Risk Level as a Service: Enhancing BGP Security through IXP-Based Risk Assessment

Stefano Servillo*, Angelo Pio Pompeo*, Pietro Spadaccino*, Francesca Cuomo*
*Department of Information engineering, Electronics and Telecommunications (DIET)
University of Rome "Sapienza", Italy, e-mail: {name.surname}@uniroma1.it

***Abstract*—The Border Gateway Protocol (BGP) is the de-facto standard for inter-domain routing, enabling Autonomous Systems (ASes) to exchange reachability information. BGP's trust-based design makes it vulnerable to configuration errors and malicious activities, disrupting global connectivity between ASes. Internet Exchange Points (IXPs), components of today's Internet enabling efficient interconnection among networks, typically implement countermeasures, enforcing filtering policies based on Internet Routing Registries (IRRs) and the Resource Public Key Infrastructure (RPKI), which protect the IXP's Route Server (RS) and members. However, IXP members may still receive malicious routes through other BGP connections and, without proper filtering, select them as best paths. A possible solution for IXP members is to deploy risk-based monitoring tools, as in compliance with the U.S. Roadmap to Enhancing Internet Routing Security. However, implementing such tools individually on each AS poses scalability challenges and adds complexity. In this work, we introduce Risk Level-as-a-Service (RLaaS), a novel security service that IXPs can provide to their members to enhance inter-domain routing resilience. RLaaS dynamically assigns a Risk Level (RL) to ASes by analyzing malicious announcements observed at the RS, potentially indicative of prefix hijacking. The RL value is distributed to IXP members, enabling them to incorporate it into their local routing policies to penalize high-risk ASes. RLaaS is intended to complement the partial deployment of RPKI, offering an additional layer of security. Through a proof-of-concept, we demonstrate that RLaaS can reduce the propagation of malicious announcements and improve the robustness of the inter-domain routing ecosystem.**

***Index Terms*—BGP, IXP, Routing security, Risk Level**

I. Introduction

The Border Gateway Protocol (BGP) serves as the de facto inter-domain routing protocol of the Internet [1]. It enables Autonomous Systems (ASes) - networks identified by unique AS Numbers (ASNs) - to exchange reachability information and enforce routing policies that reflect business agreements and operational constraints. However, BGP's original design did not incorporate security considerations, leaving it without inherent mechanisms to authenticate or validate routing information. As a result, the protocol depends on implicit trust that each AS advertises only those prefixes it legitimately owns [2].

This absence of security validation makes BGP inherently vulnerable to incorrect or malicious announcements. Two well-known threats are *prefix hijacking*, where an AS illegitimately advertises another AS's IP prefixes [3], and *route leaks*, where routes are propagated beyond intended policy boundaries [4].

ISBN 978-3-903176-82-9 © 2026 IFIP

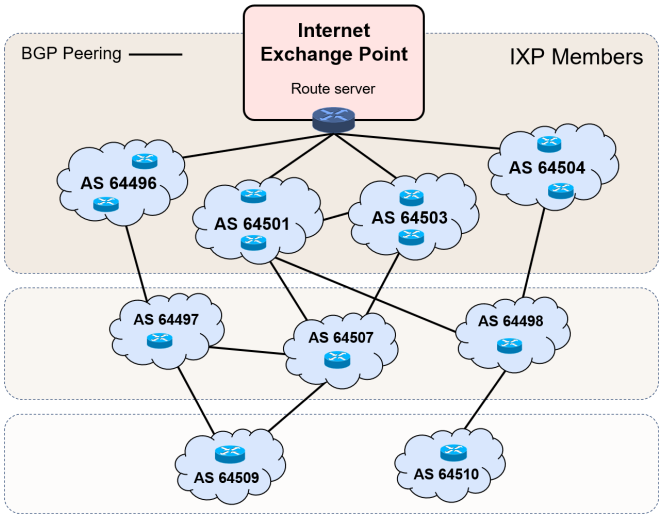

Fig. 1. Illustration of an IXP where multiple ASes interconnect via a Route Server to facilitate multilateral BGP peering.

Both can lead to misdirected, intercepted, or dropped traffic, with significant regional or even global impact [5].

These issues have repeatedly occurred in practice. In February 2008, Pakistan Telecom (PTCL) mistakenly hijacked YouTube's IP prefix while enforcing national access restrictions, resulting in a global service outage for nearly two hours [6]. More recently, in February 2022, AS 9457 advertised prefixes belonging to KaKaoTalk, a major South Korean messaging platform [7]. Attackers exploited prefix hijacking to redirect traffic to a malicious website, stealing approximately 1.9 million USD in cryptocurrency over a two-hour window.

Over the past decades, the networking community has proposed numerous solutions to mitigate these vulnerabilities. However, only two mechanisms are currently deployed at a global scale: *Internet Routing Registries (IRRs)* [8] and *Resource Public Key Infrastructure (RPKI)* [9]. IRRs allow operators to publish route objects specifying which ASes are authorized to originate specific prefixes. In contrast, RPKI introduces a cryptographic foundation for *Route Origin Validation (ROV)* using digitally signed *Route Origin Authorizations (ROAs)* that bind IP prefixes to authorized ASNs.

BGP security is critical in an Internet Exchange Point (IXP) [10], which acts as neutral interconnection hub where

multiple ASes exchange traffic via a Route Server (RS) [11]. Figure 1 illustrates a typical IXP topology in which several BGP sessions with a central RS are maintained, rather than numerous bilateral peering links. This configuration simplifies multilateral interconnection and improves scalability and operational efficiency. To preserve the integrity of routing information, IXP operators typically implement both basic filters (e.g., on prefix or AS path length) and advanced filters using IRR and RPKI data [12], thereby ensuring that only valid announcements are propagated within the exchange. However, since IXP members also connect to other ASes outside the exchange, they remain exposed to malicious or misconfigured announcements that can propagate through external links, undermining local and global routing stability.

Several studies have aimed to enhance network security. For instance, [13] applies unsupervised deep learning to jointly embed ASNs and prefixes into a vector space, detecting anomalous routing events by analyzing deviations in cosine similarity. Nevertheless, such systems rely on data from public route collectors, which can be strategically evaded by attackers to conceal malicious activity [14], thereby limiting their effectiveness in detecting real-time routing anomalies.

