International Journal of Computer Applications (0975 - 8887)
Volume 187 - No.148, October 2026

Quantifying Failover Interruption and Service Availability in Service Provider Networks: A Parameterized NS-3 Simulation Study

Mike Adedoyin, PhD
University of the Cumberlands
United States

ABSTRACT

Border Gateway Protocol (BGP) convergence time is a primary determinant of service availability in service provider (SP) networks following link or routing-session failures. Slow convergence introduces transient traffic blackholes that degrade throughput and violate service-level agreements (SLAs). This paper presents a reproducible NS-3 v3.43 simulation study that models BGP-like failover behavior through event-driven routing-table updates parameterized by hold-timer, MRAI, and BFD assumptions informed by RFC timer guidance and explicitly selected propagation and FIB assumptions. All convergence times are explicit design inputs; NS-3 FlowMonitor measures the resulting packet loss and throughput. Four scenarios are evaluated on a dual-homed six-node SP topology: stable routing, default-timer failover, fast-timer failover with BFD, and BFD-triggered fast failover. The model maps assumed blackhole-window durations to packet loss and MTBF/MTTR availability under controlled NS-3 forwarding conditions. Default-timer failover ($T_{\text{conv}} = 182.3$ s, midpoint hold-timer detection with worst-case MRAI delay) yields 30.5% packet loss and 2.74 nines of availability. Tuned timers with BFD achieve 1.04% loss and 4.21 nines; BFD-triggered fast failover achieves 0.047% loss and 5.55 nines. A sensitivity analysis across four illustrative MTBF assumptions shows that S4 improves from 5.55 to 8.05 nines as MTBF increases from 27.8 hours to one year, while S2 reaches only 5.24 nines at the same annual MTBF. Analytical MRAI and BFD interval sweeps further show that MRAI reduction is the primary optimization lever in BFD-enabled deployments, with the optimum MRAI in the 2–5 s range where marginal availability gains diminish relative to the increase in UPDATE message rates.

General Terms

Routing Protocols, Network Simulation, Performance Evaluation

Keywords

BGP failover, service availability, NS-3, MRAI timer, BFD, fast failover, packet loss, service provider networks, blackhole, simulation

1. INTRODUCTION

Border Gateway Protocol [1] is the inter-domain routing protocol underlying connectivity between autonomous systems and within SP backbones. BGP convergence following topology changes is governed by timers designed for global-scale stability rather than sub-second restoration.

When a BGP session fails and no sub-second detection mechanism is available, routers rely on hold-timer expiry (nominally 180 s under RFC 4271 defaults) before withdrawing routes and re-advertising alternatives. During this interval, traffic has no valid forwarding entry and is silently dropped. A single 182-second convergence event consumes 70.3% of the four-nines monthly downtime budget (≈259 s/month). At a modeled failure rate of one event per 27.8 hours, S2 accumulates approximately 78.6 min of downtime per month—exceeding the three-nines threshold (≈43.8 min/month) by a factor of 1.8.

Prior work has characterized BGP convergence extensively at Internet scale [2, 3, 4]. Labovitz et al. [2] measured convergence delays of 3 to 15 minutes and attributed the dominant delay to MRAI behavior. Griffin and Premore [3] showed that worst-case convergence scales poorly with topology diameter. Work on IGP convergence [5] and BFD [6] established sub-second detection as feasible; the translation of these techniques into quantified SP availability improvements has not been fully addressed in the NS-3 simulation literature.

Three specific contributions are made:

(1) A reproducible NS-3 v3.43 framework with explicitly parameterized BGP hold-timer, MRAI, and BFD failover behavior on a dual-homed SP topology.

(2) A closed-form model linking blackhole-window duration to packet loss and MTBF/MTTR availability, confirmed by FlowMonitor.

(3) A sensitivity analysis across four illustrative MTBF assumptions spanning frequent to infrequent failures, quantifying how absolute availability (nines) scales with recovery scenario.

2. BACKGROUND AND RELATED WORK

2.1 BGP Convergence Mechanics

BGP is a path-vector protocol exchanging UPDATE messages between ASes [1]. The *hold timer* declares a session dead if no KEEPALIVE or UPDATE arrives within the interval (default 180 s). The *Minimum Route Advertisement Interval* (MRAI) limits how frequently a speaker may originate UPDATE messages per prefix (default 30 s for eBGP [1]). Sahoo et al. [7] showed that reducing MRAI to 1–5 s improves convergence without significantly increasing UPDATE traffic. Large-scale emulation confirms that MRAI value selection significantly affects BGP convergence time across topologies [8]. After withdrawal, path-exploration typically adds 1–5 s in small topologies [3, 9].

