Degree Project in Computer Science and Engineering

First cycle, 15 credits

Evaluating Simulation Accuracy for Suboptimal BGP Routing in IP Anycast

A Case Study of I-root

SARAH SALEH

Evaluating Simulation Accuracy for Suboptimal BGP Routing in IP Anycast

A Case Study of I-root

SARAH SALEH

Degree Programme in Computer Engineering
Date: June 24, 2026

Supervisors: Siva Satya Sri Ganesh Seeram, Fredrik Lindeberg, Tomas Agartz
Examiner: Peter Sjödin

Host company: Netnod AB
Swedish title: Utvärdering av simuleringsnoggrannhet för suboptimal BGP-routing i IP-Anycast
Swedish subtitle: En fallstudie av I-root

© 2026 Sarah Saleh

Abstract

The I-root, managed by Netnod, achieves robust global availability through the use of Internet Protocol (IP) Anycast. Inter-domain routing, which is governed by routers employing the Border Gateway Protocol (BGP) protocol, is primarily determined by BGP attributes and administrative policies rather than by real-time performance metrics such as latency. Specifically, routing decisions are based on rules such as local preference that prioritizes certain exit points and AS-Path length that which selects the path with the fewest AS hops. Router-specific mechanisms such as longest prefix matching directs traffic to the most specific IP blocks which can also influence the routing. This can result situations with suboptimal routing, where traffic is directed to distant nodes despite the presence of closer alternatives. This thesis investigates whether publicly available data sources construct a comprehensive understanding of these routing scenarios and evaluates the accuracy of network simulations in replicating them. By utilizing RIPE Atlas measurements, Autonomous System (AS)-specific Looking Glasses (LGs), and other publicly available sources, three real-world cases were constructed based on instances of suboptimal routing in regions with high availability. These cases were subsequently reconstructed in a controlled simulation environment using Netlab and FRRouting. The findings shows that public data are generally sufficient for identifying and modeling certain routing behaviors. The simulations closely reflected real-world routing behavior, with only minor discrepancies related to undisclosed internal BGP configurations. In summary, this research demonstrates that network emulators are effective tools for operators when aiming to analyze and optimize internet infrastructures without risking live network stability.

Keywords

Routing, IP Anycast, BGP Path Selection, Anycast Catchment, Routing Instability, Latency Inflation, Peering Strategies, Internet Exchange Points, BGP Tie-breaking, BGP Attributes

Sammanfattning

I-root, som drivs av Netnod, uppnår robust global tillgänglighet genom användning av IP Anycast. Routing mellan domäner, som styrs av routrar som tillämpar BGP, avgörs främst av BGP-attribut och administrativa policyer snarare än av prestandamått i realtid såsom latens. Specifikt baseras routningsbeslut på regler såsom "local preference" som sätter prioritering för vissa exitpunkter och AS-Path-längd, som väljer den väg med minst antal AS-hopp. Routerspecifika mekanismer såsom längsta prefixmatchning styr trafiken till de mest specifika IP-blocken vilket också kan påverka routingen. Detta kan leda till situationer med suboptimal routing, där trafik dirigeras till avlägsna noder trots att närmare alternativ finns. Detta examensarbete undersöker om publikt tillgängliga datakällor möjliggör en fullständig förståelse av dessa routningsscenarier och utvärderar hur väl nätverkssimuleringar kan replikera dem. Genom att använda RIPE Atlas-mätningar, AS-specifika LGs och andra publika källor konstruerades tre verkliga fall baserade på tillfällen då suboptimal routing sker i områden med hög tillgänglighet. Dessa fall återskapades sedan i en kontrollerad simuleringsmiljö med hjälp av Netlab och FRRouting. Resultaten visar att publika data är relativt tillräckligt för att identifiera och modellera ett visst routningsbeteende. Simuleringarna återspeglade verkligt routningsbeteende väl, med endast mindre avvikelser kopplade till icke-publika interna BGP-konfigurationer. Sammanfattningsvis visar denna forskning att nätverksemulatorer är effektiva för operatörer som vill analysera och optimera internet infrastrukturer utan att riskera nätverk stabiliteten.

Nyckelord

Routing, IP-anycast, BGP-ruttval, Anycast-upptagningsområde, Routinginstabilitet, Latensinflation, Peeringstrategier, Internetknutpunkter, BGP-tie-breaking, BGP-attribut

Acknowledgments

This thesis would not have been completed without the invaluable contributions of many individuals. First and foremost, I would like to express my deepest gratitude to my supervisors at Netnod, Fredrik Lindeberg and Tomas Agartz, for their expert guidance and insightful feedback. Their expertise and knowledge gave me a deep understanding of the topic.

I am equally thankful to my supervisor at KTH, Siva Satya Sri Ganesh Seeram, for his academic mentorship and constructive advice. The support I received gave me clarity and direction to complete this thesis. I would also like to thank my examiner at KTH, Peter Sjödin, for his time and valuable feedback.

I also want to express my appreciation to the engineering and research teams at Netnod for their support and engaging discussions. A special mention goes to Arvid Zhang for going well beyond what was expected to offer his time and assistance throughout this process.

Last but certainly not least, I want to thank my family for their endless encouragement and unwavering support throughout my academic journey. I don't know where I would be without them.

Stockholm, June 2026
Sarah Saleh

Contents

1 Introduction **1**
1.1 Background . . . 2
1.2 Problem . . . 3
1.2.1 Original problem and definition . . . 4
1.2.2 Scientific and engineering issues . . . 4
1.3 Purpose . . . 5
1.4 Goals . . . 5
1.4.1 Deliverables . . . 6
1.5 Research Methodology . . . 6
1.6 Delimitations . . . 7
1.7 Structure of the thesis . . . 7

2 Background **9**
2.1 Domain Name System . . . 9
2.1.1 The root . . . 10
2.1.2 Netnod's role . . . 11
2.2 Border Gateway Protocol . . . 12
2.2.1 External vs. Interior Gateway Protocols . . . 12
2.2.2 Policy-Based Routing and Path Attributes . . . 12
2.2.3 BGP Best Path Selection Process . . . 13
2.2.3.1 BGP Communities . . . 13
2.2.4 Key Traffic Engineering Attributes . . . 13
2.2.4.1 Local Preference . . . 14
2.2.4.2 AS_PATH Prepending . . . 14
2.2.4.3 Multi-Exit Discriminator (MED) . . . 14
2.2.5 Autonomous systems . . . 14
2.2.5.1 Route reflectors and Cluster lists . . . 15
2.2.5.2 Internet Service Provider . . . 15
2.2.5.3 Internet Exchange Point . . . 15

2.2.5.4 Routing Information Base and Forwarding Information Base 16
2.3 IP Anycast and Catchments 16
2.4 Previous Research 16
2.4.1 BGP and Anycast Performance 16
2.4.1.1 The Impact of BGP Tie-breaking 17
2.5 Summary 18

3 Methods **19**
3.1 Research strategies 19
3.1.1 Cases 19
3.1.2 Simulated testing 19
3.2 Research Process 20
3.3 Research Paradigm 20
3.4 Data Collection 21
3.4.1 Sampling 22
3.4.2 Sample Size 22
3.4.3 Target Population 22
3.5 Experimental design/Planned Measurements 22
3.5.1 Test environment and Model 22
3.5.2 Software to be used 22
3.6 Assessing reliability and validity of the data collected 23
3.6.1 Validity of method 23
3.6.2 Reliability of method 24
3.6.3 Data validity 24
3.6.4 Reliability of data 24
3.7 Planned Data Analysis 25
3.7.1 Data Analysis Technique 25
3.7.2 Software Tools 25
3.8 Evaluation Framework 25
3.9 System documentation 26

4 From Collected Data to Controlled Simulation **27**
4.1 Hypothesis Formulation 27
4.2 Simulation Baseline Validation 27
4.3 Case Study 29
4.4 Simulation 29
4.4.1 Case 1 31
4.4.2 Case 2 33

4.4.3 Case 3 35

5 Results and Analysis **39**
5.1 Case 1, RIPE Atlas Probe 1007504 39
5.1.1 Data 39
5.1.2 Simulation results 43
5.2 Case 2, RIPE Atlas Probe 50288 45
5.2.1 Data 45
5.2.2 Simulation results 47
5.3 Case 3, RIPE Atlas Probe 14521 49
5.3.1 Data 49
5.3.2 Simulation 51
5.4 Reliability Analysis 54
5.5 Validity Analysis 54

6 Discussion **57**

7 Conclusions and Future work **59**
7.1 Conclusions 59
7.2 Limitations 60
7.3 Future work 60
7.3.1 What has been left undone? 60
7.3.1.1 Cost analysis 61
7.3.1.2 Security 61
7.3.2 Next obvious things to be done 61
7.4 Reflections 61

References **63**

A BGP LG **69**
A.1 Case 1 Orange Looking Glass 70
A.2 Case 3 DTAG Looking Glass 71

B YAML-Code **72**
B.1 Test-Case: Local Preference 72
B.2 Case 1 74
B.3 Case 2 76
B.4 Case 3 78

List of Figures

1.1 Illustration of a suboptimal routing scenario in which a higher latency path is selected despite the availability of closer alternatives. 3

2.1 The DNS hierarchy is structured with the root at the apex, and subdomains distributed beneath it. 9
2.2 I-root instances in the world. Map generated using Kepler.gl [17]. 11

4.1 Test topology. Generated via Netlab 28
4.2 Probe 1007504 geographical location. Source: RIPE Atlas . . 31
4.3 Netlab topology for Case 1. Generated via Netlab 32
4.4 Probe 50288 geographical location. Source: RIPE Atlas 33
4.5 Netlab topology for Case 2, rr = Route Reflector. Generated via Netlab 34
4.6 Probe 14521 geographical location. Source: RIPE Atlas 35
4.7 Netlab topology for Case 3. Generated via Netlab 36

5.1 Traceroute RIPE: Probe 1007504 Date: Mon May 11 11:15:11 CEST 2026 39
5.2 Visually presents the three signals RTT trend for each hop. . . 40
5.3 Traceroute RIPE: Probe 50288 Wed May 20 20:32:06 CEST 2026 45
5.4 Visually presents the three signals RTT trend for each hop. . . 45
5.5 Traceroute RIPE: Probe 14521 Tue May 12 10:13:49 CEST 2026 49
5.6 Visually presents the three signals RTT trend for each hop. . . 49

List of Tables

2.1 Steps in the BGP decision process, adapted from Caesar and Rexford [24]. 13

5.1 Case 1 data compiled from multiple sources 41
5.2 Case 2 data compiled from multiple sources 46
5.3 Case 3 data compiled from multiple sources 50

A.1 Orange BGP LG routing result to Netnod I-root (192.36.148.0/24) 70
A.2 DTAG AS3320 to Netnod I-Root (192.36.148.0/24) 71
A.3 DTAG AS 3320 BGP Communities [57] 71

Listings

4.1 General local preferernce test 29
5.1 Simulation result without forced connection to IIJ 43
5.2 Netnods local preference was lowered to match the output of the LG 44
5.3 Cutting the peering connection 44
5.4 The result after a the IIJ link is removed 44
5.5 Shows the observer network with 4000 local preference with correct cluster list 47
5.6 Peer to the local preference path (originator) turned of to force the test path with lower local preference [55] 47
5.7 Result of the turned off path 48
5.8 This code presents testing the specific prefix options and then testing if 51
5.9 Presents the process of turning off the two direct paths ”Main path /24” and ”Test Path /23” forcing traffic to the longer path with prefix /23 [55] 52
5.10 Output after /24 path was turned off 53
5.11 Shows the process of turning off the direct path with /24 prefix forcing traffic based on AS-path length given the two /23 options [55] 53
5.12 Output after /24 path was turned off and it chooses based on AS-path length. 54
B.1 Test-case local preference 73
B.2 Case 1 topology 75
B.3 Case 2 topology 77
B.4 Case 3 topology 79

List of acronyms and abbreviations

AS	Autonomous System
ASN	Autonomous System Number
BGP	Border Gateway Protocol
CLI	Command Line Interface
DNS	Domain Name System
DTAG	Deutsche Telekom AG
EGP	Exterior Gateway Protocol
FIB	Forwarding Information Base
FQDN	Fully Qualified Domain Name
FRR	Free Range Routing
IGP	Interior Gateway Protocol
IIJ	Internet Initiative Japan
IP	Internet Protocol
IS-IS	Intermediate System to Intermediate System
ISP	Internet Service Provider
IXP	Internet Exchange Point
LG	Looking Glass
MED	Multi-Exit Discriminator
OSPF	Open Shortest Path First
PII	Identifiable Information
RIB	Routing Information Base
RTT	Round-Trip Time
SDG	Sustainable Development Goal

TCP	Transmission Control Protocol
TLD	Top-Level Domain
UN	United Nations

Chapter 1

Introduction

The root infrastructures, responsible for initiating the DNS resolution process [1], are masterly crafted but also face many challenges in optimization, efficiency and robustness. Therefore, this thesis investigates whether collected data from scenarios involving suboptimal routing provide an accurate representation of network conditions, and whether simulations can effectively replicate these real-world scenarios.

The I-root is one of 13 root servers labeled A–M and is managed by Netnod. Therefore, it is crucial that its efficient, robust, and low-latency, given the high level of reliance on it worldwide [2].

The Domain Name System (DNS) is used for translating domain names to Internet Protocol (IP) addresses, and vice versa [3]. It features a hierarchical structure in which each domain segment is managed by its respective authoritative name server [4]. At the apex of this hierarchy sits the root; in technical notation, it is defined by a period "." at the end of a Fully Qualified Domain Name (FQDN) [3]. The root is responsible for initiating the traversal process by referring queries to the appropriate Top-Level Domain (TLD) name servers [1].

Over time, the demand for a faster and more reliable internet has increased as daily life has become increasingly dependent on it for financial, legal, and security purposes. As critical services and transactions are now primarily accessed via mobile devices and global networks, continuous availability is essential regardless of a user's geographic location [5]. To address these requirements for high availability and performance, IP Anycast was implemented.

Anycast is a routing method in which the same IP address (prefix) is advertised from multiple geographic locations [6]. This ensures that a user can

access a service from the topologically nearest node, even if the service origin is in another country. By directing traffic to the closest available instance, Anycast significantly reduces latency and enhances robustness [6].

However, latency issues still occur; while many are trivial to resolve, some remain obscured within the routing architecture. These issues frequently arise when the router, using Border Gateway Protocol (BGP), must choose between multiple identical paths that are considered to be the best options and this scenario is known as a BGP tie-break [7, 8]. Since BGP is fundamentally a policy-oriented protocol rather than one based on real-time performance, path selection may be suboptimal, resulting in increased latency [9]. BGP tie-breaks can be influenced by a hierarchy of factors including Internet Service Provider (ISP) policies, the length of the Autonomous System (AS)-Path, the arrival order of the route advertisements and more [7, 10]. The ultimate goal of this thesis is to analyze the different layers of the BGP tie-break mechanism, predict its outcomes, and propose methods to mitigate suboptimal routing decisions.

1.1 Background

BGP is the primary protocol responsible for routing and communication between AS [7]. Routers and servers utilize BGP Best Path Selection Process to determine the most efficient route through the network [11]. The BGP Best Path Selection Process is a deterministic, multi-step process that evaluates various path attributes, to select a single exit point for traffic [11].

The BGP Best Path Selection Process is designed for robustness and efficiency. However, it lacks awareness of latency, thereby risking suboptimal performance for certain paths [12]. Using BGP, a path may be selected as the "shorter" path in terms of AS-Path length but significantly longer in terms of physical distance, leading to suboptimal routing outcomes [11, 13]. These are therefore relevant to build cases on.

Netnod manages the I-root, one of the 13 logical root server identities. Netnod's primary role in this context is to ensure a resilient, secure, and robust connection to the I-root infrastructure, thereby maintaining its availability to the global internet community [2].

IP Anycast is a routing method that allows the same IP address prefix to be advertised simultaneously from multiple geographic locations [14]. Anycast enables a client to access a service via the topologically nearest node. By duplicating service IPs across various locations, Anycast provides inherent redundancy, which in turn reduces global latency, if routed correctly [6, 14]

Internet Exchange Points (IXPs) enable direct exchange of routing information, even among multi-tiered ISPs that wish to peer to provide transit to their customers [15], thereby enhancing the internet exchange process.