Contributions. This paper aims to enhance BGP security at IXPs by proposing the following contributions:

- Definition of the Risk Level as a Service (RLaaS) model for IXPs, intended to complement the current partial deployment of RPKI.
- Specification of the security actions available to IXP members based on RL metrics.
- Evaluation of the impact of RLaaS on inter-domain routing security.
- Evaluation of a Proof of Concept (PoC) demonstrating the feasibility and benefits of the proposed service.

The rest of the paper is organized as follows. Section II summarizes the main motivations underlying this work, while Section III reviews the state of the art. The RLaaS system is described and further analyzed in Section IV for its advantages and limitations. The PoC is presented in Section V, and Section VI concludes the paper.

II. Motivations

In this section, we examine the operational realities motivating our study: incomplete RPKI adoption, the role of risk-based planning, and the strategic position of IXPs in strengthening inter-domain routing resilience.

Slow and incomplete RPKI adoption. The RPKI enhances routing security by enabling validation of prefix-origin pairs through ROAs and supporting ROV. Its effectiveness, however, depends on widespread deployment.

Recent measurements from the NIST RPKI Monitor [1] show substantial progress in RPKI deployment, with approximately 60% of unique IPv4 and IPv6 prefix-origin AS pairs now covered by valid ROAs. For IPv4, this marks a clear inflection point compared to historical stagnation, while IPv6 adoption has advanced at a more gradual yet consistently positive pace. Despite this progress, the deployment of ROV remains limited. Empirical analysis [15], [16] report that 61.3% of non-stub ASes do not perform ROV at any customer-facing interfaces, and 36.2% fail to apply ROV on any interfaces within their networks. Prior works [16], [17] attribute this gap to limited awareness of routing security risks, competing operational priorities, insufficient resources, and administrative complexities. As a result, invalid or malicious announcements may still propagate despite ROA availability.

Risk-Based Planning. To address the limitations posed by the partial deployment of RPKI, the *U.S. Roadmap to Enhancing Internet Routing Security* [18] recommends that network operators *"develop, maintain, and periodically update a cybersecurity risk management plan"*. Operators are encouraged to integrate routing security and resilience into their broader cybersecurity assessments and operational practices. Consistent with this approach, the *OECD* stresses that *"an understanding and appreciation of this risk is essential for policy makers as they consider policies or measures aimed at improving the security of communication networks"* [19]. These challenges emphasize the importance of complementary mechanisms that reinforce BGP's security posture. In particular, external monitoring and detection frameworks can provide valuable visibility into suspicious routing events that evade incomplete RPKI enforcement.

Centralized security at IXPs. Deploying risk-based monitoring frameworks at the individual AS level introduces scalability and visibility challenges. IXPs offer a more efficient alternative, acting as a centralized vantage points that observe diverse BGP announcements. With hundreds of IXPs operating globally, they are well positioned to support collective security functions [20]. IXPs already provide value-added services that leverage their central position in the routing ecosystem, such as *Remote Blackhole Triggering (RBH)* [21] to mitigate *Distributed Denial-of-Service (DDoS)* attacks and selective peering mechanisms at RSes. Extending this model to include risk-level analysis represents a natural evolution of IXP functions, enabling improved security for member ASes while reducing operational overhead.

Overall, the persistent gap between ROA publication and ROV enforcement underscores the limitations of AS-level security measures. To address this, we investigate how the integration of risk-based planning with IXP-centric monitoring can complement existing RPKI mechanisms and establish a cooperative framework for enhancing global routing resilience.

III. Background and Related Work

Inter-domain routing security has been extensively studied across multiple dimensions, including prefix hijacking detection, operational best practices, and the role of IXPs in enhancing BGP security; this section reviews key contributions in these areas and situates the proposed RLaaS framework within this research landscape.

Prefix hijacking detection. Several systems have been developed to detect and mitigate prefix hijacking in inter-domain

[1]NIST RPKI Monitor available at https://rpki-monitor.antd.nist.gov/ROV

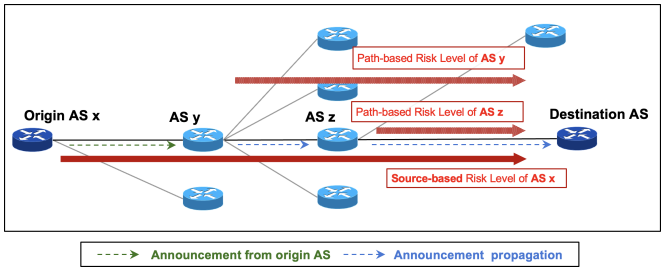

Fig. 2. Overview of the RL framework proposed in [28]. The IXP computes the source- and path-based RL metrics for an AS, based on mathematical calculations and observation of the RS RIB.

routing. Authors in [13] apply unsupervised deep learning to embed ASNs and prefixes into a 32-dimensional vector space, detecting anomalies by measuring deviations in cosine similarity. Similarly, authors in [22] integrate fast anomaly detection with automated mitigation, identifying suspicious AS links and filtering benign events before applying countermeasures. While these systems primarily protect an operator's prefixes, they are limited by their reliance on public route collectors. The work in [14] demonstrated that adversaries can evade detection by suppressing malicious announcements toward collector vantage points while continuing to propagate them across other parts of the Internet.

Security best practices. Beyond technical solutions, community initiatives promote operational best practices to enhance routing security. The *Mutually Agreed Norms for Routing Security (MANRS)* initiative defines specific actions for network operators, including route filtering, anti-spoofing, coordination, and validation [23]. Empirical evidence [24] indicates that MANRS-compliant networks experience significantly fewer routing anomalies, underscoring the initiative's practical value. Nonetheless, widespread challenges persist: many operators have yet to deploy RPKI validation, and large-scale configuration management remains complex. Consequently, even MANRS-compliant networks may inadvertently propagate misconfigurations or malicious announcements.

IXP-based Security. Several studies have examined the role of IXPs in enhancing inter-domain routing security. The work in [25] demonstrated that large IXPs, by observing traffic from a majority of global ASes, serve as ideal vantage points for detecting routing anomalies and validating control-plane information. Authors in [26] analyzed over 34,000 RTBH events at a major European IXP, showing that while RTBH can mitigate DDoS attacks, it may also cause collateral damage. The authors in [27] proposed Spoofer-IX, a methodology leveraging IXP traffic to infer spoofed packets, revealing gaps in source-address validation between 2017 and 2019. Collectively, these works highlight IXPs' potential as cooperative infrastructures for improving inter-domain routing security.