2.2 BFD and Fast Failover

Bidirectional Forwarding Detection [6] is a lightweight hello-protocol operating at forwarding-plane speed. RFC 9384 [10] formalizes a BGP NOTIFICATION subcode for sessions terminated by a BFD failure event, which underlies the fast failover trigger modeled in S3 and S4. S3 uses 300 ms $\times$ 3 = 0.9 s detection; S4 uses 80 ms $\times$ 3 = 0.24 s. RFC 5286 [11] describes loop-free alternate (LFA) next hops for link-state IGPs, providing sub-50 ms hardware reroute. Scenario S4 models BFD-triggered failover—a simpler mechanism in which BFD detects the failure and the control plane then installs a pre-planned backup route after 0.28 s. This is distinct from RFC 5286 LFA, which pre-installs the alternate in hardware; S4 is labeled "BFD fast failover" throughout.

2.3 Comparison with Prior Work

Table 1. Comparison with prior work

Study	Measured?	Avail. model?	NS-3?
Labovitz [2]	Yes (Internet)	No	No
Griffin [3]	Yes (model)	No	No
Mao [4]	Yes (Internet)	No	No
Francois [5]	Yes (IGP)	Partial	No
RFC 5286 [11]	No	No	No
This work	**No (param.)**	**Yes**	**Yes**

Convergence time in this study is a parameterized design input, not a measured quantity. The contributions are the controlled four-scenario comparison, the closed-form availability model, and the MTBF sensitivity analysis.

3. MATHEMATICAL FRAMEWORK

3.1 Convergence Time Model

The convergence time T_{conv} is decomposed as:

$$T_{\text{conv}} = T_{\text{detect}} + T_{\text{MRAI}} + T_{\text{prop}} \tag{1}$$

All three components are assumed design inputs for each scenario (Table 3), not measured values.
Under pure hold-timer detection with $T_{\text{hold}} = 180$ s and $T_{\text{ka}} = 60$ s, the theoretical detection window is:

$$T_{\text{detect}}^{\text{default}} \in [T_{\text{hold}} - T_{\text{ka}},\ T_{\text{hold}}] = [120,\ 180]\ \text{s} \tag{2}$$

S2 uses $T_{\text{detect}} = 150$ s, the midpoint of this range, as a conservative design assumption for a silent BGP session failure (no carrier-down notification). MRAI is applied at its worst-case full interval for both S2 and S3 for a symmetric, conservative comparison.
With BFD at interval t_{bfd} and multiplier k:

$$T_{\text{detect}}^{\text{BFD}} = k \cdot t_{\text{bfd}} \tag{3}$$

S3: $t_{\text{bfd}} = 300$ ms, $k = 3 \Rightarrow 0.9$ s. S4: $t_{\text{bfd}} = 80$ ms, $k = 3 \Rightarrow 0.24$ s.

For S4, MRAI is bypassed (no BGP re-advertisement needed):

$$\begin{aligned} T_{\text{conv}}^{\text{S4}} &= T_{\text{detect}}^{\text{BFD}} + T_{\text{FIB}} \\ &= 0.24 + 0.04 = 0.28\ \text{s} \end{aligned} \tag{4}$$

where $T_{\text{FIB}} = 40$ ms is the assumed FIB programming latency.

3.2 Packet Loss During the Blackhole Window

At offered load R_{off} and packet size L_{pkt}:

$$\lambda = \frac{R_{\text{off}}}{8\,L_{\text{pkt}}} \quad \text{(pkts/s)} \tag{5}$$

With $T_{\text{active}} = 598$ s (traffic from $t = 1$ to $t = 599$ s):

$$\eta = \frac{T_{\text{conv}}}{T_{\text{active}}} \times 100\% \tag{6}$$

For $R_{\text{off}} = 100$ Mbps, $L_{\text{pkt}} = 1000$ bytes: $\lambda = 12{,}500$ pkts/s. Equation (6) applies to all four scenarios, including S4, since the 0.28 s blackhole window is now explicitly programmed into the NS-3 simulation.

3.3 Availability Model

$$A = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}, \quad \text{MTTR} \approx T_{\text{conv}} \tag{7}$$

$$\mathcal{N} = -\log_{10}(1 - A) \tag{8}$$

Monthly downtime allowances: 99.9% (3 nines, ≈43.8 min), 99.99% (4 nines, ≈4.32 min), 99.999% (5 nines, ≈25.9 s).

