

Simulation of BGP Prefix Hijacking, Traffic Engineering, and Policy-Based Mitigation in a Multi-ISP Mini-Internet Using GNS3

Mohammed Shafiullah Raihan
Department of Computer Science and Engineering (CSE), International Islamic University Chittagong, Chittagong, Bangladesh.

Corresponding Author: * Mohammed Shafiullah Raihan

DOI: https://doi.org/10.5281/zenodo.24217374

Abstract

This project builds and evaluates a six–Autonomous System (AS) "mini-Internet" in GNS3 using Cisco c7200 routers to simulate realistic BGP behavior. The topology includes a Tier-1 transit provider, two ISPs, a multi-homed legitimate prefix owner (victim), a customer network, and an attacker AS. We first established normal eBGP connectivity and verified global reachability between customer and victim networks. We then simulated two BGP prefix hijacking attacks: (1) a basic origin hijack of a /24 and (2) a more-specific /25 hijack that exploits longestprefix-match forwarding behavior. Finally, we implemented layered policy-based mitigations including prefix filtering, AS-PATH filtering, and max-prefix protection, and demonstrated traffic engineering using Local Preference and AS-PATH prepending to reflect business routing decisions.

Tested results show that without filtering, the attacker's announcements can become the selected BGP best path in parts of the network, and the /25 hijack consistently diverts traffic for addresses within the more specific block. After applying customer/provider policies (prefixlists + route-maps, AS-PATH constraints, and max-prefix limits), hijacked routes were rejected and end-to-end traffic from the customer LAN to the victim LAN was restored. Traffic engineering changes successfully shifted preferred outbound and inbound paths according to policy, demonstrating how routing decisions can reflect economics and operational goals in addition to shortest-path metrics.

Manuscript Information

- **ISSN No:** 3139-9851
- **Received:** 12-09-2026
- **Accepted:** 19-10-2026
- **Published:** 14-11-2026
- **IJCATM:**1(4); 2026: 112-123
- **©2026, All Rights Reserved**
- **Plagiarism Checked:** Yes
- **Peer Review Process:** Yes
- **Paper Type:** Research paper

How to Cite

Raihan MS. Simulation of BGP Prefix Hijacking, Traffic Engineering, and Policy-Based Mitigation in a Multi-ISP Mini-Internet Using GNS3. Int. J. Current Acad. Trend. Multi. 2026;1(4): 112-123

Access this Article Online

www.ijcatmjournal.com

KEYWORDS: Global Reserve Currency, Triffin Dilemma, Savings-Investment Identity, Exchange Rate Pass-Through, Monetary Spillovers, SVAR and VECM Modeling, De-dollarization.

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

1. INTRODUCTION

The Border Gateway Protocol (BGP) is the backbone routing protocol of the global Internet, responsible for exchanging reachability information between Autonomous Systems (ASes). Despite its critical role, BGP was designed under a trust-based model and lacks built-in mechanisms to verify the authenticity of route advertisements. This architectural limitation has made BGP vulnerable to misconfigurations and malicious attacks, most notably prefix hijacking.

In recent years, numerous real-world incidents have demonstrated how incorrect or malicious BGP announcements can divert traffic, disrupt services, or enable traffic interception. These incidents highlight the importance of understanding BGP behavior and implementing operational safeguards at ISP and enterprise boundaries.

This project focuses on simulating a realistic multi-ISP environment using GNS3 to study normal BGP operation, analyze the impact of prefix hijacking attacks, and evaluate the effectiveness of commonly deployed mitigation techniques. By emulating real Cisco IOS routers, the project aims to bridge theoretical routing concepts with practical, observable behavior.

1.1 Goal of the Project

The goal of this project is to simulate, observe, and mitigate Border Gateway Protocol (BGP) prefix hijacking within a realistic multi-ISP mini-Internet environment. The project demonstrates normal eBGP route propagation across multiple interconnected Autonomous Systems and examines how routing behavior changes under adversarial conditions. Two common hijacking variants are analyzed: a same-length prefix hijack involving a /24 network and a morespecific /25 hijack that exploits longest-prefix-match forwarding behavior. In addition, the project evaluates defense-in-depth mitigation techniques deployed at ISP boundaries, including prefix filtering, AS-PATH filtering, and maximum prefix limits. Traffic engineering mechanisms using Local Preference and AS-PATH prepending are also implemented to illustrate how routing decisions are influenced by operational and business policies rather than purely technical metrics.

1.2 Summary of Work and Findings

The project begins by establishing a stable multi-AS BGP environment and verifying legitimate end-to-end connectivity between customer and victim networks. Controlled hijacking scenarios are then introduced from a malicious AS, and their impact on routing tables, path selection, and traffic forwarding is observed. Experimental results show that a basic /24 origin hijack may succeed partially depending on BGP best-path selection and network topology. In contrast, a more-specific /25 hijack consistently attracts traffic for the covered address range due to longest-prefix-match forwarding. The results further demonstrate that strict prefix filtering at provider edges is the most effective practical mitigation. AS-PATH filtering and max-prefix limits provide complementary protection and contribute to a layered security posture. Finally, traffic engineering experiments confirm that Local Preference can override AS-PATH length, allowing operators to implement routing decisions based on policy and business relationships.

2. OBJECTIVES OF THE PROJECT

The primary objective of this project is to design and implement a multi–Autonomous System mini-Internet and evaluate BGP routing behavior under both normal and adversarial conditions. The project aims to demonstrate how legitimate BGP routing is established across multiple ISPs and how insufficient routing policies can allow unauthorized or malicious route advertisements to propagate through the network. Additional objectives include simulating realistic BGP prefix hijacking attacks, analyzing their effects on routing tables, best-path selection, and reachability, and implementing layered mitigation techniques such as prefix filtering, AS-PATH filtering, and maximum prefix limits. The project also seeks to demonstrate policy-driven routing decisions through traffic engineering mechanisms that reflect real-world operational practices.

