Autonomous Systems Risk Level in the Route Server Infrastructure of an Internet Exchange Point

Stefano Servillo
DIET
Sapienza Università di Roma
Rome, Italy
stefano.servillo@uniroma1.it

Pietro Spadaccino
DIET
Sapienza Università di Roma
Rome, Italy
pietro.spadaccino@uniroma1.it

Francesca Cuomo
DIET
Sapienza Università di Roma
Rome, Italy
francesca.cuomo@uniroma1.it

Flavio Luciani
Namex (Rome IXP)
Rome, Italy
f.luciani@namex.it

***Abstract*—Internet Exchange Points (IXPs) play a fundamental role in the exchange of data between Internet Service Providers (ISPs). However, they face one of the main challenges of the Border Gateway Protocol (BGP): trust-based route sharing. This feature introduces a series of vulnerabilities that can be exploited by attackers to hijack or disrupt traffic. Despite the presence of various countermeasures such as Internet Routing Registries (IRRs) or the Resource Public Key Infrastructure (RPKI), the lack of implementation by the majority of Autonomous Systems (ASes), limits their effectiveness. In this paper, we define a tool that supports IXPs operation, enhancing the security of BGP peering in their infrastructure. The proposed approach analyses the information contained in BGP UPDATE messages received by the route-server of an IXP to identify possible prefix hijacking attacks. This set of prefixes are then used to define and compute a *Risk Level* value associated with each AS, providing network operators with an indication of anomalous behaviours. To achieve this objective, data obtained from the route-server one of the main Italian IXP, are examined, showing a real application of our tool.**

***Index Terms*—BGP, routing, security, IXP**

I. Introduction

The Border Gateway Protocol (BGP) is the inter-domain path-vector routing protocol on which Internet network traffic is based [1]. BGP allows an Autonomous System (AS), identified by an Autonomous System Number (ASN) assigned by a Regional Internet Registry (RIR), to share reachability information by announcing a set of IP prefixes. Despite being a resilient and well-established protocol, BGP has shortcomings in its security measures, leading to multiple vulnerabilities [2].

At its essence, an inherent weakness in BGP lies in the foundation of mutual trust upon which the exchange of routing information is built, where an AS is supposed to only advertise its IP prefixes - blocks of IP addresses typically assigned to the AS administrator by a RIR [3]. This is determined by the absence of several fundamental mechanisms:

- BGP does not offer tools to validate legitimacy and authorization of an AS to disseminate routing information.
- Lack of mechanisms to protect integrity and 'freshness' of messages, as well as origin authentication.

The present work is partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU.

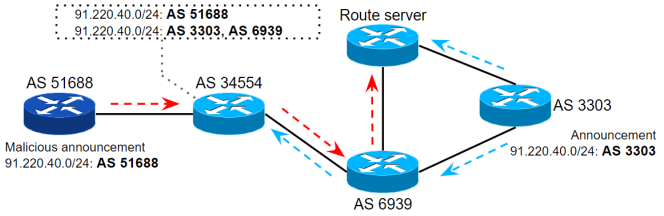

Fig. 1: A real-world case of the hijacking of '91.220.40.0/24' belonging to the AS 3303 and also announced by the AS 51688. Without correct countermeasures, the AS 34554 receiving both BGP UPDATE messages will choose the one with the shortest AS_PATH as best path.

- Absence of tools to ensure the authenticity of the AS_PATH attribute announced by an AS.

Malicious ASes can take advantage of these lack of validation and authentication mechanisms by advertising others' IP prefixes, as happened in the real case illustrated in Fig. 1. Such technique is generally recognized as 'prefix hijacking' where an AS can purposefully black-hole traffic or impersonate the receiver [4]. For example, the prefix '91.220.40.0/24', belonging to AS 3303, is announced by another AS. Without appropriate countermeasures, AS 34554 would choose the malicious path announced by AS 51688 due to a shorter path. In any case, the disruption of normal network routing occurs, as evidenced by the 2022 Twitter hijacks [5].

BGP lack of security has notably affected Internet Exchange Points (IXPs), neutral hubs where several ASes interconnect, allowing the exchange of routing information through public peering, a multilateral agreement among ASes. This exchange is facilitated by route servers, tools that reflect BGP announcements from one AS to another [6], as illustrated in Fig. 2.

To mitigate these risks, the Internet community has developed countermeasures, widely implemented by IXPs globally. The main ones comprise:

- *Internet Routing Registries (IRRs)*: These are public databases managed by RIRs, containing information regarding the ownership of Internet resources. Authorized operators of ASes can insert a route object <IP prefix; ASN>, specifying which ASN is authorized to announce

ISBN 978-3-903176-63-8 © 2024 IFIP

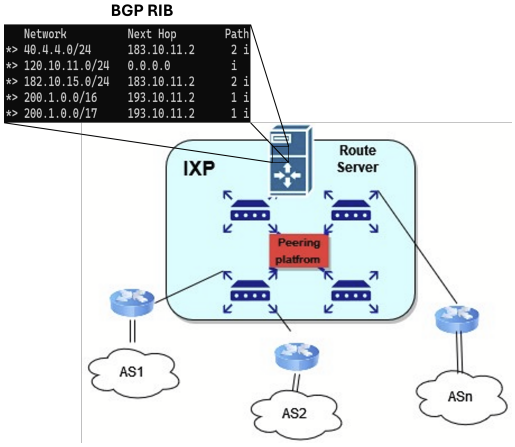

Fig. 2: Simplified network design of a typical IXP peering platform. The autonomous systems (in the scheme AS1 and AS2) perform peering with the IXP route server, which then redirects all the packets to the destination AS. An example of a BGP RIB is provided, highlighting prefixes announced by AS1 and AS2 and received by the IXP route server.

the IP prefix to the Internet [7];

- *Resource Public Key Infrastructure (RPKI)*: This framework is based on a public structure with distributed databases that include associations <IP prefix; Authorized ASN>. Each pair is associated with a Digital Certificate, allowing the verification of route origin authenticity. These pairs with the associated Digital Certificate are termed Route Origin Authorization (ROA) [8].

