Development of an Improved Routing Mechanism in Distributed Networks Using Enhanced BGP Timer Configuration

[1]Wahab Ganiyu Abolore, [2]Oyediran Mayowa Oyedepo, [2]Ipeayeda Funmilola Wumi

[1]Department of Computer Studies, The Polytechnic, Ibadan, Nigeria

[2]Faculty of Computing, Ajayi Crowther University, Oyo, Nigeria

DOI: https://doi.org/10.51584/IJRIAS.2026.11090088

Received: 25 September 2026; Accepted: 30 September 2026; Published: 15 October 2026

ABSTRACT

This study developed and evaluated an Enhanced BGP configuration with MRAI reduced from 30 to 5 seconds and Keepalive/Hold timers reduced from 60/180 to 20/60 seconds against default BGP timers, using two identical Autonomous System topologies built in GNS3 (one router per AS, joined by a single 64 kbps serial link) and subjected to identical router-shutdown and cable-disconnection failures, with packet-level data captured and analyzed via Wireshark. In the router-failure run, Enhanced BGP reduced the observed unreachable period by 25.6% (33.82 s to 25.16 s), cut unanswered ICMP requests by about 25%, and recorded a 20.3% lower pre-failure round-trip time, but recorded a higher post-recovery round-trip time than the default configuration. Convergence time, packet loss and jitter did not improve consistently across the two failure scenarios, revealing a trade-off between faster failure detection and post-recovery path stability rather than a uniform improvement. Because each scenario was run once on a minimal topology, the results are indicative rather than statistically established. The study contributes a reusable GNS3 framework for comparative inter-domain routing research.

Keywords: Border Gateway Protocol, Enhanced BGP, convergence time, MRAI, distributed networks.

INTRODUCTION

The routing protocols are the pillars of modern networking infrastructure, which dictate how routers communicate reachability and topology data and how they choose the best path to data in more and more dynamic networks (Nasir and Tariq, 2018; RamKumar and Raj Anand, 2016). Of these, the Border Gateway Protocol (BGP) is the most widely used Exterior Gateway Protocol (EGP) for negotiating routing information between, not within, Autonomous Systems (ASes). The path-vector design, TCP transport and policy-driven best-path selection mechanism allow independently administered networks to share and coordinate inter-domain routing information at Internet scale.

Even though BGP is scalable and robust, at the time of its design, its focus was more on the routing stability and scalability, rather than rapid failure detection. The conventional BGP timer parameters are Minimum Route Advertisement Interval (MRAI), Keepalive interval, and Hold time, which control how often route advertisements are made, and how BGP sessions are detected when they fail. Typical 30-second MRAI, 60-second Keepalive interval, and 180-second Hold time are typical defaults for Cisco IOS implementations. These were more conservative settings to minimize too much traffic and computation when it comes to routing updates. But they can also cause failures to be delayed in the detection process in some cases, especially when failure detection relies on the expiration of the BGP Hold timer. As a result, a network might be down for an extended period before an alternate route is chosen and normal forwarding is resumed.

For enterprise and Internet Service Provider (ISP) networks, the operation consequences of BGP convergence delay are critical, as it can impact network service availability, application performance, and network reliability. In fact, there is a balance to be struck between having a conservative timer setting to ensure routing stability and a more aggressive setting that could lead to quicker failure detection, but at the expense of

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September

workload on the control plane and routing instability. The previous works have explored adaptive or reduced configurations of the BGP timer, such as the one by Lasković and Trajković (2006), or simulation-based techniques to study the behavior and attributes of BGP by Zidan and Hamarsheh (2016). There is still a need for the controlled experimental evaluation of a standard BGP configuration against a tuned configuration under the same network conditions and failure scenarios while taking into account the multiple network-performance indicators simultaneously.

