Survey on Internet Routing Security: Stakeholder Interests, Current, and Future Research Directions

Stefano Servillo, *Graduate Student Member, IEEE*, Pietro Spadaccino, Marco Centenaro, *Member, IEEE*, Antonio Prado, Massimiliano Rossi, Flavio Luciani, Monica Scannapieco, and Francesca Cuomo, *Senior Member, IEEE*

***Abstract*—The functioning of the Internet depends on the exchange of routing information across Autonomous Systems (ASes) using the Border Gateway Protocol (BGP). The mutual trust between ASes and the complexity of routing mechanisms expose BGP to significant cyber-threats and attacks, posing risks to the broader Internet-dependent economy. This survey presents a multi-stakeholder perspective on BGP routing security, integrating academic, industrial, and institutional points of view. It reviews policy frameworks developed by both public and private organizations, outlining synergies between industry and institutions and best practices for securing inter-domain routing. The paper further examines open challenges, identifying four research and innovation directions. For each direction, we survey potential solutions and categorize them based on stakeholder interest, thereby bridging the gap between academic advances and operational deployment. Finally, we analyze potential future directions in BGP security, highlighting differentiated priorities and strategic interests across stakeholders.**

***Index Terms*—Cybersecurity, internet, network security, routing security, border gateway protocol, industry, institutions.**

I. Introduction

THE deployment of high-capacity connectivity has made the highways of the Internet available to everyone, enabling new forms of interaction, entertainment, and work. For example, cloud computing has shifted part of the application workload and storage from local, on-premise data centers to centralized, public ones. Another example is streaming services, which have become so widely available that they are making television broadcasting outdated. This has brought new players into the Internet arena, other than traditional Internet Service Providers (ISPs) and network operators, such as Cloud Service Providers (CSPs) and Content Providers (CPs).

In this context, the importance of the *cybersecurity* of the Internet is also increasing. In particular, the ability to securely exchange routing information between different parties is fundamental to the correct forwarding of packets, thus enabling the effective functioning of the global "network of networks". Since such information exchange is performed via the Border Gateway Protocol (BGP), the security of this protocol represents a critical issue for the management of networks worldwide. An incident affecting BGP can yield i) altering Internet traffic content (breach of *integrity*), ii) performing Internet traffic and metadata analysis or, even, eavesdropping Internet traffic content (breach of *confidentiality*), and iii) creating connection outages (breach of *availability*) [1].

A. Motivation

The role of public policy in BGP routing security has become increasingly prominent, particularly considering the recent geopolitical developments that have highlighted the vulnerability of critical infrastructures [2], [3]. In this regard, three main stakeholders play a significant role: **academia**, **industry**, and **institutions**. These stakeholders interact and influence each other, but each has distinct motivations, perspectives, and approaches to address BGP routing security challenges.

Academia is the most flexible stakeholder with respect to experimentation and risk tolerance, and typically operates at low-to-mid Technology Readiness Levels (TRLs) [4]. Its primary objective is the **disruption** of the existing BGP routing ecosystem, challenging assumptions to enable paradigm shifts that may lead to long-term improvements. Industry seeks mechanisms that can be incorporated into existing production environments, thus prioritizing high-TRLs solutions. Its primary objective is finding the **tradeoff** between the operational efficiency and the security of the BGP ecosystem, ensuring a risk-controlled evolution of deployed networks. Finally, institutions are concerned with **systemic security** promotion and governance. Their role consists of protecting critical infrastructures by providing normative policies, frameworks, and best practices for the security of the global routing ecosystem and assessing the compliance to such norms. Consequently, institutions typically operate at mid to high TRLs.

Given the centrality of BGP for global Internet connectivity, its security cannot be effectively addressed without taking into account the contributions and limitations of all these actors. Hence, the primary motivation of this survey is to provide a comprehensive overview of the current state of BGP routing security, consolidating in a single study the perspectives of all

Received 28 March 2026; revised 16 June 2026; accepted 15 July 2026. Date of publication 17 July 2026; date of current version 24 July 2026. This work was supported in part by SERICS project under the MUR National Recovery and Resilience Plan funded by the European Union - Next Generation EU under Grant PE00000014. *(Corresponding author: Stefano Servillo.)*

Stefano Servillo, Pietro Spadaccino, and Francesca Cuomo are with the Department of Information Engineering, Electronics and Telecommunications (DIET), University of Rome "La Sapienza," 00184 Rome, Italy (e-mail: stefano.servillo@uniroma1.it; pietro.spadaccino@uniroma1.it; francesca.cuomo@uniroma1.it).