However, both countermeasures have their limitations. In the case of IRR, route objects are untrustworthy. Due to the manual nature of the process, these objects are frequently outdated or contain inaccurate information [9]. Despite RPKI being designed to mitigate this issue, its universal adoption remains incomplete [10]. By the end of 2023, after a decade of use, only 47.66% of routes were protected through RPKI, as evidenced by NIST RPKI Monitor [1], leaving many networks vulnerable to potential attacks.

Preventing prefix hijacking is a particularly difficult operation, requiring the cooperation of multiple ASes. Therefore, the development of detection techniques has garnered attention in recent years [11], [12], [13].

Several studies use a tool, developed by BGPMon, to obtain information about past attacks, called BGPStream[2]. This is a free resource that warns about hijacks and disruptions occurring in BGP. These are several BGPMon peers around the world that analyse the correct flow of BGP announcements. Peers report possible prefix hijacking if they receive a prefix announcement from an AS different than usual. However, BGPStream only checks the Origin ASN, without further checks. For this reason, configuration errors may be included among possible attacks, making it not very accurate.

[1]NIST RPKI Monitor available at https://rpki-monitor.antd.nist.gov/
[2]BGPStream available at https://bgpstream.com/about/


```
{
  "network": "77.239.137.144/28",
  "neighbor": "193.201.28.179",
  "rpki": "!",
  "origin": "31087",
  "path":[
    "30848",
    "31087"
  ],
  "reason": "rpki invalid"
}
```


(a) *'77.239.137.144/28' prefix , announced by AS 31087 and filtered by the route server due to 'rpki invalid' reason*


```
{
  "network": "217.77.101.0/24",
  "neighbor": "193.201.28.59",
  "rpki": "N",
  "origin": "204287",
  "path": [
    "6939",
    "62874",
    "204287"
  ],
  "reason": "irrdb filtered"
}
```


(b) *'217.77.101.0/24' prefix, announced by AS 204287 and filtered by the route server due to 'irrdb filtered' reason*

Fig. 3: Structure of data extracted from the BGP RIB of a route server of an IXP.

Rater than rely on external monitoring services, our study focuses on analyzing the data within the BGP Routing Information Base (RIB) of a route server within an IXP. The BGP RIB is a dynamic data structure containing all BGP routing information. Upon receiving a BGP UPDATE message, the route server inserts the corresponding prefix into the BGP RIB, along with its associated information. The crucial information lies in whether the prefix was selected for re-transmission or filtered out and not re-transmitted, due to several filtering policies in the route server. After extracting information from the BGP RIB, a procedure is applied to the set of filtered prefixes to detect those potentially related to attacks, in order to identify the involved ASes. Through a series of computations, each AS is assigned a specific value using the 'Risk Level' metric, reflecting its potential threat to BGP routing. It is crucial to emphasize that this process is automated to minimize potential errors, enabling IXP operators to obtain precise insights into the BGP traffic reaching the route server.

The rest of the paper is organized as follows. Section II discusses the kind of data analyzed for the purposes of the work, while Section III describes the pre-processing mechanism applied to such data. Two cases of 'Risk Level' will then be presented, the first evaluate the risk coming from the Origin AS, the one announced the prefixes, denoted *Source-Based* Risk Level (Sec. IV) and the second evaluate the risk coming from ASes along the path of an attack-related prefix, denoted *Path-Based* Amplification Level (Sec. V). Subsequently, the results coming from the Risk Level calculation on the data obtained from a route server belonging to the main Italian IXP are shown and evaluated (Sec. VI). Finally, Section VII discusses future evolution of the present study.

II. BGP RIB DATA ANALYSIS

The extraction of data from the BGP RIB of the route server is a critical operation in our study. Within the BGP RIB, two distinct categories of prefixes can be identified: those selected for re-transmission and those filtered base on specific policies. As depicted in Fig. 3, prefixes in both categories share certain common attributes:

- *Network*: This attribute identifies the announced prefix, encapsulated within the BGP UPDATE message;
- *Neighbor*: Denotes the interface address of the route server where the BGP UPDATE message was received;
- *RPKI*: Indicates the result of RPKI validation, which verifies whether the Origin ASN is associated with the IP prefix within the RPKI framework. Possible values include:
  - Valid (V): The prefix is associated with the Origin ASN;
  - Invalid (!): The prefix is not associated with the Origin ASN;
  - Not found (N): The prefix is not present in the RPKI.
- *Origin*: Identifies the ASN that announced the prefix;
- *Path*: Shows the AS_PATH attribute, which includes the list of all ASes traversed by the BGP UPDATE message;
- *Reason*: Indicates the rationale for accepting the prefix for re-transmission or filtering by the route server.

Our study focuses on identifying attacks from prefixes which have been filtered by IXP route servers. Within a route server, administrators can implement filtering policies to manage the information contained in BGP UPDATE messages. The most effective approach, as illustrate in the Fig. 4, involves implementing a cascade control mechanism based on mutually exclusive criteria, prioritized in the following sequence:

1) *RPKI Validation*: The route server interfaces with the RPKI framework to authenticate information. It verifies whether a ROA associates the Origin ASN with the prefix. If the information does not align, the prefix is filtered. This validation step ensures that the Origin ASN is authorized to announce the prefix;
2) *IRR Check*: The route server verifies the presence of a route object associating the Origin ASN with the prefix in IRRs. If no such route object is found, the prefix is filtered. This verification is crucial to ensure that the Origin ASN is authorized to advertise the prefix, particularly for those ASes not yet implementing RPKI;
3) *Subnet Mask Control*: The route server enforces scrutiny on the subnet mask length. If the subnet mask exceeds '/24', the prefix is filtered. This control prevents an AS from announcing specific prefixes, thereby mitigating the risk of compromising traffic of precise devices;
4) *BGP selection process*: Prefixes that successfully pass security checks enter the BGP selection process. However, if the path is not designated as the best-path according to BGP selection criteria, the prefix is filtered.