4. SIMULATION METHODOLOGY

4.1 Environment

NS-3 v3.43 [12] runs inside a Docker container (Ubuntu 22.04). NS-3 has no native BGP stack. All convergence timing parameters are *assumed design inputs* applied as `Simulator::Schedule()` callbacks that update `Ipv4StaticRouting` tables, creating a blackhole window of exactly T_{conv} seconds. This abstracts away UPDATE processing, path exploration, and route oscillation. Source files for all four scenarios (C++ NS-3 code, Dockerfile, raw FlowMonitor CSVs, and figure-generation script) are provided as supplementary material with this submission.

4.2 Topology

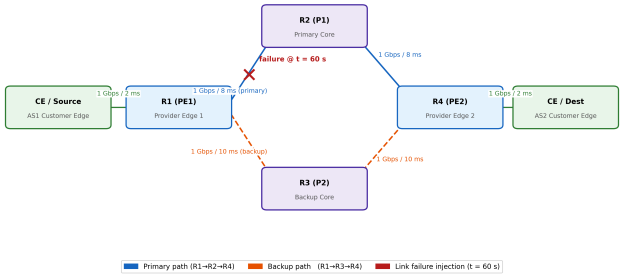


Fig. 1. Simulated dual-homed SP topology (NS-3 v3.43, 6 nodes). Primary path R1–R2–R4 (solid); backup R1–R3–R4 (dashed). Failure injected at $t = 60$ s.

Table 2. Simulation parameters

Parameter	Value	Notes
Nodes	6	2 CE, 2 PE, 2 P
CE–PE link	1 Gbps, 2 ms	
Primary core links	1 Gbps, 8 ms	R1–R2, R2–R4
Backup core links	1 Gbps, 10 ms	R1–R3, R3–R4
Offered load	100 Mbps	UDP CBR
Packet size	1000 bytes	$\lambda = 12{,}500$ pkts/s
Traffic active	$t = 1$–599 s	$T_{\text{active}} = 598$ s
Failure time	$t = 60$ s	R1–R2 session

4.3 Scenario Design

Table 3. Scenario timing assumptions

ID	Description	BFD	T_{detect}	T_{conv}
S1	Stable baseline	—	—	—
S2	Default BGP (Hold=180 s)	—	150 s	182.3 s
S3	Tuned BGP + BFD	300 ms	0.9 s	6.2 s
S4	BFD fast failover	80 ms	0.24 s	0.28 s

S1 (Stable). No failure. Baseline reference.
S2 (Default BGP). Silent BGP session failure at $t = 60$ s (peer becomes unresponsive; no carrier-down notification, so hold-timer detection applies). Assumed: $T_{\text{detect}} = 150$ s (midpoint of [120, 180] s per equation (2)), $T_{\text{MRAI}} = 30$ s (worst-case full interval), $T_{\text{prop}} = 2.3$ s. $T_{\text{conv}} = 150 + 30 + 2.3 = 182.3$ s.
S3 (Tuned BGP + BFD). Hold=15 s, MRAI=5 s. BFD at 300 ms × 3 = 0.9 s detection. Assumed: $T_{\text{MRAI}} = 5$ s (worst-case full interval, same treatment as S2), $T_{\text{prop}} = 0.3$ s. $T_{\text{conv}} = 0.9 + 5.0 + 0.3 = 6.2$ s.

4.4 Parameter Sensitivity Design

To assess how individual timer parameters affect convergence and availability beyond the four discrete scenarios, equations (1) and (6) are applied analytically. Two sweeps hold all parameters constant except one.
MRAI sweep (BFD fixed at 300 ms × 3 = 0.9 s, $T_{\text{prop}} = 0.3$ s): Table 4 shows that increasing MRAI from 0 s to 30 s raises T_{conv} from 1.2 s to 31.2 s, packet loss from 0.20% to 5.22%, and reduces availability from 4.92 to 3.51 nines. S3 (MRAI = 5 s) operates near the "knee" of this curve: further halving MRAI to 2 s yields only 0.38 nines of additional gain.