Marco Centenaro, Massimiliano Rossi, and Monica Scannapieco are with the National Cybersecurity Agency of Italy, 00198 Rome, Italy (e-mail: m.centenaro@acn.gov.it; m.rossi@acn.gov.it; m.scannapieco@acn.gov.it).

Antonio Prado and Flavio Luciani are with Namex, 00185 Rome, Italy (e-mail: antonio@prado.it; f.luciani@namex.it).

Digital Object Identifier 10.1109/COMST.2026.3714569

TABLE I
COMPARISON OF THIS WORK WITH EXISTING SURVEYS AND OTHER LITERATURE IN THE FIELD

Ref.	BGP Key Topics	BGP Security Aspects	Research Questions	Industrial Perspective	Institutional Perspective
[5]	Routing Anomalies, Machine Learning, Attack Detection	○	×	×	×
[6]	Security, Attack Visualization, Anomaly Detection	○	○	×	×
[7]	Security, Machine Learning, Attack Detection	○	○	×	×
[8]	Security, Attack Detection, Detection Techniques	●	○	×	×
[1]	Security, Attack Detection, Mitigations, Open Questions on Security	●	●	×	×
[9]	Security, Attack Detection, Mitigations, Open Questions on Security	●	●	×	×
[10]	Security, Attack Detection, Mitigations	●	●	×	×
[11]	Security, Anomaly Detection, Future Detection Needs	●	●	×	×
[12]	Security, Mitigations, Open Questions on Security	●	○	○	×
[13]	Security in Partially Deployed Frameworks, Guidelines for Industry	●	×	●	×
[14]	Security, Incidents due to Misconfigurations by Industry and Practitioners	○	×	●	×
[15]	Data Collection for Security, Monitoring Systems for Practitioners	×	○	○	×
[16]	Strategies for Governments and Industry for the Adoption of Security Frameworks	×	×	●	●
[17]	Government Strategy and Techniques to Control Routing, Censorship	×	×	○	●
	This Work	●	●	●	●

The symbols ●, ○, × indicate that the work fully covers, partially covers, or does not cover the topic, respectively.

relevant stakeholders and surveying how each state-of-the-art approach fits their specific point of view.

B. Related Work

Existing surveys and, in general, research efforts have addressed the topic of BGP routing security under various perspectives. We summarized such proposals in Table I, reporting if they provide overviews of BGP routing security aspects as well as whether they investigate open research questions and future direction from the academic, industrial, or institutional perspective. Each proposal has a different symbol and color, whether it fully covers, partially covers, or does not cover a specific field. The selection focuses on all surveys published since [1], which provides the conceptual starting point and motivation for the present work.

Authors in [5], [6], [7], and [8] provide valuable overviews of BGP attack detection techniques. They explore various approaches, including visualization methods and machine learning-based solutions, to identify and mitigate routing anomalies. However, these studies primarily focus on technical detection mechanisms and do not offer substantial insights into open research questions or future directions in the academic field. Moreover, they lack coverage of the perspectives from industry operators and institutional entities, which are critical for understanding real-world deployment challenges and policy implications. Then, works such as [1], [9], [10], and [11] provide an in-depth survey of BGP security aspects, offering a comprehensive overview of various types of attacks and corresponding detection and mitigation mechanisms. Beyond technical content, these studies also address open research questions and challenges; the authors explicitly discuss current limitations and outline potential directions for future academic work. However, despite their thorough examination of technical and theoretical aspects, these surveys do not cover the perspectives of industry operators or institutional stakeholders.

Some publications address the industry perspective only partially. In particular, while not explicitly focused on industry practices, certain studies provide valuable insights that can help infer industry-related challenges and needs. For instance, [13] investigates BGP security in contexts where security frameworks are only partially deployed. The authors show how partial adoption can still yield tangible security benefits. This is highly relevant for the industry, where full and synchronized adoption of a security standard is often slow and difficult to achieve. Analyzing how security frameworks perform under partial deployment allows operators to better assess and manage the expected security level, even in environments where widespread adoption is still incomplete. Continuing, the work in [14] discusses common incidents and misconfiguration errors made by network operators. The paper highlights that many of these errors stem from complex or unintuitive configurations, and concludes that an improved routing design

could prevent a significant number of operational mistakes. These findings are especially relevant to the industry, as they suggest actionable improvements that could enhance network reliability in real-world deployments. Authors in [15] focus on data collection systems for BGP security. These systems play a critical role in enabling monitoring and anomaly detection. From an industry perspective, such tools are particularly valuable because they can often be deployed externally with minimal setup, making them easier to adopt and integrate into existing infrastructures compared to full security frameworks adoption or redesign. Nevertheless, these studies do not offer a comprehensive survey of BGP routing security. Their scope is typically limited to specific challenges, and thus, they do not address the broader landscape of attacks, mitigation techniques, and research directions that a general security survey would encompass. Additionally, they do not consider the perspective of institutional and governmental entities. These actors differ significantly from industry players: they operate under distinct governance models, must align with public policy goals, and are often subject to international agreements and regulatory frameworks.

