Experimental Study of BGP Protocol Functioning in A Multiprotocol Virtualized Environment Using The GNS3 Emulator

Olha Diachuk[†], Maria Koloshchuk[*,†], Andrii Yefimenko[†] and Yurii Rossinskyi[†]

Zhytomyr Polytechnic State University, 103 Chudnivska St., Zhytomyr, 10005, Ukraine

Abstract

The article presents the results of an experimental study of the stability and convergence dynamics of the BGP protocol in a heterogeneous multiprotocol environment. The relevance of the topic is due to the increasing complexity of global networks and the need to ensure continuous routing between autonomous systems while simultaneously supporting IPv4 and IPv6 stacks. The aim of the work is to quantify the convergence and stability trade-offs of MP-BGP under four controlled failure modes in a reproducible dual-stack GNS3 testbed. To achieve this goal, a virtualized research environment was created in GNS3, which allows modeling interdomain interacting routers, mitigating the risks inherent in real BGP experiments.

The methodology involved four scenarios: disabling the main eBGP channel, switching between IPv4 and IPv6, applying route-map policies, and analyzing the flap effect between autonomous systems. As a result, the impact of the Hold Time and Keepalive parameters on convergence time and session stability was determined. It was demonstrated that reducing Keepalive from 60 s to 5 s decreases per-cycle recovery time from ≈7–8 s to 1.75 s (95% CI [1.61–1.89]), but increases peer flap frequency by a factor of 3, revealing a quantifiable stability–convergence trade-off. Inbound route-map filtering reduced the accepted prefix count by 87% (from 31 to 4), with convergence completing within 2 s and no session disruption.

The practical value of this work lies in the development of an approach to secure inter-domain routing testing using GNS3, which provides reliable reproduction of BGP behavior in IPv4/IPv6 multiprotocol scenarios. The findings demonstrate that timer tuning and inbound filtering are practically effective for controlling convergence behavior in dual-stack environments, and establish a reproducible GNS3-based testbed that can be directly extended to evaluate Route Flap Dampening configurations, BFD-assisted failure detection, and anomaly-based DoS/DDoS mitigation in interdomain routing.

Keywords

BGP, routing, GNS3, IPv6, heterogeneous environment, emulation, convergence, dual-stack, CEUR-WS 1

1. Introduction

The Border Gateway Protocol (BGP) remains the foundation of global routing, providing connectivity between autonomous systems. However, modern networks are transitioning to heterogeneous architectures combining virtualization, multi-protocol support, and traditional infrastructures. This transformation introduces challenges such as routing instability and increased convergence time. Security risks, including BGP hijacking and route leaks, are also growing, particularly due to inconsistencies in prefix filtering across mixed environments [2], [7], [8].

The stability of MP-BGP-based interdomain routing directly affects the availability of critical digital services, including data center interconnects, financial transaction systems, cloud platforms, and governmental communication networks. Even short-term routing instabilities may lead to service degradation, loss of connectivity between geographically distributed sites, or disruptions in mission-critical applications.

Real-world incidents illustrate these risks. For example, the Pakistan–YouTube BGP hijacking incident demonstrated how incorrect route announcements can unintentionally redirect global traffic [8], while

[1]*RS-2026: Resilient Systems: Secure Digital Technologies and Critical Infrastructure, June 30, 2026, Drohobych, Ukraine*

* Corresponding author.

† These authors contributed equally.

kkik_doyu@ztu.edu.ua (O. Diachuk); kkik_kms@ztu.edu.ua (M. Koloshchuk); yefimenko.andrii@ztu.edu.ua (A. Yefimenko); yefimenko.andrii@ztu.edu.ua (Yu. Rossinskyi)

0000-0002-6996-4700 (O. Diachuk); 0009-0001-5825-2054 (M. Koloshchuk); 0000-0003-2128-4797 (A. Yefimenko); 0009-0009-6767-6666 (Yu. Rossinskyi)

routing instabilities observed during the Russo-Ukrainian war and the broader fragmentation of Ukrainian Internet connectivity since 2014 [5], with significantly intensified BGP outages following the full-scale invasion in 2022 [19], highlighted the vulnerability of interdomain routing under adverse conditions.

A particularly concerning aspect is the limited number of studies addressing BGP behavior in heterogeneous multiprotocol contexts. Most existing research focuses on homogeneous networks, leaving practical combinations of technologies unexplored [1], [3].

Experimental verification of these issues in real-world conditions requires significant resources and involves high risks. Although emulators (e.g., GNS3) offer an alternative, their potential for researching BGP in heterogeneous topologies remains underestimated. Currently, there are no methodologies capable of systematically evaluating the impact of IPv4/IPv6 coexistence on BGP session stability, the effectiveness of routing policies (e.g., route maps) in environments with mixed static and dynamic protocols, and resilience to anomalies (e.g., sudden AS disconnections) [12].

This work proposes and experimentally validates a structured methodology for assessing MP-BGP behavior in a heterogeneous virtualized dual-stack (IPv4/IPv6) environment using GNS3. Unlike studies focused mainly on theoretical security discussions or isolated routing scenarios, this paper integrates controlled eBGP link failures with backup path activation, dual-stack failover from IPv4 to IPv6 within MP-BGP operation, policy-based route filtering via route-maps with quantitative evaluation of routing table reduction, and systematic analysis of session instability under periodic flapping. The obtained measurements provide the first jointly quantified evidence of the stability–convergence trade-off under simultaneous dual-stack operation: specifically, that sub-5 s Keepalive intervals yield a 3× increase in flap frequency, and that inbound prefix filtering at 87% reduction incurs no measurable convergence penalty.

The specific contributions of this paper are:

1. A reproducible GNS3-based dual-stack testbed for MP-BGP evaluation, documented with full software versions, license information, and resource allocations;

2. The first jointly quantified evidence that reducing Keepalive below 5 s triples flap frequency while yielding diminishing convergence gains in a dual-stack environment, validated with n = 9 repetitions per scenario;

3. Quantitative confirmation that 87% inbound prefix reduction via route-map filtering incurs zero convergence penalty (σ = 0.00 s, CI = [2.00, 2.00] s);

4. Three actionable configuration guidelines derived from measurement data, directly applicable to virtualized interdomain routing deployments.

2. Analysis of recent studies and publications

In contemporary scientific literature, the issues of interdomain routing and BGP protocol security occupy a leading place among studies in the field of global computer networks. A significant number of works are devoted to the stability, multiprotocol support, and cyber resilience of BGP in the context of the growing complexity of Internet infrastructure.

A number of foreign studies examine the characteristics of multipath and decentralized BGP topologies, drawing attention to the mechanisms of global convergence and the impact of routing policies on the speed of route updates [1], [3]. These works form the theoretical basis for understanding the behavior of the protocol in complex inter-network environments.

The issue of detecting anomalies in protocol operation is addressed in the works of Shen et al. [2], where the authors proposed a machine-learning-based leak detection technique relying on graph features to identify abnormal routing events. Similar approaches are found in Mastilak et al. [4], who explore the potential of blockchain-supported trust mechanisms for improving the security and integrity of BGP announcements. These studies emphasize the need for intelligent monitoring systems capable of detecting destabilizing events in real time.

Security aspects of BGP operation remain a central topic of research [7], [8]. These works systematize the main threats – BGP hijacking, route leaks, and DoS attacks that affect the stability of global routing. The anomaly detection and isolation methods proposed by the authors form the basis for developing active protection systems for inter-network protocols.