The observed configuration at the Rome IXP represents an optimal process employed by various IXPs to mitigate the propagation of erroneous information throughout the Internet. It is crucial to note that when a prefix is filtered by the route server, it does not necessarily indicate an attack. Instead, it often originates from configuration errors made by the AS authorized to announce the prefix.

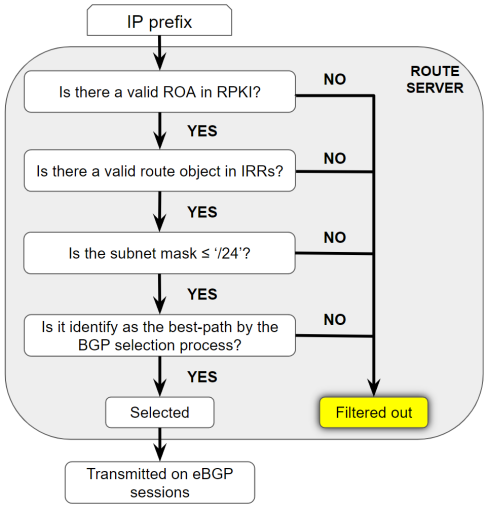

Fig. 4: Representative diagram of the cascade filtering process implemented by a route server upon the arrival of a new BGP UPDATE message. The process consists of a series of checks based on mutually exclusive criteria. If the prefix passes all checks, it will be accepted by the route server and re-propagated, otherwise it is filtered out

III. Input Data Pre-processing

The acknowledgement that the initial set of prefixes filtered by a route server $\mathbb{N}_{init}$ could include configuration errors has underscored the need for data pre-processing. The describe process consists of two stages that aims to discern the subset of possible attack-related prefixes.

A. Initial Filtering Stage

The initial filtering phase is pivotal in identifying most of configuration errors received by the route server. Its primary aim is to remove from the initial set $\mathbb{N}_{init}$ prefixes where the Origin AS, responsible for announcing the BGP UPDATE message, is the legitimate owner but has been filtered out by the route server due to incorrect data. Within this phase, two case studies are examined:

Change of owner. Typically, an IP prefix is assigned to an AS by a RIR. However, an AS may acquire an IP prefix by mutual agreement with another AS. The authorized operators of both ASes will update the relevant information in the IRR or RPKI to reflect this change. The previous AS will remove the route object or ROA associating its ASN with the prefix, while the new AS will create new entries, linking its ASN with the new IP prefix. Nevertheless, this process may not be completed. Consequently, when the route server interacts with the IRR or RPKI, it may filter out the prefix due to inconsistencies in the information provided.

Subnet mask. When an authorized operator of an AS creates an ROA, associating their ASN with an IP prefix, they are required to specify a parameter known as 'Maximum length'. This parameter denotes the maximum length of the

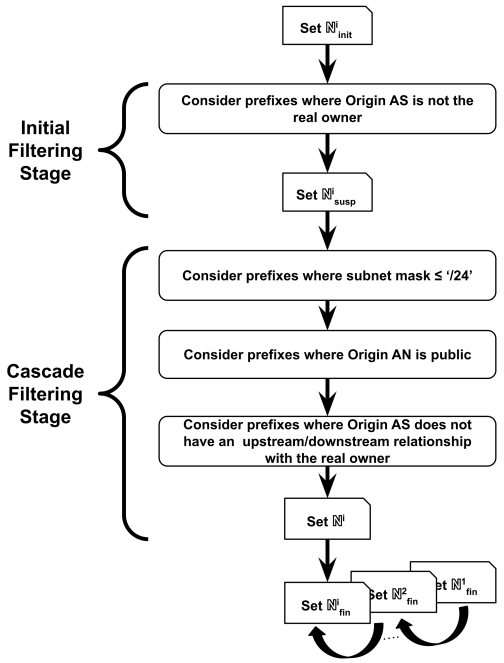

Fig. 5: Representative diagram of the entire filtering process that must be applied to the set of prefixes filtered by the route server to obtained the set of attack-related prefixes for week i. The final dataset of attack-related prefixes includes the union of all resulting sets from each week.

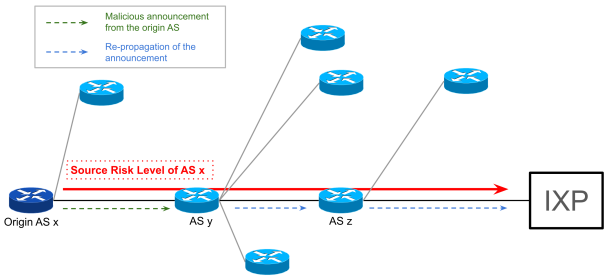

(a) Source-based Risk Level for AS x announcing an attack-related prefix.

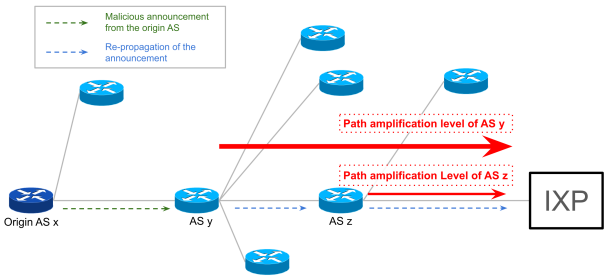

(b) Path-Based Amplification Level for AS y and AS z propagating an attack-related prefix originated by AS x.

Fig. 6: Announcement of an attack-related prefix by the origin AS x. When the announcement is received at the IXP we can then compute the Source-based Risk level for the origin AS x (a) and the different Path-Based Amplification Levels of the AS that have re-propagated the malicious announcement (b).

subnet mask permissible for the prefix, thereby delineating the allowable IP subnets that can be advertised. However, certain ASes may advertise prefixes with a subnet mask that exceeds the maximum permitted length. This results in the filtering of the prefix by the route server, as the information provided does not align with the predetermined parameters.