Another set of works, by contrast, explicitly addresses the synergies between industry and public institutions. For example, [16] discusses strategies for governments and industry to support the adoption of BGP security frameworks. The authors examine various operational and policy-driven strategies, including market-related factors, to promote secure routing practices. The work in [17] investigates mechanisms used by international governments to exert control over Internet routing, including practices such as content censorship and traffic filtering. These works adopt a broader view that incorporates institutional and governmental perspectives and touch on the operational realities faced by both public and private actors. However, they lack a focused examination of BGP routing security aspects and do not delve into academic research questions. This is largely because they are aimed at a more operational audience, prioritizing policy and deployment over technical or theoretical analysis.

While existing literature tends to focus largely on a single stakeholder point of view in isolation, to the best of our knowledge, no previous work has attempted to integrate all of them within a single, unified survey. This holistic approach is essential to develop a complete understanding of the dynamics of Internet routing security from all stakeholders' perspectives.

C. Main Contributions of Our Manuscript

The **core contribution** of our work lies in the investigation of the **full range of perspectives** from the **various stakeholders** involved in BGP routing security, namely academia, industry, and public and private institutions. By doing so, our study fills a critical gap in the existing literature, which has so far considered these viewpoints only in isolation or partially. In examining the surveyed contributions, we approach them from the perspective of each stakeholder. In the following, we report the main contributions of our manuscript.

1) We start by proposing a **structured in-depth taxonomy** of BGP routing attacks, grouping them based on their characteristics and the aspects of BGP they target. This categorization is designed to help all stakeholders to **precisely classify** and prioritize BGP routing security issues. We emphasize that several elements of the proposed taxonomy refine and standardize distinctions that are already well established in the BGP security community. In fact, our contribution is not the invention of entirely new attack notions, but the harmonization of macro- and micro-level categories under a single coherent structure and consistent acronyms.

2) We explore how the different stakeholders have **approached** the task of securing Internet routing, including the adoption of common standards, procedures, and best practices. Specifically, we survey the main directions that have been **investigated** by the academia, the main **solutions** that have been implemented by the industry and the main **policy frameworks** proposed by public and private institutions. Moreover, this section will discuss notable efforts to address BGP security, emphasizing the collaboration between governments, private companies, and industry groups to develop effective solutions. These efforts are vital not only for securing national networks but also for ensuring the overall resilience of the global Internet infrastructure.

3) Based on the proposed approaches, we **tackle open research questions** identified in prior works and **examine** the evolution of BGP routing security challenges. Specifically, we revisit the questions posed by [1] and **review the literature** addressing these issues from the **perspective of each stakeholder**. We further introduce an **interest metric** that quantifies a stakeholder's interest to pursue a given initiative, and use it to assign an interest score to each reviewed work for academia, industry, and institutions.

4) We analyze the **evolution** of the identified open questions and examine each from the perspective of **academia, industry, and institutions**. For each identified research question, we propose actions tailored to stakeholder roles and employ the interest metric to quantify their **relevance**. This allows us to explicitly characterize how stakeholder priorities align with each research direction.

D. Manuscript Organization and Definitions

The remainder of this manuscript is organized as follows. Sec. II revisits the fundamentals of Internet routing, highlighting the principal technological trends and their evolution over time. Sec. III surveys attacks targeting BGP ecosystem and introduces a structured taxonomy based on attack characteristics and targets. Building on this taxonomy, Sec. IV provides an overview of best practices across academic, industrial, and institutional domains to enhance routing security, and identifies attack classes that remain insufficiently explored or mitigated. Based on such gaps, in Sec. V we revise and extend the previous work from [1] by identifying four prominent R&I directions and investigating to what extent they have been addressed by the community by means of a quantitative evaluation of the proposed approaches. Sec. VI considers the level of interest of the different stakeholders in pursuing and improving the identified R&I directions to sketch the envisioned future trends as well as extract a set of cross-cutting insights leveraging the entire landscape of BGP

routing security surveyed in this work. Finally, we draw the conclusions of this paper in Sec. VIII.