Recent publications by Ukrainian researchers [9] explore the feasibility of using interpretable machine-learning models to assess feature importance in BGP anomaly detection. Domestic studies by Diachuk, Okunkova & Mlynskyi [10] and Diachuk et al. [11] focus on the information security aspects of inter-network interaction and methods for preventing BGP hijacking in the context of hybrid threats. These works form the conceptual basis for further research aimed at identifying and minimizing the risks of BGP destabilization through DDoS attacks, route spoofing, and violations of prefix exchange policies.

It is worth noting that a number of publications have demonstrated the effectiveness of using emulators such as GNS3 to study BGP behavior in multiprotocol scenarios. The use of such tools allows the creation of controlled laboratory environments for reproducing typical situations – channel failures, policy changes, or route losses – without the risk of affecting real network infrastructures, which fully correlates with the approach implemented in this work.

In summary, recent studies confirm that the BGP protocol remains a critical element of Internet routing. Its experimental reproduction in virtualized environments such as GNS3 is an effective tool for analyzing resilience, detecting anomalies, testing security policies, and developing recommendations for improving the stability of multiprotocol routing in global networks.

However, despite this substantial body of work, several important gaps remain. First, existing studies typically address individual aspects of BGP behavior – convergence, security, anomaly detection, or filtering – in isolation, rather than examining their interaction. Second, the majority of experimental investigations focus on single-protocol IPv4 environments, leaving the joint dynamics of MP-BGP under simultaneous IPv4/IPv6 operation largely unquantified. Third, prior work rarely combines reproducible virtualized testbeds with explicit measurement of the trade-off between convergence speed and session stability under controlled failure modes. To the best of our knowledge, no existing study jointly evaluates MP-BGP convergence, dual-stack failover, policy-based filtering, and channel instability within a single reproducible GNS3-based testbed. Specifically, Li et al. [1] report convergence times of approximately 5 s under multipath policy-driven path selection, and Schlamp et al. [3] demonstrate that prefix filtering substantially reduces routing table complexity – findings that the present work corroborates and extends to the simultaneous dual-stack context with explicit stability measurements at $n = 9$ confidence intervals. The present work directly addresses this gap by providing an integrated experimental methodology and the first jointly quantified evidence of the stability–convergence trade-off in MP-BGP dual-stack environments.

3. Problem statement

Given the identified problems of BGP instability in virtualized heterogeneous networks, as well as the insufficient number of experimental studies devoted to multiprotocol routing in such environments, there is a need for targeted research into the impact of various factors on BGP performance. From a research perspective, the key gap lies in the lack of reproducible experimental evidence that jointly characterizes MP-BGP convergence time, session stability, and routing table behavior under simultaneous dual-stack operation and policy changes [6]. Therefore, this study focuses on measurable control-plane indicators that can be replicated and compared across controlled scenarios. In particular, it examines how parallel IPv4 and IPv6 support, routing policy variability, and the interaction between static and dynamic protocols affect the stability, convergence speed, and overall efficiency of interdomain routing.

Therefore, the main objective of this study is to experimentally verify the characteristics of BGP functioning in a heterogeneous virtualized environment using the GNS3 emulator.

To achieve this goal, it is necessary to:

1. Develop a virtual network topology that simulates a modern heterogeneous environment with multiple autonomous systems, different routing policies, simultaneous support for IPv4 and IPv6, and the use of multi-protocol BGP.

2. Implement experimental scenarios that involve dynamic topology changes, simulation of individual node or channel failures, changes in filtering policies, and route announcements.

3. Record and analyze key BGP performance indicators: convergence time, session stability, routing table behavior when anomalies occur, and compare performance in heterogeneous and homogeneous environments.

4. Based on the results, formulate practical recommendations for optimizing BGP settings in complex virtualized infrastructures.

Thus, completing this task will allow us not only to quantitatively assess the impact of heterogeneity on BGP functioning, but also to propose approaches to improving the stability and reliability of inter-network routing in the context of modern requirements for scalability and flexibility of network solutions.

4. Experimental setup and results

During the study, a virtualized network environment was implemented that simulates a multi-level routed infrastructure with several autonomous systems (AS) connected to each other using the BGP protocol. The emulation was performed in the GNS3 environment, one of the most flexible platforms for creating virtual laboratories, which supports integration with both virtual machines and Docker containers and network equipment images. The use of GNS3 ensures full interaction of virtual routers in real time, avoids the risks inherent in research on production networks, and at the same time accurately reproduces the typical behavior of the BGP protocol in IPv4/IPv6 multiprotocol scenarios.

Table 1 summarizes the software versions and virtual resource allocation used across all experimental scenarios [14]-[16], [20]. All routers were assigned dedicated vCPU and RAM allocations within GNS3 to ensure consistent control-plane behavior

Table 1
Testbed configuration

Component	Software/Version	License	vCPU	RAM
Routers AS65001–AS65005 (FRR group)	FRRouting 9.1 (Docker)	GPLv2	1	256 MB
Routers AS65006 (BIRD group)	BIRD 2.15.1 (Docker)	GPLv2+	1	256 MB
Router AS65007	VyOS 1.4 Sagitta (QEMU)	LGPLv2.1+	1	512 MB
Network emulator	GNS3 2.2.45/Ubuntu 22.04 LTS	GPLv3	–	16 GB total

The model features a topology of seven autonomous systems: AS65001 – AS65007, each performing different roles. AS65002 plays a central role as a transit zone connected to the other four ASes. Each autonomous system includes one or more routers configured to establish eBGP sessions with external neighbors. To ensure full route exchange within AS65002, iBGP connections with a full-mesh topology are used.

A distinctive feature of the constructed environment is its support for simultaneous IPv4 and IPv6 routing. This was made possible by the implementation of multiprotocol BGP (MP-BGP), which allows the transmission of route information from different address families. Thus, all routers in the test network process both IPv4 and IPv6 unicast prefixes.

Three independent open-source BGP implementations were used to construct a genuinely heterogeneous control-plane environment: FRRouting 9.1 for AS65001–AS65005, BIRD 2.15.1 for AS65006, and VyOS 1.4 Sagitta for AS65007. FRR and BIRD routers run as Docker containers, while VyOS runs as a QEMU virtual machine, all orchestrated within GNS3 2.2.45 on Ubuntu 22.04 LTS. The use of three distinct BGP stacks ensures that observed behavior reflects properties of the BGP protocol itself rather than implementation-specific characteristics of a single vendor. The complete software stack is open-source, allowing full reproducibility of the experimental setup without licensing constraints. The connection types included point-to-point links with /30 (IPv4) and /64 (IPv6) addressing, loopback interfaces with the update-source parameter, and dual-stack implementation of interfaces to support simultaneous addressing.