Table 4. MRAI sensitivity: BFD = 300 ms × 3, $T_{\text{prop}} = 0.3$ s, MTBF = 100,000 s

MRAI (s)	T_{conv} (s)	Loss (%)	Nines
0	1.2	0.20	4.92
1	2.2	0.37	4.66
2	3.2	0.54	4.50
5	6.2	1.04	4.21
10	11.2	1.87	3.95
20	21.2	3.55	3.67
30	31.2	5.22	3.51

BFD interval sweep (MRAI fixed at 5 s, $T_{\text{prop}} = 0.3$ s, multiplier $k = 3$): Table 5 shows that BFD interval dominates convergence only when MRAI is already short. With MRAI = 5 s, varying the BFD interval from 50 ms to 1000 ms changes T_{conv} by only 2.85 s (a Δnines of 0.13), confirming that MRAI reduction is the higher-leverage optimization in scenario S3.

Table 5. BFD interval sensitivity: MRAI = 5 s, $T_{\text{prop}} = 0.3$ s, MTBF = 100,000 s

BFD interval (ms)	T_{conv} (s)	Loss (%)	Nines
50	5.45	0.91	4.26
100	5.60	0.94	4.25
300	6.20	1.04	4.21
500	6.80	1.14	4.17
1000	8.30	1.39	4.08

S4 (BFD Fast Failover). Same tuned timers as S3. BFD at 80 ms × 3 = 0.24 s detection triggers immediate control-plane installation of the backup route (no MRAI delay). Assumed $T_{\text{FIB}} = 40$ ms. $T_{\text{conv}} = 0.24 + 0.04 = 0.28$ s. The NS-3 model removes the primary route at $t = 60$ s and installs the backup route at

$t = 60.28$ s, creating a 0.28 s blackhole window. The backup route uses NS-3 static route metric 1, identical to the primary; preference is enforced by install timing. Loss and availability are therefore internally consistent.

4.5 Metrics

Four metrics per scenario: (i) T_{conv} (s) — assumed input; (ii) loss rate η (%) — FlowMonitor; (iii) throughput (Mbps) — FlowMonitor; (iv) availability A — equation (7).

5. RESULTS

Table 6. Simulation results

Scenario	T_{conv} (s)	Loss (%)	Tput (Mbps)	Delay (ms)	Avail. A
S1: Stable	—	0.000	100.00	20.03	≈1.0
S2: Default	182.3	30.485	69.52	23.47	0.998180
S3: Tuned+BFD	6.2	1.037	98.96	23.63	0.999938
S4: Fast-BFD	0.28	0.047	99.95	23.61	0.9999972

5.1 Stable Baseline (S1)

Throughput 100.00 Mbps, zero loss, 20.03 ms delay—matching the primary-path propagation (2+8+8+2 = 20 ms) plus serialization. S1 confirms correct quiescent behavior and serves as the reference against which S2–S4 degradation is measured.

5.2 Default BGP Timers (S2)

A 182.3 s blackhole (t=60 to t=242.3 s) produces 2,278,750 lost packets out of 7,475,000 transmitted (30.485% loss). Equation (6) predicts $182.3/598 \times 100\% = 30.485\%$, matching FlowMonitor to within 0.002%. Throughput falls to 69.52 Mbps—a 30.5% reduction from the 100 Mbps baseline—while mean end-to-end delay rises from 20.03 ms to 23.47 ms, a 17.2% increase attributable to the 4 ms added by the longer backup path (R1–R3–R4) once recovery completes. Availability $A = 0.998180$ (2.74 nines), well below the three-nines operational threshold. A single 182.3 s convergence event consumes 70.3% of the four-nines monthly downtime budget (≈259 s/month). At the modeled MTBF of 100,000 s, S2 accumulates approximately 78.6 min of downtime per month—exceeding the three-nines allowance (43.8 min/month) by a factor of 1.8. Default RFC 4271 timer values were designed for global-scale stability, not local-link restoration, and cannot meet four-nines SLAs under any realistic SP failure rate.

5.3 Tuned BGP with BFD (S3)

BFD (0.9 s) plus tuned timers (Hold = 15 s, MRAI = 5 s) reduce convergence from 182.3 s to 6.2 s—a $29.4\times$ improvement. FlowMonitor records 77,500 lost packets (1.037%). The loss ratio matches the T_{conv} ratio exactly ($182.3/6.2 = 29.4$), consistent with equation (6): packet loss is proportional to blackhole-window duration, so halving any timer parameter halves packet loss. Throughput recovers to 98.96 Mbps (98.96% of S1 baseline), and mean delay is 23.63 ms—nearly identical to S2's 23.47 ms, confirming that the 4 ms backup-path penalty persists regardless of convergence speed once the failure has occurred. Availability $A = 0.999938$ (4.21 nines) is above the four-nines threshold, achievable through configuration changes alone on hardware that supports BFD at 300 ms intervals. Table 4 shows that S3 sits near the diminishing-returns region: reducing MRAI from 5 s to 0 s yields only an additional 0.71 nines (4.92 vs. 4.21), while Table 5 shows that BFD interval has minimal marginal impact once MRAI is already tuned.