This research fills this void by carrying out a controlled comparison of default and tuned BGP timer configurations. With the help of the Graphical Network Simulator-3 (GNS3) and Cisco IOS router images, two identical Autonomous System topologies were created, and the same network failure scenarios were applied to both of them (router down, cable disconnect). The standard and adjusted BGP timer configuration was tested for each topology, with the latter configuration, known as Enhanced BGP in this study, being considered. Using Wireshark, packet-level traffic captures were analysed to measure their throughput, loss, convergence time, jitter, and round trip time (RTT). The topologies and failure conditions are kept identical for both configurations to allow a side-by-side comparison of the performance implications of BGP timer tuning. The study does not look at reduced timers in isolation, but attempts to measure the advantages and disadvantages of quicker BGP failure detection and convergence.

RELATED WORK

The behaviour of the BGP protocol when a network failure occurs has been well studied since the study by Labovitz et al. (2000) concerning the Internet's routing convergence behaviour after network failures, and they showed that the Internet could take much longer than the duration of the network failure to converge when the network fails, due to path exploration, repeated withdrawals of routes, and subsequent update cycles. They found that in addition to the physical failure, the convergence delay depends on the signalling and route selection of the routing protocol used. This observation is important in developing the foundation for studying configuration based methods such as timer tuning, to enhance failure detection and route recovery speeds.

Lasković and Trajković (2006) suggested that the value of MRAI timer be dynamically changed depending on the network condition, rather than a fixed value. The study also introduced reusable MRAI timers that limit the advertisements of individual destinations independently, and evaluated the proposal with a BGP processing-delay model based on reported measurements, with the aim of reducing BGP convergence time. The approach showed that the convergence performance can be improved in the simulated environment assessed, but it was mostly targeted to convergence and routing-update behaviour. The study did not focus specifically on a wider performance evaluation using other parameters like throughput, packet loss, jitter and post recovery network performance. The result leaves a space for studying the impact of BGP timer tunings on the overall network behavior in case of failures. Moubarak et al. (2017) likewise proposed a BGP control mechanism (BGP-NCM) driven by network performance parameters, which reflects continued interest in performance-aware BGP control.

Zidan and Hamarsheh (2016) experimented with BGP path attributes in order to test their effect on the route selection process, and showed that GNS3 can be used for controlled experimentation with BGP. The research they carried out was mainly on policies and route selection and not on the performance impact of configuration changes in BGP in the event of failure. As a result, the main metrics used for evaluation were not convergence time, or metrics related to traffic performance like latency and packet loss. Although the study does not offer an exhaustive analysis of the behavior of BGP, it does give support for the use of GNS3 as an experimental platform for studying BGP behavior.

Baggan and Panda (2019) and Bhaga (2021) were more with SDN and Multi-homed network redundancy, respectively. The results show how significant the BGP configuration is and how crucial are the redundancy mechanisms for providing resilient inter-network communication. These studies, however, were not conducted in a controlled environment with a side-by-side experimental comparison of a tuned BGP configuration to the default configuration using the same network topologies and failure scenarios. Therefore, the particular performance impact of timer-tuning with respect to various traffic and convergence measurements still hasn't been well understood.

Similar studies in other routing domains prove that a multi dimensional assessment is also beneficial. Oyediran (2024), for instance, examined the AODV based MANET system with OMNeT++ simulator and RC5 encryption of the data for security in MANET. The study offers observations on network behaviours, but does not fully quantify the performance with a set of multiple indicators like packet delivery ratio, delay, thru-put and others. Likewise, Olaniyan et al. (2020) pointed out in their survey on power-aware and secure routing methods for Mobile Ad Hoc Network (MANETs), that many of the proposed approaches have been geared towards a particular objective, such as energy efficiency, or security from a specific security threat. While different from BGP, these studies make the important point again that the methodology of examining the performance of network protocols through multiple performance metrics is important. As a part of this research, this viewpoint lends credence to the convergence time, throughput, packet loss, RTT and jitter being used as complementary metrics to determine the impact of BGP timer tuning in practice.