1.2 Problem

BGP and IP Anycast are fundamental to global internet routing, and inter-domain path selection is governed by administrative policies rather than technical performance metrics such as latency. Using BGP, routers determine the "best" path by adhering to a serial hierarchy of attributes, where local administrative preferences and routing policies primarily drive the decision-making process.

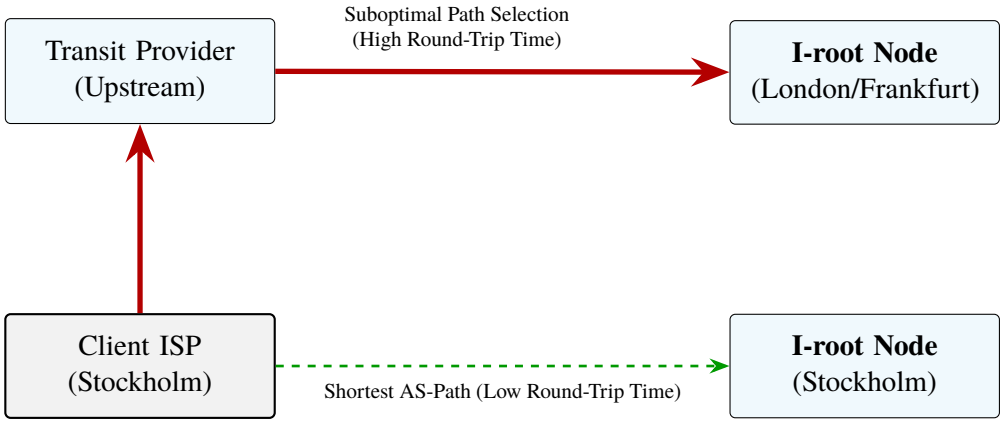

Figure 1.1: Illustration of a suboptimal routing scenario in which a higher latency path is selected despite the availability of closer alternatives.

In a global Anycast environment, such as the I-root server system, policy-driven selection can lead to suboptimal outcomes, including latency inflation, in which traffic is routed to a distant node even when a physically closer alternative is available (see Figure 1.1).

These outcomes often result from the different configurations of specific BGP attributes across ASes to manage their traffic. Unlike intra-domain protocols with full network visibility, BGP's decentralized nature means that a single policy change in one AS can have significant, and sometimes unpredictable, effects on Anycast catchments elsewhere. Therefore, when providers configure their routing policies, their perspective is limited to their own network, which can lead to consequences outside it.

Therefore for root infrastructures, analyzing these situations can appear fragmented, where information is scattered across multiple platforms. This consideration prompts the question of whether fragmented information can be integrated into a structured and comprehensive model to facilitate the troubleshooting process.

This raises two primary research questions:

1. To what extent does publicly available data provide a comprehensive understanding of network dynamics?
2. Are simulations relevant at accurately reflecting networks that include suboptimal routing?

1.2.1 Original problem and definition

Netnod has observed that even in unclear network failures, traffic still periodically shifts between different Anycast nodes, a phenomenon known as catchment shifts. While BGP is designed to ensure reachability, its reliance on a series of attributes often controls path selection based on administrative policies rather than real-time performance metrics. These shifts frequently occur when upstream ASes modify their routing policies, inadvertently forcing traffic toward suboptimal or distant Anycast nodes.

The objective of this thesis is to develop a deeper understanding of network structure by simulating scenarios involving suboptimal routing and evaluating their accuracy and relevance for future research.

1.2.2 Scientific and engineering issues

The scientific challenge of this research lies in modeling the complex relationship between decentralized BGP policy configurations and the resulting performance of a global Anycast service. From an engineering standpoint, a major issue is the lack of a direct correlation between control-plane signaling (BGP updates) and data-plane performance (latency). Because BGP routers derive their information exclusively from its immediate neighbors [9, 10, 13], it is difficult for a service provider to predict how a policy change in a neighboring AS will affect its traffic distribution.

The following research questions address these challenges:

- Can collected data from different public platforms provide a comprehensive image of the internet structure?

- To what extent can network emulation tools like Netlab be used to simulate real-world scenarios, and are simulations accurate in replicating them?

Addressing these questions, this research contributes to the broader engineering sector by developing a methodology to analyze networks and suboptimal routing.

1.3 Purpose

The purpose of this research is to bridge the gap between theoretical modeling and real-world I-root server infrastructure. The purpose is to make network performance research more accessible to future researchers, thereby leading to a more robust internet.

The purpose is to provide simulations of various network scenarios to improve traffic oversight and to understand the effects of BGP configurations on the Internet. Academically and engineering-wise, the aim is to add knowledge on how routing attributes affect Anycast, therefore supporting network operators' work.

This research also supports sustainability and social goals. More efficient and robust routing reduces unnecessary path inflation. This optimizes resource usage and maintains the internet's reliability as a vital societal service.

1.4 Goals

The primary goal of this thesis project is to properly emulate different networks with occuring suboptimal routing. This is divided into three sub-goals:

1. **Identification:** Identify and log instances of network traffic exhibiting unexpectedly high latency.

2. **Analysis and Modeling:** To construct an accurate simulation based on collected data using public sources and develop them in Netlab. This demonstrates this specific network topology and how it operates depending on which BGP attributes its configured on.

3. **Final analysis:** Analyze the network presentation accuracy of the collected data combined with Netlab simulations.

1.4.1 Deliverables

Throughout the duration of the project, the following deliverables will be produced:

- **Real-World Case Studies:** Real-world data collected from RIPE Atlas, AS-specific Looking Glasses (LGs), and other publicly available sources, compiled to document instances of suboptimal routing performance.
- **Simulation Topologies:** Pre-configured Netlab environments designed to emulate the network structures and BGP configurations of the identified real-world cases.
- **Comparative Analysis:** An evaluation of whether publicly available network data provides a comprehensive enough understanding to accurately model observed routing behaviors, and whether the resulting simulations accurately reflect these real-world scenarios.

1.5 Research Methodology

The methodology for this thesis follows a quantitative and experimental research strategy, structured as a case study focused on the I-root anycast instance. The approach is designed to bridge the gap between collected data and theoretical modeling of BGP behavior.

Philosophical Assumption: The research is grounded in pragmatism. This perspective is chosen because the primary goal is to identify functional solutions to a specific engineering problem, optimizing Anycast catchments using the most effective tools available, whether purely theoretical or purely empirical.

Research Method and Approach: This study adopts an abductive approach. It begins with an inductive phase, where collected data from RIPE Atlas, AS-specific LGs and other public sources are gathered and analyzed. This is to identify patterns of catchment shifts and latency inflation. These observations are used to construct a layout of the network flow applying the BGP attributes it has configured. This is followed by a deductive phase, in which these cases are tested through controlled experiments in a simulated Netlab environment to confirm or refute the accuracy of this simulated environment and whether BGP affects the path there as well.

Choice of Methodology and Alternatives: This approach was selected to integrate information from multiple sources and to establish a comprehensive simulated environment, rather than relying solely on numerical analysis or mathematical modeling. Mathematical models are limited in their ability to represent realistic BGP outputs, which emulators such as Netlab can effectively capture. Additionally, emulators are capable of representing non-deterministic BGP behavior, enabling researchers to study these phenomena in isolation and to configure network topologies with greater flexibility.

1.6 Delimitations

To maintain a focused and feasible scope for this research, the following delimitations have been established:

- The research will not account for server-side factors such as DNS processing time and internal router switching fabric performance.
- Data collection focuses primarily on European IXPs, as they provide a dense, structured environment for studying complex peering policies.
- The primary objective is to collect data and create simulations. If time permits, additional research on BGP behavior may be undertaken.

1.7 Structure of the thesis

The thesis is structured as follows:

- **Chapter 1: Introduction** – Provides the background, problem statement, goals, and delimitations of the project.
- **Chapter 2: Background** – Explains the fundamental concepts of BGP, IP Anycast, and the DNS root server system, as well as related work in the field.
- **Chapter 3: Methodology** – Describes the research strategy and paradigm, the tools used, and how data was collected and analyzed.
- **Chapter 4: Implementation** – Details the practical setup of the study, specifically how the collected data was used to construct and configure the isolated Netlab environments.

- **Chapter 5: Results and Analysis** – Presents the compiled data and the outcomes of the corresponding simulations.
- **Chapter 6: Discussion** – Interprets the findings, evaluates the underlying causes of the observed behaviors, and reflects on the accuracy of the simulations.
- **Chapter 7: Conclusions and Future Work** – Summarizes the findings and suggests areas for further research.

Chapter 2

Background

2.1 Domain Name System

DNS is a fundamental infrastructure component designed to translate human-readable domain names to Internet addresses [16] and is essential for modern internet functionality. It is organized as a hierarchy, often visualized as an inverted tree [16], and therefore serves as a robust and essential "phonebook" for global internet communication [16].

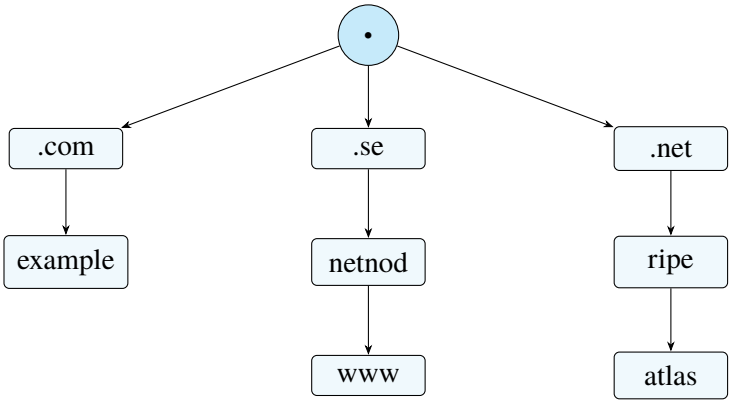

Figure 2.1: The DNS hierarchy is structured with the root at the apex, and subdomains distributed beneath it.

The DNS system consists of name servers distributed across various geographic locations, each being authoritative for specific domain levels [16]. When a client queries a domain, such as www.example.com, the query is typically processed recursively, descending the DNS hierarchy [16], as shown in Figure 2.1. The process starts at the root zone, directing the query to the

Top-Level Domain (TLD) servers (e.g., .com). The TLD, in turn, identifies the authoritative name servers for example.com, directing the client further down the hierarchy until the specific IP address is resolved [16].

2.1.1 The root

The root zone is at the apex of the DNS hierarchy and sits above the TLD such as .se, .com [1]. Root servers facilitate referrals to the appropriate top-level domains (TLDs) by using the root zone file, which includes name servers, IP addresses, and other pertinent information [1]. In its technical notation, the root is represented by a single dot (.) at the end of a Fully Qualified Domain Name (FQDN), indicating the absolute start of the hierarchy [1]. This notation is illustrated in the following example; note the trailing dot at the end of the domain name: “www.example.com.”.

There are 13 root identifiers, sorted in alphabetical order from A to M, each representing a vast network comprising hundreds of physical servers distributed globally via Anycast technology. This distribution ensures low latency, high redundancy, and high availability [1].

2.1.2 Netnod's role

Netnod operates the I-root server, ensuring a robust internet for the Nordics and Swedish organizations [5].

Figure 2.2: I-root instances in the world. Map generated using Kepler.gl [17].

As seen in Figure 2.2, the I-root instances have a large presence in Europe, but they are also widely distributed across the Americas, Asia, and Oceania, with the size of each dot on the map reflecting the number of instances in that specific city. Netnod is a Swedish-based company established in 1996 and is wholly owned by TU-stiftelsen (Stiftelsen för Telematikens utveckling) [18]. Netnod's approach to their profit model is self sustaining, reinvesting its earnings into the technical infrastructure, contributions, and expertise required to maintain a "rock-solid" standard of robustness for the internet [19].

For Swedish and Nordic organizations, it is important to maintain high control and low latency to the root zone while also contributing to global internet stability. Their contribution is essential to resilient internet infrastructure, underscoring the significance of their guiding principle: "for

the good of the internet" [5]. The infrastructure is designed such that the I-root (Autonomous System Number (ASN) 29216) is always advertised behind Netnod (ASN 8674) [20]. Consequently, the simulations in Chapter 4.4 are configured to consistently reflect this arrangement.

Netnod offered expert supervision, guided the research direction, and assisted with information verification. Additionally, they supplied a dedicated Linux server to serve as the laboratory environment for executing the demanding network simulations.

2.2 Border Gateway Protocol

BGP is a standardized path-vector protocol. It facilitates the exchange of network connectivity and reachability information between ASes on the Internet [7]. BGP uses Transmission Control Protocol (TCP) port 179 [7] and routers use BGP to update routing tables and select the best paths for data transmission [7].

2.2.1 External vs. Interior Gateway Protocols

BGP is categorized as an Exterior Gateway Protocol (EGP) [21], enabling communication between distinct ASes [7]. In contrast, Interior Gateway Protocols (IGPs), such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS), manage routing within a single administrative domain [21]. While IGPs typically maintain a comprehensive map of network topology using link-state information [22], BGP operates under the principle of "routing by rumor". As a result, it lacks a global view of the Internet's physical topology and instead relies solely on information advertised by its immediate neighbors [9, 10, 13]. Additionally, the system remains vulnerable to BGP configurations managed by independent administrative entities whose routing policies are not publicly disclosed and whose route announcements lack cryptographic verification, especially when originating from ASes without established trust relationships [23].

2.2.2 Policy-Based Routing and Path Attributes

The primary strength and complexity of BGP lie in its policy-oriented design. This allows network operators to influence route selection based on administrative and commercial objectives rather than solely on technical performance metrics [10, 11]. Instead of suggesting paths based on indicators

such as bandwidth or latency, the BGP decision process relies on Path Attributes to evaluate and filter routes [7]. These attributes are processed sequentially to select a single "best" path among the available candidates [7].

2.2.3 BGP Best Path Selection Process

The BGP Best Path Selection Process follows a sequential, step-by-step process, as outlined in Table 2.1 [24]. Vendors such as Cisco have modified this process to support custom configurations [11].

Table 2.1: Steps in the BGP decision process, adapted from Caesar and Rexford [24].

Step	Attribute	Controlled by local or neighbor AS?
1	Highest LocalPref	Local
2	Lowest AS path length	Neighbor
3	Lowest origin type	Neither
4	Lowest MED	Neighbor
5	eBGP-learned over iBGP-learned	Neither
6	Lowest IGP cost to border router	Local
7	Lowest router ID (to break ties)	Neither

2.2.3.1 BGP Communities

BGP communities are optional, transitive attributes [25] that allow network operators to group destinations and communicate administrative policies across ASes [26]. They function as numerical tags attached to route advertisements, allowing an originating AS to influence how neighboring or remote ASes treat a specific prefix [25]. In traffic engineering, communities are regularly used as signaling mechanisms to remotely trigger changes to Local Preference or to request specific AS-Path prepending behaviors in upstream networks [25]. In this study, communities are essential markers for inferring policy changes from observed BGP data.

2.2.4 Key Traffic Engineering Attributes

This thesis focuses on specific attributes that operators actively use to manage Anycast traffic. By manipulating these variables, an AS can influence the stability and latency of global catchments.

2.2.4.1 Local Preference

Local Preference, which represents the first step in the "Best Path Selection Process" (see Table 2.1), is the most powerful step for controlling outbound traffic within an AS. Because it is evaluated so early in the process, it can force traffic through a specific peer even if the AS_PATH is significantly longer [24, 27]. At this point, administrative factors, such as transit costs, frequently outweigh technical performance considerations [27].

2.2.4.2 AS_PATH Prepending

The AS_PATH attribute, which is the second step in the process (see Table 2.1), measures the number of ASes that traffic must traverse. Through prepending the practice of repeating one's own AS number multiple times in an advertisement, an operator can make a path appear artificially longer [24, 27]. This is used to discourage inbound traffic from using a specific link [24]. In the context of the I-root, operators can use prepending to "nudge" users from one Anycast node to another.

2.2.4.3 Multi-Exit Discriminator (MED)

The Multi-Exit Discriminator, located at the fourth step in the process (see Table 2.1), is an attribute sent between neighboring ASes to suggest which entry point should be preferred for inbound traffic [24, 27]. It is primarily used when multiple physical connections exist between two networks [27]. Since MED is evaluated after Local Preference and AS_PATH, it serves as a more subtle tool for fine-tuning traffic flows [27].