3. BACKGROUND AND MOTIVATION

3.1 Overview of BGP

The Border Gateway Protocol is a path-vector interdomain routing protocol responsible for exchanging reachability information between Autonomous Systems. Unlike interior routing protocols, BGP does not compute routes based on shortest-path metrics. Instead, it applies a structured decision process that prioritizes policy attributes such as Local Preference, ASPATH length, and origin type. This design allows network operators to enforce routing policies aligned with business agreements, but it also introduces complexity and security challenges.

3.2 BGP Prefix Hijacking

BGP prefix hijacking occurs when an Autonomous System advertises an IP prefix that it does not legitimately own. If other networks accept the advertisement, traffic destined for the legitimate prefix owner may be diverted, intercepted, or dropped. While many hijacking incidents result from accidental misconfigurations, deliberate attacks have also been observed. Morespecific prefix hijacks, such as advertising a /25 instead of a legitimate /24, are particularly effective because IP forwarding follows the longest-prefix-match rule, causing traffic to prefer the attacker's route even when the legitimate route remains available.

3.3 Why GNS3 (Emulation) for This Work

GNS3 was selected as the experimental platform because it supports emulation of real Cisco IOS routing behavior, including eBGP session handling, policy enforcement, and convergence dynamics. Unlike abstract simulators, GNS3 allows direct interaction with IOS configuration features such as prefix-lists, route-maps, and AS-PATH filters, making it well suited for studying operational BGP security mechanisms in a controlled environment.

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

4. LITERATURE REVIEW

Border Gateway Protocol (BGP) security has been an active area of research and operational concern since the early expansion of the global Internet. Because BGP was designed without intrinsic authentication or validation mechanisms, numerous studies and operational reports have documented its vulnerability to route leaks, misconfigurations, and deliberate prefix hijacking. This literature review summarizes relevant concepts and operational practices directly related to BGP hijacking and mitigation, which form the foundation of the experimental work in this project.

4.1 BGP Routing and Trust Model

BGP operates on a trust-based interdomain routing model in which Autonomous Systems exchange reachability information with limited inherent verification. Early studies highlight that BGP assumes cooperative behavior among network operators, relying on bilateral agreements and manual policy enforcement to maintain routing correctness. While this model enabled rapid Internet scalability, it also introduced systemic weaknesses, as routers generally accept routing announcements from neighbors unless explicitly filtered.

Research and operational analyses emphasize that BGP decision-making prioritizes policy attributes such as Local Preference and AS-PATH length over shortest-path metrics. As a result, routing behavior is highly dependent on operator-defined policies, making misconfigurations or malicious announcements capable of influencing large portions of the routing system.

4.2 BGP Prefix Hijacking

Prefix hijacking is a well-documented routing security incident in which an Autonomous System announces an IP prefix that it does not legitimately own. Studies distinguish between origin hijacks, where the attacker claims ownership of an existing prefix, and more-specific prefix hijacks, where the attacker announces a longer prefix to exploit longest-prefix-match forwarding. Operational incident analyses consistently show that more-specific hijacks are particularly effective, as forwarding behavior prefers the most specific route regardless of BGP best-path selection for less specific prefixes.

Real-world cases reported by network operators and regional Internet registries demonstrate that hijacks can lead to traffic interception, service disruption, or large-scale outages. These findings directly motivate the attack scenarios implemented in this project.

4.3 Operational Mitigation Techniques

Operational best practices for mitigating BGP hijacking focus on policy enforcement at network boundaries. Prefix filtering, based on customer or peer authorization, is widely regarded as the most effective first line of defense, as it prevents unauthorized announcements from entering the routing system. AS-PATH filtering provides additional validation by restricting acceptable origin AS values for known prefixes, though it depends on accurate knowledge of routing relationships.

Max-prefix limits are commonly deployed to protect routers from excessive route advertisements caused by configuration errors or malicious behavior. While max-prefix mechanisms do not prevent hijacking directly; they reduce operational impact by limiting routing table growth and session instability.

More recent literature introduces cryptographic solutions such as Resource Public Key Infrastructure (RPKI) and Route Origin Validation (ROV), which aim to formally verify prefix ownership. Although these mechanisms represent a significant advancement in routing security, their deployment remains partial. Consequently, traditional policy-based mitigations remain highly relevant and are the focus of this project's experimental evaluation.

5. Problem Statement

Despite widespread deployment of BGP, many networks operate with minimal filtering and validation of routing information. This project investigates how easily BGP prefix hijacking can occur in a multi-ISP environment and how effectively standard operational practices mitigate it.

5.1 Description

Design and implement a multi-AS network that:

- Demonstrates normal BGP operations across interconnected ISPs and customers.
- Simulates two hijack attacks (same-prefix and more-specific).
- Implements layered mitigation techniques.
- Demonstrates traffic engineering using policy attributes.

5.2 Calculations

Subnetting for /30 Point-to-Point Links

A /30 provides 4 addresses: network, two usable, broadcast. Example: 10.0.46.0/30

- **Network:** 10.0.46.0
- **Hosts:** 10.0.46.1, 10.0.46.2
- **Broadcast:** 10.0.46.3

/24 and /25 Prefix Behavior

- 203.0.113.0/24 provides 256 addresses (254 usable).
- 203.0.113.0/25 covers 203.0.113.0–203.0.113.127 (128 addresses).
- Forwarding uses longest-prefix-match; thus /25 is preferred over /24 for covered destina-tions.

5.3 Software and Hardware

- **Software:** GNS3 (Windows), Dynamips c7200 IOS (advipservices), VPCS.
- **Hardware:** 8 GB RAM recommended, stable CPU; Idle-PC required for c7200 performance.