The reviewed BGP literature, therefore, shows that there existed an important methodological gap. The existing research has examined BGP convergence, adaptive timer mechanisms, the behaviour of route-selection, and network redundancy, but has tended to focus on specific aspects of BGP performance in these studies, instead of offering a controlled comparison between default and tuned timer settings, over multiple metrics. First of all, there are few experimental studies published that reveal the impact of a particular strategy of tuning timers on the convergence time, throughput, packet loss, round trip time (RTT), and jitter, under the same failure events for both timers' configuration. This study aims to fill this gap by creating two identical Autonomous System (AS) topologies in GNS3, configuring one topology with the default BGP parameters and the other with the proposed tuned parameters and then simulating both topologies with the same failure scenarios for the routers and links. Finally, wireshark packet captures are then used to measure the difference in network performance, in terms of number of packets captured

Research Design

To test the effect of BGP timer setting on network performance in controlled failures a quantitative experimental research design was used. Two identical small-scale network topologies were built up in GNS3, one of them with the default BGP timer values, and the other with the suggested Enhanced BGP configuration. All other experimental parameters such as network topology, IP addressing, routing relationships, traffic patterns, router configurations and link characteristics were kept the same in order to have a controlled comparison. Thus, the configuration of the BGP timer was the main independent variable with the resulting changes in convergence time, throughput, packet loss, round-trip time (RTT) and jitter being the main performance measures. Various failure events such as router shutdown, disconnect from cables, etc. were triggered independently in both network environments and under similar conditions, resulting in controlled failure events. The resulting routing behaviour and traffic was captured and analysed for these two configurations, to identify the difference in the behaviour in case of a network failure and to quantify the performance difference between the two configurations.

Simulation Environment

The simulation environment was set up in GNS3 (Gayathri and Vadivel, 2017; Pan and Jain, 2008) with two Cisco 2691 routers with the c2691-entservicesk9-mz.124-13b_2 IOS image. The routers were connected together by a serial point-to-point link with a 64 kbps clock rate and each autonomous system (AS) was one of the routers. The Main-Router was assigned to AS 10 and used to connect to the 192.168.1.0/24 network, while the Campus-Router was assigned to AS 20 and connected to the 192.168.2.0/24 network. The inter-router connection was configured with the 10.0.0.0/8 subnet, with 10.0.0.1 and 10.0.0.2 configured on the interfaces of the two routers. Each Autonomous System (AS) was made up of Ethernet switch and four Virtual PC Simulator (VPCS) end-host nodes in order to have a controlled environment for generating and measuring network traffic. The inter-router link packet-level traffic was captured using the software WireShark, and then organised, statistically calculated and graphically analysed in Microsoft Excel for measured performance indicators during normal operation and controlled failure events. The 64 kbps clock rate represents a bandwidth-constrained link rather than a typical provider interconnect; the consequences of this choice are discussed under the limitations.