The terms and abbreviations used in the rest of this manuscript will be, as far as possible, consistent with the RFC style guide [18]. Moreover, we will distinguish between a "network operator" and an "ISP" to highlight that the former manages a telecommunication infrastructure that is typically exploited by other organizations (business-to-business model), and the latter provides access to the Internet to end users (business-to-consumer model) [19].

II. Fundamentals of Internet Routing

In this section, we will briefly describe the main technological trends relevant to the scope of this survey, and then we will recap the Internet routing architecture and BGP basics.

A. The Current Shape of Internet Technological Landscape

As a decentralized network of interconnected ASes, the Internet ecosystem is extremely fragmented. Nonetheless, ISPs are traditionally classified into three categories: i) Tier 1, which form the global backbone, ii) Tier 2, which purchase transit from Tier 1 providers while also peering with other Tier 2, offering regional or national services, and iii) Tier 3, which focus on last-mile delivery, providing access to end users.

However, the Internet traffic has changed a lot in the last decade due to two main technological trends. On the one hand, there have been increasing investments for the widespread deployment of legacy telecommunication technologies supporting broadband traffic, such as:

- *optical fiber*. According to the latest investment tracker report of the FTTH Council Europe [20], following a record-breaking peak in 2022, the FTTx financing market has seen reduced but stable activity. Although transaction volumes and deal counts remained high in recent years, activity has gradually slowed since 2022, with a more cautious outlook for 2025 amid a challenging market environment [21]. In terms of the United States, the Fiber Broadband Association [22] reported that the market continues to show substantial growth potential, emphasizing the continuing demand for fiber infrastructure from network operators and/or ISPs of all tiers. In addition, investments in fiber by private equity and government programs are expected to sustain and accelerate deployment efforts in the coming years [23];
- *submarine cables*. Investments in new digital infrastructures of this kind, which represent the backbone of global connectivity [24], are soaring, with the so-called Big Tech industries (such as, e.g., Google, Meta/Facebook, Amazon, Microsoft) that have recently become the main investors in submarine cables as opposed to Tier 1 ISPs [25]. This increase not only aims to support higher traffic, but, due to geopolitical reasons, also to increase redundancy and diversification of the backbones.

On the other hand, we have witnessed the development of several emerging technologies for digital infrastructures, like:

- *broadband cellular networks*, that is, Long-Term Evolution (LTE) [26] and fifth-generation (5G) systems [27]. With the deployment of these mobile network generations, broadband access has been granted by ISPs to billions of so-called user equipment (i.e., smartphones and tablets), enabling an unprecedented quality of experience in mobility conditions. The enhanced mobile broadband provided by 5G is also exploited for fixed wireless access, which functionally replaces optical fiber in some remote geographical areas [28];
- *Low-Earth Orbit (LEO) satellite networks* [29]. As a means to complement or speed up the broadband network coverage of optical fiber and mobile networks, both private organizations and public institutions, e.g., the European Space Agency and the U.S. Department of Defense, are investing to deploy low-orbit constellations of satellites, such as e.g., Starlink and IRIS [30];
- *cloud infrastructures* [31], deployed worldwide by CSPs [32] to empower all kinds of organizations, including enterprises [33], with flexible solutions based on virtual networking and computing resources to run applications without caring about the maintenance of the underlying hardware;
- *Content Delivery Networks (CDNs)*, which are geographically distributed networks of proxy (edge) servers providing faster access to digital content to end users. Unlike cloud computing, which focuses on scalable and on-demand delivery of applications and data, CDNs aim to deliver content as quickly as possible to end users using edge servers. This is the typical infrastructure adopted by CPs such as, e.g., Netflix and Disney+, to distribute audiovisual material – with an expected global market Compound Annual Growth Rate (CAGR) of 11% in the period 2023-2028 [34].

These two technological trends have yielded the increase of overall capacity across all network tiers, which was necessary to satisfy the demand for all new offered services, but also have been challenging the very same, ISP-centric Internet ecosystem. For example, the network infrastructure of large CSPs, also referred to as hyperscalers, are on par with (or even exceed) those of Tier 1 ISPs – with the difference that they do not fit the Tier 1 category as they do not sell transit of their own network to third parties, rather they exploit it for their own purposes. Therefore, the modern Internet architecture can be considered, with a more relaxed hierarchy, as a complex ecosystem where two types of large-scale global backbones (Tier 1 ISPs and hyperscalers) interconnect with each other and with thousands of smaller networks below them.