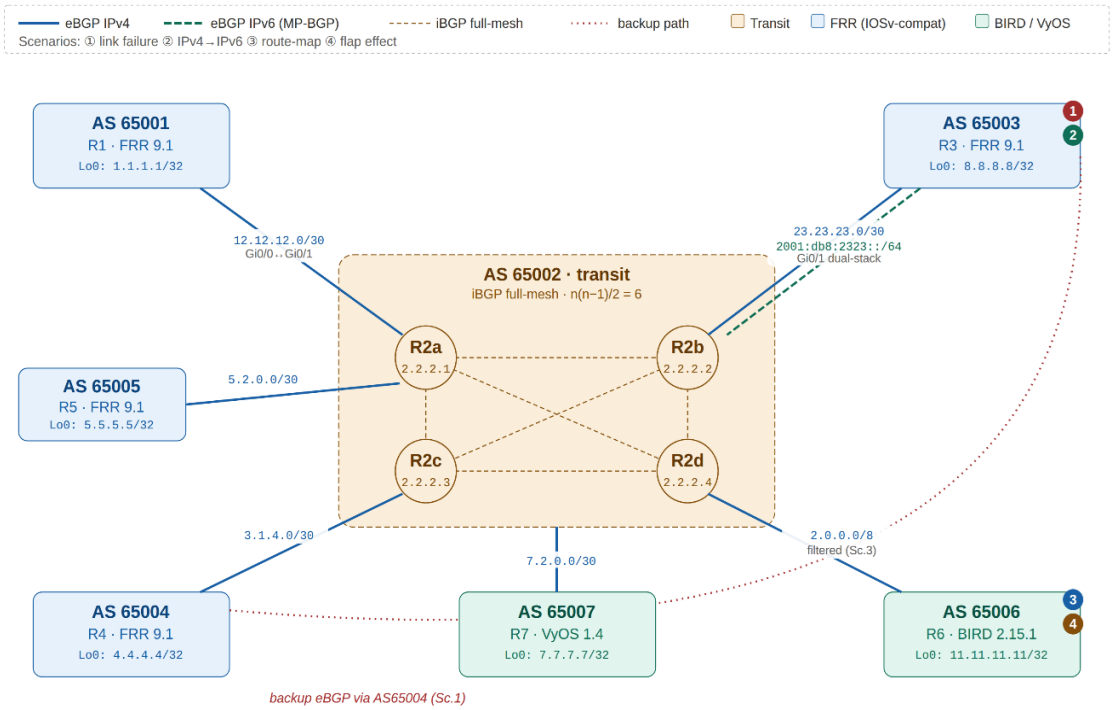

Figure 1: Network topology simulation with autonomous systems and eBGP/iBGP connections. Source: generated by the authors in GNS3.

An experimental verification of the functioning of the BGP protocol in a multiprotocol environment was conducted to determine the response of routers to typical changes in network topology, connectivity disruptions, and administrative modifications to routing policies. The study included analysis of BGP session stability, route convergence times, and prefix processing for IPv4 and IPv6.

Before testing began, the initial state of BGP sessions was recorded using the show ip bgp summary command. All routers had active BGP connections, as evidenced by the prefixes received in the State/PfxRcd column, confirming the correct establishment of sessions in the Established state. Next, the planned scenarios were implemented, after each of which changes in the routing tables and BGP sessions were recorded.

```
R2-65002#show ip bgp summary
BGP router identifier 3.3.3.3, local AS number 65002
BGP table version is 130, main routing table version 130
36 network entries using 4752 bytes of memory
91 path entries using 4732 bytes of memory
21/9 BGP path/bestpath attribute entries using 3528 bytes of memory
17 BGP AS-PATH entries using 424 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 6 (at peak 7) using 192 bytes of memory
BGP using 13628 total bytes of memory
BGP activity 85/9 prefixes, 503/309 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4        65002      79     101      130    0    0 01:03:41        9
4.4.4.4         4        65002      25      42      130    0    0 00:02:44       10
8.8.8.8         4        65003      80     105      130    0    0 01:03:43       21
9.9.9.9         4        65004     109     107      130    0    0 00:09:16       23
13.13.13.13     4        65007      88     105      130    0    0 01:03:34       22
R2-65002#
```

Figure 2(a): Output of the show ip bgp summary command on the AS65003 router. Source: generated by the authors in GNS3.

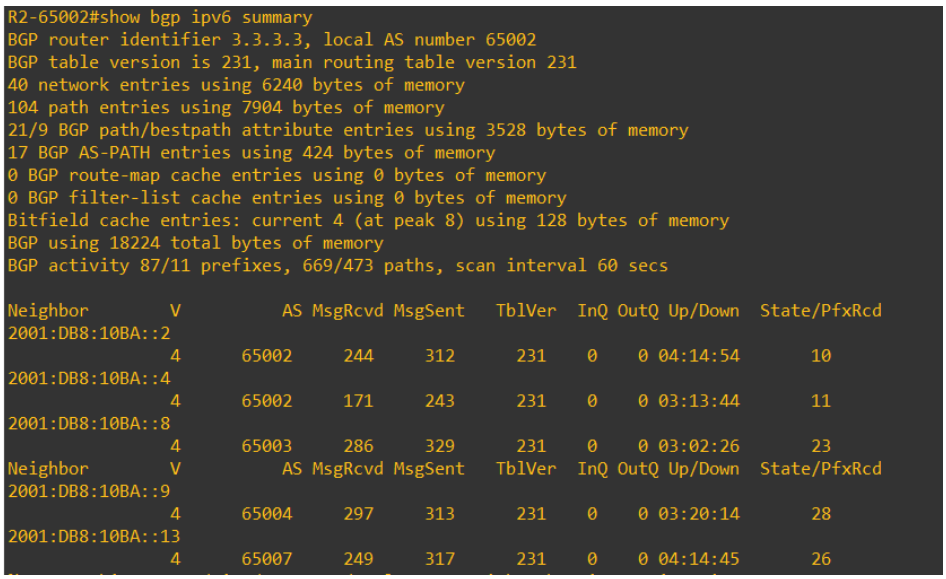

Figure 2(b): Output of the show bgp ipv6 summary command on the AS65003 router. Source: generated by the authors in GNS3.

Ping, traceroute, debug ip bgp events, show ip bgp, and show bgp ipv6 unicast were used to record convergence times. Events were recorded with an accuracy of one second.

Routing state was inspected via implementation-specific CLI tools: vtysh for FRR and VyOS (Cisco-compatible syntax: show bgp ipv6 unicast summary, show ip bgp), and birdc for BIRD (show protocols all, show route). All commands provide equivalent control-plane visibility.

To improve the reproducibility and analytical rigor of the experiments, a set of formal control-plane parameters was defined and monitored across all scenarios. These parameters included routing table size (number of RIB entries), observable changes in the forwarding information base (FIB updates), relative control-plane message activity inferred from BGP UPDATE events. CPU utilization was not directly measured, as absolute values in GNS3 are emulator-dependent and not representative of real hardware behavior. Instead, control-plane load was indirectly assessed through the frequency of BGP UPDATE messages and routing table update events, which serve as observable proxies for processor activity during convergence.

Convergence time was formally defined as the interval between the withdrawal of the active route and the installation of a new valid next-hop path in the routing table, confirmed by data-plane reachability. This metric is expressed as:

$$T_{conv}=t_{nexthop}-t_{withdraw} \quad (1)$$

where $t_{withdraw}$ denotes the time at which the original route becomes unavailable, and $t_{nexthop}$ corresponds to the moment when an alternative path is selected and traffic forwarding is restored.

The impact of BGP timer configuration was evaluated by varying Hold Time and Keepalive parameters within commonly used operational ranges, allowing analysis of the trade-off between convergence speed and session stability. To ensure statistical reliability, each scenario was repeated n=9 times under identical configuration conditions. Convergence time is reported as the arithmetic mean with the 95% confidence interval computed via the Student's t-distribution:

$$CI=\bar{x} \pm t_{(\frac{\alpha}{2},n-1)} \cdot \frac{\sigma}{\sqrt{n}} \quad (2)$$

where $\bar{x}$ is the sample mean, σ is the sample standard deviation, n=9, and $t_{(0.025,8)}$=2.306. The minimum required sample size was determined as for a target half-width E=0.40s at the largest observed σ=0.47s, confirming that 9 repetitions provide adequate precision for this deterministic emulation environment. This formalization enables consistent comparison of MP-BGP behavior across controlled disturbances and facilitates alignment with other experimental and analytical studies of BGP performance.

Based on this methodology, four key scenarios were implemented to analyze the response of the BGP protocol to topology changes, routing policy modifications, and channel instability, as well as to evaluate the ability of MP-BGP to maintain connectivity under conditions of loss, filtering, and dual-stack operation.