2.2.5 Autonomous systems

ASes are collections of IP addresses managed by network operators that share the same routing policy [28]. As mentioned previously, their communication, either within or outside the AS, is handled by different routing protocols. An important factor to note is that ASes are commonly categorized into a descriptive, multi-tier hierarchy that can be summarized as follows:

Tier 1: Tier 1 is positioned at the top of the hierarchy, often referred to as the "backbone" of the internet due to its significant global interconnectivity. It does not purchase any transit services and relies exclusively on peering [10].

Tier 2: Tier 2 typically connects the customer and internet using a combination of peering and transit [10].

Tier 3: Generally relies on purchased transit from tier 2 (sometimes tier 1), and almost no peering [10].

Although tier 1 to tier 3 are the most frequently referenced classifications, more granular hierarchies may be established with end customers generally occupying the lowest level [29].

ISPs connect using two primary forms: Transit and Peering. Transit is the method of buying internet service from an ISP, thereby gaining access to the entire global routing table. Peering, however, does not require purchasing and instead relies on direct communication between two ISPs to exchange traffic, providing a connection for their respective customers [10].

Each AS has a unique ASN used for identification [30], which is useful when BGP path-traverses using AS_PATH [10].

2.2.5.1 Route reflectors and Cluster lists

Within larger internal networks, the use of Route Reflectors acts as speakers to all neighboring nodes [31]. Therefore single network reachability messages update will be propagated through multiple parallel paths to avoid a complex full mesh structure. The Cluster ID section is used to prevent loops by recording the ordered steps a route has traversed/reflected accross the network [31].

2.2.5.2 Internet Service Provider

ISPs provide service to these different tiers of internet connections. The different levels of ISPs are located by country for tier 1, then region for tier 2 and lastly locally for tier 3. The tier 1 ISPs connect, therefore, with each other, fulfilling the global connection [32].

2.2.5.3 Internet Exchange Point

The purpose of an IXP is to facilitate the exchange of Internet traffic between ASes [33]. By providing a shared switching platform, an IXP allows participating networks (ISPs) to connect their edge routers and share routing information via BGP, enabling them to pass traffic directly between their respective networks [33].

2.2.5.4 Routing Information Base and Forwarding Information Base

Routing Information Base (RIB) stores peering information, Adj-RIB-In for adjecent peers, resulting best paths are stored in Loc-RIB then to the IP-RIB [10]. Forwarding Information Base (FIB) stores all the best path IP prefixes from the routing table enabling effective forwarding [10].

2.3 IP Anycast and Catchments

IP Anycast is a network addressing and routing method in which a single destination IP address (prefix) is shared by multiple physical endpoints (nodes) distributed across different geographic locations [6]. Through BGP, these nodes advertise the same prefix to the global internet which in turn enables the routing infrastructure to route traffic to a specific node based on BGP's path selection criteria [6].

Catchments & Catchment Leaks: Catchments are defined as the specific set of client networks or IP probes that are logically mapped to a nearby Anycast site [34]. On the other hand, a catchment leak is the opposite. The client is routed to a distant node, meaning there is a route leak occuring. BGP routing advertisements are propagated outside their intended scope, consequently violating the routing policies of the sender, receiver, or any intermediate AS [35].

2.4 Previous Research

2.4.1 BGP and Anycast Performance

Modeling BGP is inherently more complex than internal routing protocols due to its decentralized nature. A BGP speaker has knowledge only of its immediate neighbors' advertisements, making large-scale path prediction difficult [36]. This lack of global visibility often leads to suboptimal routing in Anycast environments. Research by Li et al. [9] suggests that increasing the number of Anycast nodes does not necessarily improve performance. While tier 1 instances like the C-root achieve near-optimal routing through management by a single provider, other root servers (e.g., K and L) suffer from "Anycast Path Inflation" because they rely on multiple operators whose independent policies outweigh the intended catchment goals.

A critical discovery in this context is the occurrence of "Unknown Tie-breaks," highlighting BGP's non-deterministic impact on traffic distribution.

Recent studies on Regional IP Anycast demonstrate that performance can be improved by "overriding" standard BGP preferences. Zhou et al. [14] found that by partitioning Anycast into regions, providers can bypass BGP's inherent preference for customer-based transit over shorter peering paths, which, in their study, significantly reduced tail latency, accounting for 44.1% of the observed performance gains. The feasibility of using specific attributes like Local Preference or MED to intentionally shape these catchments has been further validated through emulation-based modeling [37].

Controlling incoming inter-domain traffic remains a primary challenge in BGP traffic engineering [38]. Building on this, Zhang et al. [34] introduced AnyOpt, a system using pairwise site measurements to predict catchments. They identified that route arrival order, a non-standard BGP tie-breaker, influences up to 14% of routing decisions. By accounting for this temporal factor, AnyOpt reached a catchment prediction accuracy of 94.7%, providing a robust framework for offline configuration optimization.

2.4.1.1 The Impact of BGP Tie-breaking

The BGP selection process frequently falls through to tie-breaking rules when the primary attributes (Local Preference and AS_PATH) are equal. A significant factor in these scenarios is "Hot Potato" routing, defined by step 6 in the BGP Best Path Selection Process in Table 2.1, meaning an AS exits traffic at the closest internal point to minimize its own resource usage [8]. This mechanism prioritizes internal AS stability over global path optimality, often causing disruptive catchment shifts in response to minor internal topology changes [8].

The frequency of tie-breaking varies significantly by network tier. Simulation data from Quoitin et al. [38] shows that in the Internet core (tier 1 and tier 2), 50% of the route selection where tie-breaks due to the highly interconnected structure of these networks and therefore many with equal AS_PATH lengths [38]. In contrast, tier 3 and other lower-tier networks are less complex; 30% of their selections are based on AS_PATH length, while another 30% rely on tie-breaks. Notably, up to 45% of these smaller networks are single-homed and therefore do not invoke the decision process, as the single available route is installed by default [38].

2.5 Summary

This chapter has explored the mechanics of BGP best-path selection and the specific challenges associated with global Anycast deployments. While BGP is stable for popular destinations, its “performance-blind” nature leads to Anycast path inflation and unstable catchments.

Chapter 3

Methods

3.1 Research strategies

A dual-method approach was implemented, integrating network simulation using Netlab with real-world scenarios derived from RIPE Atlas, AS-specific LGs, and other publicly available, verified sources. This methodology was used to assess the accuracy of network simulations in relation to the precision of the collected data.

3.1.1 Cases

Each case is based around the I-root, particularly when RIPE Atlas reports high Round-Trip Time (RTT) for countries within a network where the I-root is highly available. The I-root is widely accessible in central Europe, which means there is a high number of I-root instances (see Figure 2.2) leading to low RTT. In contrast, the I-root is less available in Oceania (see Figure 2.2), so measurements with high RTT are considered typical in that region.

3.1.2 Simulated testing

The simulation provides a controlled environment for analyzing network routing scenarios. This is to confirm simulation accuracy in mirroring real-life cases.

Environment

The simulation employs Netlab and Free Range Routing (FRR) to emulate ASes, their corresponding BGP routers, and configurations.

Topology

Each topology is constructed using real-world cases extracted from RIPE Atlas, all of which exhibit higher latency than expected. The topology reflects the AS structure of each case's traceroute, and every relevant BGP configuration influencing the routing is included and tested. Additional modifications involve introducing test scenarios to demonstrate that the simulation accurately reproduces real-world results.

3.2 Research Process

Real-world data and traceroute measurements to the I-root are analyzed, and simulations are developed based on the identified suboptimal routing scenarios.

1. **Anomaly Identification:** Elevated RTT measurements from global probes to I-root nodes were identified using RIPE Atlas and documented as cases for further investigation. Notably, higher RTT values were observed in regions with a high density of I-root nodes, which makes such RTT spikes atypical. To understand the routing movement, the RIPE Atlas Traceroute Command Line Interface (CLI) was utilized exclusively.

2. **BGP Analysis:** Each documented case was studied using AS-specific LGs, PeeringDB, bgp.tools, and the RIPE NCC and APNIC Whois databases. This allowed for further understanding of the BGP configurations and their effect on the suboptimal routing.

3. **Simulation:** Each case was reconstructed within a simulated environment using Netlab, and the suboptimal routing was replicated. The accuracy of the simulation was evaluated by comparing the results to the collected data.

3.3 Research Paradigm

This study adopts, to a certain extent, a post-positivist research paradigm, which acknowledges that observations are inherently imperfect and that knowledge is provisional rather than absolute. The analysis assumes that routing behavior can be observed through passive measurements and that BGP behavior can be inferred from publicly available data. Since some BGP

configurations are not fully visible to entities outside the respective AS, certain interpretations were necessary.

An abductive reasoning approach is employed, utilizing observations from both RTT measurements and BGP-specific data sources to infer the underlying causes of routing behavior. These inferences are subsequently tested in a simulated environment (Netlab) to determine whether they reproduce the observed phenomena.

Additionally, this study integrates both quantitative and qualitative methodologies. The analysis of RTT measurements to detect anomalous routing behavior represents the quantitative component, while the examination of BGP policy configurations constitutes the qualitative aspect. Each anomaly is treated as an individual case, emphasizing the relevance of network dependency and variability in inter-domain routing.

3.4 Data Collection

For this study, data will be collected from the simulated environment, passive measurements and traceroutes, as well as BGP specific configurations.

Passive Measurement Data & BGP specific configuraitons: Passive latency measurement data were collected from RIPE Atlas using globally distributed probes to observe RTT directed toward I-root instances. To analyze the specific routing paths, traceroute data was obtained via the RIPE Atlas CLI. Furthermore, additional BGP configuration and routing policy information was gathered from AS-specific LGs, PeeringDB, bgp.tools, and both the RIPE NCC and APNIC Whois databases.

Simulation Data: Simulation data are collected from Netlab, specifically from the output generated after configuration. These data are then compared to the collected data to assess their validity.

Ethical Considerations: The data utilized is public and does not contain Personally Identifiable Information (PII). Measurement probes in RIPE Atlas are hosted by volunteers, and the research adheres to the RIPE community's ethical guidelines regarding data usage and credit.

The simulation environment is implemented using Netlab, an open-source and released under the MIT license. All experiments are conducted locally and do not impact or interact with production networks, preventing impacts from reaching operators or the global network.

3.4.1 Sampling

This study samples based on the selected probe within certain geographically dense areas where I-root instances are available to a high degree. Highlighting the abnormality of high RTT measurements within those areas.

3.4.2 Sample Size

The number of RIPE Atlas probes utilized depends on the traceroute results and whether their inclusion significantly affects the observed route. Private nodes or IP addresses are excluded, as these typically represent clients exiting their own networks.

3.4.3 Target Population

The target population comprises selected probes and the corresponding data for each significant segment of their traceroute. All data includes AS-specific information and details regarding their interactions within the global network.

3.5 Experimental design and Planned Measurements

3.5.1 Test environment and Model

The experimental testbed is built with Netlab, a virtual network laboratory tool that uses infrastructure-as-code principles, with YAML files to construct topologies [39]. The model replicates real-life network systems and their BGP configurations. Specifically for this study, FRR was used to run the actual BGP daemon code. This ensures that the simulation behaves identically to real-world routing hardware.

3.5.2 Software to be used

- **Virtual network tool:** Netlab is the core simulation framework used to define and deploy the network topologies [39]. It is built on ContainerLab, which runs network nodes as Docker containers [40]
- **Topology Visualization:** The visual representations of Netlab topology graphs were generated using Netlab's internal tool. These graphs were subsequently edited with custom designs to enhance clarity.

- **Routing Engine:** FRR provides the BGP protocol stack within the emulated ASes. [41].
- **RIPE Atlas:** A global measurement platform used to collect and present real-world latency data, specifically utilized in this thesis to observe RTT inflation [42].
- **RIPE Atlas traceroute CLI:** This tool offers multiple functionalities and in the context of this thesis, it is used to map reachability to the destination in a hop-by-hop manner [43].
- **LGs:** Tools provided by the specific AS providers that present their BGP configurations [44]. The third-party LG bgp-lg.ping.pe, which is a diagnostic tool to study network connectivity problems at the IP transit level [45], was consulted if the AS-specific ones were not available.
- **PeeringDB:** Is a non-profit, community-driven database of networks. PeeringDB is a freely available, user-maintained database that documents the global interconnection of networks at Internet Exchange Points (IXPs), data centers, and other interconnection facilities [46].
- **I-root Specifics:** Using official I-root metadata and their node locations [47].
- **bgp.tools:** Have both free and paid features. In this thesis, the free tools where used which are, the BGP routing information as well as AS-specific prefixes and connectivities [48].
- **Whois from RIPE NCC and APNIC Databases:** A registry is used to identify properties, and in this thesis, BGP communities of specific AS were identified [49, 50].

3.6 Assessing reliability and validity of the data collected

3.6.1 Validity of method

The validity of the method is demonstrated by employing FRRouting in the simulation, which utilizes the same BGP protocol stack as enterprise-grade network hardware. Administrative policies of real-world Internet Service Providers are defined using data from AS-specific LGs, PeeringDB, bgp.tools,

the RIPE NCC and APNIC Whois databases, and AS-specific articles. The approach is further validated by replicating observed scenarios within a simulated environment, confirming the relationship between BGP routing policies and Anycast performance.

3.6.2 Reliability of method

Methodological reliability is achieved through the application of infrastructure-as-code principles. The Netlab environment is configured using YAML files to maintain consistency. Furthermore, the RIPE Atlas measurement methodology utilizes a scheduled polling interval, which renders the data collection process systematic and reproducible by other researchers employing identical parameters. This approach also ensures that the data remains current. Reliability of AS-specific information is actively verified using multiple sources to support various hypotheses. LG tools offer direct visibility into BGP configurations for each ISP, which are compared and analyzed to assess the realism of the results. Similarly, PeeringDB and bgp.tools provide metadata for each AS, contributing to the verification process. Additionally, consultations with Netnod and the use of their internal tools further corroborate the validity of the findings.

3.6.3 Data validity

The well-established reputations of the RIPE Atlas, AS-respective LGs as well as other publicly available information such as PeeringDB and bgp.tools underpin the credibility of the collected data. RIPE Atlas operates a globally distributed network of hardware and software probes with standardized configurations, enabling precise measurements of RTT.

3.6.4 Reliability of data

Reliability is ensured by cross-verifying data across multiple validated sources. The probable cause of the elevated RTT observed via RIPE Atlas is identified by comparing these measurements with information from LGs and other public sources. This approach confirms that high-latency spikes are a consistent issue, likely attributable to BGP policy configuration, rather than sporadic anomalies.

3.7 Planned Data Analysis

This analysis employs an abductive reasoning framework. Observed routing anomalies are examined to infer the most probable BGP policy explanation, which is subsequently tested within a controlled simulation environment. Each case is addressed by identifying the relevant BGP policy attribute associated with the observed routing behavior. The hierarchical structure of the BGP decision process requires that attributes are evaluated in a fixed order of precedence, which is essential for determining the consequences of BGP policy settings.

3.7.1 Data Analysis Technique

The main analytical methodology combines direct analysis of BGP configurations with policy inference and consists of the following procedures:

- **Event-Performance Correlation:** Correlation of BGP events with observed RIPE Atlas RTT spikes to identify unexpected routing shifts.
- **Attribute Correlation:** Identification of BGP routing-specific behaviors using AS-specific LGs and other public sources to enhance the accuracy of formulated hypotheses.
- **Comparative Validation:** Comparison of mirrored simulation results in Netlab with real-world data to assess the validity of the simulation approach.

3.7.2 Software Tools

As mentioned in previous Section 3.5.2, the same tools used for simulation are used for data analysis. Specifically, RIPE Atlas and the BGP policy oriented tools to analyze the routing behavior. Simulation tools data are used to be analyzed whether its an accurate representation.

3.8 Evaluation Framework

The evaluation of the thesis is based on two primary criteria:

1. First, the extent to which data, combined with information from BGP specific sources, contributes to a comprehensive understanding that enables effective simulation is assessed.

2. If the simulation can mirror real-world data as well as mirror the results, proving that it can be a useful tool for providers.