A. Risk level

The concept of RL in the context of BGP security was introduced by Servillo et al. [28]. Their framework analyzes the *Routing Information Base (RIB)* of an RS to detect prefix hijacking and to estimate the threat posed by ASes to inter-domain routing from the perspective of IXPs. Two complementary metrics were defined: the *Source-Based Risk Level* RL_{SB} and the *Path-Based Amplification Level* RL_{PB}. The first quantifies the risk of Origin ASes announcing attack-related prefixes, whereas the second evaluates the amplification potential of transit ASes that re-propagate them, as summarized in Fig. 2. An AS with extensive connectivity that redistributes a malicious prefix can significantly amplify the scope and impact of an attack. For completeness, we briefly restate the definition introduced in [28].

The RL_{SB} metric is derived from five factors: Time (T, Frequency (F), Path (P), Filtering (G), and Owner (O), each ranging within [0, 1]. These factors are combined through a weighted linear aggregation that reflects their importance from the IXP perspective. The aggregation is given by:

$$RL_{SB} = T \cdot \alpha + F \cdot \beta + P \cdot \gamma + G \cdot \delta + O \cdot \epsilon$$

The weights are obtained using the Analytic Hierarchy Process (AHP), which relies on expert-driven pairwise comparisons among the five categories to derive a consistent priority vector [28].

The RL_{PB} metric, on the other hand, assesses the role of ASes along the AS-PATH based on their connectivity relative to the Origin AS. It is calculated by comparing the number of connections of an AS along the AS-PATH (C_a), and the number of connections of the Origin AS (C_o):

$$RL_{PB} = \frac{C_a}{C_o} \cdot RL_{SB}$$

Unlike RL_{SB}, the RL_{PB} value is not bounded within [0, 1], allowing it to capture the amplification effect of large, well-connected networks in disseminating high-risk prefixes. For further information about the calculation of RL_{SB} and RL_{PB} we refer the reader to [28].

From a computational perspective, RL metrics are derived from a snapshot of the RS RIB, which is significantly smaller than the stream of BGP UPDATE messages, as multiple updates can correspond to the same prefix. The RL computation relies on fixed attribute sets and external datasets accessed via constant-time lookups. The computation of the RL metrics is applied only to the subset of prefixes classified as potentially attack-related. Even in the worst case, where all prefixes are deemed suspicious, the computation involves a linear scan over the RIB snapshot, avoiding nested iterations across prefixes or ASes. As a result, the approach is computationally efficient and suitable for deployment in large-scale IXPs.

Although these metrics provide valuable insight into routing behavior, authors in [28] did not propose any operational use for them. The RL values are computed and maintained internally by IXPs solely for awareness and record-keeping. In this work, we extend the utility of the RL framework by integrating its values into an automated system that actively strengthens the security of inter-domain routing and complements the existing partial deployment of RPKI.

IV. Risk Level as a Service

In this section, we introduce RLaaS, a security service enabling IXP members to enhance their local routing defense mechanisms using RL metrics computed at the IXP, without violating net neutrality. RLaaS provides authenticated risk information on high-risk ASes in near real time, allowing members to make informed routing decisions. The service requires only minimal client-side configuration changes and does not modify the RS architecture or its operational behavior.

To describe the functioning of the RLaaS, we refer to the classical scenario where an IXP operates an RS and multiple ASes establish BGP peering sessions with it. The RS collects BGP `UPDATE` messages from all peers, propagating valid announcements and filtering invalid ones according to its policy. As outlined in [28], the IXP periodically computes the RL metric for each AS using data extracted from the RS RIB. This computation yields two floating-point values (RL_{SB}, RL_{PB}) associated with the corresponding AS.

In the remainder of this paper, we refer to a *subscriber* AS (sAS) as an IXP member using the RLaaS service, and a *non-subscriber* AS (nsAS) as one that does not. For the sake of simplicity, we consider the case in which an AS is connected to only one IXP. The scenario where an AS is connected to multiple IXPs, thus receiving different values (RL_{SB}, RL_{PB}) from each IXP, will be considered as future work.

RLaaS exchanges information using tuples of the form $(ASN, MetricType, RL)$, where each tuple identifies the evaluated AS, the risk metric type $(MetricType \in (RL_{SB}, RL_{PB}))$, and the corresponding RL value. Normalization is performed at the IXP before dissemination: RL_{SB}, bounded in [0,1], is linearly mapped to a discrete integer scale in [0,10], while RL_{PB}, which is unbounded, is normalized using a transformation function (e.g., logarithmic or min–max normalization) to the same range, ensuring comparability across subscribers. Each normalized tuple is securely transmitted from the IXP to sASes via mutually authenticated TLS sessions, with additional digital signatures to guarantee integrity and origin authenticity.

Figure 3 illustrates this architecture. The RS detects a prefix hijack originating from AS1 and computes the two RL values. These values are transmitted to the sASes (AS2 and AS3) through the RLaaS interface, allowing them to enforce local mitigation policies. In contrast, the nsAS (AS4), which is not subscribed to the service, does not receive RL updates and may incorrectly accept the hijacked prefix as valid.

The IXP maintains full operational neutrality throughout this process. It does not alter routing policies or traffic flows; its sole role is the computation and dissemination of mathematically derived RL values. All subsequent mitigation actions are independently executed by the sASes. This separation ensures compliance with the principle of network neutrality while improving the overall resilience of inter-domain routing.