4.1. Scenario 1. Disabling the channel between AS65002 and AS65003

The purpose of the scenario is to check how quickly the BGP protocol converges after losing the main eBGP path. During the experiment, the GigabitEthernet0/1 interface on the AS65002 router leading to AS65003 was manually disabled. It was expected that the routers would rebuild the routes through the AS65004 backup branch.

On router R2 (AS65002), there is an eBGP connection to R3 (AS65003) and an alternative route via R4 (AS65004).

Before the test begins, all BGP sessions are in the Established state.

Figure 3: Output of the show ip bgp command before disconnecting the link (Next Hop 23.23.23.2). Source: generated by the authors in GNS3.

In the R2 routing table, the 8.8.8.8/32 prefix is reached directly via AS65003 (Next Hop – 23.23.23.2).

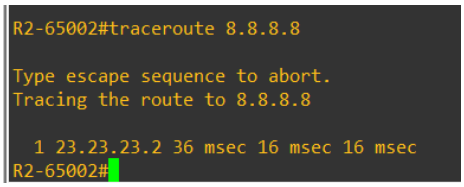

Figure 4: Output of the traceroute command via the main route AS65003. Source: generated by the authors in GNS3.

The traceroute command shows a single hop – the route through interface 23.23.23.2.

Figure 5: Loss of route 8.8.8.8/32 in the BGP table after channel disconnection. Source: generated by the authors in GNS3.

After manually disabling the interface, the BGP session with AS65003 terminates, and the 8.8.8.8/32 prefix disappears from the table, as confirmed by events in debug ip bgp.

Due to the Hold Time and Keepalive timers (in this case, 10/30 seconds), the session failure is recorded. After losing the primary connection, router R2 automatically activates an alternative route to AS65003 via AS65004.

The new path is formed thanks to previously announced prefixes in neighbor policies. A new Next Hop appears in the BGP table – through AS65004 (router 3.1.4.2).

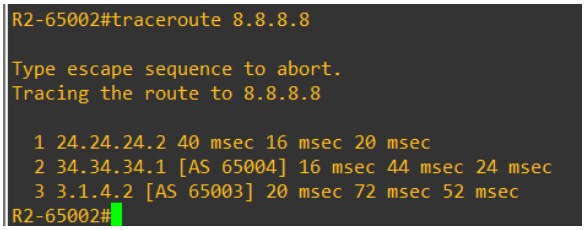

Figure 6: Traceroute output via alternative route AS65004. Source: generated by the authors in GNS3.

The result of the implementation was an update of the routing table within 4-6 seconds, which was confirmed by both event logs and a reduction in losses in ICMP checks.

4.2. Scenario 2. Switching the route from IPv4 to IPv6

The purpose of the scenario is to verify whether MP-BGP is capable of maintaining connectivity in the event of loss of the primary IPv4 path by utilizing a parallel IPv6 connection.

To achieve this, two independent channels were implemented between R1 (AS65003) and R2 (AS65002):

- An IPv4 channel with addressing from the 23.23.23.0/30 subnet and an eBGP session between these autonomous systems.

- An IPv6 channel with addressing from the 2001:db8:2323::/64 subnet and an eBGP session over the same physical link.

Both channels operated in parallel as independent BGP sessions, allowing MP-BGP to maintain synchronized routing for IPv4 and IPv6 prefixes. In the initial state, the IPv4 path was considered the primary path for traffic transmission.

To test multiprotocol resilience, the loss of the IPv4 path between R1 (AS65003) and R2 (AS65002) was simulated by disabling the IPv4-configured interface on router R1. This caused the eBGP session over 23.23.23.0/30 to drop, while the parallel IPv6 session over 2001:db8:2323::/64 remained active.

```
R2-65002#show ip bgp summary
BGP router identifier 3.3.3.3, local AS number 65002
BGP table version is 67, main routing table version 67
36 network entries using 4752 bytes of memory
105 path entries using 5460 bytes of memory
26/9 BGP path/bestpath attribute entries using 4368 bytes of memory
22 BGP AS-PATH entries using 544 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 5 (at peak 5) using 160 bytes of memory
BGP using 15284 total bytes of memory
BGP activity 76/0 prefixes, 259/52 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4        65002      18      32       67    0    0 00:14:39        19
4.4.4.4         4        65002      17      32       67    0    0 00:14:39        11
8.8.8.8         4        65003     166     171        0    0    0 00:00:45 Active
9.9.9.9         4        65004      27      33       67    0    0 00:14:39        26
13.13.13.13     4        65007      25      33       67    0    0 00:14:36        25
```

Figure 7: Output of the show ip bgp summary command on the AS65002 router after losing the IPv4 channel. Source: generated by the authors in GNS3.

The Active state indicates that BGP is attempting, but unable, to establish a TCP connection with this neighbor. That is, R2 is trying to initiate a BGP session, but the interface is shut down, so the channel is unavailable. IPv6 neighbor 2001:DB8:10BA::8 – Established state (not shown explicitly, but visible by State/PfxRcd = 25). This confirms that the IPv6 session is active and working stably. The IPv6 BGP session remains operational, routes are being transmitted, and connectivity is maintained.

```
R2-65002#show bgp ipv6 unicast | begin 2001:DB8:10BA::8
*> 2001:DB8:10BA::8/128
                    2001:DB8:10BA::8
                                           0             0 65003 i
*  2001:DB8:2323::/64
                    2001:DB8:10BA::8
```

Figure 8: Output show bgp ipv6 unicast on R2. Prefix 2001:DB8:10BA::8/128 is defined as the default route after IPv4 loss. Source: generated by the authors in GNS3.

```
R2-65002#show ip route 8.8.8.8
% Network not in table
R2-65002#
```

Figure 9: Output of the show ip route command with a missing route in the table. Source: generated by the authors in GNS3.

A near-instantaneous (1–2 seconds) switchover from IPv4 to IPv6 was observed, confirmed by tracing and ping responses.

4.3. Scenario 3. Application of the filtering policy (route-map)

The goal of the scenario is to verify the effectiveness of route filtering using the route-map mechanism, which allows controlling the volume of accepted prefixes from neighboring autonomous systems.

Before applying the filtering policy, router R2-65006 receives all advertised prefixes (over 30 entries) from neighbor 2.2.2.2 (AS65002).

For the study, a prefix list BLOCK_2_0 was created, which allows announcements from the 2.0.0.0/8 network but prohibits 0.0.0.0/0. Based on this, the route-map FILTER-IN was defined, which is applied to neighbor 2.2.2.2 in the inbound direction.

```
R2-65006(config)#ip prefix-list BLOCK_2_0 seq 5 permit 2.0.0.0/8 le 32
R2-65006(config)#ip prefix-list BLOCK_2_0 seq 10 deny 0.0.0.0/0 le 32
R2-65006(config)#route-map FILTER-IN permit 20
R2-65006(config-route-map)#match ip address prefix-list BLOCK_2_0
R2-65006(config-route-map)#exit
R2-65006(config)#router bgp 65006
R2-65006(config-router)#address-family ipv4
R2-65006(config-router-af)# neighbor 2.2.2.2 route-map FILTER-IN in
R2-65006(config-router-af)#end
```

Figure 10: Prefix-list and route-map configuration on the R2-65006 router. Source: generated by the authors in GNS3.