3.9 System documentation

The topologies are constructed using the collected data and their corresponding levels of detail. As the data specifies each AS hop and the IXP or certain router traversed, only these elements were modeled. Neighboring nodes not directly involved were excluded to prevent unnecessary system complexity. This approach demonstrates that networks can be effectively simulated and analyzed in isolation.

Chapter 4

From Collected Data to Controlled Simulation

The research was conducted primarily through network simulation to emulate real-world I-root instances, thereby simplifying the complex network topology. The methodology adopted a case-specific approach by minimizing topology sizes and including only the relevant ASes and routing mechanisms. This approach enables isolated replication of the observed BGP anomalies and eliminates unnecessary network noise from uninvolved parties.

4.1 Hypothesis Formulation

The research process began with a hypothesis informed by data collected primarily from RIPE Atlas, AS-specific LGs, and the RIPE NCC and APNIC Whois databases. Where necessary, supplementary data was sourced from PeeringDB and bgp.tools. Instances of high RTT were identified within the RIPE Atlas data. Specific cases were then selected and classified for further study if routing inefficiencies were expected to be minimal, based on their geographical location and proximity to nearby I-root instances.

4.2 Simulation Baseline Validation

The initial methodology focused on analyzing BGP behavior to determine whether specific BGP configurations could lead to inaccurate routing. The plan involved collecting case studies, mapping them, including Netnod's neighboring nodes, and abstracting these into a general scenario with a fixed

number of nodes. The identities of the specific nodes were not prioritized; instead, the focus was on the topology size and conducting multiple tests, see Figure 4.1 and Listing B.1. This phase was discontinued upon recognizing that BGP routers when operating in isolation within its AS, functions as documented. Consequently, rather than testing all BGP attributes, it was always as established in official documentation that the attribute with the highest priority in the BGP Best Path Selection Process (see Table 2.1) hierarchy dictates the preferred route. If all attributes are equal, route selection defaults to the arrival order.

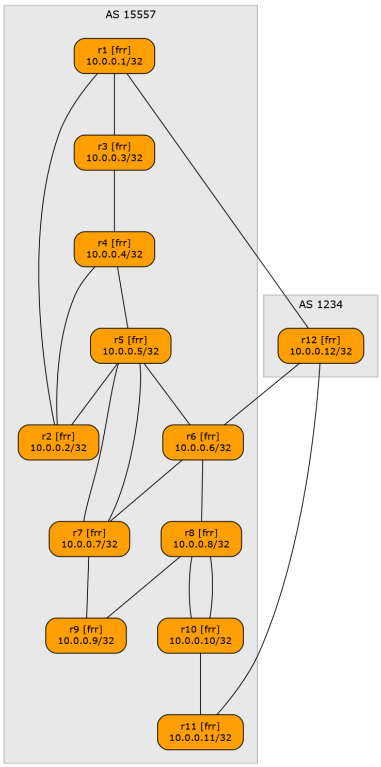

Figure 4.1: Test topology. Generated via Netlab

```
r11# show ip bgp 1.2.3.0/24
BGP routing table entry for 1.2.3.0/24, version 4
Paths: (2 available, best #1, table default)
  Advertised to peers:
  r12(10.1.0.66)
  1234
    10.0.0.1(r1) (metric 60) from r1(10.0.0.1) (10.0.0.1)
      Origin IGP, metric 0, localpref 101, valid, internal, bestpath-from-AS 1234, best (Local Pref)
      Last update: Mon May 18 11:50:04 2026
  1234
    10.1.0.66(r12) from r12(10.1.0.66) (10.0.0.12)
      Origin IGP, metric 0, valid, external
      Last update: Mon May 18 11:49:51 2026
```

Listing 4.1: General local preferernce test

In this specific test case (see Listing 4.1), the advertising is done from `r11`, which is closest to the destination `r12`. This shows that the simulation forces the packet to travel to the node with the highest local preference rather than taking the shortest topological path.

4.3 Case Study

This thesis presents a case study of various suboptimal routing scenarios. For each case, publicly available tools are utilized. The analysis begins by examining each traceroute to determine the network path. Subsequently, the respective AS LGs are consulted to investigate the providers' BGP configurations. This process includes the use of bgp.tools and whois.ripe.net to access all publicly available BGP configuration data. If aspects of the network remain unclear, additional sources such as PeeringDB, Hurricane Electric, or AS-specific articles are consulted.

Providers, in general can't identify if outside AS are transit or peer. Both are technically the same but transit open the entire internet to the client instead of a direct one on one connection.

4.4 Simulation

Based on the collected information, a simulation will be constructed for each case. Each faulty BGP configuration will be tested and presented to demonstrate that simulations can accurately replicate both the presentation and behavior of a network.

MIX-IXP and DIX-IXP: In the simulation, the IXP is modeled as an intermediate node in the topology. In practice, an IXP acts as a transparent intermediary where each connected endpoint perceives the destination on the opposite side directly, rather than the IXP's own IP address. This behavior

is accurately reflected in the simulation, where endpoints communicate as though physically connected end-to-end. To further reflect real-world conventions, direct links use /31 prefixes while IXP connections use /24 prefixes.

BGP communities: These are specified as AS-specific and, unless manually modified, do not affect traffic independently. Instead, they serve to enhance visual accuracy and facilitate effective debugging.

RIPE NCC traceroute: The most reliable traceroute results were obtained using the RIPE NCC CLI, which provides raw data. In contrast, the TraceMon visualization tended to misrepresent certain endpoints, resulting in it being excluded from the final analysis.

Simulation: In the current simulation, traffic is directed exclusively toward the I-root. As a result, specific details, such as the prefixes for other nodes, were not visible. To enhance the accuracy of the simulation representation, implementing tracerouting in multiple directions would provide more comprehensive results. This was however not included in the final result as direction toward the I-root was prioritized.

Lastly, it is important to note that internet traffic is inherently volatile and may fluctuate over time, which renders these results indicative rather than conclusive. The primary objective was to simulate a realistic scenario and evaluate path selection strategies.

Netlab is designed to reflect real-world network functionality. As mentioned previously, the primary exception involves BGP communities, as some AS-specific BGP communities cannot be directly replicated. Instead, these communities are emulated as specific paths, with BGP communities represented in the output to preserve visual accuracy. This method aligns the simulation's behavior with real-world scenarios; however, accuracy may decrease if the network structure is altered. Unless explicitly specified, IP addresses are assigned arbitrarily. The syntax and configuration, including BGP policy implementation, were derived from Netlab's official GitHub documentation. [51], as well the FRR CLI [52]

4.4.1 Case 1

Figure 4.2: Probe 1007504 geographical location. Source: RIPE Atlas

Figure 4.2 shows that the probe is located north of France. The world map in Figure 2.2 indicates that this region has high I-root availability.

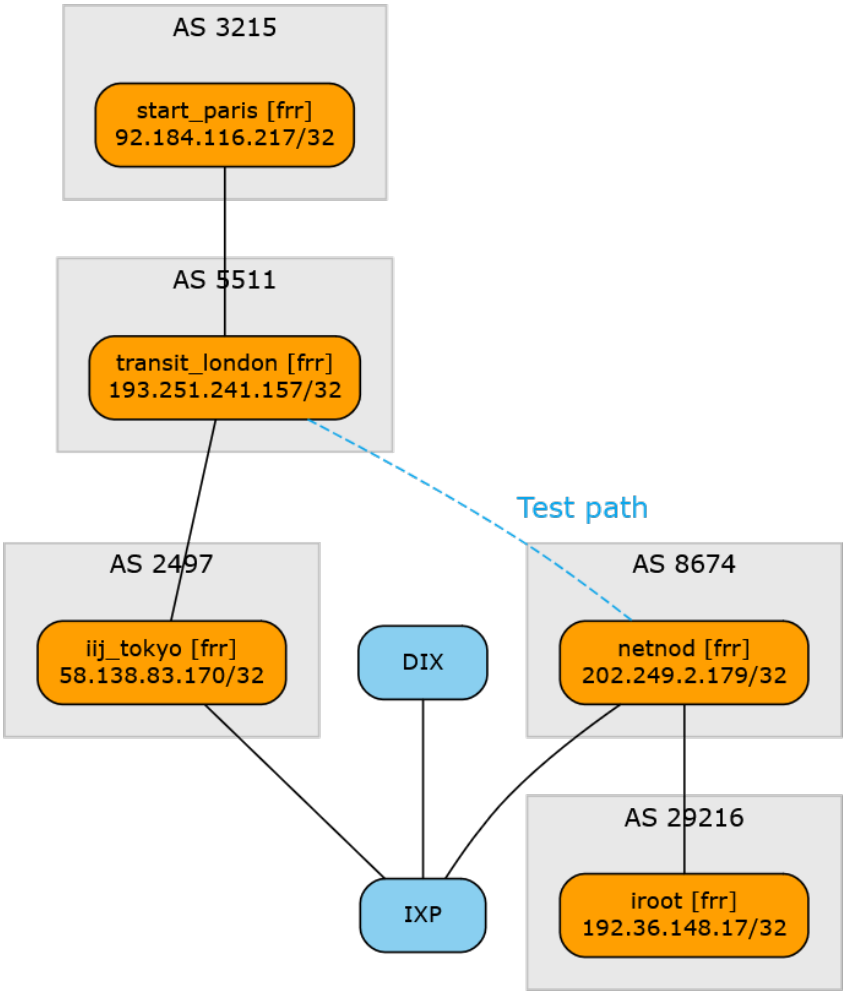

Figure 4.3: Netlab topology for Case 1. Generated via Netlab

The primary focus of this topology (see Figure 4.3 and Listing B.2) is the AS-path and its influence on routing behavior. A test path demonstrates that, under default conditions, routing will prefer the shorter path. However, introducing a lower local preference to the shorter path results in it becoming less preferable. This adjustment is made because the compiled data does not explicitly specify all attributes, suggesting the possibility of hidden attributes not visible in the LG. Therefore, a lower local preference is assumed for the analysis. This topology is structured based on the compiled data presented in Chapter 5 and reflects the current network configuration for Orange transit from Paris (see Table A.1).

4.4.2 Case 2

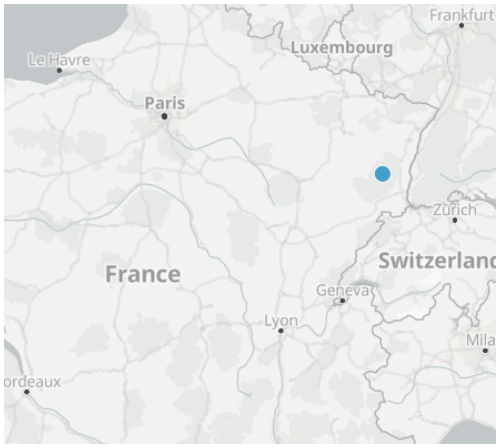

Figure 4.4: Probe 50288 geographical location. Source: RIPE Atlas

Figure 4.4 indicates that the probe is located in eastern France. As illustrated in Figure 2.2, I-root demonstrates high availability in this region.

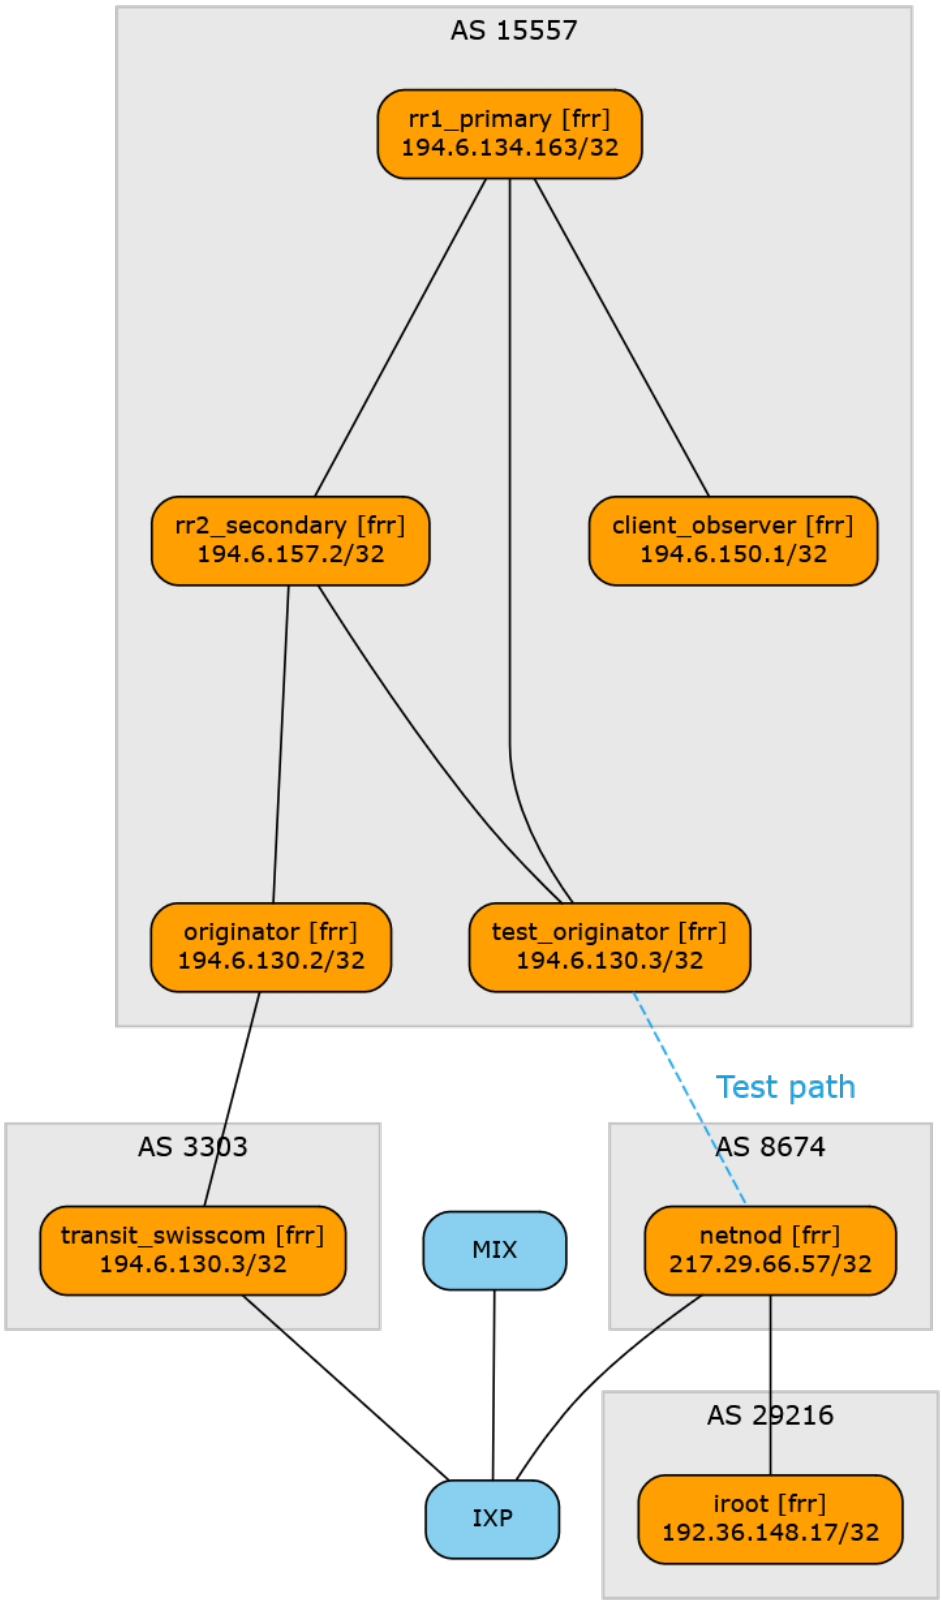

Figure 4.5: Netlab topology for Case 2, rr = Route Reflector. Generated via Netlab

This topology (see Figure 4.5 and Listing B.3) is designed to demonstrate the impact of high local preference. A test path with lower local preference was incorporated to validate this effect. Route Reflectors were included to maintain consistency with the LG presented in Table 5.2. Furthermore, because IXPs primarily serve as extensions rather than active routing units, the IXP was included in this simulation solely for visual representation.