A. Subscriber-side actions

Upon receiving a tuple $(ASN, MetricType, RL)$ from the IXP, a sAS retains full autonomy in determining how to utilize

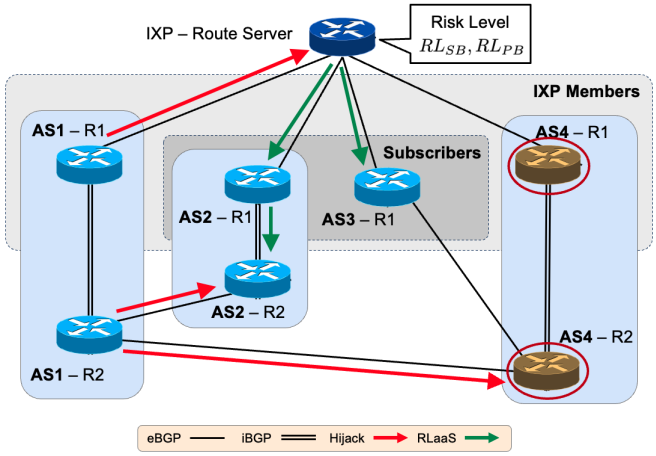

Fig. 3. The RS detects a prefix hijack from AS1 and computes two RL values. The sASes (AS2, AS3) obtain these values to mitigate the attack, while the nsAS (AS4), lacking RL data, may incorrectly accept the hijacked prefix.

the information. The simplest use case involves passively logging the received RL values for situational awareness, incident correlation, and post-event forensic analysis, consistent with the U.S. national cybersecurity roadmap [18]. During routing anomalies, such records provide diagnostic evidence that facilitates the identification of potentially implicated ASes, thereby supporting timely mitigation and recovery. This mechanism enables operators to iteratively refine their routing policies and enhance network resilience.

A more proactive mitigation approach allows a sAS to penalize high-risk ASes directly within the BGP decision process. BGP route selection is governed by a sequence of configurable attributes, including Local Preference, AS Path length, Multi-Exit Discriminator (MED), and Community values. By manipulating these attributes, operators can deprioritize routes originated or propagated by ASes exhibiting high RL values, thus reducing their likelihood of being selected as best paths.

Penalization is preferred over complete filtering, as rejecting all routes from an AS risks collateral impact on legitimate prefixes. In contrast, penalization selectively lowers route preference while maintaining global reachability.

The RL metric can be incorporated into the local routing policy to adjust relevant attributes dynamically—for instance, decreasing the Local Preference for prefixes originated by high-risk ASes or inflating their AS Path through prepending. These actions reduce the possibilities of routes coming from misbehaving ASes to be selected as best paths without altering the BGP path selection process.

Given its precedence in the BGP decision hierarchy, Local Preference is the most suitable attribute for applying RL-based penalties. Routes with higher Local Preference values are prioritized; thus, decreasing this attribute for high-risk ASes effectively lowers their route ranking. Rather than applying a static penalty, the sAS can use the received RL value as a dynamic decrement factor, ensuring proportional penalization across ASes with varying RL values.

For a practical deployment, two considerations must be

addressed. First, the penalization must be applied consistently across all BGP-enabled routers and across every active BGP session within the sAS, not only to the session established with the RS. Applying the adjustment network-wide ensures that malicious routes are consistently deprioritized. Second, the default Local Preference value (commonly 100) may differ across operator configurations. Therefore, penalization should be computed relative to the actual baseline Local Preference configured in each session, avoiding cases where a fixed decrement produces no tangible reduction.

This mechanism applies to both RL metrics. For RL_{SB}, penalization targets the origin AS (the rightmost AS in the AS path), reflecting the source of the announcement. For RL_{PB}, it targets intermediate ASes (from the penultimate to the leftmost AS), reflecting those that propagate potentially invalid routes. In the case of RL_{PB}, a stronger countermeasure can be applied: when the involved AS is a peer connected through a private link, the sAS can de-peer from it, thereby terminating the BGP session and halting route exchange, while maintaining connectivity through the more controlled ans secure IXP. This approach reduces exposure through less constrained private links while preserving routing through the IXP.

Finally, since both RL metrics may produce multiple values over time, the sAS can define a threshold parameter T_{RL} to control the activation of automatic mitigation. This threshold allows fine-tuning of sensitivity to RL fluctuations, balancing responsiveness against stability in routing policy adjustments.

B. Technical Requirements

This section describes the technical prerequisites for deploying and operating the RLaaS framework, distinguishing between IXP-side and subscriber-side infrastructures.

IXP-side. To provide RLaaS, the IXP must deploy a dedicated infrastructure for risk computation and dissemination, separate from the operational RS environment. This infrastructure must periodically acquire RIB snapshots from the RSes, compute RL values, and generate normalized risk indicators. Additionally, the IXP must also offer a publicly accessible distribution endpoint for RL information, protected by authentication, authorization, and secure transport protocols.

Subscriber-Side. Each sAS must deploy a client infrastructure to securely receive and process RL updates. To enforce consistent policies, RL information must be available at all border routers involved in external routing, either through local client components or internal distribution. The infrastructure must also support automated integration of RL values into border router configurations, enabling dynamic adjustments to local routing policies based on risk indicators.

Although RLaaS requires client-side integration at sASes, the deployment burden is minimal compared to existing BGP security mechanisms. RLaaS does not require protocol extensions, object creation, or changes to inter-domain routing semantics. Adoption involves deploying a lightweight client component and integrating risk indicators into existing policy frameworks, commonly used for traffic engineering and routing control. RLaaS enables incremental deployment and immediate benefit even under partial adoption.

C. RLaaS Implications

This section discusses the main implications, outcomes, and benefits derived from the use of the RLaaS.

A key aspect of RLaaS is its design to **complement**, rather than replace RPKI, which remains the long-term solution for securing BGP routing. However, as outlined in Section II, RPKI adoption is far from universal. RLaaS was therefore developed to address gaps in areas not yet covered by RPKI. For instance, when a prefix is received without an associated ROA, RPKI validation will not be able to make a definitive decision and therefore filter it. In contrast, RLaaS can step in to penalize the route derived from such an announcement without sacrificing reachability.

One primary implication of RLaaS is the **centralization** of RL computation within the IXP. This approach removes the need for each AS to deploy and maintain its own monitoring or processing systems, reducing operational complexity. In addition, computing RL values directly from the RS's RIB provides a broader and more accurate picture of routing dynamics, as the RS aggregates BGP updates from hundreds of ASes, including those that might be invisible to IXP members. Thus, RLaaS turns the IXP into a collective observation point for routing risk, enabling more scalable and reliable assessments than any single AS could achieve.