B. Routing Architecture: Evolution Over Time

As shown in Fig. 1, the routing of traffic across the Internet can be represented at two different levels, depending on the relationships within and outside an AS. The AS is an essential element of the previously described Internet ecosystem and is defined [35] as an interconnected set of one or more Internet Protocol (IP) prefixes managed by one or more network operators that has a *unique* and *clearly defined* routing policy. The two mentioned plans are as follows:

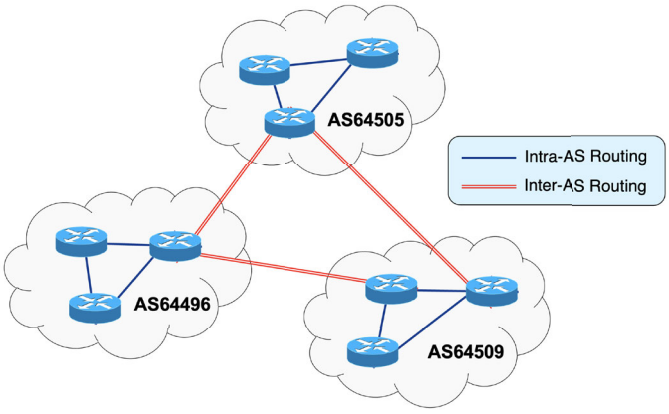

Fig. 1. Sample Internet routing architecture. We distinguish among three ASes (denoted as "AS64496", "AS64505", and "AS64509"), whereby the solid blue line and the red line represent the intra-AS routing domain and inter-AS routing domain, respectively.

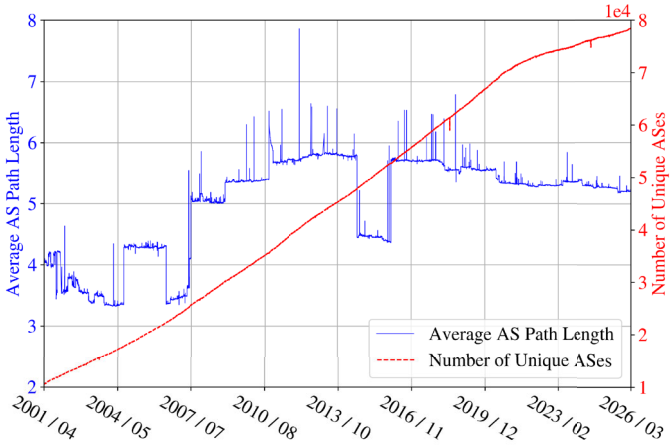

Fig. 2. Average AS_PATH length (blue) and number of unique ASes (red).

- *intra-AS routing* is composed of distinct collections of routing elements. According to the Number Resource Organization (NRO) reports more than 130,000 ASes assigned by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIRs),[1] under the supervision of the Internet Corporation for Assigned Names and Numbers (ICANN) [36]. In Q4 2025, around 118,000 ASNs were assigned by RIRs, with a total of 78,270 ASes active on the Internet, defined as networks that have made at least one announcement [37]. Note that the ASes' owners may not only be ISPs, but also CSPs, CPs, large enterprises, etc., thus making the panorama of Internet routing administration and management extremely heterogeneous. Each of these ASes is identified by an AS Number (ASN) issued by the respective RIR and is assigned a set of IP address prefixes. At this level, different Interior Gateway Protocols (IGPs) can be used for routing within a single AS, such as, e.g., Open Shortest Path First (OSPF) [38] and Intermediate System to Intermediate System (IS-IS) [39];
- *inter-AS routing*, allowing the establishment of connections among different ASes. BGP serves *de facto* as the only External Gateway Protocol (EGP), allowing different ASes to announce and select routes based on policies, economic agreements, and network performance. Other than for exchanging routing information between ASes, *exterior* BGP (eBGP), the protocol can be used to extend the reach of routing information from the edge of the AS inward, *interior* BGP (iBGP).[2]

In this paper, we will focus on inter-AS routing, since it involves routing packets across ASes belonging to different network operators and, thus, necessitates the verification of BGP routing information exchanged by ASes themselves. Depending on the topology of inter-AS connections, the amount of ASes a packet traverses, i.e., the length of the AS_PATH attribute, varies – Fig. 2 shows the evolution in time of the average AS_PATH length [36]. The graph should be interpreted with consideration for the fluctuations in the maximum number of ASes operating within the network over time. During the early 2000s, the number of active ASes was approximately 10,000, whereas presently it has grown to just over 70,000, as mentioned previously. However, as shown in the figure, as the number of ASes increases, the average length of the AS_PATH does not increase proportionally. This phenomenon can be interpreted as an indication of the augmented rate of interconnection among ASes, a process that effectively mitigates the average length of the AS_PATH.