Listing B.3 also illustrates the use of `module:`, which restricts the permitted routing protocols. In both cases 2 and 3, the IXP was restricted to none to prevent incorrect functionality.

4.4.3 Case 3

Figure 4.6: Probe 14521 geographical location. Source: RIPE Atlas

Figure 4.6 presents the probe location in central Germany, where, according to Figure 2.2, I-root demonstrates high availability.

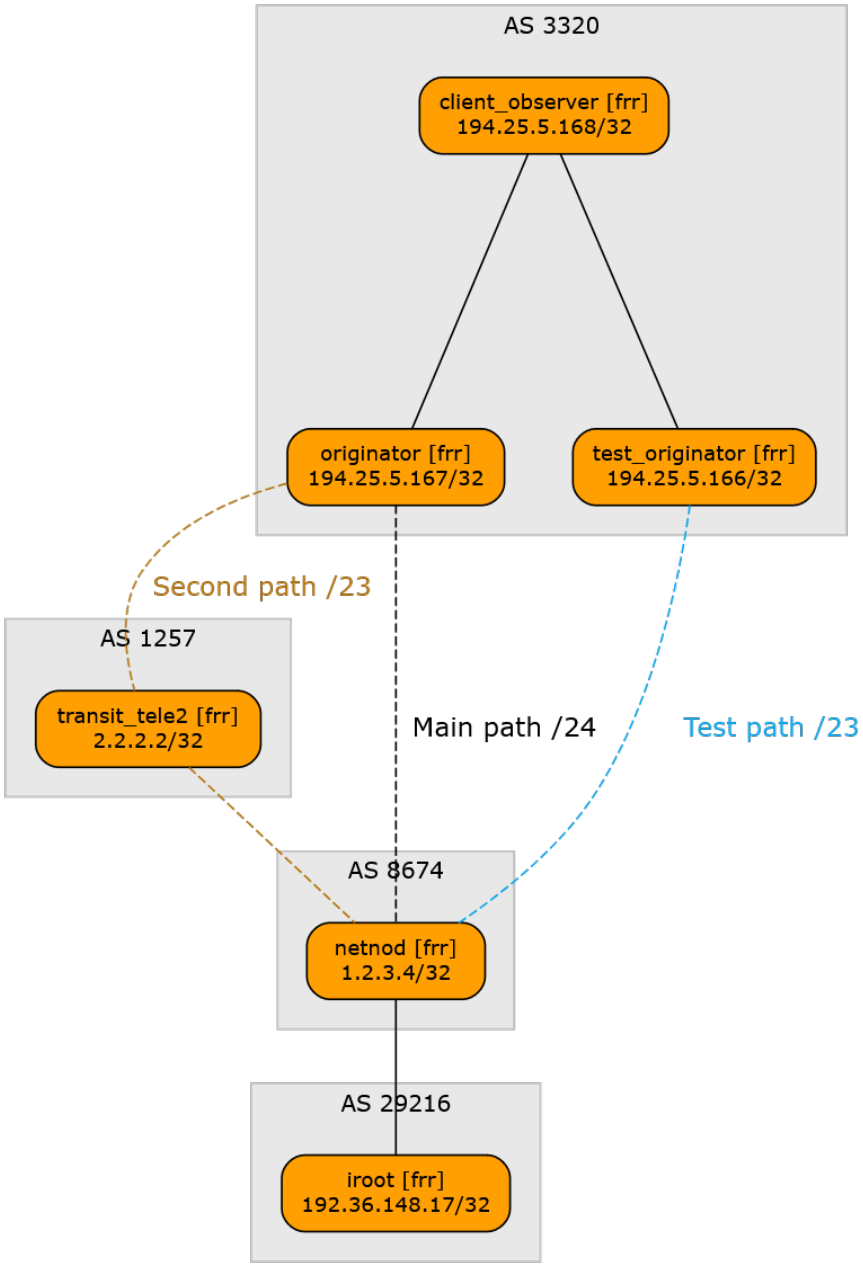

Figure 4.7: Netlab topology for Case 3. Generated via Netlab

This topology (see Figure 4.7 and Listing B.4) the main focuse revolved around the different prefixes available to iroot. Therefore it has another test path with prefix /23 the same amount of AS as the Main path with prefix /24. This is to test if it prefer the longest matching prefix when other policies are

the same. The ordinarie Second path is also added to mirror the real-world case with its correct prefix setting.

Chapter 5

Results and Analysis

5.1 Case 1, RIPE Atlas Probe 1007504

5.1.1 Data

Hop	IP	Country
1	172.16.27.1	
2	10.164.83.64	
3	10.164.82.62	
4	10.164.82.65	
5	193.251.111.125	France (Orange)
6	*	
7	*	
8	*	
9	193.251.241.157	(Orange)
10	*	
11	58.138.83.178	Japan (IIJ)
12	58.138.88.85	
13	58.138.102.206	
14	202.249.2.179	Tokyo, Japan (DIX-IXP)
15	192.36.148.17	I-root
16	192.36.148.17	I-root

Figure 5.1: Traceroute RIPE: Probe 1007504 Date: Mon May 11 11:15:11 CEST 2026

Figure 5.1 presents the traceroute result using RIPE CLI tool. Every row represents as hop and every significant traffic shifting hop is marked in blue while green represents the final destination.

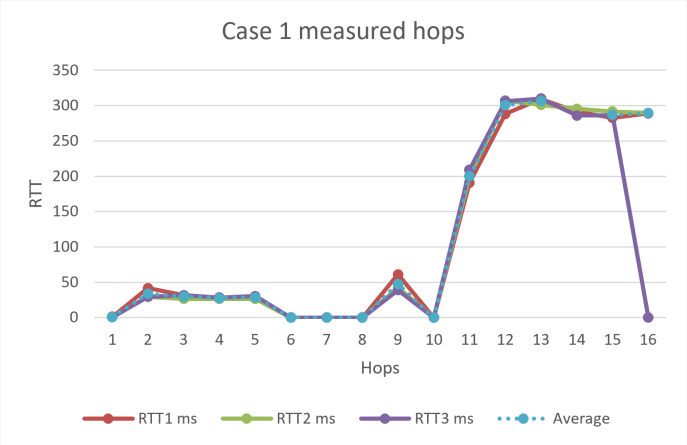

Figure 5.2: Visually presents the three signals RTT trend for each hop.

The traceroute yields notable results, as illustrated in Figure 5.2 and Figure 5.1. The route proceeds from Paris to the Orange transit, a tier 1 ISP, through London Telehouse (see Table 5.1). It then continues to the Tokyo IXP known as DIX before reaching the I-root. The minor increase in RTT between hops 5 and 9 supports the conclusion that the transition occurs toward the United Kingdom. Additionally, the RIPE Atlas traceroute in Figure 5.1 demonstrates a substantial RTT increase between hops 10 and 12, indicating that the packet traversed a considerable distance, which supports the inference of a Japanese destination. Moreover, the minimal RTT increases at hops 12 and 13 further confirm that the network traffic remains within Japan and in close proximity. The third IP address in the CLI output (58.138.102.206) is registered to Internet Initiative Japan (IIJ), as indicated by bgp.tools in Table 5.1.

Table 5.1: Case 1 data compiled from multiple sources

Case 1 Compiled Data	
RIPE Atlas data[1]	Probe 1007504 · ASN 3215 · France · Min RTT 289.4 ms · 16 hops · 2026-05-18 07:15:08
Traceroute path[2]	Paris (AS3215) → London (AS5511) → Tokyo (AS2497) [5] → Netnod (8674) → I-root (AS29216)

Looking Glass[3]	`Prefix: 192.36.148.0/24`
	`Next hop: 172.25.4.81`
	`AS_PATH: 2497 8674 29216`
	`Local Pref: 100`

AS path[2][3]	`3215 → 5511 → 2497 → 8674 → 29216`
BGP communities[3][4]	`5511:500` European Customer (Deprecated) `5511:515` UK and N.Ireland (Deprecated) `5511:999` BGP Comm. append to all customer routes `5511:30228` N. Europe `5511:30232` UK and N. Ireland `5511:30535` London Telehouse `5511:30594`

[1] RIPE Atlas [2] RIPE Atlas traceroute [3] Orange transit BGP Looking Glass A.1
[4] `whois` (RIPE NCC & APNIC) [5] `bgp.tools` [20]

All BGP communities, including Orange LG, are listed in the appendix (see Table A.1). The results presented are limited to those for which public information is available. Important to note: Subsequent analysis with bgp.tools identified an alternative /23 prefix path routed via AS 5511 through the Arelion AS 1299. This alternative was excluded from the final simulation because routing behaviors associated with prefix-length preference are already comprehensively studied in Case 3. Including this scenario would have resulted in redundant findings.

Upon further research, the reason it moved to Japan becomes clearer. Orange is a provider for IIJ, therefore clarifying why the traffic is influenced to travel in this specific direction [53]. AS5511 is Orange transit, and according to the data gathered, the probe is a customer of Orange, which is noted by the BGP communities 5511:500 and 5511:999 Table 5.1. The 5511:500

community is marked as deprecated, which creates ambiguity regarding its current applicability or potential removal. Its redundancy may indicate a transition from legacy to updated BGP communities.

Client ISP needed to route through AS5511 in London, also proved by the BGP Community 5511:30535 shown in Table 5.1 and to traverse the network. To reach the I-root, it could simply move to an I-root instance in London, but it still sends it to Japan. Although it is known that Orange and IIJ have an agreement, it remains unclear which BGP configuration or community has forced the specific connection to Japan. Therefore the rest of the BGP communities 5511:1000, 5511:2000, 5511:3000, 5511:10009, 5511:40176 in Table A.1 seem to be hidden for public use and therefore not detailing a comprehensive image. Another possible case for the direct path to I-root be less preferred could be that the LGs don't provide all the information and don't show all the results, and therefore the local preference could have been set to a much lower value without it being visible.

5.1.2 Simulation results

The simulation in Figure 4.3 shows an accurate depiction of the compiled data.

```
1
2  transit_london# show ip bgp 192.36.148.0/24
3  BGP routing table entry for 192.36.148.0/24, version 7
4  Paths: (2 available, best #1, table default)
5    Advertised to peers:
6    netnod(10.0.0.1) start_paris(10.0.0.2) iij_tokyo(10.0.0.5)
7    8674 29216
8      10.0.0.1(netnod) from netnod(10.0.0.1) (202.249.2.179)
9        Origin IGP, valid, external, bestpath-from-AS 8674, best (AS Path)
10       Last update: Thu May 21 09:43:53 2026
11   2497 8674 29216
12     10.0.0.5(iij_tokyo) from iij_tokyo(10.0.0.5) (58.138.83.170)
13       Origin IGP, valid, external, bestpath-from-AS 2497
14       Community: 5511:500 5511:515 5511:999 5511:1000 5511:2000 5511:3000 5511:10009 5511:30228
               5511:30232 5511:30535 5511:30594 5511:40176
15       Last update: Thu May 21 09:43:53 2026
```

Listing 5.1: Simulation result without forced connection to IIJ

Without forcing any path preference toward IIJ, it will naturally gravitate toward the path with the shortest AS hops. Which is the direct Test Path from transit_london to Netnod (see Figure 4.3).

```
transit_london# show ip bgp 192.36.148.0/24
BGP routing table entry for 192.36.148.0/24, version 4
Paths: (2 available, best #1, table default)
  Advertised to peers:
  netnod(10.0.0.1) start_paris(10.0.0.2) iij_tokyo(10.0.0.5)
  2497 8674 29216
    10.0.0.5(iij_tokyo) from iij_tokyo(10.0.0.5) (58.138.83.170)
      Origin IGP, valid, external, bestpath-from-AS 2497, best (Local Pref)
      Community: 5511:500 5511:515 5511:999 5511:1000 5511:2000 5511:3000 5511:10009 5511:30228
           5511:30232 5511:30535 5511:30594 5511:40176
      Last update: Thu May 21 09:37:14 2026
  8674 29216
    10.0.0.1(netnod) from netnod(10.0.0.1) (202.249.2.179)
      Origin IGP, localpref 50, valid, external, bestpath-from-AS 8674
      Last update: Thu May 21 09:37:14 2026
```

Listing 5.2: Netnods local preference was lowered to match the output of the LG

The path through IIJ is chosen and portrays the correct output. It appears the BGP Communties do not affect the path selection since its custom to Orange AS. However, the topology structure is already connected to London, therefore fulfilling the BGP community purpose. Nevertheless, there is no indication as to which BGP configuration is routing it to Japan. Therefore it was an edit in this research to make the local preference of Netnod set to 50. Resulting in the traffic seen in the LG 5.2. Setting local preference to 50 is based on the knowledge that LGs don't show all possible results, meaning, that there could be configurations not shown but still affect the routing.

```
netlab connect transit_london
transit_london(bash)# vtysh
transit_london# conf
transit_london(config)# router bgp 5511
transit_london(config-router)#  neighbor 10.0.0.5 shutdown
transit_london(config-router)# end
```

Listing 5.3: Cutting the peering connection

Providers are able to remove peer connections within the FRR environment [41], enabling them to activate or deactivate links as necessary. In this instance, the link was deactivated.

```
transit_london# show ip bgp 192.36.148.0/24
BGP routing table entry for 192.36.148.0/24, version 8
Paths: (1 available, best #1, table default)
  Advertised to peers:
  netnod(10.0.0.1) start_paris(10.0.0.2)
  8674 29216
    10.0.0.1(netnod) from netnod(10.0.0.1) (202.249.2.179)
      Origin IGP, localpref 50, valid, external, bestpath-from-AS 8674, best (First path received)
      Last update: Thu May 21 09:37:13 2026
```

Listing 5.4: The result after a the IIJ link is removed

In real life, Netnod can force back the traffic by turning of the peering connection through IIJ (see Listing 5.3). When running the test again, we get the expected result, as seen in Listing 5.4, which returns the traffic back to the direct path toward the I-root.

5.2 Case 2, RIPE Atlas Probe 50288

5.2.1 Data

Hop	IP	Country
1	192.168.8.1	
2	*	
3	10.2.9.114	
4	10.203.116.1	
5	194.6.149.123	France (SRF)
6	194.6.146.71	
7	193.5.122.102	
8	138.187.129.98	Switzerland (Swisscom)
9	138.187.129.27	Switzerland (Swisscom)
10	138.187.129.128	Switzerland (Swisscom)
11	138.187.129.203	Switzerland (Swisscom)
12	138.187.131.1	Switzerland (Swisscom)
13	217.29.66.57	Milan, Italy (MIX-IXP)
14	192.36.148.17	I-root

Figure 5.3: Traceroute RIPE: Probe 50288 Wed May 20 20:32:06 CEST 2026

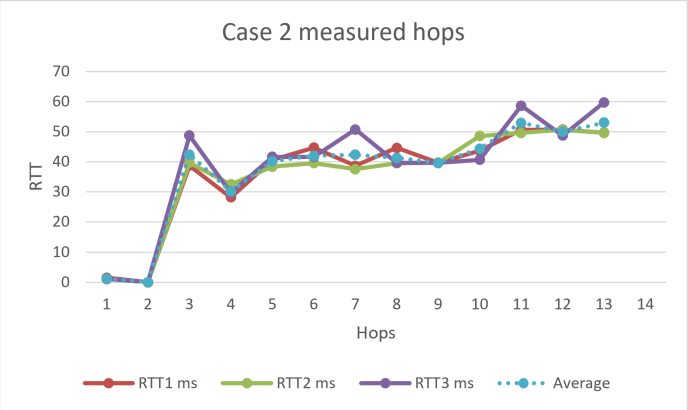

Figure 5.4: Visually presents the three signals RTT trend for each hop.

Looking at the RIPE Atlas traceroute 5.2 the ip right before the Netnod IP is an IP located in Milan (217.29.66.57), specifically the IXP MIX-IT. This indicates that the route went from Milan to reach I-root.