6. Designs

6.1 AS Roles and Prefix Ownership

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

Table 6.1: Autonomous System Roles

AS Number	Role	Notes
AS100	Customer network	Originates 198.51.100.0/24
AS200	Attacker AS	Hijacks 203.0.113.0/24 and 203.0.113.0/25
AS300	Regional ISP	Aggregation + policy enforcement point
AS400	Victim / Legit owner	Originates 203.0.113.0/24 (multi-homed)
AS500	Tier-1 transit	Transit between regions
AS600	Upstream ISP of AS400	Provider path for AS400

Table 6.2: Prefix Ownership

AS	Prefix
AS400	203.0.113.0/24 (victim/legitimate)
AS100	198.51.100.0/24 (customer)
AS200	Unauthorized advertisements: 203.0.113.0/24, 203.0.113.0/25

IP Addressing Plan

Table 6.3: Inter-AS Links (Point-to-Point)

Link	Subnet	Purpose
R400–R600	10.0.46.0/30	AS400–AS600
R600–R500	10.0.56.0/30	AS600–AS500
R500–R300	10.0.53.0/30	AS500–AS300
R300–R100	10.0.31.0/30	AS300–AS100
R300–R200	10.0.32.0/30	AS300–AS200
R300–R400	10.0.34.0/30	AS300–AS400

Table 6.4: LAN Networks

LAN	Prefix	Gateway
AS100 LAN	198.51.100.0/24	198.51.100.1 (R100 f0/0)
AS400 LAN	203.0.113.0/24	203.0.113.1 (R400 f0/0)

6.3 BGP Design

All inter-AS BGP sessions are eBGP. AS400 is multi-homer to AS300 and AS600. The attacker AS200 peers to AS300.

6.4 Policy/Mitigation Design

- Prefix filtering: Accept only authorized prefixes from neighbors (especially customers).
- AS-PATH filtering: Ensure expected origin AS for critical prefixes (lab policy).
- Max-prefix: Limit number of accepted prefixes to reduce blast radius of leaks/floods.

7. Network Design and Topology

7.1 Overall Architecture

The simulated network represents a small-scale but realistic interdomain Internet environment composed of six Autonomous Systems interconnected through eBGP sessions.

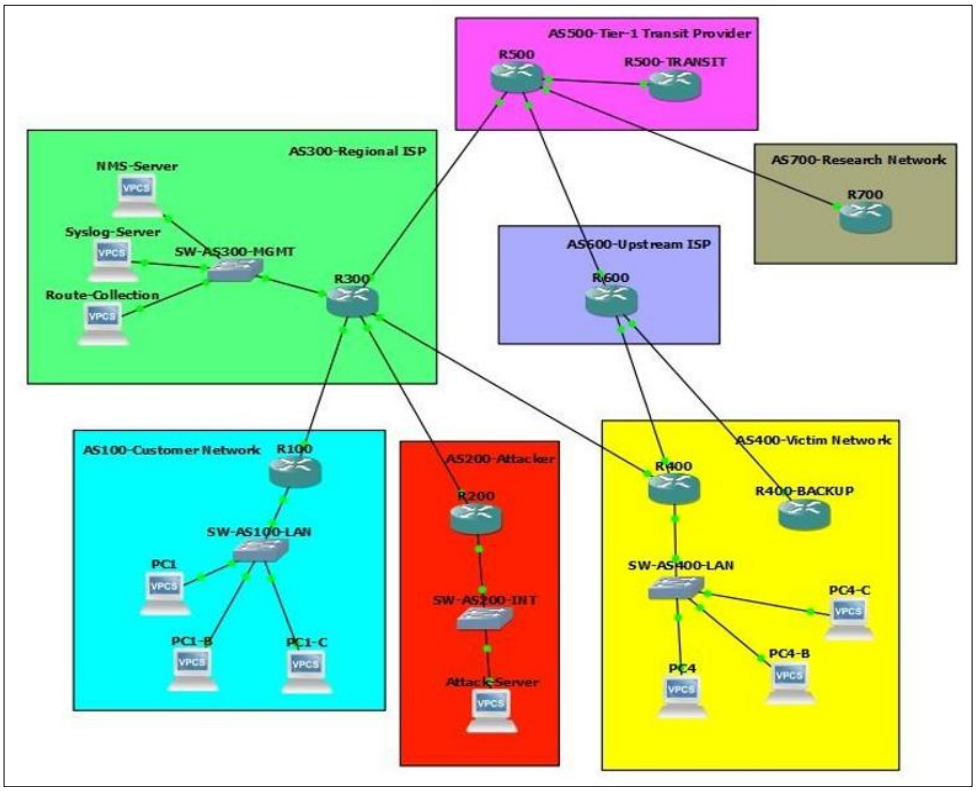

Fig 7.1: Multi-AS Mini-Internet Topology in GNS3

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

7.2 Roles of Autonomous Systems

Table 7.1: Autonomous System Roles and Functions

AS Number	Role	Description
AS100	Customer Network	Originates customer prefix
AS200	Attacker	Announces illegitimate prefixes
AS300	Regional ISP	Aggregation and policy enforcement
AS400	Victim	Legitimate prefix owner
AS500	Tier-1 Provider	Internet transit
AS600	Upstream ISP	Primary provider for AS400

7.3 IP Addressing Scheme

Table 7.2: Summary of IP Address Allocation

Link / Network	IP Subnet
AS400 – AS600	10.0.46.0/30
AS600 – AS500	10.0.56.0/30
AS500 – AS300	10.0.53.0/30
AS300 – AS100	10.0.31.0/30
AS300 – AS200	10.0.32.0/30
AS300 – AS400	10.0.34.0/30
AS100 LAN	198.51.100.0/24
AS400 LAN	203.0.113.0/24

8. Implementation

8.1 Simulation Environment

The network was implemented using GNS3 with Cisco IOS c7200 routers running the advipservices image and VPCS hosts. An appropriate Idle-PC value was applied to each c7200 router to prevent excessive CPU utilization and ensure stable simulation performance. All routers were configured with IP routing enabled and basic interface configurations before establishing BGP sessions.