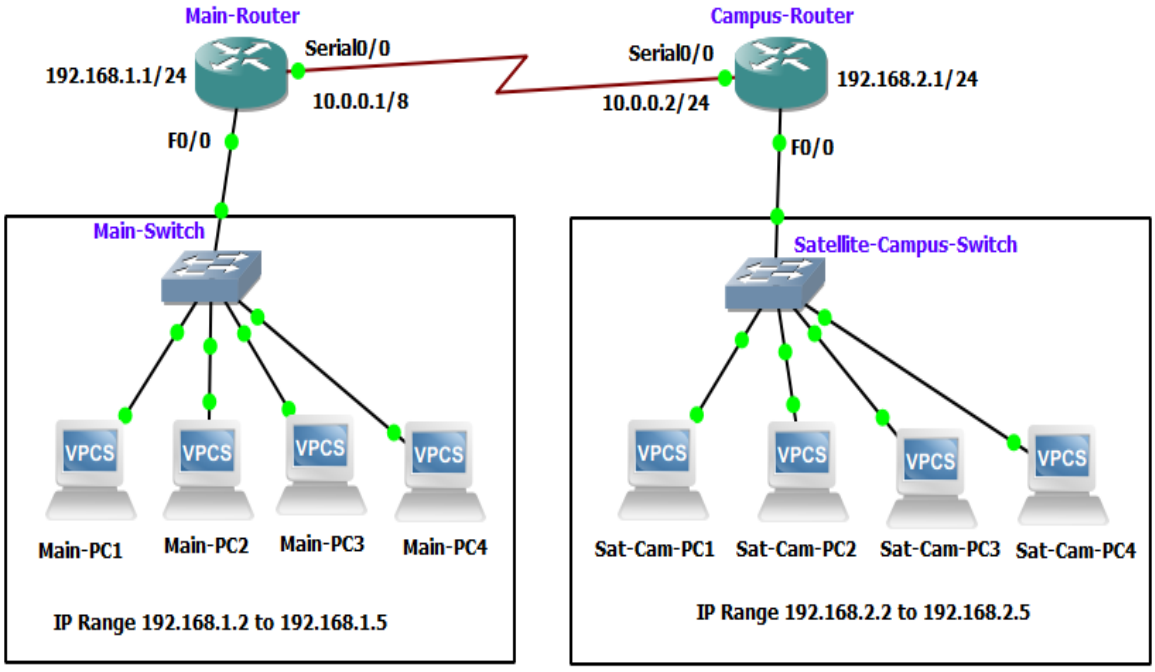

Figure 1: Network topology implemented in GNS3, showing Main-Router (AS 10) and Campus-Router (AS 20) connected via eBGP over a serial link, with their respective switches and VPCS end hosts.

BGP Configurations

Both routers were set up to peer eBGP between the two Autonomous Systems, and advertise the local networks of both ASes and the inter-router network. The routers were configured to use the following default BGP timers for the baseline configuration: Minimum Route Advertisement Interval (MRAI) is 30 seconds; Keepalive is 60 seconds; Hold timer is 180 seconds. With the Enhanced BGP configuration, the MRAI was set to 5 seconds, and both BGP neighbor relationships' Keepalive and Hold timers were set to 20 and 60 seconds, respectively. The other BGP parameters, the routing configuration, the network topology, addressing schemes, procedures for generating traffic and link characteristics were kept unchanged in both experimental configurations. The controlled configuration allowed the differences in the behaviour of the networks to be studied mainly against the BGP timer changes specified.

Failure Scenarios and Data Collection

Two failure scenarios have been tested in each of the BGP setups: controlled router shutdown and physical disconnection of the inter-AS link. In each experiment, the end-hosts in AS 10 and 20 generated continuous ICMP traffic and the traffic passing through the inter-router link was captured by using Wireshark. The data captured were then analysed to obtain the key performance metrics. The number of bytes captured, multiplied by 8, and divided by the capture time was used as throughput. The percentage of ICMP Echo Requests that did not receive an Echo Reply was calculated as a percentage of all the Echo Requests sent. Two related but different recovery measures were computed. The observed unreachable period (Table 1) was taken as the elapsed time between the first and the last ICMP Destination Unreachable message in the capture, that is, the interval during which a router reported that it had no route to the destination. The convergence time (Table 2) was taken as the largest gap between successive successful ICMP Echo Replies that was caused by the failure event and ended when BGP re-convergence restored reachability, that is, the longest interval without working end-to-end connectivity. It can therefore include periods in which packets were lost without a Destination Unreachable message being generated, and so it is expected to be longer than the unreachable period. The jitter was computed as the average of the absolute differences between the successive Round-Trip Time (RTT) readings. Connectivity and routing-path correctness was checked before each timed experiment by pinging and tracerouting to establish a baseline before the start of the experiment and to determine that the possibly intended routes were active. This process helped ensure that the changes in performance measured were

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September

correlated with the failures and subsequent convergence of BGP, and not with pre-existing connectivity or routing problems.

RESULTS

Table 1 summarizes the primary packet-level measurements captured in the router-failure scenario for both configurations (one run per configuration).

Table 1: Main packet-level measurements captured through Wireshark in the router-shutdown scenario (one run per configuration)

Measure	**Enhanced BGP**	**Default BGP**
Total packets captured	461	401
Capture duration (s)	237.08	243.16
ICMP Destination Unreachable messages	25	33
Observed unreachable ~~period~~period, first to last Destination Unreachable message (s)	25.16	33.82
Unanswered ICMP requests	24	32
Time from last unreachable to first success (s)	1.05	1.02
Average RTT before failure (ms)	63.7	80.0
Average RTT after failure (ms)	137.4	103.8

Table 2 reports throughput, packet loss, convergence time, and jitter across both failure scenarios (router shutdown, cable disconnection) for each configuration.