Moreover, RLaaS influences route preference in a **non-destructive** way: when a sAS penalizes an involved AS, two outcomes are possible, depending on the available routes. If only an invalid route is available, the sAS continues to use it, as RLaaS is designed to penalize rather than block announcements. In contrast, when both valid and invalid routes are present, *without RLaaS*, best-path selection follows the local routing policy and the standard BGP decision process; absent strict filtering, the invalid route participates in the selection and can be chosen, for example, due to a shorter AS-path. *With RLaaS*, the penalization lowers the invalid route's local preference, ensuring selection of the valid route regardless of update order or tie-break conditions. Moreover, as sASes propagate their selected routes to their peers according to the valley-free rule, the overall system suppresses the amplification of attacks, reinforcing secure and consistent route dissemination across the Internet.

While the penalization strategy improves routing security, it may introduce side effects. RLaaS inherits limitations from the RL methodology in [28], particularly its reliance on inferred AS relationships, which can lead to misclassifications and false positives, potentially penalizing ASes not involved in malicious activity. Moreover, since penalization operates at the AS level, both malicious and legitimate announcements may be affected. Under RL_{SB}-based penalization, a sAS de-preferences all routes originated by a specific AS, whereas RL_{PB}-based penalization extends this effect to all routes propagated by that AS, including those originated by its downstream customers. Consequently, legitimate traffic may

TABLE I
OVERVIEW OF RLAAS IMPLICATIONS AND ITS IMPACT ON ROUTING SECURITY

Discussion Area	Key Idea	Implications / Outcomes
Complementing RPKI	RLaaS complements RPKI by addressing deployment gaps	Provides an additional security layer by penalizing invalid routes not filtered by RPKI
Centralized computation	IXP computes RL values from the RS view, avoiding per-AS computation.	Provides more accurate and scalable risk assessment; improves visibility across multiple peering relationships.
Selection as best-path	Penalization influences BGP best-path selection without blocking announcements.	Ensures valid routes are preferred; limits propagation of invalid routes; mitigates attack amplification.
Penalized routes	Penalization may affect all prefixes of an AS.	Can induce longer but safer paths on legitimate traffic
Prevention	Repeated malicious AS behavior increases RL, lowering route preference automatically.	Prevents future exploitation by deprioritizing malicious sources; cumulative RL effects strengthen defense.
Recovery	RL-based penalization aids post-incident mitigation by reducing the attractiveness of hijacked routes.	Decreases impact of hijacks; may reduce need for more-specific prefix announcements (e.g., /25).
IXP Net Neutrality	RLaaS computes and disseminates RL without enforcing routing policies.	Maintains IXP neutrality; no traffic prioritization or discrimination introduced.
Economic	IXPs can monetize RLaaS through subscription fees.	Supports IXP operations with a sustainable revenue model.

be impacted, potentially causing suboptimal path selection or increased latency. However, connectivity is preserved, as RLaaS influences route preference rather than blocking routes. Therefore, **traffic remains deliverable**, while routing decisions progressively shift toward ASes adopting stricter filtering and validation practices.

Although RLaaS is primarily designed for post-attack mitigation, it also exhibits **prevention** capabilities, particularly against *serial* malicious ASes that conduct repeated attacks over time [29]. Once an attack is detected, the IXP disseminates an elevated RL value for involving AS. SASes then reduce the local preference of routes associated with that AS, ensuring that subsequent attacks, whether targeting the same or different prefixes, are automatically deprioritized. As detailed in [28], this cumulative adjustment progressively increases the RL metric, discouraging future malicious activity and transforming RLaaS into a lightweight, adaptive reputation mechanism integrated within the BGP decision process.

Beyond prevention, RLaaS also contributes to post-incident **recovery**, which refers to the process of regaining traffic after a routing security incident. During a prefix hijack, both the legitimate owner and the hijacker announce the same prefix, leading to a split in the global AS selection. *Without the RLaaS*, the legitimate AS often faces a topological disadvantage due to its route typically having a longer AS path. To compensate, the victim commonly announces more specific prefixes (e.g., /25s) to exploit the longest-prefix-match rule and reestablish reachability. For instance, during the 2008 YouTube BGP hijack, two /25 prefixes were used to recover the hijacked /24 [6]. However, this approach is limited, since many IXPs and transit providers filter prefixes longer than /24, preventing such sub-prefixes from being widely propagated [28]. *With RLaaS*, most sASes already deprioritize the hijacker's routes because of the high RL assigned to that AS. Consequently, the majority of the ASes converge toward the legitimate route even before the victim issues more specific announcements. The legitimate AS may still choose to advertise sub-prefixes, but only to recover ASes topologically closer to the attacker, where the false route remains temporarily preferred.

From the point of view of the IXP, we make two main considerations. Firstly, despite its influence on routing decisions, RLaaS fully preserves **IXP neutrality**. Since penalization decisions are executed by IXP members, the IXP itself remains passive in terms of traffic manipulation. The service does not introduce any mechanism of discrimination, prioritization, or interference with participants' routing policies. The IXP's role is limited to computing RL values based on mathematical formulas and observed BGP updates received from its members. These values are then published to all sASes, which independently determine how to apply them. Consequently, the IXP maintains a strictly neutral position, consistent with the principles traditionally upheld by the majority of IXPs.

Secondly, RLaaS can be offered by an IXP as an optional, value-added service, potentially generating an **economic interest**. The adoption of any associated economic model is not inherent to RLaaS itself and remains entirely at the discretion of the IXP, aligned with its principles of neutrality. From the perspective of sASes, RLaaS can reduce the need for deploying and maintaining loval risk assessment mechanisms. This structure aligns the interests of IXPs and their members, ensuring the long-term viability of RLaaS and incentivizing continuous improvement in routing security.

V. PROOF OF CONCEPT

This section presents the PoC implementation of the RLaaS framework, aimed at assessing its practical feasibility and effectiveness in reducing prefix hijack propagation, complementing existing RPKI validation. The PoC reproduces a realistic inter-domain routing scenario where multiple ASes interconnect through an IXP RS. The testbed was deployed using the Kathara framework [30], an open-source container-based network emulation platform that enables the deployment of multiple interconnected routers on a single host.