8.2 Normal BGP Operation (Verification)

Key commands used: show ip bgp summary show ip bgp show ip bgp 203.0.113.0 ping 203.0.113.10 traceroute 203.0.113.10

```
R300#show ip bgp
BGP table version is 4, local router ID is 10.0.53.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.31.2                0             0 100 i
 *>  203.0.113.0      10.0.34.2                0             0 400 i
 *                    10.0.53.1                              0 500 600 400 i
R300#show ip bgp 203.0.113.0
BGP routing table entry for 203.0.113.0/24, version 3
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 2
  400
    10.0.34.2 from 10.0.34.2 (203.0.113.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 2
  500 600 400
    10.0.53.1 from 10.0.53.1 (10.0.56.2)
      Origin IGP, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
R300#
```

Fig 8.1: show ip bgp summary on AS300 During Normal Operation

```
R100#show ip bgp
BGP table version is 3, local router ID is 198.51.100.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     0.0.0.0                  0         32768 i
 *>  203.0.113.0      10.0.31.1                              0 300 400 i
R100#show ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     203.0.113.0/24 [20/0] via 10.0.31.1, 00:12:12
R100#
```

Fig 8.2: show ip bgp summary on AS100 During Normal Operation

```
R500#show ip bgp
BGP table version is 3, local router ID is 10.0.56.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.53.2                              0 300 100 i
 *   203.0.113.0      10.0.53.2                              0 300 400 i
 *>                   10.0.56.1                              0 600 400 i
R500#
```

Fig 8.3: show ip bgp summary on AS500 During Normal Operation

8.1 Attack Scenarios

8.1.1 Attack 1: Basic Prefix Hijack (/24)

The attacker (AS200) originates the victim's prefix 203.0.113.0/24 using BGP. A static route to Null0 is added so IOS can advertise the prefix via the network statement.

```
R300#show ip bgp 203.0.113.0
BGP routing table entry for 203.0.113.0/24, version 3
Paths: (3 available, best #2, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  200
    10.0.32.2 from 10.0.32.2 (10.0.32.2)
      Origin IGP, metric 0, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 2
  400
    10.0.34.2 from 10.0.34.2 (203.0.113.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 2
  500 600 400
    10.0.53.1 from 10.0.53.1 (10.0.56.2)
      Origin IGP, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
R300#
```

Fig 8.4: Evidence of Basic Prefix Hijack in BGP Table in AS300

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

```
R100#show ip bgp 203.0.113.0
BGP routing table entry for 203.0.113.0/24, version 3
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  300 400
    10.0.31.1 from 10.0.31.1 (10.0.53.2)
      Origin IGP, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
R100#show ip route 203.0.113.0
Routing entry for 203.0.113.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 300, type external
  Last update from 10.0.31.1 00:26:29 ago
  Routing Descriptor Blocks:
  * 10.0.31.1, from 10.0.31.1, 00:26:29 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 300
      MPLS label: none
R100#
```

Fig 8.5: Evidence of Basic Prefix Hijack in BGP Table in AS100

```
R200#show ip bgp
BGP table version is 5, local router ID is 10.0.32.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.32.1                              0 300 100 i
 *>  203.0.113.0      0.0.0.0                  0         32768 i
 *                    10.0.32.1                              0 300 400 i
R200#
```

Fig 8.6: Evidence of Basic Prefix Hijack in BGP Table in AS200

```
PC1> ping 203.0.113.10
84 bytes from 203.0.113.10 icmp_seq=1 ttl=61 time=90.559 ms
84 bytes from 203.0.113.10 icmp_seq=2 ttl=61 time=91.281 ms
84 bytes from 203.0.113.10 icmp_seq=3 ttl=61 time=111.314 ms
84 bytes from 203.0.113.10 icmp_seq=4 ttl=61 time=103.447 ms
84 bytes from 203.0.113.10 icmp_seq=5 ttl=61 time=83.068 ms
```

Fig 8.7: Evidence of Basic Prefix Hijack in BGP Table in PC1

8.1.2 Attack 2: More-Specific Prefix Hijack (/25)

The attacker announces 203.0.113.0/25. Due to longest-prefix-match, traffic destined to addresses within 203.0.113.0–203.0.113.127 is forwarded using the /25 route if present, making this attack highly effective.

```
R300#show ip bgp
BGP table version is 5, local router ID is 10.0.53.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.31.2                0             0 100 i
 *>  203.0.113.0/25   10.0.32.2                0             0 200 i
 *   203.0.113.0      10.0.32.2                0             0 200 i
 *>                   10.0.34.2                0             0 400 i
 *                    10.0.53.1                              0 500 600 400 i
R300#
```

Fig 8.8: More-Specific Hijack and Forwarding Impact-AS300

```
R400#show ip bgp
BGP table version is 4, local router ID is 203.0.113.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *   198.51.100.0     10.0.46.1                              0 600 500 300 100 i
 *>                   10.0.34.1                              0 300 100 i
 *   203.0.113.0/25   10.0.46.1                              0 600 500 300 200 i
 *>                   10.0.34.1                              0 300 200 i
 *>  203.0.113.0      0.0.0.0                  0         32768 i
R400#
```

Fig 8.9: More-Specific Hijack and Forwarding Impact-AS400

```
R200#show ip bgp
BGP table version is 6, local router ID is 10.0.32.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.32.1                              0 300 100 i
 *>  203.0.113.0/25   0.0.0.0                  0         32768 i
 *>  203.0.113.0      0.0.0.0                  0         32768 i
 *                    10.0.32.1                              0 300 400 i
R200#
```