TABLE II
TOP 5 IXPS WORLDWIDE BY PEAK EXCHANGED TRAFFIC IN 2026

Name	Reference	Country	Peak Exchanged Traffic
IX.br	[44]	Brasil	32 Tbps
DE-CIX	[45]	Germany	18 Tbps
Pit-Chile	[46]	Chile	15 Tbps
AMS-IX	[47]	Netherlands	14.22 Tbps
LINX	[48]	United Kingdom	11 Tbps

Moreover, to improve interconnection, AS owners often rely on hubs known as IXPs or Network Access Points (NAPs), which are typically regarded as neutral entities [40]. Each of these infrastructures, which are spread all over the world [41], provides a peering platform that facilitates bilateral and multilateral agreements, promotes efficient traffic exchange, and improves overall network resilience. Currently, there are around 80 IXPs worldwide that have reached a traffic peak of over 1Tbps [42]. Tab. II presents the top 5 IXPs based on the peak exchanged traffic in 2026, with a combined total of 90.22 Tbps. Among these, the South American IX.br in Brazil and the European DE-CIX in Germany account for the majority of the exchanged traffic. As noted in [43], despite the significant traffic exchanged at these IXPs, the amount is one order of magnitude lower than that of a large ISP, even though AS adjacencies have improved and the number of datacenters associated with IXPs has increased, but the shape is the same.

[1]Five RIRs operate worldwide: RIPE NCC (Europe, Middle East, and Central Asia); ARIN (North America); APNIC (Asia Pacific); AFRINIC (Africa and Indian Ocean region); LACNIC (Latin America and Caribbean).

[2]BGP is built on top of the Transmission Control Protocol (TCP). However, while eBGP neighbors shall be directly connected, iBGP can leverage the IGP in use within the AS to reach a neighbor for its peering connection.

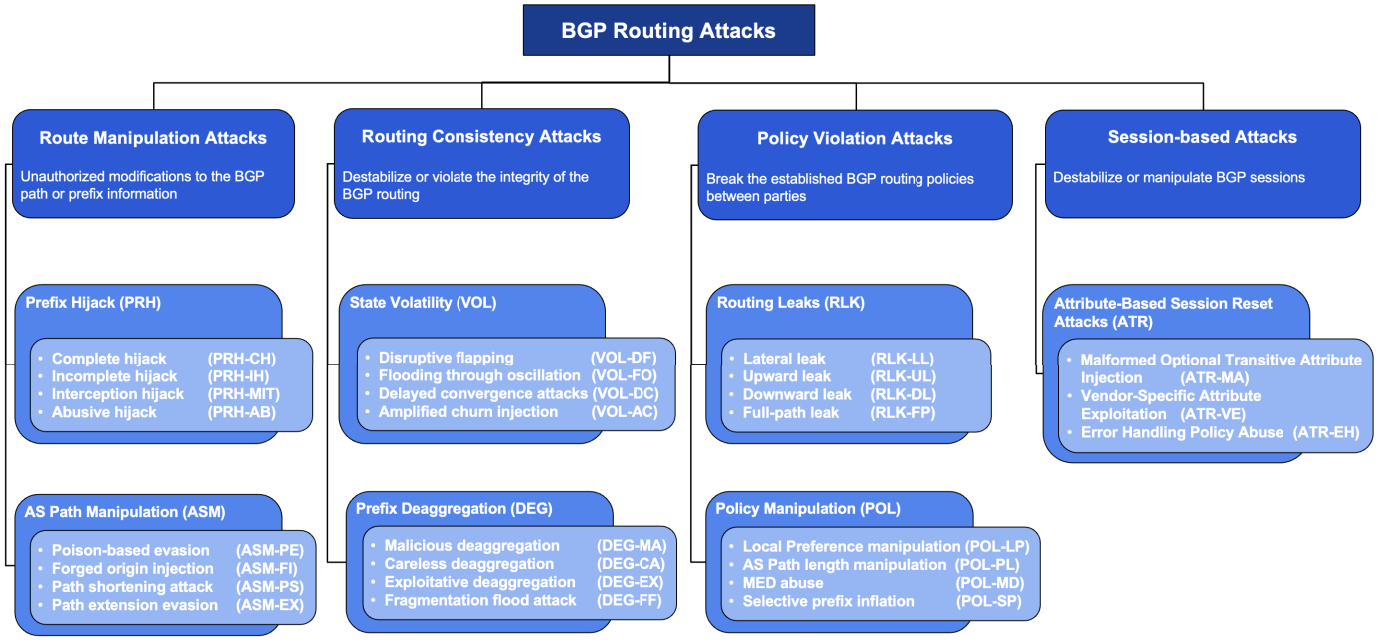

Fig. 3. Macro and micro taxonomy of BGP routing attacks used throughout this survey.