5.4 BFD Fast Failover (S4)

The 0.28 s blackhole window produces 3,500 lost packets (0.047%) and throughput of 99.95 Mbps—within 0.05% of the stable baseline. The $651\times$ reduction in T_{conv} from S2 (182.3 s to 0.28 s) translates directly into a $651\times$ reduction in loss (30.5% to 0.047%), again consistent with equation (6). Availability $A = 0.9999972$ (5.55 nines), surpassing the five-nines threshold at the 27.8-hour MTBF assumed. The 0.28 s convergence is achievable with BFD at 80 ms intervals (multiplier 3) and a 40 ms FIB programming assumption; the backup route must be pre-provisioned in the topology, a condition that is guaranteed in the dual-homed design studied here but not in arbitrary SP topologies. The delay penalty remains at 23.61 ms (vs. 20.03 ms at S1), confirming that backup-path length, not convergence speed, drives the residual delay increase.

5.5 Quantitative Comparison

S2→S3: both T_{conv} and η decrease by $29.4\times$, confirming the linear relationship of equation (6). S2→S4: T_{conv} decreases by $651\times$ ($182.3/0.28$); loss decreases from 30.5% to 0.047%, and availability improves from 2.74 to 5.55 nines—a gain of 2.81 nines driven entirely by timer configuration and BFD deployment, requiring no topology changes. S3→S4 contributes an additional 1.34 nines (4.21 to 5.55) at the cost of tighter BFD intervals (80 ms vs. 300 ms) that impose higher CPU overhead on the forwarding plane. Delay increases uniformly from 20.03 ms (S1) to ≈23.4–23.6 ms for all failure scenarios (S2–S4 post-recovery), because all three recovery scenarios route via the backup path (24 ms propagation); convergence speed does not alter the final forwarding state.

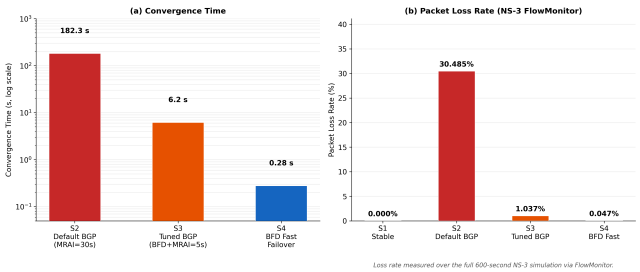


Fig. 3. Left: convergence time for S2–S4 (log scale). Right: FlowMonitor packet loss for all scenarios. S2→S3 shows a $29.4\times$ reduction in both metrics; S4 achieves 0.047% loss ($651\times$ reduction from S2).