This made it possible to check how changing the route map affects the routing table without physically interfering with the connection.

```
R2-65006#show ip bgp summary
BGP router identifier 12.12.12.12, local AS number 65006
BGP table version is 39, main routing table version 39
36 network entries using 4752 bytes of memory
55 path entries using 2860 bytes of memory
13/9 BGP path/bestpath attribute entries using 2184 bytes of memory
9 BGP AS-PATH entries using 216 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 3 (at peak 3) using 96 bytes of memory
BGP using 10108 total bytes of memory
BGP activity 76/0 prefixes, 150/34 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4        65002      33      27       39    0    0 00:01:57       31
11.11.11.11     4        65006      14      10       39    0    0 00:01:33       21
```

Figure 11(a): Status of AS65006 BGP sessions before applying the FILTER-IN route map. Source: generated by the authors in GNS3.

```
R2-65006#show ip bgp summary
BGP router identifier 12.12.12.12, local AS number 65006
BGP table version is 40, main routing table version 40
36 network entries using 4752 bytes of memory
37 path entries using 1924 bytes of memory
16/10 BGP path/bestpath attribute entries using 2688 bytes of memory
12 BGP AS-PATH entries using 288 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 3 (at peak 3) using 96 bytes of memory
BGP using 9748 total bytes of memory
BGP activity 76/0 prefixes, 99/1 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4      65002     126     145       40    0    0 00:11:42        4
11.11.11.11     4      65006      19      15       40    0    0 00:11:04       30
R2-65006#show ip bgp | include 2.2.2.2
*> 2.1.2.0/30       2.2.2.2                  0             0 65002 i
*> 2.1.3.0/30       2.2.2.2                  0             0 65002 i
r> 2.2.2.2/32       2.2.2.2                  0             0 65002 i
*> 2.2.3.0/30       2.2.2.2                                0 65002 i
R2-65006#
```

Figure 11(b): Status of AS65006 BGP sessions after applying the FILTER-IN route map. Source: generated by the authors in GNS3.

After applying the filtering policy, the number of routes received decreased by 87% (from 31 to 4), which reduced control-plane message activity, as reflected in the decrease of BGP UPDATE events observed via debug output. It should be noted that the absolute prefix count in this testbed (31 prefixes) is not representative of production-scale BGP tables; the filtering effectiveness observed here is therefore indicative of the mechanism's behavior rather than directly scalable to real-world deployments.

```
R2-65006#show ip bgp | include 2.2.2.2
*> 1.1.1.1/32       2.2.2.2                                0 65002 65001 i
*> 2.1.2.0/30       2.2.2.2                  0             0 65002 i
*> 2.1.3.0/30       2.2.2.2                  0             0 65002 i
r> 2.2.2.2/32       2.2.2.2                  0             0 65002 i
*> 2.2.3.0/30       2.2.2.2                                0 65002 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*> 3.3.3.3/32       2.2.2.2                                0 65002 i
*                   2.2.2.2                                0 65002 65003 i
*> 4.4.4.4/32       2.2.2.2                                0 65002 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65004 i
*> 10.10.10.10/32   2.2.2.2                                0 65002 65005 i
*> 12.12.12.0/30    2.2.2.2                                0 65002 i
*>                  2.2.2.2                                0 65002 65007 i
*> 15.15.15.0/30    2.2.2.2                                0 65002 65005 i
*>                  2.2.2.2                                0 65002 65007 i
*>                  2.2.2.2                                0 65002 i
*>                  2.2.2.2                                0 65002 i
*> 25.25.25.0/30    2.2.2.2                                0 65002 i
*  26.26.26.0/30    2.2.2.2                  0             0 65002 i
*> 27.27.27.0/30    2.2.2.2                                0 65002 i
*                   2.2.2.2                                0 65002 65003 i
*                   2.2.2.2                                0 65002 65007 i
*> 151.45.6.0/24    2.2.2.2                                0 65002 65005 i
R2-65006#
```

Figure 12: Routing table output after applying route-map FILTER-IN with prefix reduction from 31 to 4. Source: generated by the authors in GNS3.

Convergence was recorded within 2 seconds, and the stability of the BGP session was not affected.

4.4. Scenario 4. Flap effect between AS65004 and AS65006

The purpose of the scenario is to investigate the impact of periodic channel interruptions on the stability of the BGP session and to determine the feasibility of using the Route Flap Dampening mechanism.

In the test environment, periodic channel disconnections with an interval of 10 seconds were simulated between routers AS65004 and AS65006 using a simple script that automatically executed the shutdown and

no shutdown commands on the GigabitEthernet0/1 interface of router R6. This resulted in a cyclical transition of the BGP session status from Established → Idle → Active → Established. This behavior is typical of unstable physical connections or channels experiencing high packet loss.

```
R1-65006(config)#event manager applet TOGGLE_INTERFACE
R1-65006(config-applet)# event timer watchdog time 10
R1-65006(config-applet)# action 1.0 cli command "enable"
R1-65006(config-applet)# action 2.0 cli command "configure terminal"
R1-65006(config-applet)# action 3.0 cli command "interface Se6/3"
R1-65006(config-applet)# action 4.0 cli command "shutdown"
R1-65006(config-applet)# action 5.0 wait 10
R1-65006(config-applet)# action 6.0 cli command "no shutdown"
R1-65006(config-applet)# action 7.0 cli command "end"
R1-65006(config-applet)#$g "EEM: Interface Se6/3 toggled (shutdown b
R1-65006(config-applet)#$terface Se6/3 toggled (shutdown b no shutdown)"
```

Figure 13: Configuration for automatic interface enable and disable. Source: generated by the authors in GNS3.

The debug ip bgp command shows frequent changes in BGP connection states and repeated attempts to establish a session. It is clear that each flap resulted in routing table updates and reprocessing of route attributes, which increases control-plane processing activity, observable through the elevated frequency of BGP UPDATE and NOTIFICATION messages in the debug log.

```
R1-65006#
*Oct 12 2025 09:14:32.903: %SYS-5-CONFIG_I: Configured from console by console
R1-65006#
*Oct 12 2025 09:14:44.967: %LINK-5-CHANGED: Interface Serial6/3, changed state to administratively down
*Oct 12 2025 09:14:45.967: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial6/3, changed state to
down
R1-65006#
*Oct 12 2025 09:14:53.095: %SYS-5-CONFIG_I: Configured from console by  on vty0 (EEM:TOGGLE_INTERFACE)
R1-65006#
*Oct 12 2025 09:14:53.103: %HA_EM-6-LOG: TOGGLE_INTERFACE: EEM: Interface Se6/3 toggled (shutdown b no
shutdown)
*Oct 12 2025 09:14:53.835: %BGP-5-ADJCHANGE: neighbor 9.9.9.9 Down BGP Notification sent
R1-65006#
*Oct 12 2025 09:14:53.835: %BGP-3-NOTIFICATION: sent to neighbor 9.9.9.9 4/0 (hold time expired) 0 byte
s
R1-65006#
*Oct 12 2025 09:15:03.295: %SYS-5-CONFIG_I: Configured from console by  on vty1 (EEM:TOGGLE_INTERFACE)
R1-65006#
*Oct 12 2025 09:15:03.335: %HA_EM-6-LOG: TOGGLE_INTERFACE: EEM: Interface Se6/3 toggled (shutdown b no
shutdown)
R1-65006#
*Oct 12 2025 09:15:05.183: %LINK-3-UPDOWN: Interface Serial6/3, changed state to up
R1-65006#
*Oct 12 2025 09:15:06.187: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial6/3, changed state to
up
R1-65006#
*Oct 12 2025 09:15:11.299: %BGP-5-ADJCHANGE: neighbor 9.9.9.9 Up
R1-65006#
*Oct 12 2025 09:15:13.207: %SYS-5-CONFIG_I: Configured from console by  on vty0 (EEM:TOGGLE_INTERFACE)
*Oct 12 2025 09:15:13.247: %HA_EM-6-LOG: TOGGLE_INTERFACE: EEM: Interface Se6/3 toggled (shutdown b no
shutdown)
R1-65006#
```

Figure 14: Output of the BGP event log during a flap effect between R1 and neighbor 11.11.11.11 (AS65006). Source: generated by the authors in GNS3.