Table 5.2: Case 2 data compiled from multiple sources

Case 2 Compiled Data	
RIPE Atlas data[1]	Probe 50288 · ASN 15557 · France · Min RTT 49.6 ms · 14 hops · 2026-05-18 07:32:05
Traceroute path[2]	Paris (AS15557) → Switzerland (AS3303) → Milan (IXP) → Netnod (8674) → I-root (AS29216)

Looking Glass[3]	`Prefix: 192.36.148.0/24`
	`Next hop (Originator): 194.6.130.2`
	`Cluster list: 194.6.134.163 194.6.157.2`
	`AS_PATH: 3303 8674 29216`
	`Local Pref: 4000`

AS path[3]	`15557 → 3303 → 8674 → 29216`
BGP communities[4][5]	`15557:90 15557:300 15557:500` (SFR, unknown meaning) `3303:1000` customer route (Swisscom) `3303:1006` European route (Swisscom) `3303:3055` route received in Milan (Swisscom)

[1] RIPE Atlas [2] RIPE Atlas traceroute [3] SFR AS15557 BGP Looking Glass
[4] Swisscom AS3303 BGP Looking Glass (`bgp-lg.ping.pe`) [5] `whois` (RIPE NCC)

Compared to previous case, the SFR LG 5.2 data reveals a clear finding: the route's Local Preference is set at 4000, which is exceptionally high compared to the default value, 100. Further research also reveals that the publicly known BGP communities direct the traffic to Milan due to Swisscoms (AS 3303) peering with SFR (15557), and the BGP community if its a European customer 5.2.

The local preference of 4000 does bring up the simple question of how long it stretches within its AS and listing 4.1 demonstrates that, within the same AS, geographical distance is irrelevant and the entity with the highest local preference is selected.

5.2.2 Simulation results

```
client_observer# show ip bgp 192.36.148.0/24
BGP routing table entry for 192.36.148.0/24, version 2
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  3303 8674 29216
    194.6.130.2(rr2_secondary) (metric 30) from rr2_secondary(194.6.157.2) (194.6.130.2)
      Origin IGP, localpref 4000, valid, internal
      Community: 15557:90 15557:300 15557:500
      Originator: 194.6.130.2, Cluster list: 194.6.157.2
      Last update: Tue May 26 12:08:06 2026
  3303 8674 29216
    194.6.130.2(rr1_primary) (metric 30) from rr1_primary(194.6.134.163) (194.6.130.2)
      Origin IGP, localpref 4000, valid, internal, bestpath-from-AS 3303, best (First path received)
      Community: 15557:90 15557:300 15557:500
      Originator: 194.6.130.2, Cluster list: 194.6.134.163
      Last update: Tue May 26 12:08:05 2026
```

Listing 5.5: Shows the observer network with 4000 local preference with correct cluster list

Similarly to real-world scenarios the local preference and BGP communities in the simulation were configured at the AS exit point, designating the originator to attract network traffic and subsequently forward it along the intended route. This configuration redirects traffic away from the test_originator, thereby achieving the simulation objective of demonstrating the impact of local preference, even when more optimal paths exist.

The test path, as seen in Figure 4.5, does not appear in the result output (see Listing 5.5) because it becomes invalid immediately once the path with local preference 4000 is known within the AS.

Using the netlab standard syntax to set up route reflectors [54] correctly defines their purpose. However it doesn't completely mirror the LG (see Table 5.2). Netlab mirrors the behavior of cluster lists correctly to the standard of cluster lists, meaning it works in parallel [54]. Each route reflector advertises their own path to the client_observer, which is why we see two pathways with one cluster list each, as seen in Listing 5.5. However, this does not fully replicate the SFR LG output (see Table 5.2), indicating that additional, unidentified configurations may be present. Achieving a strictly sequential cluster list on NetLab requires further custom configurations, but these were omitted because they are based on speculation rather than established facts.

```
netlab connect originator
originator(bash)# vtysh
originator# conf
originator(config)# router bgp 15557
originator(config-router)# neighbor 10.0.0.11 shutdown
originator(config-router)# end
```

Listing 5.6: Peer to the local preference path (originator) turned of to force the test path with lower local preference [55]

To verify the remediation, the peering connection to the originator was shut down (see Listing 5.6), causing the router to fall back to the path with fewer hops.

```
client_observer# show ip bgp 192.36.148.0/24
BGP routing table entry for 192.36.148.0/24, version 18
Paths: (2 available, best #1, table default)
  Not advertised to any peer
  8674 29216
    194.6.130.3(rr1_primary) (metric 20) from rr1_primary(194.6.134.163) (194.6.130.3)
      Origin IGP, localpref 100, valid, internal, multipath, bestpath-from-AS 8674, best (Neighbor IP)
      Originator: 194.6.130.3, Cluster list: 194.6.134.163
      Last update: Tue May 26 12:24:09 2026
  8674 29216
    194.6.130.3(rr2_secondary) (metric 20) from rr2_secondary(194.6.157.2) (194.6.130.3)
      Origin IGP, localpref 100, valid, internal, multipath
      Originator: 194.6.130.3, Cluster list: 194.6.157.2
      Last update: Tue May 26 12:24:09 2026
```

Listing 5.7: Result of the turned off path

Listing 5.7 confirms that shutting down the peer connection successfully redirected traffic to the shorter AS-path with lower local preference.

5.3 Case 3, RIPE Atlas Probe 14521

5.3.1 Data

Hop	IP	Country
1	192.168.1.1	
2	192.0.0.1	
3	*	
4	*	
5	*	
6	*	
7	*	
255	192.36.148.17	I-root

Figure 5.5: Traceroute RIPE: Probe 14521 Tue May 12 10:13:49 CEST 2026

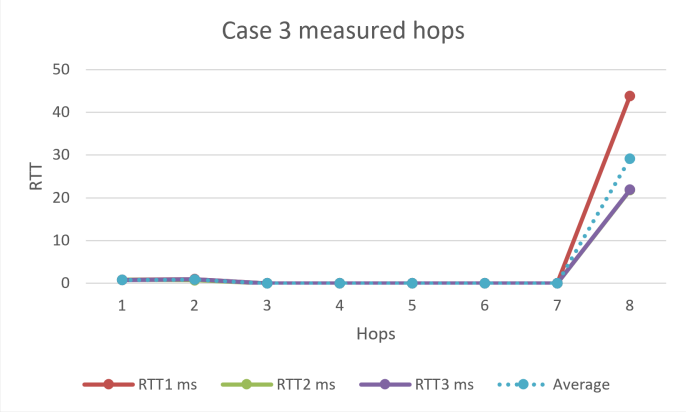

Figure 5.6: Visually presents the three signals RTT trend for each hop.

The traceroute in Figure 5.5 returns a hop count of 255, which is interpreted as the next valid hop and is therefore represented as hop 8 in Figure 5.6. The final hop shows a sharp RTT spike directly after the non-responsive routers along the path. What makes this case noteworthy is that the probe's AS, Deutsche Telekom AG (DTAG), is a tier 1 ISP. As such, it relies exclusively on peering

rather than transit and typically maintains direct access to infrastructure like the I-root [10]. Eight hops and elevated RTT are therefore unexpected, a tier 1 network operating within its primary region of central Europe should, in principle, reach the I-root with considerably fewer.

Table 5.3: Case 3 data compiled from multiple sources

Case 3 Compiled Data	
RIPE Atlas data[1]	Probe 14521 · ASN 3320 · Germany · Min RTT 21.8 ms · 8 hops · 2026-05-15 09:13:49
Traceroute path[2]	Germany (AS3320) → Netnod (8674) → I-root (AS29216) *(Possibly crossing DE-CIX)* [4]
Looking Glass	`Primary Prefix: 192.36.148.0/24` (Direct)[3][4] `Backup Prefix: 192.36.148.0/23` (Via AS1257)[4] `Next hop (Originator): 194.25.5.167`[3] `Local Pref: 100` [3]
AS path	`3320` → `8674` → `29216` (for /24)[3][4] `3320` → `1257` → `8674` → `29216` (for /23)[4]
BGP communities[5]	`3320:1276`[3] Imported in Germany (DTAG)[5] `3320:2010`[3] Imported in Europe (DTAG)[5] `3320:9010`[3] Imported from customer (DTAG)[5] `1257:50 1257:52 1257:3752 1257:9000` (Tele2, unknown meaning)[4] `64512:100` (Private AS, unknown meaning)[3]

[1] RIPE Atlas [2] Ripe Atlas CLI Traceroute [3] DTAG AS3320 BGP LG Table A.2
[4] `bgp.tools` [5] DTAG BGP Community documentation see Table A.3

The probe is located in Germany (see Figure 4.6). Based on the available data in Table 5.3, the network path connects directly to the I-root, which is corroborated by both the DTAG LG (Table A.2) and the traceroute (Figure 5.5). However, the non-responsive hops visible in Figure 5.5 introduce uncertainty regarding the precise traversal path. Additionally, bgp.tools

identified a possible IXP along the route (Table 5.3); as this does not influence the routing outcome, it was excluded from the simulation, though it could be added in future work to improve topological accuracy.

It is worth noting that bgp.tools presents different results because it uses different vantage points, reading from its nearest neighbors compared to DTAG LG with its own internal view, see Table A.2.

BGP communities were extracted from DTAG's official documentation [56] and are listed in the compiled data Table 5.3. These indicate that DTAG routes the traffic solely based on it being in Germany, within Europe, and from a customer. Analysis of the LG output indicates that additional BGP communities influence routing, as shown in Table A.2. However, the subsequent AS BGP communities are not publicly accessible. This lack of transparency obscures the BGP analysis, preventing definitive conclusions.

A further examination of the bgp.tools LG output Table 5.3 reveals additional routing details. It is observed that the system selects the more specific /24 prefix rather than the /23 prefix, which aligns with standard routing expectations. Consequently, traffic originating from Germany connects to DTAG and reaches the I-root directly. The /23 prefix, associated with Tele2 as specified by the ASN in the Table 5.3, is subsequently routed to I-root and may serve as a backup route if the primary path becomes unavailable. Another thing worth noting is that bgp.tools identified another IXP that the traffic passes through called Deutsche Commercial Internet Exchange (DE-CIX) which is likely that it has.

The comprehensive data collection indicates that although part of the path is identifiable, the complete route remains indeterminate. Traceroute results don't indicate any definitive path, while LG data provide a more definitive outcome. Therefore, the simulations will be based on the information extracted from bgp.tools and DTAG LG, specifically testing the routing based on the the /23 prefix and AS path.

5.3.2 Simulation

The simulation accurately represents network traffic in isolation. It has all the possible scenario available based on the bgp.tools LG, and it accurately chooses the best path. The DE-CIX was, however, not added, but could be for future simulations. Similar to previous cases, it serves as a visual help and does not affect the routing result.

```
client_observer# show bgp ipv4 unicast 192.36.148.0/24
BGP routing table entry for 192.36.148.0/24, version 4
Paths: (1 available, best #1, table default)
  Not advertised to any peer
```

```
  8674 29216
    194.25.5.167(originator) (metric 10) from originator(194.25.5.167) (194.25.5.167)
      Origin IGP, localpref 100, valid, internal, bestpath-from-AS 8674, best (First path received)
      Community: 3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000
      Last update: Thu May 21 12:05:12 2026
client_observer# show bgp ipv4 unicast 192.36.148.0/23
BGP routing table entry for 192.36.148.0/23, version 2
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  8674 29216
    194.25.5.166(test_originator) (metric 10) from test_originator(194.25.5.166) (194.25.5.166)
      Origin IGP, localpref 100, valid, internal, bestpath-from-AS 8674, best (First path received)
      Community: 3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000
      Last update: Thu May 21 12:05:12 2026
client_observer# show bgp ipv4 unicast 192.36.148.17
BGP routing table entry for 192.36.148.0/24, version 4
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  8674 29216
    194.25.5.167(originator) (metric 10) from originator(194.25.5.167) (194.25.5.167)
      Origin IGP, localpref 100, valid, internal, bestpath-from-AS 8674, best (First path received)
      Community: 3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000
      Last update: Thu May 21 12:05:12 2026
```

Listing 5.8: This code presents testing the specific prefix options and then testing if

As seen, if asking for /24 we get the path with the /24 prefix, if we only want the /23 it chooses the one with fewer hops and lastly if we just enter the address to I-root (without specified prefix) it will choose the longest matching prefix naturally as shown in Listing 5.8.

```
netnod# clear ip bgp 10.0.0.6
netnod# clear ip bgp 10.0.0.4
netnod# conf
netnod(config)# router bgp 8674
netnod(config-router)# neighbor 10.0.0.6 shutdown
netnod(config-router)# neighbor 10.0.0.4 shutdown
netnod(config-router)# exit
netnod(config)# exit
netnod# write
```

Listing 5.9: Presents the process of turning off the two direct paths ”Main path /24” and ”Test Path /23” forcing traffic to the longer path with prefix /23 [55]

To test how it will be forced to use the path with longer AS path and with /23 prefix. To start, all other peering options removed (see Listing 5.9), it will consequently be forced to peer with the /23 prefix path which is through Tele2 as shown in Listing 5.10.

```
originator# show ip bgp 192.36.148.17
BGP routing table entry for 192.36.148.0/23, version 10
Paths: (1 available, best #1, table default)
  Advertised to peers:
  transit_tele2(10.0.0.11) test_originator(194.25.5.166) client_observer(194.25.5.168)
  1257 8674 29216
    10.0.0.11(transit_tele2) from transit_tele2(10.0.0.11) (2.2.2.2)
      Origin IGP, valid, external, bestpath-from-AS 1257, best (First path received)
      Community: 3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000
      Last update: Thu May 21 12:05:01 2026
```

Listing 5.10: Output after /24 path was turned off

The output returns the longer path with prefix /23, which is to be expected.

```
netnod#  clear ip bgp 10.0.0.4
netnod# conf
netnod(config)#  router bgp 8674
netnod(config-router)# neighbor 10.0.0.4 shutdown
netnod(config-router)# exit
netnod(config)# exit
netnod# write
```

Listing 5.11: Shows the process of turning off the direct path with /24 prefix forcing traffic based on AS-path length given the two /23 options [55]

```
client_observer# show ip bgp 192.36.148.17
BGP routing table entry for 192.36.148.0/23, version 4
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  8674 29216
    194.25.5.166(test_originator) (metric 10) from test_originator(194.25.5.166) (194.25.5.166)
      Origin IGP, localpref 100, valid, internal, bestpath-from-AS 8674, best (First path received)
      Community: 3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000
      Last update: Thu May 21 12:16:37 2026
```

Listing 5.12: Output after /24 path was turned off and it chooses based on AS-path length.

When eliminating the direct contact between the Originator and Netnod, as demonstrated in Listing 5.11. This configuration compels selection between two /23 paths. After reverting previous changes in Listing 5.9 and applying those specified in Listing 5.11, the system selects either the Test Path /23, which has a shorter AS path (see Figure 4.7), or the Second Path /23, which routes through Tele2 (see Figure 4.7). Testing confirms that the path with the shorter AS-path is consistently chosen as shown in Listing 5.12.

5.4 Reliability Analysis

This research uses Netlab, a network simulation tool licensed under MIT. The project is implemented using YAML files and Free Range Routing (FRR) as the routing software, and therefore the program behaves deterministically and does not yield unpredictable results.

Although RIPE Atlas measurements can exhibit noise such as jitter, the data is generally stable and any notable fluctuations were treated as points of interest rather than errors. Other sources such as PeeringDB is more static in nature and therefore inherently more consistent.

5.5 Validity Analysis

This research has high validity because it uses FRR routing stack, which implements BGP that is RFC-compliant. The environment was further confirmed through baseline validation testing, which demonstrated that the simulated routers strictly adhere to the standard BGP Best Path Selection process.

Comparative Validity: Validity is further strengthened by grounding each simulation in real-world case studies, allowing the accuracy of the Netlab environment to be evaluated against observed routing behavior.

Abstraction: To keep the simulation manageable, private hops such as local routers were excluded from the topology. These are not relevant to the simulation as they are specific to the individual user's access point. Only globally significant hops were retained, as these are common across multiple users within the same region. **Transit vs. Peer Uncertainty:** The distinction between transit and peering is not always discernible from LG outputs, meaning some assumptions were made in the topology design. This introduces uncertainty regarding the commercial nature of certain network relationships, though it does not materially affect the logical routing flow, which remains the primary focus of the simulation.