C. BGP Basics

As a path vector routing protocol for the Internet, BGP determines the best path for data packets based on a selection process in which one of the key elements is *distance*, in terms of the number of ASes. The protocol employs five kinds of messages: i) `OPEN`, for setup and establishment of a session with a neighboring BGP router; ii) `UPDATE`, for advertisement, update, or withdrawal of routes; iii) `NOTIFICATION`, for indication of an error condition to a BGP neighbor; iv) `KEEPALIVE`, to ensure that BGP neighbors are still alive; v) `ROUTE-REFRESH`, for requesting the resend of routing information. When multiple paths to the same IP prefix are available – each received from a different BGP neighbor – the router applies a hierarchical decision process to select the best path [49]. The first criterion is the Local Preference, a configurable value that indicates route priority. If Local Preference values are equal, BGP then considers the `AS_PATH` length, preferring the shortest path. However, if multiple paths still tie, other attributes such as origin type, Multi-Exit Discriminator (MED), or BGP identifier are evaluated.[3]

III. Main Attacks Against Internet Routing

The reliability of the information exchanged via BGP is entirely dependent on the trustworthiness of the peering ASes. Both accidental misconfigurations and malicious actors can compromise the integrity of inter-AS routing. In such scenarios, the information received from a given AS via BGP may be unreliable. In addition, BGP lacks built-in verification mechanisms, making it impossible to ensure the accuracy and integrity of received routing information *at runtime* – but it may only be detected *a posteriori* with different means such as, e.g., via data science methods [11], [60]. All these vulnerabilities can lead to suppression of routing information, alteration of the route object that is being passed, or the generation of spurious routing objects – cf. [1, § IV-C]. Such attacks are highly diverse, originating from different sources and producing a wide range of effects on routing behavior. Hence, it becomes essential to categorize them systematically.

In the scientific and standardization literature, there exist various taxonomies of BGP routing attacks. However, they are either at a very high level, like [61, cf. Security challenges of the Border Gateway Protocol], [62, cf. Possible attacks on BGP] or they treat very sectorial aspects, such as [51]. To the best of our knowledge, existing taxonomies remain incomplete with respect to providing an integrated view of research, innovation, and institutional dimensions. Thus, here we propose a harmonized taxonomy of BGP routing attacks, aimed at thoroughly categorizing the various types of attacks based on their scope and the underlying mechanisms they exploit. In particular, as shown in Fig. 3, we first distinguish four macro-categories of attacks, namely i) route manipulation attacks, ii) routing consistency attacks, iii) policy violation attacks, iv) session-based attacks. Each macro-category is then refined into multiple micro-categories. Within each macro-category, we describe routing incidents encompassing both illustrative scenarios and real-world cases. Each incident is sequentially indexed as [A*n*], where *n* denotes the progressive identifier. A summary of the presented incidents is then provided in Table IV, including source classification, affected autonomous systems, and corresponding attack categories.

Relation to prior taxonomies. The proposed taxonomy (Fig. 3) harmonizes categories that prior surveys and RFCs either treat at a high level or in a specialized manner. To clarify this relationship, Table III provides a crosswalk from common attack labels in the literature to our macro/micro taxonomy, highlighting where prior notions are i) split into finer micro-classes, ii) merge from scattered terminology,

[3]As mentioned in the introduction, for further information about the design and operation of BGP per se, we refer the reader to [1, § 3].

TABLE III

CROSSWALK BETWEEN BGP ATTACK CATEGORIES USED IN PRIOR LITERATURE AND THE MACRO-MICRO TAXONOMY ADOPTED IN THIS SURVEY (FIG. 3). THE "OP." COLUMN DENOTES THE MAPPING OPERATION: UNCHANGED CATEGORIES (SAME), REFINED INTO SUBCATEGORIES (SPLIT), COMBINED FROM MULTIPLE NOTIONS (MERGE), OR REASSIGNED TO A DIFFERENT MACRO-CATEGORY (RECLASS)