After filtering out these cases, we obtain the set $\mathbb{N}_{susp}$ comprising prefixes where the Origin ASes responsible for announcing them are not the real owners. These prefixes are deemed suspicious. Nonetheless, even within this refined set, there remains the possibility of additional configuration errors. Consequently, a second filtering step becomes necessary.

B. Cascade Filtering Stage

The second part stems from observations regarding the behaviours of ASes. This has led to the implementation of a cascading filtering mechanism on the set $\mathbb{N}_{susp}$:

1) *Subnet mask*: In BGP, it is a best practice for an AS to use prefixes with subnet masks greater than '/24' for internal network purposes and advertise aggregate prefixes with subnet masks of '/24' or less to the global Internet. As a result, ASes enforce control of the subnet mask within their filtering mechanisms. This consensus ensures that malicious ASes cannot effectively launch attacks by announcing prefixes with subnet masks greater than '/24', as it would be blocked by neighboring ASes. Therefore, the receipt of this type of prefix indicates a misconfiguration of the announcing AS, which failed to advertise the appropriate aggregate prefix. Consequently, prefixes with subnet masks greater than '/24' are eliminated from the set;
2) *AS Type*: Within the network infrastructure, an AS may use ASes for private use or documentation purposes. As outlined in RFC 6996 [14], certain ranges of ASN values are reserved to them. These ASes should not be able to announce prefixes to the public Internet. However, due to misconfigurations, many ASes may re-announce prefixes originating from private ASNs. Consequently, these prefixes cannot be associated with potential attacks and are eliminated from the set;
3) *AS Relationship*: In the Internet, AS relationships can take three distinct forms: upstream/downstream, peer-to-peer, or no established relationship. The upstream/downstream relationship exists between a provider AS and its client AS. It is common for a client to advertise prefixes that actually belong to its provider, or vice versa. This behavior is often driven by configuration decisions and contractual agreements between the involved ASes, rather than being indicative of malicious intent. Consequently, prefixes where the

Origin AS has an upstream/downstream relationship with the real owner are eliminated from the set. This task is often complicated by the confidentially of AS relationships.

At the end of the process, we have arrived at a set $\mathbb{N}$ comprising prefixes exhibiting well-defined characteristics associated with potential attacks. To be included in the set, a prefix must meet the following criteria:

- The Origin AS announcing the prefix does not correspond to the real owner;
- The prefix is announced by a public Origin AS;
- The Origin AS has either a peering relationship or no established relationship with the real owner;
- The prefix has a subnet mask that does not exceed '/24'.

For a correct execution, the entire two-steps filtering process is applied to the set of prefixes filtered by the route server each week. As can be seen from Fig. 5, the resulting set at a specific week i is denoted as $\mathbb{N}^i$. The set of attack-related prefixes up to week i is represented by $\mathbb{N}^i_{fin}$ and can be computed as $\mathbb{N}^i_{fin} = \mathbb{N}^1 \cup \mathbb{N}^2 \cup ... \cup \mathbb{N}^i$, with the understanding that $\mathbb{N}^1_{fin} \subseteq ... \subseteq \mathbb{N}^{i-1}_{fin} \subseteq \mathbb{N}^i_{fin}$.

This dataset serves as input for calculating the Risk Level, a metric designed to assess the risk associated with an AS to the Internet routing. This study defines two distinct Risk Levels: the *Source-Based* Risk Level (RL_{SB}) and the *Path-Based* Amplification Level (RL_{PB}). The first focuses on assigning a risk value to Origin ASes, which are responsible for announcing attack-related prefixes, as shown in Fig. 6a. Meanwhile, the *Path-Based* Amplification Level aims to evaluate the risk posed by ASes along the path of an attack-related prefix, as shown in Fig. 6b. An AS re-propagating a prefix, could have significantly more connections than the Origin AS, therefore amplifying the spread of an attack.

IV. Source-Based Risk Level

The *Source-Based* Risk Level (RL_{SB}) is designed to asses the risk posed by each Origin AS that has announced attack-related prefixes within the $\mathbb{N}^i_{fin}$ set. Calculating the score for an Origin AS necessitates deriving the set $\mathbb{C}_t \subseteq \mathbb{N}^i_{fin}$ denoting the subset of attack-related prefixes announced by the t-th Origin AS. Let $C_t = |\mathbb{C}_t|$. Information from five distinct categories is extracted, each assigning a score ranging from 0 to 1. These scores are then combined using a weighted formula, where each category is allocated a specific weight reflecting the importance of information analyzed. This process yields the RL_{SB} score associated with the t-th Origin AS.

The expression defined for the calculation of the RL_{SB}, along with the corresponding weight coefficients, is provided below:

$$RL_{SB} = T \cdot \alpha + F \cdot \beta + P \cdot \gamma + G \cdot \delta + O \cdot \epsilon \quad (1)$$

where $\{\alpha, \beta, \gamma, \delta, \epsilon\}$ are tunable parameters by the IXP operators.

The following subsections will define the formulas used to compute the score for each category using the variable described in Table I.