Chapter 6

Discussion

Across all cases, the simulations accurately and adequately represented real-world scenarios, with outputs aligning with expected results based on the data collected. Custom configurations were occasionally required to achieve correct outputs, such as AS-specific BGP communities. The data provides a relatively comprehensive perspective, although certain aspects remain undisclosed and had to be estimated. The simulation proved to be a valuable tool, closely mirroring real-world cases and enabling safe testing in isolated environments. The overall process was effective but the primary challenges involved compiling accurate information, and additional relevant data emerged during the research process.

Although the simulations were generally accurate, certain cases remain unresolved, resulting in incomplete or insufficient simulation outcomes for those instances. For instance in Case 1, assumptions were necessary to determine as to why certain paths were less preferable. By assigning the local preference to 50 to the more latency-efficient path with fewer AS-hops as seen in Figure 4.3. LGs do not provide a complete view, instead offering only a limited selection of optimal paths, necessitating further estimations. Another example in Case 2, the rationale for including two Route Reflectors in the cluster list is not publicly specified, nor is the reason for the route propagating through both instead of just one. This ambiguity reduces the simulation's accuracy. This situation is expected because most visibility occurs within the operator's own AS, and the specifications of other ASes are typically not fully accessible.

This thesis aligns with previous research, confirming that BGP relies on "routing by rumor" and that BGP configurations significantly influence Anycast stability. The results corroborate previous research discussed in

Section 2.4, which found that multiple providers can negatively impact Anycast catchments [9]. This was demonstrated through the relationship between Orange and IIJ. Certain routes are prioritized for commercial reasons, sometimes at the expense of more effective options. Prior research also indicates that tier 1 providers generally achieve near-optimal results [9]; which is mostly true, but Case 3 demonstrated a less optimal outcome. It is hypothesized that the second path with the /23 was used, possibly as a result of operator shutdowns or route flapping, which led to the less preferred outcome.

This thesis also aligns with previous research in Section 2.4, which emphasized the importance of analyzing routing situations rather than attempting large-scale predictions. Given BGP's limited visibility and the challenges in predicting its behavior on a broader scale, it was necessary to reconstruct each step of the routing process incrementally. Consequently, the simulations served to provide a more complete representation.

However, maintaining the simulations presents challenges, as routing is dynamic and does not constitute a static structure.

Chapter 7

Conclusions and Future work

7.1 Conclusions

This thesis addressed two primary research questions: the extent to which publicly available data enables a comprehensive understanding of network dynamics, and the accuracy of simulations in representing real-world scenarios displaying suboptimal routing.

In response to the first question, sources such as RIPE Atlas, Autonomous System (AS)-specific Looking Glasses (LG), PeeringDB, and bgp.tools were, in these cases, sufficient though incomplete to compile a comprehensive understanding and to identify the Border Gateway Protocol (BGP) attributes influencing network dynamics. However, these sources do not provide a complete representation. Certain AS-specific configurations, including some BGP Communities, remain undisclosed and must be inferred. Despite these limitations, the combined use of these data sources was adequate for constructing well-founded hypotheses and simulation topologies.

In response to the second question, the simulations exhibited high accuracy in both visual and functional representations of each scenario. Local preference, longest-prefix match, and AS-Path length, were accurately modeled in terms of their influence on the route selection. Exceptions were necessary for custom provider configurations, such as hard-coded connections defined by BGP Communities. Additionally, the cluster list could not be fully replicated due to undisclosed Interior Gateway Protocol configurations. These discrepancies are attributable to data limitations rather than inherent simulation constraints.

The main contribution of this thesis is the demonstration that fragmented, publicly available data is sufficient to reconstruct a comprehensive represen-

tation suitable for simulating suboptimal routing scenarios. This validates network emulation as both a theoretical and practical tool for operators managing global Anycast catchments, enabling the analysis and testing of remediation strategies without impacting live infrastructure.

7.2 Limitations

An exclusive focus on BGP and inter-domain routing necessarily excludes other factors that influence network behavior, underscoring the need to study specific scenarios in a controlled environment. The research scope also shifted from examining BGP in strict isolation to analyzing real-world suboptimal routing cases, as isolated BGP testing yielded insufficiently meaningful results (see Figure 4.1). To summarize, the resulting data should be considered indicative rather than definitive because the network is continuously changing. Verification against live network conditions would require accounting for additional influencing factors, both hardware-related performance and accurate external ISP-specific policies.

7.3 Future work

For future scope, it would be beneficial to study BGP, and networks uisng simulations. Since Netlab proved to be extremely useful in concretizing the network, allowiwng engineers to research without affecting their network. It would also therefore be interesting to study Anycast catchments on a larger scale and examine how different BGP configurations affects overall catchment stability.

7.3.1 What has been left undone?

The study did not address potential improvements in the simulation that could enhance outcomes for each scenario. Furthermore, comprehensive research on the impact of BGP configurations on large-scale networks was not conducted. Such research would have been more feasible and less risky through simulation. Investigation into BGP vulnerabilities and the behavior of malicious actors within simulated environments also remains an area for future work.

7.3.1.1 Cost analysis

The simulations require significant computational resources, including both CPU and memory, depending on the hardware utilized. Access to Linux servers mitigates these performance challenges. In this research, no RIPE credits were used; however, future work may incur costs if such credits become necessary. Another challenge involves the development of each network topology, which can be time-consuming when structured from the ground up, particularly in assigning IP addresses. Establishing a base standard would streamline this process.

7.3.1.2 Security

As mentioned previously in possible future work, this thesis did not address BGP vulnerabilities and tested different malicious actors. In addition to the lack of research on BGP security, vulnerabilities in networks with undisclosed topology information may pose significant security risks. This underscores the critical importance of addressing the fragmented information structure.

7.3.2 Next obvious things to be done

The most immediate next steps would be to examine the impact of malicious actors on the various simulation cases. Subsequently, the study should focus on expanding the network and analyzing its behavior at larger scales.

7.4 Reflections

Evaluating modifications to multiple networks and BGP configurations in a controlled environment prior to deployment mitigates potential adverse impacts on the live network. This strategy reduces the likelihood of costly damages and minimizes the need for extensive troubleshooting. Reducing suboptimal routing can decrease the number of long geographical distance hops, which in turn lowers energy consumption and reduces the associated carbon footprint. From a social and ethical standpoint, it is important to have stable internet, especially when modern societies rely heavily on it. Simulations tools like Netlab allows for safety troubleshooting and preventions of damage. This work contributes to the United Nations (UN) Sustainable Development Goals (SDGs) number 9 by by enhancing the effectiveness of infrastructure research, with a particular focus on network systems. Future

efforts aim to achieve more stable, reliable, and resilient networks to further progress toward Goal 9.

References

[1] D. Conrad, "Brief overview of the root server system," ICANN, Tech. Rep. OCTO-010, May 2020, accessed: Apr. 8, 2026. [Online]. Available: https://icann-hamster.nl/ham/icann/octo/pub/octo-010-en.pdf.

[2] Netnod. (2026) i.root-servers.net. Netnod AB. Accessed: Apr. 12, 2026. [Online]. Available: https://www.netnod.se/i-root/i.root-servers.net

[3] "Domain names - concepts and facilities," RFC 1034, Nov. 1987. [Online]. Available: https://www.rfc-editor.org/info/rfc1034

[4] "Domain names - implementation and specification," RFC 1035, Nov. 1987. [Online]. Available: https://www.rfc-editor.org/info/rfc1035

[5] R. Allen. (2021) 25 years working for the good of the internet. Netnod AB. Accessed: 2026-04-10. [Online]. Available: https://www.netnod.se/blog/25-years-working-good-internet

[6] J. Abley and K. Lindqvist, "Operation of Anycast Services," Internet Requests for Comments, RFC 4786, December 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4786

[7] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," Internet Requests for Comments, RFC 4271, January 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4271

[8] R. Teixeira, T. G. Griffin, M. G. C. Resende, and J. Rexford, "Tie breaking: tunable interdomain egress selection," in *Proceedings of the 2005 ACM Conference on Emerging Network Experiment and Technology*, ser. CoNEXT '05. New York, NY, USA: Association for Computing Machinery, 2005. doi: 10.1145/1095921.1095935. ISBN 159593197X p. 93–104. [Online]. Available: https://doi.org/10.1145/1095921.1095935

[9] Z. Li, D. Levin, N. Spring, and B. Bhattacharjee, "Internet anycast: performance, problems, & potential," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '18. New York, NY, USA: Association for Computing Machinery, 2018. doi: 10.1145/3230543.3230547. ISBN 9781450355674 p. 59–73. [Online]. Available: https://doi.org/10.1145/3230543.3230547

[10] R. Zhang and M. Bartell, *BGP Design and Implementation*. Indianapolis, IN: Cisco Press, 2004. ISBN 978-1587051098

[11] V. Jain and B. Edgeworth, *Troubleshooting BGP: A Practical Guide to Understanding and Troubleshooting BGP*. Indianapolis, IN: Cisco Press, 2017. ISBN 978-1587144646

[12] M. Apostolaki, A. Singla, and L. Vanbever, "Performance-driven internet path selection," in *Proceedings of the ACM SIGCOMM Symposium on SDN Research (SOSR)*, ser. SOSR '21. New York, NY, USA: Association for Computing Machinery, 2021. doi: 10.1145/3482898.3483366. ISBN 9781450390842 p. 41–53. [Online]. Available: https://doi.org/10.1145/3482898.3483366

[13] A. Soliman, "Analysis of geography-oriented internet routing," Master's thesis, Carleton University, Ottawa, Ontario, 2024. [Online]. Available: https://carleton.scholaris.ca/server/api/core/bitstreams/6849764d-15c8-4845-9e3f-582459a25b74/content

[14] M. Zhou, X. Zhang, S. Hao, X. Yang, J. Zheng, G. Chen, and W. Dou, "Regional ip anycast: Deployments, performance, and potentials," in *Proceedings of the ACM SIGCOMM 2023 Conference*, ser. ACM SIGCOMM '23. New York, NY, USA: Association for Computing Machinery, 2023. doi: 10.1145/3603269.3604846. ISBN 9798400702365 p. 917–931. [Online]. Available: https://doi.org/10.1145/3603269.3604846

[15] M. Z. Ahmad and R. Guha, "Internet exchange points and internet routing," in *2011 19th IEEE International Conference on Network Protocols*, 2011. doi: 10.1109/ICNP.2011.6089065 pp. 292–294.

[16] C. Liu and P. Albitz, *DNS and BIND*, 5th ed. Sebastopol, CA, USA: O'Reilly, May 2006.

[17] Contributors to the kepler.gl project, "Kepler.gl: a powerful open source geospatial analysis tool for large-scale data sets," 2026, accessed: 2026-05-27. [Online]. Available: https://kepler.gl/

[18] Netnod. (2026) Supporting critical infrastructure at the core of the internet. Netnod AB. Accessed: 2026-04-10. [Online]. Available: https://www.netnod.se/about/Netnod

[19] Netnod, "We provide robustness for digitalised organisations," Official Website, 2026, accessed: 2026-04-10. [Online]. Available: https://www.netnod.se/

[20] "AS29216 Netnod - I-root - bgp.tools." [Online]. Available: https://bgp.tools/as/29216#connectivity

[21] N. Karamela and D. A. Karras, "A comparative analysis of ospf and eigrp routing protocol evaluation," *Journal of Transactions in Systems Engineering*, vol. 1, no. 2, p. 73–103, Jul. 2023. doi: 10.15157/JTSE.2023.1.2.73-103. [Online]. Available: https://journals.tultech.eu/index.php/jtse/article/view/116

[22] H. Smit and N. Shen, "Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels," RFC 3906, Oct. 2004. [Online]. Available: https://www.rfc-editor.org/info/rfc3906

[23] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A survey of bgp security issues and solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010. doi: 10.1109/JPROC.2009.2034031

[24] M. Caesar and J. Rexford, "Bgp routing policies in isp networks," *IEEE Network*, vol. 19, no. 6, pp. 5–11, 2005. doi: 10.1109/MNET.2005.1541715

[25] F. Streibelt, F. Lichtblau, R. Beverly, A. Feldmann, C. Pelsser, G. Smaragdakis, and R. Bush, "Bgp communities: Even more worms in the routing can," in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC '18. New York, NY, USA: Association for Computing Machinery, 2018. doi: 10.1145/3278532.3278557. ISBN 9781450356190 p. 279–292. [Online]. Available: https://doi.org/10.1145/3278532.3278557

[26] T. Li, R. Chandra, and P. S. Traina, "BGP Communities Attribute," RFC 1997, Aug. 1996. [Online]. Available: https://www.rfc-editor.org/info/rfc1997

[27] I. van Beijnum, *BGP*. O'Reilly Media, Inc., September 2002.

[28] J. A. Hawkinson and T. J. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)," RFC 1930, Mar. 1996. [Online]. Available: https://www.rfc-editor.org/info/rfc1930

[29] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the internet hierarchy from multiple vantage points," in *Proceedings.Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 2, 2002. doi: 10.1109/INFCOM.2002.1019307 pp. 618–627 vol.2.

[30] RIPE NCC, "What is an AS number?" n.d., accessed: Apr. 8, 2026. [Online]. Available: https://www.ripe.net/manage-ips-and-asns/as-numbers/.

[31] E. Chen, T. J. Bates, and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)," RFC 4456, Apr. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4456

[32] N. F. Mir, *Computer and Communication Networks*. Upper Saddle River, NJ, USA: Pearson, 2006.

[33] N. Chatzis, G. Smaragdakis, and A. Feldmann, "On the importance of internet exchange points for today's internet ecosystem," 2013. [Online]. Available: https://arxiv.org/abs/1307.5264

[34] X. Zhang, T. Sen, Z. Zhang, T. April, B. Chandrasekaran, D. Choffnes, B. M. Maggs, H. Shen, R. K. Sitaraman, and X. Yang, "Anyopt: predicting and optimizing ip anycast performance," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, ser. SIGCOMM '21. New York, NY, USA: Association for Computing Machinery, 2021. doi: 10.1145/3452296.3472935. ISBN 9781450383837 p. 447–462. [Online]. Available: https://doi.org/10.1145/3452296.3472935

[35] K. Sriram, D. Montgomery, D. R. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks," RFC 7908, Jun. 2016. [Online]. Available: https://www.rfc-editor.org/info/rfc7908

[36] N. Feamster, J. Winick, and J. Rexford, "A model of bgp routing for network engineering," *SIGMETRICS Perform. Eval. Rev.*, vol. 32, no. 1, p. 331–342, Jun. 2004. doi: 10.1145/1012888.1005726. [Online]. Available: https://doi.org/10.1145/1012888.1005726

[37] ——, "A model of bgp routing for network engineering," in *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '04/Performance '04. New York, NY, USA: Association for Computing Machinery, 2004. doi: 10.1145/1005686.1005726. ISBN 1581138733 p. 331–342. [Online]. Available: https://doi.org/10.1145/1005686.1005726

[38] B. Quoitin, C. Pelsser, O. Bonaventure, and S. Uhlig, "A performance evaluation of bgp-based traffic engineering," *International Journal of Network Management*, vol. 15, no. 3, pp. 177–191, 2005. doi: https://doi.org/10.1002/nem.559. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/nem.559

[39] netlab: a virtual networking labbing tool. [Online]. Available: https://netlab.tools/

[40] I. Pepelnjak, "ipspace/netlab," May 2026, original-date: 2020-12-09T18:22:22Z. [Online]. Available: https://github.com/ipspace/netlab

[41] "FRRouting." [Online]. Available: https://frrouting.org/

[42] "RIPE Atlas." [Online]. Available: https://www.ripe.net/analyse/internet-measurements/ripe-atlas/

[43] The RIPE Atlas Team, "RIPE Atlas Tools (Magellan) 3.3.1 documentation." [Online]. Available: https://ripe-atlas-tools.readthedocs.io/en/latest/

[44] M. Stubbig, "Looking Glass Command Set," RFC 8522, Feb. 2019. [Online]. Available: https://www.rfc-editor.org/info/rfc8522

[45] ping.pe, "About ping.pe." [Online]. Available: https://bgp-lg.ping.pe/about

[46] "PeeringDB." [Online]. Available: https://www.peeringdb.com/

[47] "Root Server Technical Operations Association." [Online]. Available: https://root-servers.org/

[48] “BGP.Tools.” [Online]. Available: https://bgp.tools/

[49] “Webupdates.” [Online]. Available: https://www.ripe.net/manage-ips-and-asns/db/webupdates

[50] “Whois search | APNIC.” [Online]. Available: https://www.apnic.net/about-apnic/whois_search/

[51] “netlab/tests/integration/routing at dev · ipspace/netlab.” [Online]. Available: https://github.com/ipspace/netlab/tree/dev/tests/integration/routing

[52] “Command Line Interface — FRR latest documentation.” [Online]. Available: https://docs.frrouting.org/projects/dev-guide/en/latest/cli.html

[53] “IIJ America reinforce their global services through the Dunant cable,” https://wholesale.orange.com/international/en/news/iij-america-reinforce-their-global-services-through-the-dunant-cable.html, jan 2022, accessed: 2026-05-17.