The reference topology, shown in Fig. 4, consists of 20 ASes and one RS deployed within a single Kathara lab. Each AS is modeled by a border router running the FRRouting daemon [31], establishing BGP sessions with neighbors and/or with the RS. Every AS originates a unique prefix and exchanges BGP UPDATE messages according to the valley-free rule, emulating realistic inter-domain behavior. The figure

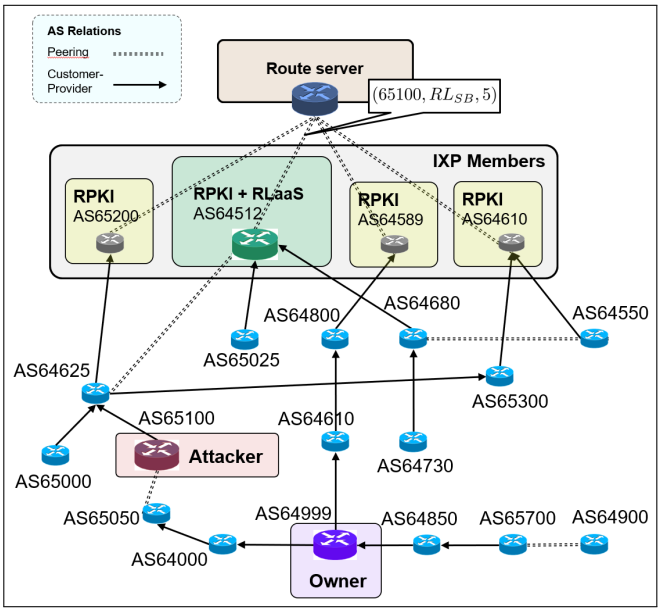

Fig. 4. Proof of concept topology showing interconnected ASes. The attacker, AS65100, hijacks the prefix 13.43.1.0/24, while AS64999, as an RLaaS subscriber, receives the tuple (65100, RL_{SB}, 5) and penalizes the route.

also depicts AS business relationships: solid arrows represent customer-provider links (from customer to provider), while dotted lines denote peer-to-peer connections.

Within this setup, the IXP instance collects BGP UPDATE messages from all members via the RS and computes the two RL metrics, RL_{SB} and RL_{PB}, as described in Section IV. Each normalized tuple $(ASN, MetricType, RL)$ is then securely distributed to sASes. Upon receipt, a sAS applies penalization to routes either originated or propagated by the corresponding AS, depending on the specific RL metric. As discussed in Section IV-C, an action is triggered when the RL exceeds a predefined threshold T_{RL}. In this PoC, $T_{RL} = 0$, meaning that all RL values activate a subscriber's action.

All IXP members in the testbed implement RPKI-based ROV to secure BGP announcements. The scenario involves the following main entities:

- **AS64999**–*owner*: legitimate holder of prefix `13.97.1.0/24`, without a signed ROA.
- **AS65100**–*attacker*. malicious AS that illegitimately announces the same prefix.

For analytical purposes, an AS is considered *compromised* if it selects the malicious announcement as the best path. The evaluation test was structured into three consecutive phases:

- **Without RLaaS**, which characterizes baseline routing behavior in the absence of RLaaS.
- **With RLaaS**, which evaluates the impact of RL_{SB} and RL_{PB}-based penalization.
- **Recovery**, which assesses RLaaS's role in post-incident recovery.

A. Without RLaaS

Initially, no IXP member subscribes to RLaaS, but all perform RPKI validation. The *owner* announces prefix `13.97.1.0/24`, which is correctly selected by all ASes. When the *attacker* advertises the same prefix, IXP members do not filter the attacker's announcement, as RPKI validation yields no result due to the absence of a ROA. As a result, the network undergoes route convergence based on standard BGP path selection. As shown in Fig. 4, ASes closer to the attacker prefer the illegitimate route due to its shorter AS-path. Although the RS enforces filtering and discards the attacker's announcements, IXP members still receive and may select the hijacked route via other BGP sessions.

TABLE II
BGP RIB OF AS64512 BEFORE RLAAS. AS64512 SELECTS THE MALICIOUS PATH (MARKED BY '>') FOR PREFIX `13.97.1.0/24`.

Network	Next Hop	LocPrf	Path
* 13.97.1.0/24	11.4.1.1	100	64589 64800 64610 64999 i
*>	***10.40.1.2***	***100***	***64625 65100 i***
*> 13.96.0.0/16	10.50.8.1	100	64610 64550 i
*> 26.0.10.0/24	11.61.1.2	100	65025 i
*> 42.0.0.0/24	10.40.1.2	100	64625 65000 i
*	11.4.1.16	100	65200 64625 65000 i
*	11.4.1.20	100	64610 65300 64625 65000 i

TABLE III
BGP RIB OF AS64512 AFTER RL_{SB} PENALIZATION. AS64512 SELECTS THE CORRECT PATH (MARKED BY '>') FOR PREFIX `13.97.1.0/24`.

Network	Next Hop	LocPrf	Path
*> 13.97.1.0/24	11.4.1.1	100	64589 64800 64610 64999 i
*	10.40.1.2	***95***	64625 65100 i
*> 13.96.0.0/16	10.50.8.1	100	64610 64550 i
*> 26.0.10.0/24	11.61.1.2	100	65025 i
*> 42.0.0.0/24	10.40.1.2	100	64625 65000 i
*	11.4.1.16	100	65200 64625 65000 i
*	11.4.1.20	100	64610 65300 64625 65000 i

Table II presents the RIB of AS64512, connected to the RS, which incorrectly selects the malicious route.

Under these conditions, 11 of the 19 ASes (attacker excluded) ($\approx 57.9\%$) were compromised and forwarded traffic to the attacker. The remaining 8 ASes ($\approx 42.1\%$) remained unaffected, typically those topologically closer to the legitimate origin. These results serve as the *baseline* against which subsequent analyses are compared.

B. RLaaS subscribers using RL_{SB}

In the second phase, AS64512 subscribes to RLaaS and applies penalization based on the RL_{SB} metric. Upon detecting an illegitimate announcement from the *attacker*, the IXP computes a raw RL_{SB} value of 0.5 according to [28], which is then normalized to the integer 5 before dissemination. The tuple $(65100, RL_{SB}, 5)$ is then distributed to AS64512.