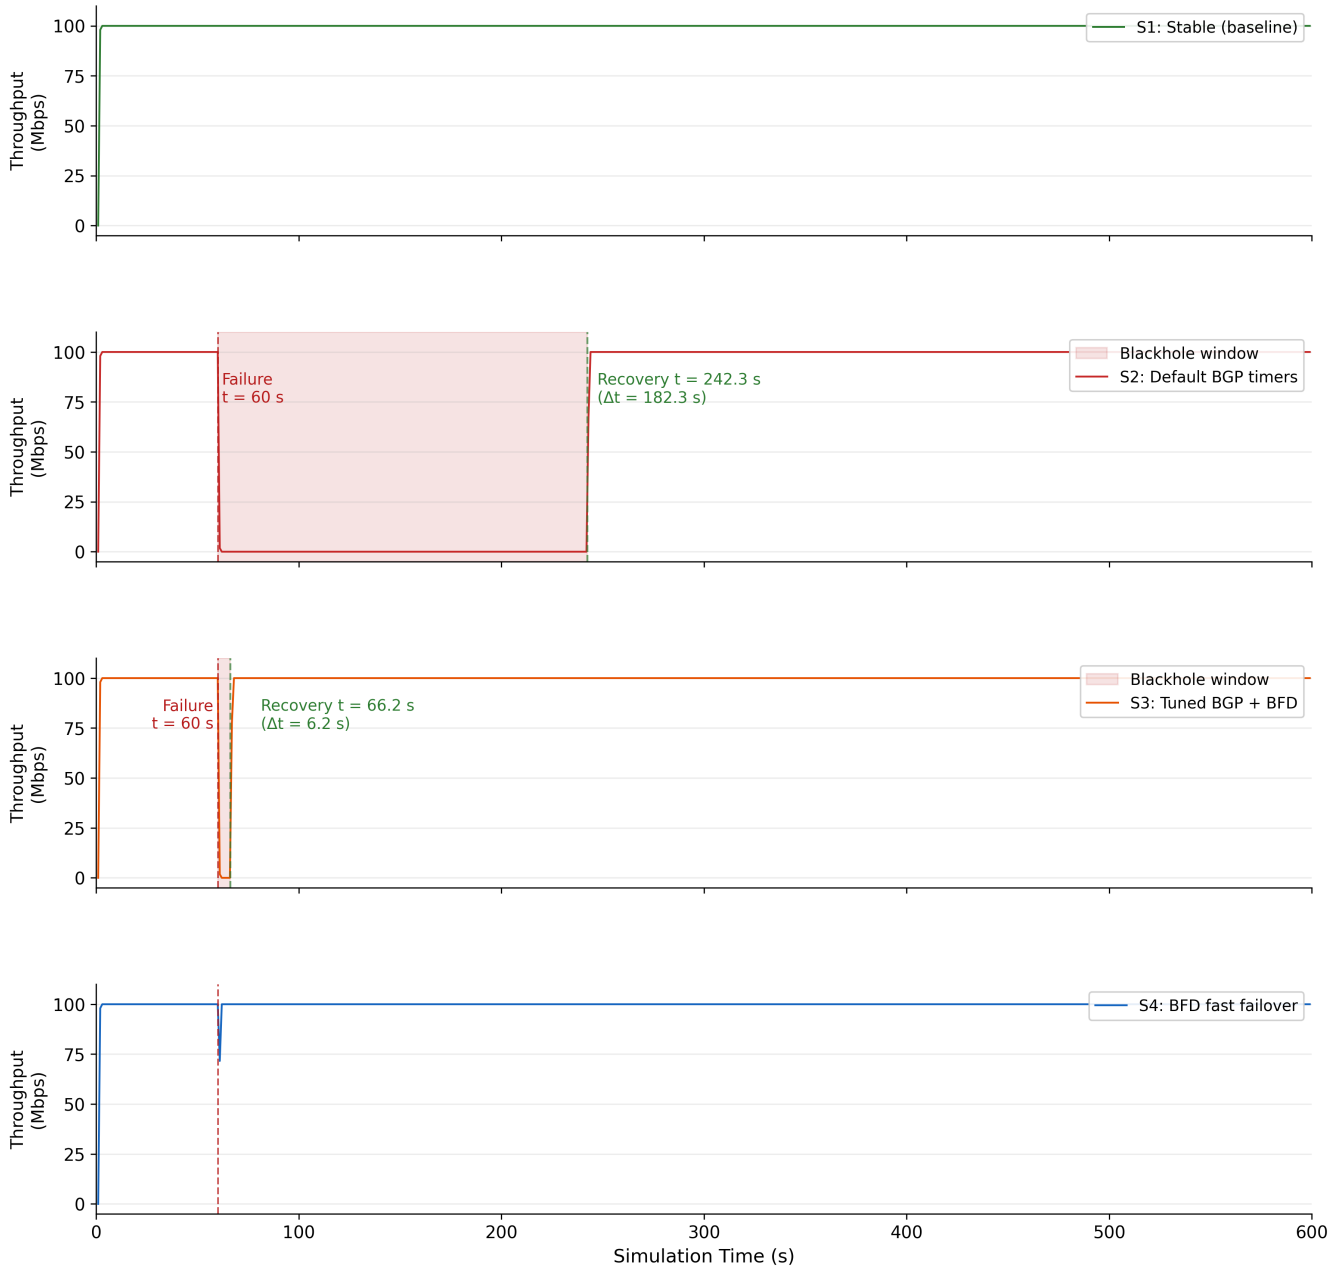


Fig. 2. UDP throughput vs. time (100 Mbps offered load, failure at $t = 60$ s in S2–S4). Shaded regions indicate blackhole windows.