Fig 8.10: More-Specific Hijack and Forwarding Impact-AS200

```
R100#show ip route 203.0.113.10
Routing entry for 203.0.113.0/25
  Known via "bgp 100", distance 20, metric 0
  Tag 300, type external
  Last update from 10.0.31.1 00:02:32 ago
  Routing Descriptor Blocks:
  * 10.0.31.1, from 10.0.31.1, 00:02:32 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 300
      MPLS label: none
R100#show ip route 203.0.113.200
Routing entry for 203.0.113.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 300, type external
  Last update from 10.0.31.1 00:37:41 ago
  Routing Descriptor Blocks:
  * 10.0.31.1, from 10.0.31.1, 00:37:41 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 300
      MPLS label: none
R100#
```

Fig 8.11: More-Specific Hijack and Forwarding Impact Observed at AS100

```
PC1> ping 203.0.113.10
*10.0.32.2 icmp_seq=1 ttl=253 time=62.097 ms (ICMP type:3, code:1, Destination host unreachable)
*10.0.32.2 icmp_seq=2 ttl=253 time=77.154 ms (ICMP type:3, code:1, Destination host unreachable)
*10.0.32.2 icmp_seq=3 ttl=253 time=76.237 ms (ICMP type:3, code:1, Destination host unreachable)
*10.0.32.2 icmp_seq=4 ttl=253 time=75.961 ms (ICMP type:3, code:1, Destination host unreachable)
*10.0.32.2 icmp_seq=5 ttl=253 time=75.741 ms (ICMP type:3, code:1, Destination host unreachable)

PC1>
```

Fig 8.12: More-Specific Hijack and Forwarding Impact-PC1

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

9. Experimental and Theoretical Results

9.1 Sequence Chart (Control-Plane Events)

Normal sequence (simplified):

1. eBGP sessions established over TCP.
2. AS100 announces 198.51.100.0/24 to AS300.
3. AS400 announces 203.0.113.0/24 to AS300 and AS600.
4. Routes propagate via AS500 and across the topology.
5. PC1 reaches PC4 with successful ping/traceroute.

9.2 Outcome Summary Table

Table 9.1: Routing Behavior Under Different Scenarios

Scenario	Selected/Installed Route (Typical)	Outcome
Normal Operation	Legit route from AS400 present globally	Legitimate routing; PC1→PC4 works
Basic /24 Hijack	AS200 route may become best in some ASes	Partial diversion/blackholing possible
More-Specific/25 Hijack	AS200/25 installed for covered IPs	Strong diversion for 203.0.113.0–127
After Filtering	Only authorized routes accepted	Correct routing restored; hijacks blocked
After Traffic Engineering	Policy-driven preference enforced	Path preferences align with business goals

9.2 Discussion

9.2.1 Why /25 Hijack Works Reliably

Forwarding uses longest-prefix-match. Even if the victim's /24 remains installed, the /25 is chosen for destinations it covers. This is a key reason more-specific hijacks are widely used.

9.2.2 Why Filtering is Effective

Provider edges act as trust boundaries. If AS300 and AS600 only accept authorized prefixes, malicious announcements are rejected early and cannot propagate widely.

10. Mitigation Techniques

10.1 Layer 1: Prefix Filtering

Prefix filtering permits only expected prefixes from a neighbor. This is the strongest operational control in many ISP environments.

Table 10.1: Prefix Filtering Policy Summary

Device	Neighbor	Policy
R300	AS400 (10.0.34.2)	Accept only 203.0.113.0/24
R300	AS100 (10.0.31.2)	Accept only 198.51.100.0/24
R300	AS200 (10.0.32.2)	Deny all (lab attacker)
R600	AS400 (10.0.46.2)	Accept only 203.0.113.0/24

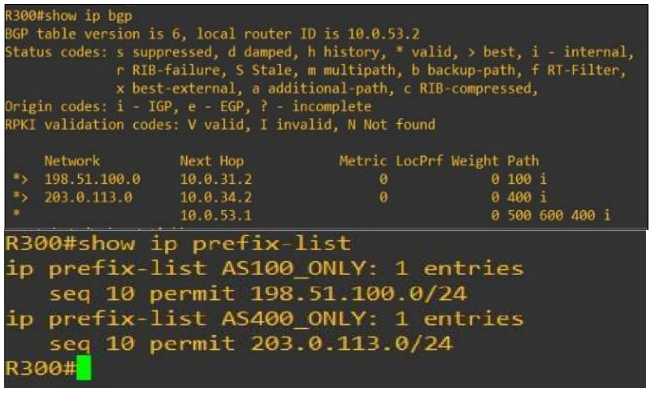

Fig 10.1: Prefix Filtering on Routers

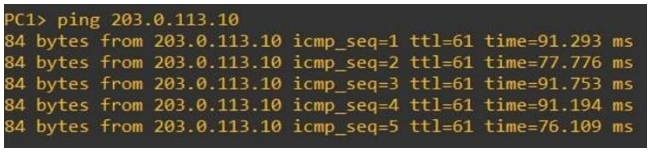

Fig 10.2: Output After Prefix Filtering on PC1

10.2 Layer 2: AS-PATH Filtering

AS-PATH filtering enforces expected origin ASes (lab policy). Example: accept 203.0.113.0/24 only if it originates in AS400.

10.3 Layer 3: Max-Prefix Protection

Max-prefix prevents route floods and limits session abuse by capping how many prefixes a neighbor can send.