The sAS reacts by lowering the local preference of all routes originated by AS65100 to 95 across all its BGP sessions, using the configuration command shown below.

```
as-path access-list ORIGIN-65100 permit _65100$
route-map FILTER permit 10
 match as-path ORIGIN-65100
 set local-preference 95
neighbor <peer-address> route-map FILTER in
```

TABLE IV
OVERVIEW OF COMPROMISED AND UNAFFECTED ASES ACROSS PHASES.

Scenario	Compromised ASes	Unaffected ASes
Before RLaaS	10 (55.6%)	8 (44.4%)
After RLaaS (one subscriber)	6 (33.3%)	12 (66.6%)
Full adoption	3 (16.6%)	15 (83.4%)

TABLE V
BGP RIB OF AS64512 AFTER RL_{PB} PENALIZATION. AS64512 SELECTS THE CORRECT PATH (MARKED BY '>') FOR PREFIX `13.97.1.0/24`.

Network	Next Hop	LocPrf	Path
*> 13.97.1.0/24	11.4.1.1	100	64589 64800 64610 64999 i
*	10.40.1.2	***95***	64625 65100 i
*> 13.96.0.0/16	10.50.8.1	100	64610 64550 i
*> 26.0.10.0/24	11.61.1.2	100	65025 i
*> 42.0.0.0/24	10.40.1.2	***95***	64625 65000 i
*	11.4.1.16	***95***	65200 64625 65000 i
*	11.4.1.20	***95***	64610 65300 64625 65000 i

As a result, the legitimate route from AS64999 regains preference and becomes the selected path. Table III shows the resulting RIB of AS64512, where the valid route now has a higher local preference and is chosen as best.

This local adjustment triggers a cascade effect, as AS64512 begins propagating the valid route to all its peers, prompting further best-path recalculations across the topology. Some neighboring ASes, such as AS65025, subsequently replace the malicious path with the legitimate one, as their sole route to the prefix depends on AS64512's selection.

Following this intervention, the overall security of the topology improves: the number of compromised ASes decreases to 7 ($\approx 36.8\%$), while unaffected ASes rise to 12 ($\approx 63.2\%$). However, some IXP members, such as AS65200, still rely solely on RPKI validation, which does not filter the malicious announcement, and remain compromised.

The cumulative effect becomes more significant as additional ASes adopt RLaaS. For example, when AS64610 adopts RLaaS, the number of compromised ASes further decreases to 5 ($\approx 26.3\%$). Under full adoption, only four ASes ($\approx 21.0\%$) remain compromised, with 15 unaffected ($\approx 78.9\%$). This corresponds to an overall improvement in network security of $\approx 36.8\%$ compared to the *baseline*.

Table IV summarizes the evolution of compromised and unaffected ASes across the experimental phases, highlighting the security improvement achieved through RLaaS adoption.

C. RLaaS subscribers using RL_{PB}

In addition to the RL_{SB}, the IXP distributes the RL_{PB} values corresponding to intermediate ASes, such as AS64625. This provides AS64512 with an alternative mitigation strategy.

Rather than penalizing the malicious origin, AS64512 may choose to penalize the intermediate AS64625, directly connected to it. When the illegitimate announcement from the *attacker* is received, the IXP had already computed the raw RL_{PB} value of 12.5, based on the formulas presented in [28]. Following this, a normalization step is applied using the logarithmic function: $(10 * log(1 + RL_{PB}))/(1 + log(1 + RL_{PB}))$, and the result is rounded to the nearest integer. Subsequently, the IXP sends the tuple $(64625, RL_{PB}, 7)$, where AS64512 applies a local preference reduction to all announcements propagated by AS64625 using the following configuration:

```
as-path access-list PATH-64625 permit _64625_
```

Table V presents the resulting RIB of AS64512. All routes containing AS64625 in their AS path are assigned a lower local preference, causing AS64512 to select the legitimate route announced by the *owner*. Quantitatively, the protection level matches that of the RL_{SB} approach, reducing the number of compromised ASes from 11 to 7. However, it proves more effective when AS64625 propagates multiple malicious routes, since it affects all the announcements it forwards, regardless of their specific origin. In this way, AS64512 mitigates not only the attack from AS65100 but also any additional malicious prefixes that AS64625 may forward.

TABLE VI
BGP RIB OF AS64512 AFTER DE-PEERING FROM AS64625. THE MALICIOUS PATH FOR PREFIX `13.97.1.0/24` IS NO LONGER RECEIVED.

Network	Next Hop	LocPrf	Path
*> 13.97.1.0/24	11.4.1.1	100	64589 64800 64610 64999 i
*> 13.96.0.0/16	10.50.8.1	100	64610 64550 i
*> 26.0.10.0/24	11.61.1.2	100	65025 i
*> 42.0.0.0/24	11.4.1.16	100	65200 64625 65000 i
*	11.4.1.20	100	64610 65300 64625 65000 i

Given their direct peering relationship, AS64512 may alternatively decide to de-peer from AS64625, thereby terminating the corresponding BGP session. Table VI shows the RIB state of AS64512 after de-peering, where no routes with AS64625 as the left-most AS remain. Consequently, the malicious path for prefix `13.97.1.0/24` is no longer received, as AS64512 now obtains routes exclusively through the RS, which performs filtering. The resulting protection level is equivalent to that achieved through local-preference penalization.

Importantly, AS64512 preserves full prefix reachability under both penalization and de-peering. In the former case, connectivity is maintained through deprioritized routes, while in the latter, traffic continues through unaffected peers, ensuring secure and uninterrupted operation. For example, the prefix `13.101.0.0/16`, which was previously learned via AS64625, is still reachable through an alternative path (`65200 64625 65000`) received from the RS.

D. Recovery

Once the hijack was detected, the legitimate owner initiated a conventional recovery procedure by announcing two /25 sub-prefixes of the hijacked block `13.97.1.0/24`. *Without RLaaS*, these announcements were propagated through the network but filtered by the RS due to its prefix-length policy, which discards routes longer than /24. Consequently, the /25 prefixes were not redistributed to other IXP members, and most ASes continued to select the hijacked /24 route as their best path. As a result, 10 ASes ($\approx 52.6\%$) remained compromised, while the other 9 ($\approx 47.4\%$) continued to route correctly toward the legitimate origin.