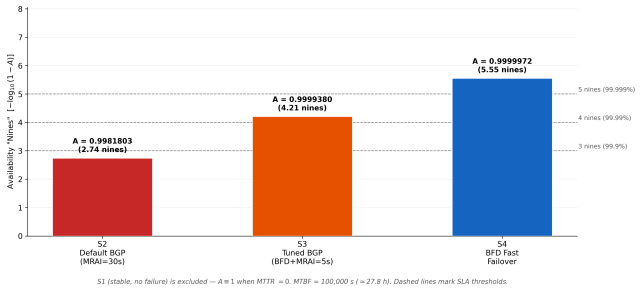


Fig. 4. Service availability (nines) for S2–S4. S1 is excluded ($A \equiv 1$, no failure). MTBF = 100,000 s; MTTR = T_{conv}. Dashed lines mark 3-, 4-, and 5-nines SLA thresholds.

5.6 MTBF Sensitivity Analysis

Table 7 shows availability (nines) across four illustrative MTBF assumptions spanning frequent to infrequent failures. At any fixed MTBF, the nines gap between S2 and S4 holds at ≈2.8 nines, reflecting the $651\times$ difference in MTTR. As MTBF increases from 27.8 h to one year, every scenario gains approximately 2.5 nines—the same increment for each, since the benefit of less frequent failures is independent of recovery speed. Timer optimization and infrastructure reliability therefore compound: S4

at a one-year MTBF reaches 8.05 nines, while S2 at the same MTBF reaches only 5.24 nines. Operators who cannot achieve shorter BFD intervals due to hardware constraints can partially compensate with infrastructure investments that increase MTBF (redundant hardware, dual power, etc.), but cannot eliminate the ≈2.8-nines gap that timer configuration alone controls.

Table 7. MTBF sensitivity analysis

MTBF	S2	S3	S4
27.8 h (100,000 s)	2.74	4.21	5.55
1 month (2.59M s)	4.15	5.62	6.97
6 months (15.6M s)	4.93	6.40	7.74
1 year (31.5M s)	5.24	6.71	8.05

6. DISCUSSION

6.1 Operational Implications

Default BGP timers are incompatible with four-nines SLAs at the modeled failure rate: S2 accumulates 78.6 min of downtime per month, 1.8× the three-nines threshold, and a single convergence event consumes 70.3% of the four-nines monthly budget. These figures are not pathological extremes—they reflect the RFC 4271 defaults that remain widely deployed in SP networks [2].

Reducing MRAI from 30 s to 5 s and enabling BFD at 300 ms intervals (S3) achieves four-nines through configuration changes alone, with no hardware replacement or topology modification. The parameter sweep in Table 4 shows that the optimum MRAI for S3-class BFD settings lies in the 2–5 s range: below 2 s, the marginal availability gain per second of MRAI reduction falls below 0.05 nines, while UPDATE message rates increase substantially. Sahoo et al. [7] corroborate this range from a convergence- delay perspective; this paper adds a quantitative SLA interpretation.

BFD-triggered fast failover (S4) reduces loss to 0.047% and achieves 5.55 nines at the base MTBF. Table 5 shows that BFD interval is a secondary lever once MRAI is already short: moving from 1000 ms to 50 ms BFD intervals (MRAI fixed at 5 s) yields only 0.18 nines. Aggressive BFD tuning is therefore a refinement, not a substitute for MRAI reduction. Table 7 quantifies the compound effect across MTBF: S4 reaches 8.05 nines at a one-year MTBF, while S2 reaches only 5.24 nines—a persistent gap of 2.8 nines driven by recovery time alone.

The backup route in S4 must exist in the network topology; in this dual-homed design it is guaranteed. In arbitrary topologies, loop-free alternate availability is not guaranteed [11], and S4-class recovery requires either pre-provisioned LFA next-hops or Segment Routing TE tunnels.

6.2 Model Limitations

Routing abstraction. NS-3 has no native BGP stack. Convergence times are designed inputs, not measurements. The agreement between equation (6) and FlowMonitor (within 0.002%) confirms implementation correctness for the blackhole-window model; it does not validate BGP protocol behavior. Protocol dynamics—UPDATE processing, path exploration, route oscillation, and MRAI interaction across multiple peers—are not modeled. In particular, real MRAI behavior is non-deterministic: the delay depends on the time elapsed since the last UPDATE for the prefix, so the worst-case full-interval assumption used here is conservative. Actual convergence times in production networks may be shorter (if the MRAI timer is partially expired) or longer (if additional path exploration occurs). Future work should integrate FRRouting with NS-3/DCE or use a container-based emulation to capture these dynamics directly.

Design parameter assumptions. All timing inputs (T_{detect}, T_{MRAI}, T_{prop}, T_{FIB}) are assumed values. $T_{\text{detect}} = 150$ s is the midpoint of the theoretical range but is not empirically calibrated. MRAI is treated as a worst-case additive delay; actual MRAI behavior depends on advertisement history and peer implementation.