Table 2: Summary of metrics for both failure scenarios (one run per scenario)

Scenario	**Throughput (kbps)**	**Packet Loss (%)**	**~~Convergence Time (s)~~Convergence Time, longest gap in ICMP replies (s)**	**Jitter (ms)**
Default BGP – Router Off	1.17	30.15	93.89	83.90
Default BGP – Cable Disconnect	1.22	58.06	151.66	112.65
Enhanced BGP – Router Off	1.43	33.73	146.95	109.19
Enhanced BGP – Cable ~~Off~~Disconnect	1.90	32.79	126.05	130.10

In the router-failure run (Table 1), the Enhanced BGP configuration showed a smaller observed unreachable period of 25.16 s than the default configuration (33.82 s), a reduction of 25.6%. It also left fewer ICMP requests unanswered (24 compared with 32, about 25.0% fewer). The measured pre-failure RTT was lower under Enhanced BGP (63.7 ms) than under the default configuration (80.0 ms), a difference of about 20.3%.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September

These are descriptive differences from a single run per configuration. The unreachable period is derived from probe timestamps, so its resolution is limited by the probing interval, and no measure of variability or significance test is reported. The results were not consistently favourable to Enhanced BGP, however. The default configuration had a shorter post-failure RTT (103.8 ms) than Enhanced BGP (137.4 ms), and the time between the last unreachable message and the first successful reply was slightly shorter for the default configuration (1.02 s) than for Enhanced BGP (1.05 s). Table 2 shows a mixed pattern across the two failure scenarios. In the cable-disconnection scenario, Enhanced BGP had the lower packet loss (32.79% versus 58.06%) and the shorter convergence time (126.05 s versus 151.66 s). In the router-shutdown scenario it had the higher packet loss (33.73% versus 30.15%) and the longer convergence time (146.95 s versus 93.89 s). Enhanced BGP recorded higher throughput in both scenarios (1.43 and 1.90 kbps versus 1.17 and 1.22 kbps) and higher jitter in both (109.19 and 130.10 ms versus 83.90 and 112.65 ms). The router-shutdown convergence time in Table 2 therefore points in the opposite direction to the unreachable period in Table 1 for the same scenario. The two measures are not synonymous (see the definitions under Failure Scenarios and Data Collection): the unreachable period covers only the interval in which a router returned Destination Unreachable messages, whereas the convergence time, as defined, spans the longest gap in successful replies and can include intervals in which packets were lost without such a message. The present data do not show which phase of recovery accounts for the difference between the two measures, so the discrepancy is reported as a limitation of the dataset rather than resolved.