[54] “BGP Configuration Module — netlab documentation.” [Online]. Available: https://netlab.tools/module/bgp/

[55] “BGP — FRR latest documentation.” [Online]. Available: https://docs.frrouting.org/en/latest/bgp.html#clicmd-neighbor-PEER-shutdown-message-MSG

[56] Deutsche Telekom Technik GmbH, “Bgp signaling for as3320 ip-transit / ip-transit light,” September 2025, version 2.2. Accessed: 2026-05-16. [Online]. Available: https://wholesale.telekom.com/fileadmin/user_upload/documents/global/downloads/BGP_Signaling_for_AS3320.pdf

[57] D. T. T. GmbH, “Bgp signaling for as3320,” Sep. 2025, accessed: 2026-06-02. [Online]. Available: https://wholesale.telekom.com/fileadmin/user_upload/documents/global/downloads/BGP_Signaling_for_AS3320.pdf

Appendix A

BGP LG

A.1 Case 1 Orange Looking Glass

Table A.1: Orange BGP LG routing result to Netnod I-root (192.36.148.0/24)

Path	BGP Nexthop	BGP Prefix	AS Path	Communities	RPKI	Metrics	LocPref
1 ⋆	172.25.4.81	192.36.148.0/24	2497 8674 29216	5511:500, 5511:515, 5511:999, 5511:1000, 5511:2000, 5511:3000, 5511:10009, 5511:30228, 5511:30232, 5511:30535, 5511:30594, 5511:40176	none	IGP	100
2	172.25.4.81	192.36.148.0/24	2497 8674 29216	5511:500, 5511:515, 5511:999, 5511:1000, 5511:2000, 5511:3000, 5511:10009, 5511:30228, 5511:30232, 5511:30535, 5511:30594, 5511:40176	none	IGP	100
3	172.25.4.81	192.36.148.0/24	2497 8674 29216	5511:500, 5511:515, 5511:999, 5511:1000, 5511:2000, 5511:3000, 5511:10009, 5511:30228, 5511:30232, 5511:30535, 5511:30594, 5511:40176	none	IGP	100

A.2 Case 3 DTAG Looking Glass

Table A.2: DTAG AS3320 to Netnod I-Root (192.36.148.0/24)

Vantage Point	Prefix	RD	AS Path	Local Pref	MED	BGP Next Hop	BGP Originator	BGP Communities
Berlin, DE	192.36.148.0/24	–	8674 29216	100	0	194.25.5.167	194.25.5.167	(3320,1276) (3320,2010) (3320,9010) (8674,200) (64512,100) (64900,31201) (64901,10000)
Düsseldorf, DE	192.36.148.0/24	–	8674 29216	100	0	194.25.5.167	194.25.5.167	(3320,1276) (3320,2010) (3320,9010) (8674,200) (64512,100) (64900,31201) (64901,10000)
Frankfurt, DE	192.36.148.0/24	–	8674 29216	100	0	194.25.5.167	–	(3320,1276) (3320,2010) (3320,9010) (8674,200) (64512,100) (64900,31201) (64901,10000)
Munich, DE	192.36.148.0/24	–	8674 29216	100	0	194.25.5.167	194.25.5.167	(3320,1276) (3320,2010) (3320,9010) (8674,200) (64512,100) (64900,31201) (64901,10000)
Hamburg, DE	192.36.148.0/24	–	8674 29216	100	0	194.25.5.167	194.25.5.167	(3320,1276) (3320,2010) (3320,9010) (8674,200) (64512,100) (64900,31201) (64901,10000)

Table A.3: DTAG AS 3320 BGP Communities [57]

Code	Description	Tag
3320:1276	Imported in Germany	tag.origin.country.de
3320:2010	Imported in Europe	tag.origin.region.eu
3320:9010	Imported from customer	tag.origin.type.customer

Appendix B

YAML-Code

B.1 Test-Case: Local Preference

```
provider: clab

defaults.device: frr



module: [ bgp, ospf, routing ]
plugin: [ bgp.policy ]

defaults.attributes.node.bgp:

routing.policy:
  locpref_policy:
    locpref: 101


nodes:
  r1:
    bgp.rr: true
  r2:
    bgp.rr: true
  r3:
    bgp.rr: true
  r4:

    bgp.rr: true
  r5:
    bgp.rr: true
  r6:

    bgp.rr: true
  r7:

    bgp.rr: true
  r8:

    bgp.rr: true
  r9:

    bgp.rr: true
  r10:

    bgp.rr: true
  r11:

    bgp.rr: true
  r12:
    bgp.rr: true
    bgp.as: 1234
```

```
50      bgp.originate:
51        - ipv4: 1.2.3.0/24
52
53  bgp:
54    as: 5555
55
56  links:
57    - r1:
58      r2:
59
60    - r1:
61      r3:
62
63    - r3:
64      r4:
65    - r4:
66      r5:
67    - r4:
68      r2:
69
70
71    - r5:
72      r2:
73    - r5:
74      r7:
75    - r5:
76      r6:
77
78    - r6:
79      r7:
80    - r6:
81      r8:
82
83    - r7:
84      r5:
85    - r7:
86      r9:
87
88    - r8:
89      r9:
90    - r8:
91      r10:
92
93    - r10:
94      r8:
95    - r10:
96      r11:
97
98    - r11:
99      r12:
100
101   - r12:
102     r6:
103
104   - r12:
105     r1:
106       bgp.policy.in: locpref_policy
```

Listing B.1: Test-case local preference

B.2 Case 1

```
version: 2.0.0
provider: clab

defaults.device: frr
defaults.attributes.node.bgp:

module: [ bgp, ospf, routing ]
plugin: [ bgp.policy ]

routing.policy:
  in_5511:
    set.community.standard: '5511:500 5511:515 5511:999 5511:1000 5511:2000 5511:3000 5511:10009
         5511:30228 5511:30232 5511:30535 5511:30594 5511:40176'
  depref_netnod:
    locpref: 50

nodes:

#---AS 3215 Orange France---

  start_paris:
    module: [ bgp, ospf ]
    bgp.as: 3215
    bgp.rr: true
    bgp.router_id: 92.184.116.217
    loopback:
      ipv4: 92.184.116.217/32

#---AS 5511---Orange transit London

  transit_london:
    module: [ bgp, ospf, routing ]
    bgp.as: 5511
    bgp.rr: true
    bgp.router_id: 193.251.241.157
    loopback:
      ipv4: 193.251.241.157/32



#---AS 2497--- Internet Initiative Japan
  iij_tokyo:
    module: [ bgp, ospf, routing ]
    bgp.as: 2497
    bgp.router_id: 58.138.83.170
    loopback:
      ipv4: 58.138.83.170/32


 #----



 # Netnod AS 8674
  netnod:
    module: [ bgp ]
    bgp.as: 8674
    bgp.router_id: 202.249.2.179
    loopback:
      ipv4:  202.249.2.179/32




  # I-root server - AS 29216
  iroot:
    module: [ bgp ]
    bgp.as: 29216
    bgp.router_id: 192.36.148.17
    loopback:
      ipv4: 192.36.148.17/32
    bgp.originate:
      - ipv4: 192.36.148.0/24


```

```
  # IXP Tokyo - IXP marker (visual representation only)
  ixp_tokyo:
    module: [ ]
    loopback: false
 
 
 
 
 
links:
 
 
 # Alternative: direct London -> netnod
  - transit_london:
      ipv4: 10.0.0.0/31
      bgp.policy.in: depref_netnod
    netnod:
      ipv4: 10.0.0.1/31
 
 
  # Paris -> London -> Tokyo -> IXP Tokyo -> netnod -> iroot
 
  - start_paris:
      ipv4: 10.0.0.2/31
    transit_london:
      ipv4: 10.0.0.3/31
 
  - transit_london:
      ipv4: 10.0.0.4/31
      bgp.policy.in: in_5511
    iij_tokyo:
      ipv4: 10.0.0.5/31
 
  # IXP Tokyo fabric between Tokyo and Netnod
  - iij_tokyo:
      ipv4: 10.0.1.1/24
    netnod:
      ipv4: 10.0.1.2/24
    ixp_tokyo:
      ipv4: false
      ipv6: false
 
  - netnod:
      ipv4: 10.0.0.6/31
    iroot:
      ipv4: 10.0.0.7/31
```

Listing B.2: Case 1 topology

B.3 Case 2

```yaml
version: 2.0.0
provider: clab

defaults.device: frr

module: [ bgp, ospf, routing ]
plugin: [ bgp.policy ]

routing.policy:
  in_3303:
    locpref: 4000
    set.community.standard: '15557:90 15557:300 15557:500'
  drop_all:
    action: deny

nodes:
  # Route Reflector Cluster 1
  rr1_primary:
    module: [ bgp, ospf ]
    bgp.as: 15557
    bgp.rr: true
    bgp.rr_cluster_id: 194.6.134.163
    loopback:
      ipv4: 194.6.134.163/32

  # Route Reflector Cluster 2
  rr2_secondary:
    module: [ bgp, ospf ]
    bgp.as: 15557
    bgp.rr: true
    bgp.rr_cluster_id: 194.6.157.2
    loopback:
      ipv4: 194.6.157.2/32

  originator:
    bgp.as: 15557
    bgp.router_id: 194.6.130.2
    loopback:
      ipv4: 194.6.130.2/32


  test_originator:
    bgp.as: 15557
    bgp.router_id: 194.6.130.3
    loopback:
      ipv4: 194.6.130.3/32


 # Netnod AS 8674
  netnod:
    module: [ bgp ]
    bgp.as: 8674
    bgp.router_id: 217.29.66.57
    loopback:
      ipv4: 217.29.66.57/32


  # Transit AS 3303
  transit_swisscom:
    module: [ bgp ]
    bgp.as: 3303
    bgp.router_id: 194.6.130.3
    loopback:
      ipv4: 194.6.130.3/32


  # I-root server AS 29216
  iroot:
    module: [ bgp ]
    bgp.as: 29216
    bgp.router_id: 192.36.148.17
    loopback:
      ipv4: 192.36.148.17/32

    bgp.originate:
      - ipv4: 192.36.148.0/24
```

```yaml

  # MIX-IT Milan IXP marker
  mix:
    module: [ ]
    loopback: false

  # Client router receiving the route
  client_observer:
    module: [ bgp, ospf ]
    bgp.as: 15557
    bgp.router_id: 194.6.150.1
    loopback:
      ipv4: 194.6.150.1/32


links:

  - rr1_primary:
      ipv4: 10.0.0.0/31
    rr2_secondary:
      ipv4: 10.0.0.1/31

  # RR to originator
  - rr2_secondary:
      ipv4: 10.0.0.2/31
    originator:
      ipv4: 10.0.0.3/31

  # RR to test_originator
  - rr1_primary:
      ipv4: 10.0.0.4/31
    test_originator:
      ipv4: 10.0.0.5/31
  - rr2_secondary:
      ipv4: 10.0.0.6/31
    test_originator:
      ipv4: 10.0.0.7/31

  # RR to  client_observer
  - rr1_primary:
      ipv4: 10.0.0.8/31
    client_observer:
      ipv4: 10.0.0.9/31

  # originator to AS 3303, with localpref 4000
  - originator:
      ipv4: 10.0.0.10/31
      bgp.policy.in: in_3303
    transit_swisscom:
      ipv4: 10.0.0.11/31

  # MIX-IT Milan - shared L2 IXP
  - transit_swisscom:
      ipv4: 10.0.1.1/24
    netnod:
      ipv4: 10.0.1.2/24
    mix:
      ipv4: false
      ipv6: false

  # test_originator peers with AS 8674 directly, localpref 100
  - test_originator:
      ipv4: 10.0.0.14/31
      bgp.policy.out: drop_all # to avoid advertising to the AS8674, otherwise it risks loops.
    netnod:
      ipv4: 10.0.0.15/31


  # AS 8674 to AS 29216
  - netnod:
      ipv4: 10.0.0.16/31
    iroot:
      ipv4: 10.0.0.17/31
```

Listing B.3: Case 2 topology

B.4 Case 3

```
version: 2.0.0
provider: clab

defaults.device: frr

module: [ bgp, ospf, routing ]
plugin: [ bgp.policy ]

routing.policy:
  allow_23:
    - match.prefix: root_23
      action: permit
    - action: deny
  allow_24:
    - match.prefix: root_24
      action: permit
    - action: deny

  # BGP communities
  in_3320:
    set.community.standard: '3320:1276 3320:2010 3320:9010 8674:200 64512:100 64900:31201 64901:10000'

routing.prefix:
  root_23:
    - ipv4: 192.36.148.0/23
  root_24:
    - ipv4: 192.36.148.0/24



nodes:


# ---AS3320---

  # observer node within AS3320
  client_observer:
    bgp.as: 3320
    bgp.router_id: 194.25.5.168
    loopback:
      ipv4: 194.25.5.168/32


# Exit point within AS3320
  originator:
    bgp.as: 3320
    bgp.router_id: 194.25.5.167
    loopback:
      ipv4: 194.25.5.167/32


  # test exit point within AS3320
  test_originator:
    bgp.as: 3320
    bgp.router_id: 194.25.5.166 # test ip
    loopback:
      ipv4: 194.25.5.166/32 # test ip


  # Transit via Tele2 (roundabout path)
  transit_tele2:
    bgp.as: 1257
    bgp.router_id: 2.2.2.2 #placeholder ip
    loopback:
      ipv4: 2.2.2.2/32 #placeholder ip



 # Netnod
  netnod:
    bgp.as: 8674
    bgp.router_id: 1.2.3.4 #placeholder ip
    loopback:
      ipv4: 1.2.3.4/32 #placeholder ip


```

```
77
78    # I-root server
79    iroot:
80      bgp.as: 29216
81      bgp.router_id: 192.36.148.17
82      loopback:
83        ipv4: 192.36.148.17/32
84      bgp.originate:
85        - ipv4: 192.36.148.0/24
86        - ipv4: 192.36.148.0/23
87
88
89
90
91
92
93  links:
94    # Internal AS3320 links
95    - client_observer:
96        ipv4: 10.0.0.0/31
97      originator:
98        ipv4: 10.0.0.1/31
99
100   - client_observer:
101       ipv4: 10.0.0.2/31
102     test_originator:
103       ipv4: 10.0.0.3/31
104
105   # Direct connection from originator to Netnod /24
106   - originator:
107       ipv4: 10.0.0.4/31
108       bgp.policy.in: in_3320
109     netnod:
110       ipv4: 10.0.0.5/31
111       bgp.policy.out: allow_24
112
113   # Roundabout path from AS3320 to AS8674 via Tele2
114   - originator:
115       ipv4: 10.0.0.10/31
116       bgp.policy.in: in_3320
117     transit_tele2:
118       ipv4: 10.0.0.11/31
119       bgp.policy.in: allow_23
120
121   # Tele2 to Netnod
122   - transit_tele2:
123       ipv4: 10.0.0.12/31
124       bgp.policy.in: allow_23
125     netnod:
126       ipv4: 10.0.0.13/31
127       bgp.policy.out: allow_23
128
129   # Netnod advertises /23 to test_originator
130   - test_originator:
131       ipv4: 10.0.0.6/31
132       bgp.policy.in: in_3320
133     netnod:
134       ipv4: 10.0.0.7/31
135       bgp.policy.out: allow_23
136
137   # Netnod to iroot
138   - netnod:
139       ipv4: 10.0.0.8/31
140     iroot:
141       ipv4: 10.0.0.9/31
```

Listing B.4: Case 3 topology

TRITA-EECS-EX-2026:397
Stockholm, Sweden 2026

www.kth.se