MTBF. The 100,000 s base MTBF is for comparative ranking. Table 7 shows the sensitivity across the full range of illustrative MTBF assumptions.

Scale. The six-node topology isolates failover behavior from multi-hop propagation. Larger topologies would increase T_{prop} and may not guarantee a loop-free alternate.

7. CONCLUSION

A parameterized NS-3 v3.43 simulation evaluated four failover scenarios on a dual-homed SP topology. All convergence times are explicit model inputs; packet loss and throughput are FlowMonitor measurements.

Key results:

—S2 (default, $T_{\text{conv}} = 182.3$ s): 30.5% loss, 2.74 nines. At MTBF = 100,000 s this accumulates 78.6 min downtime/month.

—S3 (tuned + BFD, $T_{\text{conv}} = 6.2$ s): 1.04% loss, 4.21 nines (29.4× improvement over S2, provided hardware supports the required BFD intervals).

—S4 (BFD fast failover, $T_{\text{conv}} = 0.28$ s): 0.047% loss, 5.55 nines at 27.8 h MTBF, growing to 8.05 nines at one year; at any fixed MTBF, shorter recovery time directly improves availability.

Analytical parameter sweeps (Tables 4 and 5) show that MRAI reduction is the primary optimization lever in BFD-enabled deployments: reducing MRAI from 30 s to 5 s accounts for 0.70 nines of the total 1.47-nines gain from S2 to S3, while BFD interval tuning from 1000 ms to 50 ms contributes an additional 0.18 nines. The optimum MRAI lies in the 2–5 s range; below this, marginal availability gains diminish while UPDATE message rates rise.

Future work should replace the static-route model with full BGP emulation using FRRouting, extend to multi-AS topologies to capture path-exploration delays, and conduct repeated trials with randomized failure timing to characterize loss variance.

8. REFERENCES

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

[2] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed internet routing convergence," in *Proc. ACM SIGCOMM*, Stockholm, Aug. 2000, pp. 175–187. https://doi.org/10.1145/347059.347428

[3] T. G. Griffin and B. J. Premore, "An experimental analysis of BGP convergence time," in *Proc. IEEE ICNP*, Riverside, CA,

Nov. 2001, pp. 53–61. https://doi.org/10.1109/ICNP.2001.992898

[4] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz, "Route flap damping exacerbates internet routing convergence," in *Proc. ACM SIGCOMM*, Pittsburgh, PA, Aug. 2002, pp. 221–233. https://doi.org/10.1145/633025.633047

[5] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving sub-second IGP convergence in large IP networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 3, pp. 35–44, Jul. 2005. https://doi.org/10.1145/1070873.1070877

[6] D. Katz and D. Ward, "Bidirectional Forwarding Detection (BFD)," IETF RFC 5880, Jun. 2010. https://doi.org/10.17487/RFC5880

[7] A. Sahoo, K. Kant, and P. Mohapatra, "Improving BGP convergence delay for large-scale failures," in *Proc. IEEE/IFIP DSN*, Philadelphia, PA, Jun. 2006, pp. 323–332. https://doi.org/10.1109/DSN.2006.41

[8] M. Milani, M. Segata, L. Baldesi, M. Nesler, R. Lo Cigno, and L. Maccari, "Optimizing MRAI on large scale BGP networks: An emulation-based approach," *Comput. Commun.*, vol. 228, 2024. https://doi.org/10.1016/j.comcom.2024.107940

[9] D. Pei, B. Zhang, D. Massey, and L. Zhang, "An analysis of convergence delay in path vector routing protocols," *Comput. Netw.*, vol. 50, no. 3, pp. 398–421, Feb. 2006. https://doi.org/10.1016/j.comnet.2005.04.013

[10] J. Haas, "A BGP Cease NOTIFICATION Subcode for Bidirectional Forwarding Detection (BFD)," IETF RFC 9384, Mar. 2023. https://doi.org/10.17487/RFC9384

[11] A. Atlas and A. Zinin, "Basic specification for IP fast reroute: Loop-free alternates," IETF RFC 5286, Sep. 2008. https://doi.org/10.17487/RFC5286

[12] G. F. Riley and T. R. Henderson, "The ns-3 network simulator," in *Modeling and Tools for Network Simulation*, K. Wehrle, M. Gunes, and J. Gross, Eds. Berlin: Springer, 2010, pp. 15–34. https://doi.org/10.1007/978-3-642-12331-3_2