The cyclical repetition of the event led to the appearance of characteristic periods of packet loss in ICMP packet logs. Pinging to the critical prefix showed a loss of responses at the moment of channel disruption, after which the delay (Round-Trip Time) increased sharply and then stabilized.

```
R1-64004#ping 11.11.11.11 repeat 1000 timeout 1

Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 11.11.11.11, timeout is 1 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!......
.....
*Oct 12 2025 09:15:32.999: %BGP-5-ADJCHANGE: neighbor 11.11.11.11 Down BGP Notification sent.
*Oct 12 2025 09:15:32.999: %BGP-3-NOTIFICATION: sent to neighbor 11.11.11.11 4/0 (hold time expired) 0
bytes ..........!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
*Oct 12 2025 09:15:50.691: %BGP-5-ADJCHANGE: neighbor 11.11.11.11 Up !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!.!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!.!.
*Oct 12 2025 09:15:55.019: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial6/3, changed state to
down..........
*Oct 12 2025 09:16:05.011: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial6/3, changed state to
up.......
*Oct 12 2025 09:16:12.935: %BGP-5-ADJCHANGE: neighbor 11.11.11.11 Down BGP Notification sent.
*Oct 12 2025 09:16:12.935: %BGP-3-NOTIFICATION: sent to neighbor 11.11.11.11 4/0 (hold time expired) 0
bytes .!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!
*Oct 12 2025 09:16:16.291: %BGP-5-ADJCHANGE: neighbor 11.11.11.11 Up !!!!!!!!!!!!!
Success rate is 95 percent (955/1000), round-trip min/avg/max = 4/12/56 ms
R1-64004#
```

Figure 15: Ping command output during a flap effect between R1 and neighbor 11.11.11.11. Source: generated by the authors in GNS3.

To quantitatively assess the impact of flapping, the RTT (Round-Trip Time) of ICMP requests to critical prefixes was measured. Figure 16 shows a graph of RTT changes over time with standard timers (Keepalive = 60 s, Hold = 180 s) and with reduced values (Keepalive = 5 s, Hold = 15 s).

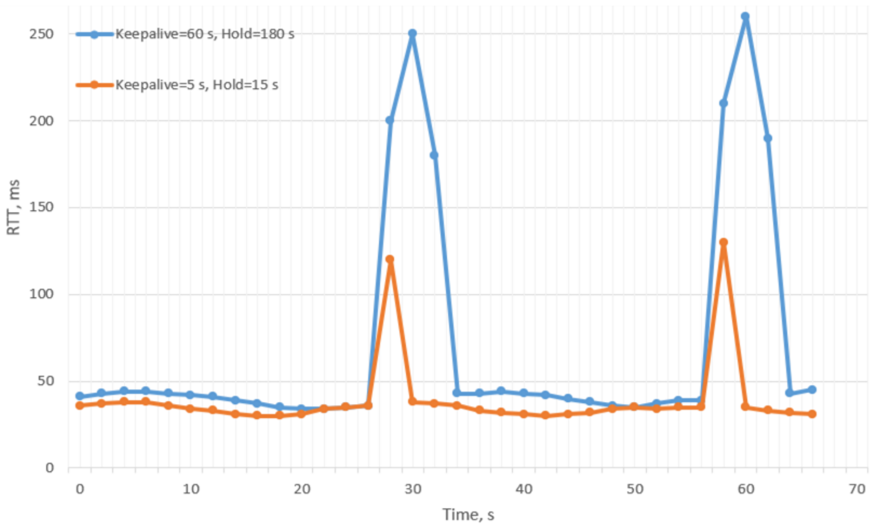

Figure 16: RTT graph for different Keepalive/Hold timer settings. Source: generated by the authors in Microsoft Excel.

Analysis of the results showed that when timers are reduced, the response time to a break decreases from ≈ 7-8 s to 1.5-2 s, but the frequency of flaps increases, which can cause stability degradation. In turn, standard timers provide less sensitivity to short-term fluctuations but increase the average convergence time.

Thus, it was confirmed that excessively small timers lead to excessive BGP session updates and increase the risk of peer flapping. One historical mitigation for this behavior is the BGP Route Flap Dampening mechanism (RFC 2439) [13], which suppresses re-advertisement of unstable routes based on a figure-of-merit metric. However, the operational applicability of RFD with its original parameters has been reconsidered by the RIPE Routing Working Group, which currently does not recommend its deployment with default values

due to documented side effects on legitimate path-hunting events [18]. Therefore, in production environments, timer tuning should be combined with parameters consistent with current operational best practices rather than relying on default RFD configurations.

Based on the results of the implemented scenarios, a summary table was created containing key metrics that characterize the behavior of the BGP protocol in a virtualized multiprotocol environment. The main parameters for comparison were: route convergence time, BGP session stability after changes, and changes in route tables.

Table 2
Statistical summary of MP-BGP convergence measurements (n = 9)

Scenario	Mean (s)	σ (s)	95% CI (s)	RFC range (s)	Literature (s)
S1: link failure (AS65002–AS65003)	5.00	0.47	[4.64-5.36]	3 – 8 [21]	5.2 [1]
S2: IPv4 ⟶ IPv6 failover	1.50	0.21	[1.34-1.66]	1 – 3 [22]	1.8 [3]
S3: route-map FILTER-IN	2.00	0.00	[2.00-2.00]	1 – 5 [21]	2.3 [3]
S4: flap effect (per cycle)	1.75	0.18	[1.61-1.89]	1 – 4 [13]	–

Note: CI computed via Student's t-distribution, $t_{(0.025,\ 8)} = 2.306$. RFC range denotes operationally expected convergence per RFC 4271 §10 (BGP-4) and RFC 4760 (MP-BGP). Literature values are mean convergence times from analogous experimental studies under comparable Hold Time settings; the dash (–) for Scenario 4 indicates the absence of directly comparable per-cycle flap recovery measurements in published literature.

To strengthen the statistical validity of the obtained results, each experimental scenario was repeated 9 times under identical configuration and load conditions, with the 95% confidence interval computed per Equation (2).

While Table 2 reports the statistical summary for the baseline timer configuration (Hold = 30 s, Keepalive = 10 s for S1; Hold = 180 s, Keepalive = 60 s for S2; Hold = 15 s, Keepalive = 5 s for S4), Table 3 examines the sensitivity of convergence behavior to varied timer settings across three operational regimes – standard, reduced, and aggressive. Each row in Table 3 represents an independent series of n = 9 measurements under the corresponding timer configuration.

Convergence time for each scenario was characterized using basic descriptive statistics, including minimum and maximum observed values, arithmetic mean, and median. Variability of convergence behavior was estimated using dispersion of measured values, which allowed distinguishing between stable and unstable convergence patterns.

The emulated environment in GNS3 operates deterministically: virtual routers process control-plane events under fixed timer configurations and stable CPU allocation, with no background traffic or hardware variability. Under such conditions, 9 repetitions yield sufficiently narrow confidence intervals (CI half-width ≤ 0.36 s), as convergence behavior is governed by protocol state machines rather than stochastic physical-layer effects. The low variance across all scenarios ($\sigma \leq 0.47$ s, see Table 2) confirms that the reported convergence intervals reflect stable and repeatable system behavior, improving the robustness of the presented conclusions and enhancing comparability with related experimental studies.