TABLE I: Variables for Computing Risk Levels

Parameters	Definitions
$\mathbb{C}_t$	Set of attack-related prefixes announced by Origin AS t
C_t	Cardinality of the set $\mathbb{C}_t$
W^d_k	Exponential weight for the kth week
N^d_k	Number of attack-related prefixes in the kth week
M	Number of months considered in the analysis
A	Monthly average of attack-related prefixes
N^m_k	Number of attack-related prefixes in the kth month
m_k	Binary variable indicating if attack-related prefixes exceed A in the kth month
W^a_k	Exponential cumulative weight for AS_PATH length k
L	Maximum AS_PATH length for an attack-related prefix
N^a_k	Number of attack-related prefixes for AS_PATH length k
$\mathbb{S}_t$	Set of suspicious prefixes announced by t-th Origin AS
S_t	Cardinality of the set $\mathbb{S}_t$
U_o	Count of ASes included in the IXP_{AS}
V_k	AS owner of the kth attack-related prefix
R_k	Ranking position of V_k in IXP_{AS}
W^o_k	Ranking weight for V_k
$\mathbb{I}_p$	Set of distinct attack-related prefixes traversing p-th AS
I_p	Cardinality of the set $\mathbb{I}_p$
B^k_o	number of connections of the Origin AS of the k-th prefix
B^k_p	count of connections of p-th AS in prefix k's AS_PATH
W^p_k	Weight of the k-th prefix within the set $\mathbb{I}_p$

A. Time

The *Time* category is based on the notion that the risk associated with each Origin AS decreases over time. This premise aligns with the objective of assigning greater importance to prefixes linked to more recent attacks.

To determine the relative score, it is necessary to ascertain the week in which each prefix within the $\mathbb{C}_t$ set was announced by the t-th Origin AS. This step allows the calculation of the count of prefixes announced for each week. An exponential distribution, parameterized by $\lambda = \frac{1}{7}$, is used to assign weights to each week. Notably, the abscissa 0 corresponds to the week in which the RL_{SB} is computed. The weight allocated to the k-th week through the exponential distribution W^d_k is computed as follows:

$$W^d_k = \frac{e^{-\frac{k}{7}}}{7} \quad (2)$$

Let D denote the number of weeks considered. The score T of the *Time* category is expressed as follows:

$$T = \frac{\sum_{k=0}^{D}(N^d_k \cdot W^d_k)}{C_t} \quad (3)$$

where N^d_k represents the number of prefixes announced in the k-th week.

B. Frequency

The *Frequency* category operates on the premise that the risk associated with each Origin AS increases proportionally with the frequency of its announcement of attack-related prefixes.

Let M denote the number of months under consideration (with a value equal to a maximum of 12). To compute the corresponding score, we first identify the monthly average of prefixes A for the t-th Origin AS, represented as $\frac{C_t}{M}$. Subsequently, for each prefix within the $\mathbb{C}_t$ set, we determine

the month of its announcement by the t-th Origin AS. This enables the calculation of the count of prefixes announced in the k-th month N_k^m. We introduce the binary variable m_k, which signifies whether the number of prefixes announced in the k-th month exceeds A, defined as:

$$m_k = \begin{cases} 1 & \textit{if } N_k^m \geq A, \\ 0 & \textit{otherwise} \end{cases} \tag{4}$$

The score for the *Frequency* category F is expressed as follows:

$$F = \frac{\sum_{k=0}^{M} m_k}{M} \tag{5}$$

C. Path

The *Path* category is based on the notion that the risk associated with each Origin AS should increase proportionately with the number of ASes traversed by an attack-related prefix before reaching the IXP route server. The length of the AS_PATH attribute serves as a metric in this assessment: the longer the AS_PATH is, the more ASes have accepted the malicious prefix as best-path and modified their routing tables accordingly. This leads the involved ASes to prefer the path through the malicious AS.

To compute the corresponding score, it is essential to extract the length of the AS_PATH attribute. Note that the length should be determined by ignoring any manipulation such as 'prepending', a common technique employed to alter the incoming path. Through prepending, an AS inserts its ASN multiple times into the AS_PATH, influencing the routing decisions of other ASes. By eliminating prepending, the real number of ASes traversed by the prefix can be ascertain.

A cumulative distribution is employed to assign weights to different lengths of the AS_PATH attribute. The length of an AS_PATH connecting two different ASes is usually lower than four. Therefore, we decided to assign a non-zero weight W_k^a only to AS_PATHs having a length k greater than four as follows:

$$W_k^a = \begin{cases} 1 - e^{4-k} & \textit{if } k \geq 4, \\ 0 & \textit{if } k < 4 \end{cases} \tag{6}$$

Let L denote the maximum length of the AS_PATH of a prefix belonging to the set $\mathbb{C}_t$, and let k such that $0 < k \leq L$. The *Path* category score P is expressed as follows:

$$P = \frac{\sum_{k=0}^{L}(N_k^a \cdot W_k^a)}{C_t} \tag{7}$$

where N_k^a represents the number of prefixes announced with an AS_PATH length of k.

D. Filtering

The *Filtering* category operates under the premise that the risk associated with each Origin AS increases in relation to the number of attack-related prefixes announced.

To compute the relative score, two precise sets will be compared for each AS: the set of suspicious prefixes and the set of attack-related prefixes. More precisely, the number of prefixes belonging to these two sets announced by an Origin AS will be compared. Let $\mathbb{S}^i$ represent the set of all suspect prefixes up to week i, denoted as $\mathbb{S}^i = \mathbb{N}^1_{susp} \cup ... \cup \mathbb{N}^i_{susp}$. It is important to derive the set $\mathbb{S}_t$, which denotes the subset of suspicious prefixes announced by the t-th Origin AS, with the understanding that $\mathbb{S}_t \subseteq \mathbb{S}^i$.

Let $S_t = |\mathbb{S}_t|$. The score of the *Filtering* category G is expressed as follows:

$$G = \frac{C_t}{S_t} \tag{8}$$

E. Owner

The *Owner* category hold paramount importance as it revolves around the notion that the risk associated with each Origin AS escalates concerning the victim of its attack, the actual owner of the advertised prefix. In this category, a non-zero score is assigned if the victim of an attack is connected to the IXP. A singular consideration governs this: ASes connected to the IXP will have a level of significance commensurate with the services they provide.

Consequently, the concept of risk has been tailored to accommodate the IXP perspective, with a focus on the likelihood of an attack occurring, particularly those involving ASes connected to it, thereby refining the concept of risk.

To compute the corresponding score, a special ranking termed IXP_{AS} has been devised. This ranking includes all ASes connected to the IXP, classified based on the relevance of the service offered. Subsequently, it is imperative to identify the AS victim of each attack-related prefix in the C_t set. Let V_k represent the AS owner of the k-th attack-related prefix in C_t. A weight between 0 and 1 is then assigned to the prefix based on the position of V_k in the ranking.