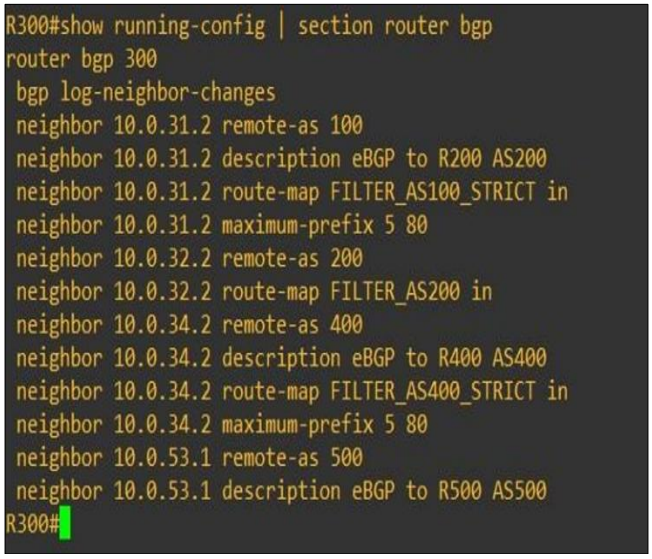

Fig 10.3: Routers After Mitigation Techniques

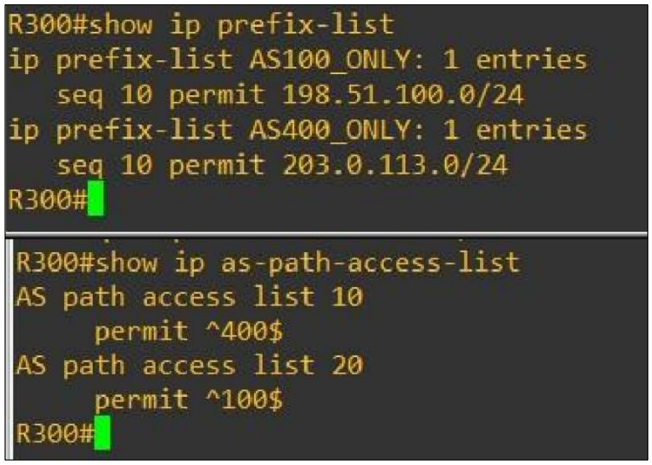

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

```
R300#show ip bgp summary
BGP router identifier 10.0.53.2, local AS number 300
BGP table version is 6, main routing table version 6
2 network entries using 288 bytes of memory
3 path entries using 240 bytes of memory
3/2 BGP path/bestpath attribute entries using 408 bytes of memory
3 BGP AS-PATH entries using 72 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
BGP using 1008 total bytes of memory
BGP activity 3/1 prefixes, 5/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.31.2       4          100      93      95       6    0    0 01:18:08        1
10.0.32.2       4          200      95      93       6    0    0 01:16:13        0
10.0.34.2       4          400     109     102       6    0    0 01:24:32        1
10.0.53.1       4          500     109     103       6    0    0 01:28:26        1
R300#
```

```
R300#show ip bgp
BGP table version is 6, local router ID is 10.0.53.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  198.51.100.0     10.0.31.2                0             0 100 i
 *>  203.0.113.0      10.0.34.2                0             0 400 i
 *                    10.0.53.1                               0 500 600 400 i
R300#
```

Fig 10.4: Routers After Mitigation Techniques

```
PC1> ping 203.0.113.10
84 bytes from 203.0.113.10 icmp_seq=1 ttl=61 time=92.002 ms
84 bytes from 203.0.113.10 icmp_seq=2 ttl=61 time=90.620 ms
84 bytes from 203.0.113.10 icmp_seq=3 ttl=61 time=93.250 ms
84 bytes from 203.0.113.10 icmp_seq=4 ttl=61 time=91.893 ms
84 bytes from 203.0.113.10 icmp_seq=5 ttl=61 time=76.984 ms
```

Fig 10.5: PC1 After Mitigation Techniques Traffic Engineering

11.1 Local Preference (Outbound Control)

Local Preference is used to influence outbound traffic selection. In this project, AS400 prefers AS600 as primary by setting higher local preference on routes learned from AS600.

```
R400#show ip bgp 198.51.100.0
BGP routing table entry for 198.51.100.0/24, version 8
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1           2
  Refresh Epoch 2
  600 500 300 100
    10.0.46.1 from 10.0.46.1 (10.0.56.1)
      Origin IGP, localpref 200, valid, external, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 3
  300 100
    10.0.34.1 from 10.0.34.1 (10.0.53.2)
      Origin IGP, localpref 50, valid, external
      rx pathid: 0, tx pathid: 0
R400#
```

```
R400#show ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     198.51.100.0/24 [20/0] via 10.0.46.1, 00:15:27
R400#
```

Fig 11.1: Local Preference

11.2 AS-PATH Prepending (Inbound Influence)

AS400 prepends its AS multiple times when advertising to AS300 to make the AS300 path less attractive for inbound selection.

```
R300#show ip bgp 203.0.113.0
BGP routing table entry for 203.0.113.0/24, version 7
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 3
  500 600 400
    10.0.53.1 from 10.0.53.1 (10.0.56.2)
      Origin IGP, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
R300#
```

Fig 11.2: AS-PATH Prepending

11.3 BGP Decision Process (Short Reference)

Common decision order:

1. Weight (higher wins
2. Local Preference (higher wins)
3. Locally originated rout
4. AS-PATH length (shorter wins)
5. Origin type (IGP > EGP > Incomplete)
6. MED (lower wins)
7. eBGP over Ibg
8. Lowest IGP metric to next ho
9. Oldest path
10. Lowest router ID / lowest neighbor IP

12. CONCLUSION

This project demonstrated the vulnerability of BGP to prefix hijacking and the effectiveness of policy-based mitigations in a realistic multi-AS mini-Internet. A basic /24 hijack can partially succeed depending on best-path selection, while a more-specific /25 hijack is highly effective due to longest-prefix-match forwarding. Implementing defense-in-depth (prefix filtering, AS-PATH filtering, max-prefix) blocked malicious routes and restored correct routing. Traffic engineering using Local Preference and AS-PATH prepending illustrated how routing decisions can reflect business policy rather than shortest paths.