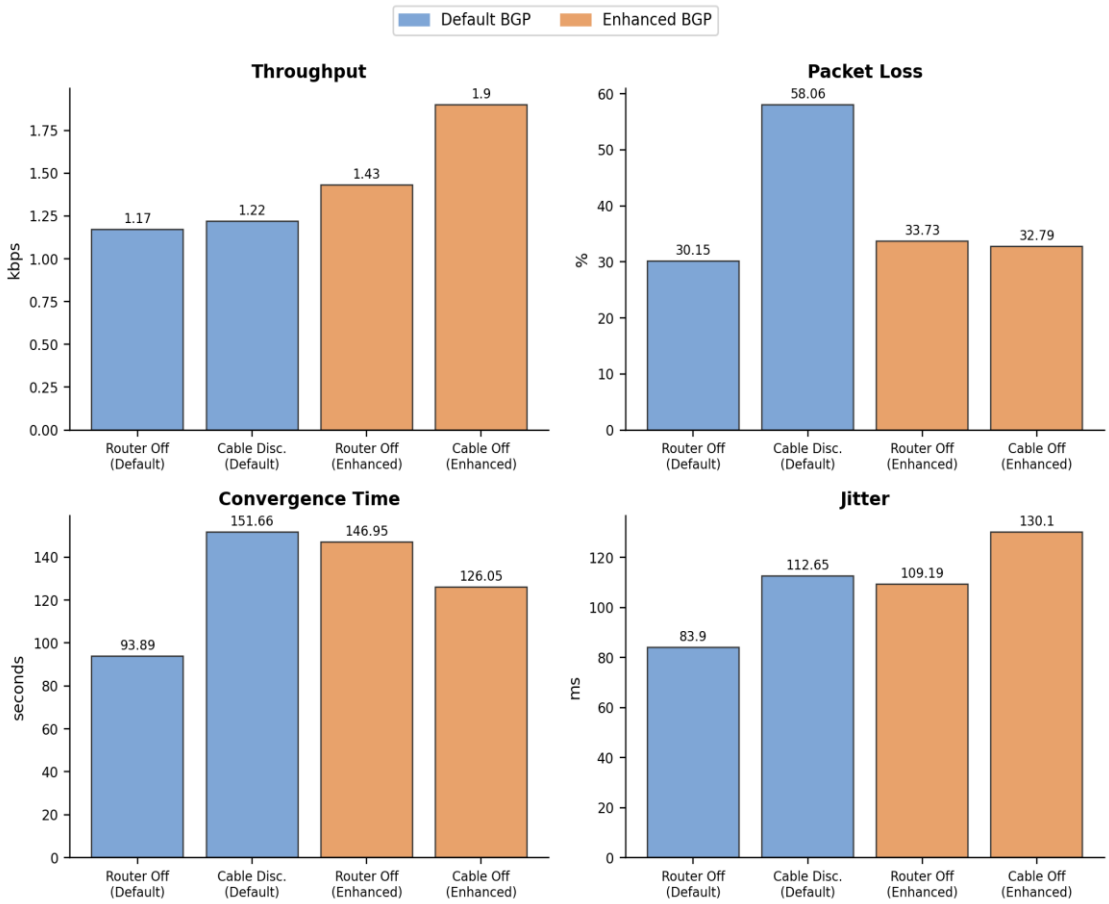

Figure 2. Throughput, packet loss, convergence time, and jitter for Default and Enhanced BGP across the router-off and cable-disconnect scenarios.

DISCUSSION

The results are not a uniform improvement; the pattern is mixed, which is informative in itself. The shorter unreachable period and the smaller number of unanswered requests under Enhanced BGP are in the direction expected from a shorter Hold timer: with the Hold time reduced from 180 s to 60 s, a silent neighbour is declared down sooner than with the longer timer. However, the measured difference (about 8.7 s) is much

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September

smaller than the 120 s difference between the two Hold times, so the Hold timer cannot be the only factor. One possibility is that some failures on a directly connected link are detected through the interface state (for example, through fast external fallover) before the Hold timer expires. This was not isolated in the experiment, and the present data cannot show which detection mechanism operated in each run. The lower pre-failure RTT (63.7 ms vs. 80.0 ms) is not readily explained by the timer change, because more frequent Keepalives would be expected to add, not remove, control traffic on the link. With a single run and a 64 kbps link on which queueing delay is sensitive to small differences in load, this difference is best treated as unexplained and possibly due to run-to-run variation.

The downside is that the default configuration fared better than Enhanced BGP on post-failure RTT (103.8 ms vs. 137.4 ms). One possible explanation is control-plane contention on the 64 kbps serial link between the routers: with the MRAI reduced from 30 s to 5 s, updates can be advertised up to six times as often (once every 5 s instead of once every 30 s), and immediately after a failure, when withdrawal and re-advertisement traffic is already dense, this higher rate of signalling competes for the same bandwidth as the measured data traffic. This is an inference from the timer configuration and link characteristics, not a measured mechanism, because link utilisation was not recorded in this study. It is offered as a plausible explanation and not a confirmed cause. The same caution applies to the link itself: 64 kbps is far below typical inter-AS capacities, so the size of any contention effect in a production setting cannot be inferred from these results.

Once the failure had been detected, the time from the last unreachable message to the first successful reply was similar for the two configurations (1.05 s for Enhanced BGP and 1.02 s for Default BGP). This suggests that the tuned timers mainly change when recovery begins rather than how quickly the network recovers once it has begun, although with one run per configuration and a difference of 0.03 s this should be read as similarity, not as evidence of a difference. The topology also has a single inter-AS path, so there is no alternate route to fail over to: recovery here means restoration of the failed component and re-establishment of the BGP session and routes, not re-routing around the failure.