With RLaaS, the network exhibited markedly different behavior. Under full adoption, the dissemination of RL tuples had already caused sASes to apply local-preference penalization, leading to 15 ASes to route correctly toward the legitimate origin, already outperforming the results achieved through the conventional recovery mechanism involving the announcement of /25 prefixes. When the legitimate owner additionally announced the two /25 sub-prefixes, the situation improved marginally, reaching 16 unaffected ASes ($\approx 84.2\%$) and only 3 compromised ($\approx 15.8\%$).

VI. Conclusion and Future Work

This paper introduced Risk Level-as-a-Service (RLaaS), a framework designed to enhance inter-domain routing security within IXP. RLaaS leverages RL metrics derived from RS observations to assess AS trustworthiness and provide actionable intelligence. By integrating RL values into routing policies, ASes can penalize high-risk neighbors, reducing attack propagation while maintaining IXP neutrality. A proof-of-concept demonstrated that RLaaS effectively mitigates prefix hijack propagation and improves routing stability.

Future work will explore inter-IXP collaboration to address scenarios where an AS receives risk values from multiple IXPs. By computing a *Global Risk Level* that aggregates intelligence from multiple IXPs, detection accuracy and resilience can be enhanced. Additionally, we will investigate adversarial scenarios in which risk values may be maliciously manipulated or strategically biased, to ensure robustness and trustworthiness of the proposed framework under such conditions. In parallel, a large-scale longitudinal evaluation using real-world BGP data will assess RLaaS scalability, performance, and long-term impact on global routing security.

Acknowledgements

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART"), and by SERICS under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU (PE00000014).

References

[1] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, Jan. 2006.

[2] T. Tofoni, F. Luciani, and A. Prado, *BGP from theory to practice*. Reiss Romoli, Nov. 2023.

[3] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "Bgp hijacking classification," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2019, pp. 25–32.

[4] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, "Rfc 7908: Problem definition and classification of bgp route leaks."

[5] P. Spadaccino, S. Bruzzese, F. Cuomo, and F. Luciani, "Analysis and emulation of bgp hijacking events," in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*, 2023, pp. 1–4.

[6] RIPE, "Youtube hijacking: A ripe ncc ris case study." [Online]. Available: https://www.ripe.net/about-us/news/youtube-hijacking-a-ripe-ncc-ris-case-study/

[7] A. Siddiqui, "Klayswap – another bgp hijack targeting crypto wallets." [Online]. Available: https://manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/

[8] B. Du, K. Izhikevich, S. Rao, G. Akiwate, C. Testart, A. C. Snoeren, and K. Claffy, "Irregularities in the internet routing registry," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023.

[9] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," RFC 8210, Sep. 2017.

[10] N. Hilliard, E. Jasinska, R. Raszuk, and N. Bakker, "Internet Exchange BGP Route Server Operations," RFC 7948, Sep. 2016.

[11] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger, "Peering at peerings: On the role of ixp route servers," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014.

[12] Flavio Luciani, "Checking prefix filtering in ixps with bird and openbgpd." [Online]. Available: https://blog.apnic.net/2021/11/15/checking-prefix-filtering-in-ixps-with-bird-and-openbgpd/

[13] T. Shapira and Y. Shavitt, "Ap2vec: an unsupervised approach for bgp hijacking detection," *IEEE Transactions on Network and Service Management*, 2022.

[14] H. Birge-Lee, M. Apostolaki, and J. Rexford, "Global bgp attacks that evade route monitoring," in *International Conference on Passive and Active Network Measurement*. Springer, 2025, pp. 335–357.

[15] L. Qin, L. Chen, D. Li, H. Ye, and Y. Wang, "Understanding route origin validation (rov) deployment in the real world and why manrs action 1 is not followed," in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. NDSS, 2024.

[16] W. Li, Z. Lin, M. I. Ashiq, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "Rovista: Measuring and analyzing the route origin validation (rov) in rpki," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023.

[17] C. Testart, J. Wolff, D. Gouda, and R. Fontugne, "Identifying current barriers in rpki adoption," *Available at SSRN*, 2024.

[18] The White House Office of the National Cyber Director, "Roadmap to enhancing internet routing security." [Online]. Available: https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/09/Roadmap-to-Enhancing-Internet-Routing-Security.pdf

[19] Organisation for Economic Co-operation and Development (OECD), "Routing security: Bgp incidents, mitigation techniques and policy actions," OECD Publishing, Tech. Rep., 2022.

[20] Euro-ix, "Internet exchange points - 2021 report." [Online]. Available: https://www.euro-ix.net/media/filer_public/35/73/3573f355-c90a-4b31-ae83-851b76cfa36b/ixp_report_2021.pdf

[21] W. Kumari and D. R. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)," RFC 5635.

[22] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "Artemis: Neutralizing bgp hijacking within a minute," *IEEE/ACM transactions on networking*.

[23] MANRS, "Actions." [Online]. Available: https://manrs.org/netops/actions/

[24] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and K. Claffy, "Mind your manrs: measuring the manrs ecosystem," in *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022.

[25] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann, "On the benefits of using a large ixp as an internet vantage point," in *Proceedings of the 2013 conference on Internet measurement conference*.

[26] M. Nawrocki, J. Blendin, C. Dietzel, T. C. Schmidt, and M. Wählisch, "Down the black hole: dismantling operational practices of bgp blackholing at ixps," in *Proceedings of the Internet Measurement conference*.

[27] L. Mueller, M. Luckie, B. Huffaker, M. Barcellos *et al.*, "Spoofed traffic inference at ixps: Challenges, methods and analysis," *Computer Networks*, vol. 182, p. 107452, 2020.

[28] S. Servillo, P. Spadaccino, F. Luciani, and F. Cuomo, "Estimating autonomous system risk levels by analyzing ixp route server rib," *Computer Communications*, vol. 237, p. 108154, 2025.

[29] E. Jaw, M. Müller, C. Hesselman, and L. Nieuwenhuis, "Serial bgp hijackers: A reproducibility study and assessment of current dynamics," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, 2024.

[30] G. Bonofiglio, V. Iovinella, G. Lospoto, and G. Di Battista, "Kathará: A container-based framework for implementing network function virtualization and software defined networks," in *NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2018.

[31] FRRouting, "Frrouting (frr)." [Online]. Available: https://frrouting.org/