13. FUTURE WORK

Future extensions:

- Implement RPKI/ROV conceptually and compare with policy filtering.
- Add route leak prevention policies (customer-to-provider export rules).
- Add monitoring/alerting (syslog-based detection of unexpected origin changes).
- Expand topology to include more ASes and simulate route leaks vs hijacks.
- Add BGP communities for advanced policy tagging and blackholing.

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

14. REFERENCES

1. Rekhter Y, Li T, Hares S. A border gateway protocol 4 (BGP-4). RFC 4271. Internet Engineering Task Force; 2006.
2. Huston G, Michaelson G, Loomans R. BGP prefix origin validation. RFC 6811. Internet Engineering Task Force; 2013.
3. Cisco Systems. BGP configuration guide. Cisco IOS Documentation. Topics: prefix-lists, route-maps, AS-PATH access-lists, maximum-prefix limits.
4. Mutually Agreed Norms for Routing Security (MANRS). Internet Society. Available from: https://www.manrs.org
5. RIPE Network Coordination Centre. Routing security and BGP hijacking case studies. RIPE NCC.

Appendix A
Addressing and Interface Assignment

Table A.1: Interfaces and IP Assignments

Device	Interface	IP/Mask	Peer/Use
R400	e1/0	10.0.46.2/30	to R600 e1/0
R400	e1/1	10.0.34.2/30	to R300 e1/3
R400	f0/0	203.0.113.1/24	Victim LAN gateway
R600	e1/0	10.0.46.1/30	to R400 e1/0
R600	e1/1	10.0.56.1/30	to R500 e1/0
R500	e1/0	10.0.56.2/30	to R600 e1/1
R500	e1/1	10.0.53.1/30	to R300 e1/0
R300	e1/0	10.0.53.2/30	to R500 e1/1
R300	e1/1	10.0.31.1/30	to R100 e1/0
R300	e1/2	10.0.32.1/30	to R200 e1/0
R300	e1/3	10.0.34.1/30	to R400 e1/1
R100	e1/0	10.0.31.2/30	to R300 e1/1
R100	f0/0	198.51.100.1/24	Customer LAN gateway
R200	e1/0	10.0.32.2/30	to R300 e1/2
PC1	e0	198.51.100.10/24	GW 198.51.100.1
PC4	e0	203.0.113.10/24	GW 203.0.113.1

Appendix B
Verification Commands
Interface and IP Verification
show ip interface brief
show ip route connected
ping <neighbor-ip>

BGP Verification
show ip bgp summary
show ip bgp
show ip bgp 203.0.113.0
show ip bgp 203.0.113.0/25

Policy Verification
show ip prefix-list
show route-map
show ip as-path-access-list
show running-config | section router bgp
show ip bgp neighbors <neighbor-ip> routes

Traffic Tests
ping 203.0.113.10
traceroute 203.0.113.10
ping 198.51.100.10
traceroute 198.51.100.10

Appendix C
Full Router and VPCS Configurations (Divided by Phases)
Important Note on IOS BGP network Statements
For IOS, a network X mask Y statement requires that the exact prefix exists in the routing table. Legitimate LANs are connected routes. Hijacked prefixes on AS200 are created using static routes to Null0 so they can be originated.

Phase 3: Base Configuration (All Routers)
enable
configure terminal
no ip domain-lookup
! hostname RXXX (set per router)
end
write memory

Phase 2: IP Addressing
R100
enable
configure terminal
hostname R100
no ip domain-lookup
interface e1/0
description Link to R300
ip address 10.0.31.2 255.255.255.252
no shutdown
interface f0/0
description AS100 LAN
ip address 198.51.100.1 255.255.255.0
no shutdown
end
write memory

R200
enable
configure terminal
hostname R200
no ip domain-lookup
interface e1/0
description Link to R300
ip address 10.0.32.2 255.255.255.252
no shutdown
end
write memory

```
R300#show ip bgp 203.0.113.0
BGP routing table entry for 203.0.113.0/24, version 7
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 3
  500 600 400
    10.0.53.1 from 10.0.53.1 (10.0.56.2)
      Origin IGP, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
R300#
```


© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

R300
enable
configure terminal
hostname R300
no ip domain-lookup
interface e1/0
description Link to R500
ip address 10.0.53.2 255.255.255.252
no shutdown
interface e1/1
description Link to R100
ip address 10.0.31.1 255.255.255.252
no shutdown
interface e1/2
description Link to R200
ip address 10.0.32.1 255.255.255.252
no shutdown
interface e1/3
description Link to R400
ip address 10.0.34.1 255.255.255.252
no shutdown
end
write memory

R400
enable
configure terminal
hostname R400
no ip domain-lookup
interface e1/0
description Link to R600
ip address 10.0.46.2 255.255.255.252
no shutdown
interface e1/1
description Link to R300
ip address 10.0.34.2 255.255.255.252
no shutdown
interface f0/0
description AS400 LAN
ip address 203.0.113.1 255.255.255.0
no shutdown
end
write memory

R500
enable
configure terminal
hostname R500
no ip domain-lookup
interface e1/0
description Link to R600
ip address 10.0.56.2 255.255.255.252
no shutdown
interface e1/1
description Link to R300
ip address 10.0.53.1 255.255.255.252
no shutdown
end
write memory

R600
enable
configure terminal
hostname R600
no ip domain-lookup
interface e1/0
description Link to R400
ip address 10.0.46.1 255.255.255.252
no shutdown
interface e1/1
description Link to R500
ip address 10.0.56.1 255.255.255.252
no shutdown
end
write memory