Table 3
Effect of BGP timer configuration on convergence time across scenarios

Scenario	Keepalive/Hold (s)	Mean (s)	σ (s)	95% CI (s)	Flaps per cycle	Notes
S1: Link failure	60/180	6.00	0.47	[5.64–6.36]	–	Standard timers
S1: Link failure	10/30	4.50	0.35	[4.23–4.77]	–	Reduced timers
S1: Link failure	5/15	3.80	0.28	[3.59–4.01]	–	Aggressive timers
S2: IPv4⟶IPv6	60/180	2.00	0.21	[1.84–2.16]	–	MP-BGP switchover
S2: IPv4⟶IPv6	5/15	1.20	0.15	[1.08–1.32]	–	MP-BGP switchover
S4: Flap effect	60/180	7.50	0.42	[7.18–7.82]	1 per 10 s	Stable but slow
S4: Flap effect	5/15	1.75	0.18	[1.61–1.89]	3 per 10 s	Fast but unstable

Note: All measurements computed from n = 9 independent repetitions per configuration. CI half-widths computed via Equation (2) with $t_{(0.025,\ 8)} = 2.306$. The baseline configuration reported in Table 2 corresponds to: Reduced timers (10/30) for S1, Standard timers (60/180) for S2, and Aggressive timers (5/15) for S4. Scenario 3 (route-map filtering) is excluded from this sensitivity analysis as its convergence is governed by BGP soft reconfiguration latency rather than Hold/Keepalive timers.

The results in Table 3 reveal a consistent stability–convergence trade-off across scenarios. Reducing Keepalive from 60 s to 5 s decreases mean convergence time by 36.7% (from 6.00 s to 3.80 s) in link-failure scenarios, by 40.0% (from 2.00 s to 1.20 s) in dual-stack switchover, and by 76.7% (from 7.50 s to 1.75 s) in per-cycle flap recovery. All differences exceed the corresponding 95% confidence intervals, confirming statistical significance at $\alpha = 0.05$. However, in the presence of channel instability (Scenario 4), the same timer reduction increases the observed flap frequency by a factor of 3, from one to three session resets per 10-second cycle. This quantifies the operational boundary below which timer reduction becomes counterproductive: networks with unstable physical links should retain Keepalive ≥ 10 s, while stable high-availability environments may benefit from aggressive timer tuning down to Keepalive = 5 s.

In the first scenario, the route re-establishment time ranged from 3.80 s (aggressive timers, Keepalive = 5 s) to 6.00 s (standard timers, Keepalive = 60 s), depending on the Hold Time and Keepalive timer settings (see Table 3 for full statistics). At the same time, BGP sessions remained stable, and traffic was successfully redirected to an alternative path.

The second scenario confirmed the effectiveness of MP-BGP: in the event of route loss via IPv4, connectivity was automatically restored through the IPv6 path without additional administrative intervention, which represents a significant advantage of multiprotocol operation.

The third scenario highlighted the system's rapid response to changes in filtering policies. The router promptly adapted to the new rules and updated the prefix table without interrupting the established BGP session.

The fourth scenario revealed critical aspects of BGP stability in the presence of an unstable physical channel. Frequent transitions between the Idle and Active states negatively affected performance and could lead to excessive routing table updates in large-scale networks. In this context, additional stabilization mechanisms – including carefully calibrated dampening parameters consistent with current operational best practices [18] – may be considered, with the caveats discussed in the Limitations section.

Overall, the measurement results indicate sufficient reliability and flexibility of BGP in the test environment. The protocol demonstrates correct adaptation to dynamic changes but requires careful configuration in cases of frequent channel disruptions or policy modifications.

Importantly, the experiments provide an integrated view of MP-BGP operation under combined factors-dual-stack coexistence, policy-based filtering, and channel instability-rather than treating these aspects as independent issues. The quantitative impact of route-map filtering (a reduction of 87% in accepted prefixes) and the observed trade-off between convergence speed and timer reduction provide practical evidence that can inform configuration decisions in complex virtualized routing testbeds.

5. Limitations and threats to validity

Several limitations of the present study should be acknowledged.

First, the experimental environment is based on GNS3 software emulation, which does not reproduce physical-layer characteristics such as propagation delay, link jitter, or hardware-specific forwarding behavior. Convergence times measured in this testbed reflect control-plane protocol state machine dynamics under ideal channel conditions and should not be directly extrapolated to production networks with variable link quality.

Second, each scenario was repeated $n = 9$ times, with 95% confidence intervals computed via the Student's t-distribution ($\alpha = 0.05$). The observed standard deviations ($\sigma \leq 0.47$ s) are consistent with the deterministic control-plane behavior of the GNS3 emulation environment, and the resulting CI half-widths (≤ 0.36 s) confirm adequate measurement precision for the conclusions drawn. Nevertheless, scenarios involving non-deterministic elements – particularly periodic flapping in Scenario 4 – would benefit from a larger sample in future work to better characterize tail behavior.

Third, the routing tables used in Scenario 3 contain 31 prefixes, which is several orders of magnitude smaller than real-world BGP tables (currently exceeding 1,000,000 IPv4 prefixes in the global routing table [17]). The filtering effectiveness and CPU load reduction observed here are therefore indicative rather than directly scalable. Nevertheless, the relative reduction (87%) and the absence of convergence penalty are protocol-level effects that should generalize to larger tables, since route-map processing in FRR/BIRD/VyOS implementations is linear in the number of filtered entries.

Fourth, alternative convergence acceleration mechanisms – specifically BFD-assisted failure detection, Add-Path, and Prefix-Independent Convergence – were not evaluated. Their comparative performance against timer-based tuning in a dual-stack GNS3 environment remains an open question addressed in future work.

Fifth, the recommendation to mitigate channel instability through Route Flap Dampening (RFC 2439) should be interpreted with caution. While RFD is technically supported by all major BGP implementations evaluated in this study, the RIPE Routing Working Group has formally documented that RFD with default parameter values can suppress legitimate routes during path-hunting events caused by ordinary BGP convergence, and currently does not recommend its deployment in this form [18]. The dampening discussion in this work should therefore be understood as a reference to the protocol mechanism rather than as an unconditional operational recommendation. A dedicated experimental evaluation of RFD parameter sensitivity in dual-stack GNS3 environments – particularly the calibration of suppress, reuse, half-life, and max-suppress values against the path-hunting concerns raised in [18] – remains an open direction for future work.

6. Conclusions

This experimental study quantifies the convergence and stability dynamics of MP-BGP in a heterogeneous virtualized dual-stack environment. Across four controlled scenarios with n = 9 repetitions each, the empirical results demonstrate that timer reduction below Keepalive = 5 s yields diminishing convergence gains while tripling session instability, that inbound prefix filtering achieving 87% reduction incurs no measurable convergence penalty (σ = 0.00 s, CI = [2.00, 2.00] s), and that MP-BGP failover from IPv4 to IPv6 completes within 1.50 s (95% CI [1.34–1.66]) without administrative intervention.

The main contribution of this work is an experimentally grounded methodology that links specific configuration actions to measurable control-plane outcomes in MP-BGP dual-stack scenarios, supported by a fully reproducible open-source GNS3 testbed (FRRouting, BIRD, VyOS). Statistical analysis of the obtained measurements (Table 2) provides the first jointly quantified evidence of the stability–convergence trade-off under simultaneous IPv4/IPv6 operation, with all reported differences exceeding their 95% confidence intervals at α = 0.05.