These results support a more precise claim than "Enhanced BGP is faster". In this experiment, shortening the MRAI and Hold/Keepalive timers was associated with a shorter observed unreachable period and fewer unanswered requests in the router-failure run, but also with a higher post-recovery RTT and with mixed convergence-time, packet-loss and jitter results across the two scenarios. A possible increase in control-plane overhead is a hypothesis that this study did not measure. On this evidence, the tuned timers may be attractive where minimising outage time matters most, but the conclusion is tentative until the experiment is repeated on larger topologies with replicated runs. Standard timers may remain preferable where path stability during the recovery window matters more, and on slow links, higher-bandwidth inter-AS connections or rate-limiting should be considered to limit the signalling burden.

Practical Implications and Deployment Considerations

Timer tuning of this kind is most relevant to small and enterprise networks with one or a few eBGP peerings, where the operator controls both ends of the session and faster detection of a silent failure is worth a modest increase in keepalive and update traffic. It is less clearly suitable for large provider networks. Reducing the timers carries risks. A shorter Hold time makes a session more sensitive to transient congestion or packet loss, so a busy or low-capacity link may drop a healthy session and trigger withdrawals and re-advertisements (session flapping). A shorter MRAI increases the number of update messages and the processing and bandwidth they consume, and in larger topologies it can increase the number of intermediate updates sent during path exploration (Labovitz et al., 2000). Interoperability also matters: the Hold time used on a session is the smaller of the values proposed by the two peers (Rekhter et al., 2006), so the effective timers depend on the neighbour's configuration, and the Keepalive interval is normally set to one third of the Hold time, as in the 20/60 s setting used here. A neighbouring AS that keeps the default values will therefore limit the benefit.

Timer tuning should also be compared with other mechanisms for faster recovery. Bidirectional Forwarding Detection (BFD; Katz and Ward, 2010) detects forwarding-path failures independently of the BGP timers and can signal BGP to bring the session down as soon as the path is declared failed. It avoids lowering the BGP Hold time and the associated false-positive risk on the BGP session itself, although BFD has its own

configuration and processing costs. BGP Prefix Independent Convergence (PIC) addresses a different part of the problem: it pre-computes and installs backup paths in the forwarding table so that traffic can be switched to an alternate path quickly after a failure is detected, independently of the number of prefixes. PIC requires an alternate path to exist, which the single-path topology used here does not provide, so it could not be evaluated. Tuned timers, BFD and PIC are complementary rather than competing: BFD shortens detection, PIC shortens repair after detection, and tuned timers are a configuration-only option where neither is available. A comparison of Enhanced BGP with BFD-assisted failure detection is a natural next experiment.

Finally, this study used only two ASes, so it says nothing about how tuned timers affect the propagation of routes across multiple ASes. Each AS along a path applies its own MRAI, and repeated session resets could cause route flapping and, where route flap damping is used (Villamizar et al., 1998), suppression of the affected prefixes at other ASes. These effects would need a multi-AS topology to evaluate.

CONCLUSION AND LIMITATIONS

In this study, we designed and tested an Enhanced BGP configuration that reduces the MRAI, Keepalive and Hold timers relative to the BGP defaults, and compared it with the default configuration on two identical AS topologies in GNS3 in which the router and the link were failed in the same way. In the router-failure run, Enhanced BGP shortened the observed unreachable period and reduced the number of unanswered requests, but it did not improve round-trip time after recovery, and its convergence-time, packet-loss and jitter results were mixed across the two failure scenarios. The findings therefore suggest a trade-off between failure-detection performance and path performance after recovery, rather than an across-the-board improvement, and they should be treated as indicative because of the small topology and the single run per scenario.