VPCS
PC1 ip 198.51.100.10 255.255.255.0 198.51.100.1 save
PC4 ip 203.0.113.10 255.255.255.0 203.0.113.1 save

Phase 4: Legitimate BGP (Normal Operation)
R100 (AS100)
enable
configure terminal
router bgp 100
bgp log-neighbor-changes
network 198.51.100.0 mask 255.255.255.0
neighbor 10.0.31.1 remote-as 300
end
write memory

R200 (AS200) – No advertisements yet
enable
configure terminal
router bgp 200
bgp log-neighbor-changes
neighbor 10.0.32.1 remote-as 300
end
write memory

R300 (AS300)
enable
configure terminal
router bgp 300
bgp log-neighbor-changes
neighbor 10.0.53.1 remote-as 500
neighbor 10.0.31.2 remote-as 100
neighbor 10.0.32.2 remote-as 200
neighbor 10.0.34.2 remote-as 400
end
write memory

R400 (AS400) – Victim / Legit Owner
enable
configure terminal
router bgp 400

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

```
bgp log-neighbor-changes
network 203.0.113.0 mask 255.255.255.0
neighbor 10.0.46.1 remote-as 600
neighbor 10.0.34.1 remote-as 300
end
write memory
```

R500 (AS500)

```
enable
configure terminal
router bgp 500
bgp log-neighbor-changes
neighbor 10.0.56.1 remote-as 600
neighbor 10.0.53.2 remote-as 300
end
write memory
```

R600 (AS600)

```
enable
configure terminal
router bgp 600
bgp log-neighbor-changes
neighbor 10.0.46.2 remote-as 400
neighbor 10.0.56.2 remote-as 500
end
write memory
```

Phase 5: Attack Scenario 1 (Basic /24 Hijack) – On R200

```
enable
configure terminal
ip route 203.0.113.0 255.255.255.0 Null0
router bgp 200
network 203.0.113.0 mask 255.255.255.0
end
write memory
```

Phase 6: Attack Scenario 2 (More-Specific /25 Hijack) – On R200

```
enable
configure terminal
ip route 203.0.113.0 255.255.255.128 Null0
router bgp 200
network 203.0.113.0 mask 255.255.255.128
end
write memory
```

Phase 7: Mitigation Layer 1 (Prefix Filtering)
On R300

```
enable
configure terminal
ip prefix-list AS100_ONLY seq 10 permit 198.51.100.0/24
ip prefix-list AS400_ONLY seq 10 permit 203.0.113.0/24
route-map FROM_AS100 permit 10 match ip address prefix-list AS100_ONLY
route-map FROM_AS400 permit 10 match ip address prefix-list AS400_ONLY
route-map FROM_AS200 deny 10
router bgp 300
neighbor 10.0.31.2 route-map FROM_AS100 in
neighbor 10.0.34.2 route-map FROM_AS400 in
neighbor 10.0.32.1 route-map FROM_AS200 in
end
write memory
clear ip bgp 10.0.31.2 soft in
clear ip bgp 10.0.34.2 soft in
clear ip bgp 10.0.32.2 soft in
```

On R600 (extra protection from AS400)

```
enable
configure terminal
ip prefix-list AS400_LEGIT seq 10 permit 203.0.113.0/24
route-map FROM_AS400_UP permit 10 match ip address prefix-list AS400_LEGIT
router bgp 600
neighbor 10.0.46.2 route-map FROM_AS400_UP in
end
write memory
clear ip bgp 10.0.46.2 soft in
```

Phase 8: Mitigation Layer 2 (AS-PATH Filtering) – On R300

```
enable
configure terminal
ip as-path access-list 10 permit ^100$
ip as-path access-list 20 permit ^400$
route-map FROM_AS100_STRICT permit 10
match ip address prefix-list AS100_ONLY
match as-path 10
route-map FROM_AS400_STRICT permit 10
match ip address prefix-list AS400_ONLY
match as-path 20
router bgp 300
neighbor 10.0.31.2 route-map FROM_AS100_STRICT in
neighbor 10.0.34.2 route-map FROM_AS400_STRICT in
end
write memory
clear ip bgp * soft in
```

Phase 9: Mitigation Layer 3 (Max-Prefix Protection) – On R300

```
enable
configure terminal
router bgp 300
neighbor 10.0.32.2 maximum-prefix 1 warning-only
neighbor 10.0.31.2 maximum-prefix 5 80
neighbor 10.0.34.2 maximum-prefix 5 80
end
write memory
```

Phase 10: Traffic Engineering
Local Preference (Outbound Preference) – On R400

```
enable
configure terminal
route-map PREFER_AS600 permit 10 set local-preference 200
route-map BACKUP_AS300 permit 10 set local-preference 50
router bgp 400
```


© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/

```
neighbor 10.0.46.1 route-map PREFER_AS600 in
neighbor 10.0.34.1 route-map BACKUP_AS300 in
end
write memory
clear ip bgp * soft in
```

AS-PATH Prepending (Inbound Influence) – On R400

```
enable
configure terminal
route-map PREPEND_TO_AS300 permit 10 set as-path prepend 400 400 400
router bgp 400
neighbor 10.0.34.1 route-map PREPEND_TO_AS300 out
end
write memory
clear ip bgp 10.0.34.1 soft out
```

Creative Commons (CC) License

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution–Non-Commercial–No Derivatives 4.0 International (CC BY-NC-ND 4.0) license. This license permits sharing and redistribution of the article in any medium or format for non-commercial purposes only, provided that appropriate credit is given to the original author(s) and source. No modifications, adaptations, or derivative works are permitted under this license.

© 2026 Mohammed Shafiullah Raihan. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY NC ND).https://creativecommons.org/licenses/by/4.0/