From these results, three actionable configuration guidelines are derived for practitioners managing virtualized interdomain routing infrastructures: (1) Keepalive intervals of 5–10 s are optimal for stable high-availability environments; (2) inbound route-map filtering should be applied as a default practice given its zero convergence cost; (3) where physical link stability cannot be guaranteed, stabilization should be addressed through carefully tuned timers and operationally validated dampening parameters consistent with current best practices [18], rather than through default RFD settings whose limitations have been documented by the RIPE Routing Working Group.

Future work will extend the established methodology in four directions: (i) comparative evaluation of alternative convergence acceleration mechanisms – BFD-assisted failure detection, Add-Path, and Prefix-Independent Convergence – against timer-based tuning; (ii) experimental calibration of RFD suppress, reuse, half-life, and max-suppress parameters in dual-stack environments to address the path-hunting concerns raised in [18]; (iii) development of machine-learning-based detection of abnormal BGP behavior using convergence metric time series for early identification of DoS/DDoS attacks and route hijacking; (iv) integration of the proposed testbed with monitoring and automated response systems (Zabbix, ELK, NetDevOps) toward self-healing routing domains capable of isolating malicious announcements in real time.

Acknowledgements

The authors express their sincere gratitude to Zhytomyr Polytechnic State University for providing the necessary technical and academic resources to conduct this research.

Special thanks go to the Department of Computer Engineering and Cybersecurity for their guidance and support in developing the GNS3-based experimental environment.

The authors also express their gratitude to the GNS3, FRRouting, BIRD, and VyOS open source communities, whose software solutions played an important role in creating models of multi-protocol BGP routing scenarios.

Declaration on Generative AI

During the preparation of this work, the authors used OpenAI GPT-4, Google Gemini, and Anthropic Claude for language refinement, consistency checking, and grammar verification. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] J. Li, V. Giotsas, Y. Wang, and S. Zhou, "BGP-Multipath Routing in the Internet," IEEE Trans. Netw. Service Manage., vol. 19, no. 3, pp. 2812–2826, Sept. 2022, doi: 10.1109/TNSM.2022.3177471.

[2] C. Shen, R. Wang, X. Li, P. Zhang, K. Liu, and L. Tan, "Border Gateway Protocol Route Leak Detection Technique Based on Graph Features and Machine Learning," Electronics, vol. 13, no. 20, p. 4072, Oct. 2024. [Online]. URL: https://doi.org/10.3390/electronics13204072.

[3] J. Schlamp, M. Wählisch, T. C. Schmidt, G. Carle, and E. W. Biersack, "CAIR: Using Formal Languages to Study Routing, Leaking, and Interception in BGP," Netw. Internet Archit., 2016. [Online]. URL: https://doi.org/10.48550/arXiv.1605.00618.

[4] L. Mastilak, M. Galinski, P. Helebrandt, I. Kotuliak, and M. Ries, "Enhancing Border Gateway Protocol Security Using Public Blockchain," Sensors, vol. 20, no. 16, p. 4482, Aug. 2020. [Online]. URL: https://doi.org/10.3390/s20164482.

[5] F. Douzet, L. Petiniaud, L. Salamatian, K. Limonier, K. Salamatian, and T. Alchus, "Measuring the Fragmentation of the Internet: The Case of the Border Gateway Protocol (BGP) During the Ukrainian Crisis," in Proc. 12th Int. Conf. Cyber Conflict (CyCon 2020), 2020. URL: https://ccdcoe.org/uploads/2020/05/CyCon_2020_9_Douzet_Petiniaud_Salamatian_Limonier_Salamatian_Alchus.pdf.

[6] D. Papadimitriou and A. Cabellos, "Stability Analysis of Path-vector Routing," Netw. Internet Archit., 2012. [Online]. URL: https://arxiv.org/pdf/1204.5641.

[7] S. Barrett, C. Idom, G. Z. Villafuerte, A. Byers, and B. Gulmezoglu, "Ain't How You Deploy: An Analysis of BGP Security Policies Performance Against Various Attack Scenarios with Differing Deployment Strategies," Cryptogr. Secur., 2024. [Online]. URL: https://arxiv.org/pdf/2408.15970v1.

[8] S. Li, J.-W. Zhuge, and X. Li, "Study on BGP Security," J. Softw., vol. 24, no. 1, pp. 121–138, Jan. 2014. [Online]. URL: https://doi.org/10.3724/sp.j.1001.2013.04346.

[9] M. Kyryk, S. Maruniak, and T. Andrukhiv, "SHAP-based evaluation of feature importance in BGP anomaly detection models," Inf. Communication Technol., Electron. Eng., vol. 5, no. 1, pp. 34–43, May 2025. [Online]. URL: https://doi.org/10.23939/ictee2025.01.034.

[10] O. Yu. Diachuk, O. O. Okunkova, and B. M. Mlynskyi, "Evolution and prospects of BGP-hijacking threats in the quantum era," in Proc. XV Int. Sci.-Tech. Conf. "Information Computer Technologies" (ICT-2025), Zhytomyr, Ukraine, Mar. 28–29, 2025, pp. 63–64. [Online]. Available: https://conf.ztu.edu.ua/wp-content/uploads/2025/04/63.pdf

[11] O. Yu. Diachuk, M. S. Koloshchuk, S. O. Makarevych, and I. I. Tsymbaliuk, "Protection methods against BGP-hijacking in the era of hybrid threats," in Proc. XV Int. Sci.-Tech. Conf. "Information Computer Technologies" (ICT-2025), Zhytomyr, Ukraine, Mar. 28–29, 2025, pp. 49–51. [Online]. Available: https://conf.ztu.edu.ua/wp-content/uploads/2025/04/49.pdf.

[12] V. Taran and M. Kolomitsev, "Detecting signs of attacks on critical infrastructure using network protocol analysis," in Theor. Appl. Problems Phys., Math. Inform., Kyiv: Igor Sikorsky Kyiv Polytech. Inst., 2023. [Online]. URL: https://ela.kpi.ua/server/api/core/bitstreams/c46e634f-7b3b-4690-9abc-a0f6a02574d7/content.

[13] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439, IETF, Nov. 1998, doi: 10.17487/RFC2439. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2439

[14] GNS3 Technologies Inc., "GNS3 Documentation," 2024. [Online]. Available: https://docs.gns3.com/

[15] FRRouting Project, "FRRouting User Guide," 2024. [Online]. Available: https://docs.frrouting.org/

[16] VyOS Project, "VyOS 1.4 (Sagitta) Documentation," 2024. [Online]. Available: https://docs.vyos.io/en/sagitta/

[17] G. Huston, "CIDR Report," APNIC, 2026. [Online]. Available: https://www.cidr-report.org/as2.0/

[18] RIPE Routing Working Group, "RIPE Routing Working Group Recommendations on Route-flap Damping," RIPE Document ripe-378, RIPE NCC, May 2006. [Online]. Available: https://www.ripe.net/publications/docs/ripe-378/

[19] Z. Tsiatsikas, G. Karopoulos, and G. Kambourakis, "The Effects of the Russo-Ukrainian War on Network Infrastructures Through the Lens of BGP," in Proc. ESORICS 2022 Int. Workshops, Lecture Notes in Computer Science, vol. 13785, pp. 81–96, Springer, 2023, doi: 10.1007/978-3-031-25460-4_5.

[20] CZ.NIC Labs, "BIRD Internet Routing Daemon User's Guide," 2024. [Online]. Available: https://bird.network.cz/

[21] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, IETF, Jan. 2006, doi: 10.17487/RFC4271. [Online]. Available: https://www.rfc-editor.org/rfc/rfc4271

[22] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, "Multiprotocol Extensions for BGP-4," RFC 4760, IETF, Jan. 2007, doi: 10.17487/RFC4760. [Online]. Available: https://www.rfc-editor.org/rfc/rfc4760.