Let U_o denote the number of ASes included in $\textit{IXP}_{AS}$, and R_k denote the ranking position of V_k. The corresponding weight W_k^o is computed using the following expression:

$$W_k^o = \begin{cases} \frac{[U_o - R_k + 1]}{U_o} & \textit{if } V_k \textit{ in } \textit{IXP}_{AS} \\ 0 & \textit{otherwise} \end{cases} \tag{9}$$

The *Owner* category score O is expressed as follows:

$$O = \frac{\sum_{k=0}^{C_t} W_k^o}{C_t} \tag{10}$$

V. Path-Based Amplification Level

The *Path-Based* Amplification Level (RL_{PB}) serves to evaluate the risk posed by an AS present along the path of an attack-related prefix within the designated set $\mathbb{N}^i_{fin}$. This assessment is based on the role such ASes play in either mitigating or amplifying the impact of an attack.

Upon receiving an attack-related prefix, ASes along the path may filter it out and not transmitting it further, thereby limiting the spread of the attack. Conversely, accepting the prefix as legitimate leads these ASes to propagate it across their eBGP connections, thereby amplifying the attack. The RL_{PB} score is computed for all ASes present in the AS_PATH attribute, excluding the Origin AS, which occupies the final position.

The calculation of the score associated with an AS starts with the identification of the set $\mathbb{I}_p \subseteq \mathbb{N}^i_{fin}$, representing the set of distinct attack-related prefixes traversing the p-th AS. Let $I_p = |\mathbb{I}_p|$. Next, we consider the number B^k_o of connections of the Origin AS for the k-th prefix and the number B^k_p of connections of the p-th AS in the AS_PATH of the k-th prefix. Consequently, the k-th prefix within the set $\mathbb{I}_p$ is assigned a weight W^p_k using the formula:

$$W^p_k = \frac{B^k_p}{B^k_o} \cdot RL^k_{SB} \tag{11}$$

where RL^k_{SB} denotes the *Source-Based* Risk Level associated with the Origin AS of the k-th prefix.

The *Path-Based* Amplification Level score RL_{PB} is expressed as follows:

$$RL_{PB} = \frac{\sum_{k=0}^{I_p} W^p_k}{I_p} \tag{12}$$

This equation aims to characterize the risk associated to an AS by computing the average of the weights associated with all the prefixes that have traversed it. It is important to note a particular scenario that may arise: when an AS re-transmits a substantial number of prefixes, a high weight value W^p_k assigned to the k-th prefix might be attenuated by numerous prefixes with very low weights. To address this, we introduce RL^{max}_{PB} and RL^{sum}_{PB}, which represents respectively the maximum value and the sum of the weights W^p_k assigned to the k-th re-transmitted prefix:

$$RL^{max}_{PB} = \max_k W^p_k \tag{13}$$

$$RL^{sum}_{PB} = \sum_{k=0}^{I_p} W^p_k \tag{14}$$

We note that the *Path-Based* Amplification Level, contrary to the *Source-Based* Risk Level, is not contained in the $[0, 1]$ interval. Indeed, with the RL_{PB} we want to emphasize the effect of a large AS reflecting and amplifying erroneous announcements of a smaller AS having a high RL_{SB}.

VI. Implementation and result

To evaluate the Risk Levels, a Python software is developed to interface with a route server. The files containing the data of the RIB will be extracted from it, containing all the information that will be given as input to the software for the analysis. Once the RIB has been obtained, the information relating to the filtered prefixes need to be extracted, which will undergo the filtering process described in Section III.

The initial phase, focused on identifying suspicious prefixes, presented challenges due to lack of correct information. To address the issue, the software uses different tools:

- *BGPView*[3]: Given an AS number as input, this tool returns the prefixes associated with it;
- *Whois*: A tool that compares information with that contained in the IRRs [15].

[3]BGPView available at https://bgpview.io/

TABLE II: Values of the Parameters for the *Source-Based* Risk Level (RL_{SB})

Weight	Category Name	Value
α	Time	0.2
β	Frequency	0.1
γ	Path	0.2
δ	Filtering	0.1
ϵ	Owner	0.4

- *Routinator*[4]: given re prefix-AS number pair as input it is possible to validate it through the RPKI framework;

After identifying the attack-related prefixes via cascade filtering, the pertinent information is stores within a PostgreSQL database. This repository serves as a crucial resource for computing the RL_{SB} and RL_{PB}.

In Table II we report the numerical values of the weights for the different categories of the RL_{SB} computation, as shown in 1. The weights were determined through data tuning and detailed consultation with IXP operators. The main idea is to better protect the networks connected to the IXP.

A. Results

In this section, we present the outcomes derived from a comprehensive test carried out on the route server of the main Italian IXP. This particular route server has connectivity with 137 ASes and specializes in the exchange of IPv4 prefixes. Commencing on June 5, 2023, the evaluation spans over 8 months, concluding on January 28, 2024. We note that the assessment period culminates with the calculation of Risk Levels on the final day. Consequently, throughout a 34-week time frame, our software actively engages with the route server, systematically extracting data from its RIB.

In assessing the *Source-Based* Risk Level, our objective is to allocate a high score to Origin ASes announcing prefixes capable of compromise routing integrity. Conversely, the *Path-Based* Amplification Level focuses on regulating the conduct of ASes encountered along the path of potentially attacks. ASes exhibiting sub-optimal filtering policies are assigned elevated scores, as they amplify the dissemination of an attack.