There are several limitations of this study. First, the numbers reported for each scenario are not means from multiple trials but come from a single run per scenario. No measures of variability or statistical tests are reported, so the differences between configurations are descriptive and cannot be shown to be statistically significant; future work should repeat each condition, report means and standard deviations, and apply appropriate tests. Second, the topology had only two ASes, one router per AS, no redundant path and a single 64 kbps serial link, which is far below the capacity of typical inter-AS links and could have exaggerated control-plane contention; a larger number of ASes, alternate paths and higher-bandwidth links would show whether the effects persist at scale, and the results should not be generalised to production inter-domain routing. Third, link utilisation, CPU load and BGP message counts were not recorded, so the explanations offered for the post-recovery RTT and for the pre-failure RTT difference are untested hypotheses. Fourth, the two recovery measures (observed unreachable period and convergence time) are defined differently and gave conflicting indications for the router-failure scenario, and the unreachable period was reported for that scenario only. Fifth, the study compares one tuned configuration against the defaults and does not separate the effect of the MRAI change from that of the Keepalive/Hold change, or compare with mechanisms such as BFD. Lastly, only two types of failure were tested (router shutdown and cable disconnection); other failures in the original scope were not evaluated and are a direction for future research.

REFERENCES

1. Baggan, V., & Panda, S. N. (2019). Enhancing border gateway routing protocol with software defined networking. International Journal of Innovative Technology and Exploring Engineering, 8(8).
2. Bhaga, 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).
3. Gayathri, C., & Vadivel, R. (2017). Basic concept of network simulation tools. International Journal of Advanced Research in Computer and Communication Engineering, 6(1).
4. Katz, D., & Ward, D. (2010). Bidirectional Forwarding Detection (BFD) (RFC 5880). Internet Engineering Task Force. https://www.rfc-editor.org/info/rfc5880
5. Labovitz, C., Ahuja, A., Bose, A., & Jahanian, F. (2000). Delayed internet routing convergence. In Proceedings of the ACM SIGCOMM 2000 Conference (pp. 175–187). Association for Computing Machinery.

6. Lasković, N., & Trajković, L. (2006). BGP with an adaptive minimal route advertisement interval. In Proceedings of the 25th IEEE International Performance, Computing, and Communications Conference (IPCCC) (pp. 135–142). IEEE.
7. Moubarak, M. T., Elbayoumy, A. D., & Megahed, M. H. (2017). Design and implementation of BGP novel control mechanism (BGP-NCM) based on network performance parameters. Ain Shams Engineering Journal. https://doi.org/10.1016/j.asej.2017.02.008
8. Nasir, A., & Tariq, U. (2018). A comparative study of routing protocols including RIP, OSPF and BGP. LGU Research Journal of Computer Science and Information Technology.
9. Olaniyan, O. M., Olusesi, A. T., Omodunbi, B. A., Esan, A., Wahab, W. B., Oyediran, M. O., & Adanigbo, O. O. (2020). Power aware and secured routing protocol in mobile ad-hoc network: A survey. International Journal of Advanced Research in Engineering and Technology, 11(7), 706–717.
10. Oyediran, M. O. (2024). Development of an improved routing protocol in Wanet. Ajayi Crowther Journal of Pure and Applied Sciences, 3(4), 1–8. https://doi.org/10.56534/acjpas.v3i4.124
11. Pan, J., & Jain, R. (2008). A survey of network simulation tools: Current status and future developments. Washington University in St. Louis. http://www.cse.wustl.edu/~jain/cse567-08/index.html
12. RamKumar, K., & Raj Anand, S. (2016). Optimal performance analysis enabling OSPF and BGP in internal and external WAN. International Journal on Computational Sciences & Applications, 6(1). https://doi.org/10.5121/ijcacs.2016.1101
13. Rekhter, Y., Li, T., & Hares, S. (Eds.). (2006). A Border Gateway Protocol 4 (BGP-4) (RFC 4271). Internet Engineering Task Force. https://www.rfc-editor.org/info/rfc4271
14. Villamizar, C., Chandra, R., & Govindan, R. (1998). BGP route flap damping (RFC 2439). Internet Engineering Task Force. https://www.rfc-editor.org/info/rfc2439
15. Zidan, N., & Hamarsheh, M. (2016). Implementation of Border Gateway Protocol (BGP) attributes. International Journal of Application or Innovation in Engineering & Management, 5(2).