As depicted in Table III, the route server filtered the majority of prefixes, amounting to 99.58%, categorized as configuration errors, thus exempt from the Risk Level calculation. A notable case involves AS 15720, as illustrated in Table IV, where four of its prefixes are propagated by four distinct ASes, totaling 120 instances. Examination of ASN reveals the presence of four private ASes associated with AS 15720. Due to inadequate configuration of filtering policies, AS 15720 lacks control over the Origin AS of received BGP UPDATE messages. This instance underscores the significance of adhering to proper configuration practices outlined in RFC 8212 [16], as it elucidates how a prefix might undergo filtering by the route server despite lacking any connection to an attack.

Over an eight-month period, the software successfully identified 35 potential attack-related prefixes, collectively announced 118 times. It's noteworthy that among the 2000-

[4]Routinator available at https://rpki-validator.ripe.net/ui/

TABLE III: Results of the Proposed Approach

Parameters	Results
Filtered Prefixes	8359
Configuration Error Rate	99.58%
Possible Attack-Related Prefixes	35
Announcement Frequency	118
Analysed Origin ASes	2011
Possible Attack-related Origin ASes	22
Possible Victim Owner ASes	25
Possible Attack-related ASes along the path	23

TABLE IV: Real Case of Configuration Errors made by AS 15720 which does not check the Origin ASN of its prefixes

Network	Owner AS	Origin AS	Number of times
62.241.17.0/24	15720	65515	30
62.241.19.0/24	15720	65517	30
62.241.20.0/24	15720	65510	30
62.241.21.0/24	15720	65506	30

plus Origin ASes analyzed, 22 of them were implicated in the dissemination of potential attacks, impacting the traffic flow of other 25 ASes. The reliability of the obtained results depends on the accurate identification of the relationship between the Origin AS announcing a prefix and its actual owner. As already discussed, this operation is complicated by the confidentially of AS relationships, thus increasing the possibility of encountering false positives.

Table V shows the importance of Owner and Path category scores in determining the *Source-Based* Risk Level. AS 51688 emerges as the most critical entity, being the only one AS with an Owner score exceeding 0. Fig. 1 visually illustrates AS 51688's announcement of '91.220.40.0/24' prefix, a subnet belonging to AS 3303, an AS directly connected to the IXP route server. Notably, AS 34554's routing decision towards AS 51688 leads to the loss of services offered by AS 3303. Securing the second to fifth positions are ASes with non-zero Path category scores. Among these, AS 37142 stands out due to two key factors. Firstly, its high Path category score stems from the announcement of a prefix with an AS_PATH length of 6, traversing through 5 other ASes. Secondly, the announcement of the attack-related prefix two weeks before the Risk Level calculation garnered a notable score in the Time category, consequently bolstering AS 37142's position.

Table VI presents the results for the *Path-Based* Amplification Level. The high values reported in the table show that a malicious prefix, originally announced by a small AS, was re-propagated by a larger AS. As an example, AS 6939 has RL_{PB} exceeding 1000, meaning that, on average, it re-propagated malicious announcements of ASes which have a thousand times less connections, therefore amplifying the attack. We can observe also ASes which have a value of RL_{PB} lower than 1, such as AS 29802 and AS 203639. Such ASes have re-propagated an attack-related prefix of a much larger AS. Therefore, the risk posed by the original announcement, indicated by the *Source-Based* Risk Level, should be considered more critical than the risk posed by its re-propagation by an intermediate AS. For instance, AS 29802, with only 3 connections, re-propagates an announcement originally broadcasted by AS 61317, having 91 connections. The re-propagation by AS 29802 is less critical than the original propagation by the source AS 61317, therefore AS 29802 gets a low value of RL_{PB}. In our analysis, the majority of ASes only re-propagate one attack-related prefix. This is shown in Table VI in the column I_p, which is the number of attack-related prefixes re-propagated by an AS. There are some notable exceptions, such as AS 6939 with $I_p = 17$. As a case study, in Fig. 7 we have reported the distribution of the weights W_k^p for the *Path-Based* Amplification Level of AS 6939, providing insights into the contributions of different ASes to the RL_{PB}. In our experiments all the weights W_k^p of AS 6939 have values greater than one, suggesting that even the

TABLE V: Ranking of Origin ASes Based on *Source-Based* Risk Level (RL_{SB})

Origin ASN	Time	Frequency	Path	Filtering	Owner	RL_{SB}
51688	0.867	0.125	0	1	0.796	0.604
37124	0.867	0.125	0.865	1	0	0.459
142519	0.368	0.125	0.950	1	0	0.376
9009	0.490	0.125	0.632	1	0	0.337
15731	0.018	0.125	0.865	1	0	0.289
6843	0.867	0.125	0	1	0	0.286
28716	0.570	0.375	0	1	0	0.252
216145	0.490	0.125	0	1	0	0.211
396982	0.490	0.125	0	1	0	0.211
198361	0.490	0.125	0	1	0	0.211
216172	0.240	0.125	0	1	0	0.160
197706	0.083	0.250	0	1	0	0.142
203639	0.082	0.250	0	1	0	0.141
57142	0.068	0.875	0	0.261	0	0.127
211886	0.012	0.250	0	1	0	0.127
33924	0.066	0.125	0	1	0	0.126
7155	0.062	0.125	0	1	0	0.125
57454	0.058	0.125	0	1	0	0.124
59741	0.021	0.125	0	1	0	0.117
39120	0.020	0.125	0	1	0	0.116
61317	0.009	0.125	0	1	0	0.114

TABLE VI: Ranking of ASes Ordered by Decreasing Values of *Path-Based* Amplification Level (RL_{PB})

ASN	I_p	RL_{PB}	RL_{PB}^{max}	RL_{PB}^{sum}
3356	1	1896.71	1896.71	1896.71
6939	17	1219.83	5883.56	20737.11
60501	1	545.92	545.92	545.92
37662	1	322.68	322.68	322.68
4755	1	293.66	293.68	293.66
6453	1	242.90	242.90	242.90
6830	2	157.60	178.22	315.21
5398	5	146.97	374.87	734.85
15169	1	40.83	40.83	40.83
12732	1	36.12	36.12	36.12
327708	1	33.05	33.05	33.05
34554	1	25.37	25.37	25.37
44901	4	23.47	37.56	93.89
132779	1	7.14	7.14	7.14
36926	1	6.88	6.88	6.88
37075	1	6.43	6.43	6.43
51395	2	3.04	3.04	6.08
56327	1	2.23	2.23	2.23
3257	2	2.15	2.15	4.31
56550	1	1.51	1.51	1.51
9002	2	1.36	1.36	2.72
203639	4	0.33	0.75	1.33
29802	1	0.13	0.13	0.13

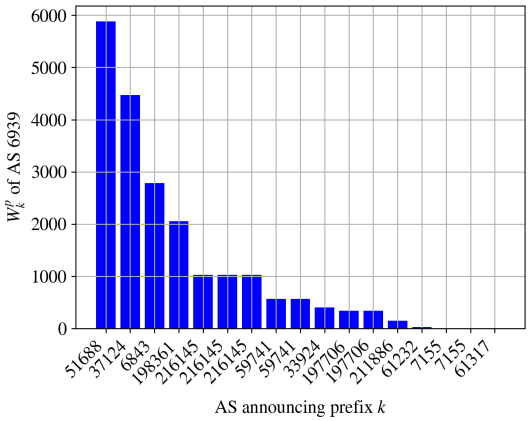

Fig. 7: Distribution of Path-Based Amplification Level weights W_k^p of AS 6939 for each source AS announcing the $k-$prefix. The same AS can appear multiple times since it may have originated multiple attack-related prefixes.

re-propagation of attack-related prefixes with the lowest value of W_k^p, which are not visible in the figure, are considerably reflected and amplified by AS 6939.

VII. Future Extensions

Our study currently focuses on defining and evaluating Risk Levels through data obtained from the RIB of a single route server within an IXP. However, it is important to note that an IXP typically operates multiple route servers, dedicated to exchange reachability information for both IPv4 and IPv6 prefixes. By expanding our software, we aim to extract and analyze data from the RIB of all route servers within an IXP. Since a prefix can traverse various paths and reach different IXPs by crossing different ASes, our next phase involves equipping a greater number of IXPs with this software. This expansion will enable us to identify the complete path taken by potential attacks, thereby obtaining all involved ASes.

Another aspect to consider is AS_PATH prepending. As highlighted in Section V, the scoring process for the *Path* category within the *Source-Based* Risk Level involves counting the exact number of ASes traversed by possible attacks. This was due to the elimination of any possible prepending manipulation within the AS_PATH. While BGP prepending is a legitimate tool for ASes to manage their routes, its excessive use, whether by the Origin AS or any AS along the route, can pose security risks [17]. Integrating this consideration into our software opens up the possibility of addressing new types of attacks. Moreover, it presents an opportunity to develop functionality aimed at identifying the optimal prepending value an AS can employ while maintaining benign behavior. Such scenario would empower Origin ASes to exercise greater control over their route paths. This proactive approach would help prevent ASes from deviating routes through excessive prepending or exploiting excessively prepended paths to favor shorter AS_PATH announcements.

VIII. Conclusions

We have introduces the Risk Level, a new metric designed to precisely quantify the threat posed by each AS to the entire Internet, based on the analysis of the attacks carried out by them. Given the current limitations in existing countermeasures against BGP attacks, the primary aim of the Risk Level is to help IXP operators in identifying ASes that could compromise routing, thus facilitating the identification of optimal countermeasures to implement. Furthermore, the Risk Level aims to promote awareness and widespread adoption of already existing countermeasures across all ASes. It is crucial to recognize that optimal protection requires the simultaneous development of all countermeasures to defend against each type of attack. Since attacks can be executed with real ease, we propose extending the software to a larger number of IXPs, providing an overall overview of BGP routing and enhancing its security through software tools operated by IXPs.

References

[1] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)." RFC 4271, Jan. 2006.
[2] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How secure are secure interdomain routing protocols," in *Proceedings of the ACM SIGCOMM 2010 Conference*, SIGCOMM '10, (New York, NY, USA), p. 87–98, Association for Computing Machinery, 2010.
[3] T. Tofoni, F. Luciani, and A. Prado, *BGP from theory to practice*. Reiss Romoli, Nov. 2023.
[4] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "Bgp hijacking classification," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 25–32, 2019.
[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*, pp. 1–4, 2023.
[6] E. Jasinska, N. Hilliard, R. Raszuk, and N. Bakker, "Internet Exchange BGP Route Server." RFC 7947, Sept. 2016.
[7] "Representation of IP Routing Policies in a Routing Registry (ripe-81++)." RFC 1786, Mar. 1995.
[8] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1." RFC 8210, Sept. 2017.
[9] B. Kuerbis and M. Mueller, "Internet routing registries, data governance, and security," *Journal of Cyber Policy*, vol. 2, no. 1, pp. 64–81, 2017.
[10] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark, "To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today," in *Passive and Active Measurement Conference (PAM)*, Mar 2020.
[11] T. Shapira and Y. Shavitt, "Ap2vec: An unsupervised approach for bgp hijacking detection," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2255–2268, 2022.
[12] K. McGlynn, H. B. Acharya, and M. Kwon, "Detecting bgp route anomalies with deep learning," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1039–1040, 2019.
[13] J. Zhang, D. Li, and B. Zhao, "A prefix hijacking detection model based on the immune network theory," *IEEE Access*, vol. 7, pp. 132384–132394, 2019.
[14] J. Mitchell, "Autonomous System (AS) Reservation for Private Use." RFC 6996, July 2013.
[15] T. Vinni, *WHOIS*, ch. 9, pp. 175–199. John Wiley & Sons, Inc., 2020.
[16] J. Mauch, J. Snijders, and G. Hankins, "Default External BGP (EBGP) Route Propagation Behavior without Policies." RFC 8212, July 2017.
[17] K. Wang, "Investigating the effectiveness of stealthy hijacks against public route collectors: Is as-path prepending enough to hide from public route collectors?," 2023.