References	Typical label / notion in prior work	Proposed macro-category(ies)	Proposed micro-category(ies)	Op.
[9], [10], [12], [50]	Exact-prefix/more-specific hijacks; complete/incomplete/interception variants	Route Manipulation	PRH-CH, PRH-IH, PRH-MIT (+ PRH-AB)	Split
[10], [12]	Poisoning, forged-origin injection, path shortening/extension	Route Manipulation	ASM-PE, ASM-FI, ASM-PS, ASM-EX	Split
[51]	Policy-violating propagation	Policy Violation	RLK-UL, RLK-LL, RLK-DL, RLK-FP	Reclass
[10], [12]	LocalPref/MED manipulation, selective export, community-based steering	Policy Violation	POL-LP, POL-MD, POL-PL, POL-SP	Merge
[10], [12], [52]	Route flapping, oscillations, amplified churn, delayed convergence	Routing Consistency	VOL-DF, VOL-FO, VOL-AC, VOL-DC	Reclass
[10], [12]	Table growth, fragmentation/flooding via many specifics	Routing Consistency	DEG-CA, DEG-MA, DEG-EX, DEG-FF	Reclass/Split
[53]	Malformed optional transitive attributes; vendor-specific handling; session resets	Session-based	ATR-MA, ATR-VE, ATR-EH	Reclass

TABLE IV

SUMMARY OF BGP ROUTING INCIDENTS AND ILLUSTRATIVE EXAMPLES

ID	Attack Class	Description	Source Type	Affected AS
[A1]	PRH-CH/IH	2022 Twitter/X prefix diversion	[54] Academic Literature	RTComm, Twitter (X), Global Downstream ASes
[A2]	PRH	2020 Rostelecom accidental prefix announcements	[55] Industry Report	AS 12389, Multiple Global ASes (Downstream)
[A3]	ASM-PE	Legitimate traffic engineering via AS_PATH poisoning	Illustrative Example	Traffic Originator, Bypassed Intermediate AS, Destination AS
[A4]	ASM-FI	Forged origin injection to appear directly connected	Illustrative Example	Malicious AS, Target Prefix Owner, Receiving Networks
[A5]	ASM-PS	Path shortening by stripping intermediate ASNs	Illustrative Example	Stripping AS, Stripped Intermediate AS, Legitimate Destination
[A6]	ASM-EX	Path extension evasion via fake ASN prepending	Illustrative Example	Attacker AS, Upstream Provider, Competing Network
[A7]	DEG-EX	IPv6 /29 to multiple /32 deaggregation for traffic control	Illustrative Example	Originating ISP, Downstream Customer Networks
[A8]	DEG-FF	IPv6 /29 to 524,288 /48 fragmentation flood	Illustrative Example	Malicious Actor, Global BGP Resolvers, IXP Route Servers
[A9]	RLK-FP	2019 Safe Host misconfiguration propagating 70k routes	[56] Industry Report	Safe Host, China Telecom, Global Downstream/Peer ASes
[A10]	RLK-UL	2025 AS 52863 upward route leak (20k routes)	[57] Industry Report	AS 52863, Provider AS, Downstream ASes
[A11]	POL-SP	Surgical interception via selective community propagation	[58] Academic Literature	Attacker AS, Selected Upstream Providers, Target Regional ASes
[A12]	ATR-EH/MA/VE	May 20, 2025 malformed Prefix-SID cascade	[59] Industry Report	Asia-Pacific Origin AS, Starlink, Disney, Zscaler, ByteDance, Global IXPs

or iii) reclassified under the "Routing Consistency" macro-category. Several micro-classes follow established distinctions in the BGP security literature (e.g., hijack subtypes and route-leak definitions), while the main contribution of our taxonomy is the integrated macro/micro structure and consistent acronyms used throughout the survey, enabling systematic linkage to defenses and open challenges discussed later.

In the following, we describe the identified macro- and micro-categories; please note that the taxonomy presented in the following provides acronyms for BGP attacks that will be referenced throughout the remainder of the paper.

A. Attack Macro-Class #1: Route Manipulation Attacks

These attacks involve unauthorized modifications to BGP path or prefix information by altering the `AS_PATH` or advertising prefixes. Such manipulations can hijack traffic, impersonate legitimate routes, or inject false prefixes into the BGP system. These actions undermine routing integrity and may lead to traffic interception, blackholing, or other security issues. In this category, we find attacks belonging to the micro-categories *Prefix Hijack* and *AS_PATH Manipulation*.

1) Prefix Hijack: BGP hijacking, also known as *Prefix Hijacking (PRH)*, is one of the most dangerous attacks against BGP. An AS, intentionally or through malfeasance, advertises one or more IP prefixes belonging to another AS. This causes traffic to be redirected to the wrong AS, with serious consequences for service security and availability [63].

Certain types of hijacking are classified as *Complete Hijack (PRH-CH)*, where the attacker announces the exact IP prefix belonging to the legitimate AS. In *Incomplete Hijack (PRH-IH)*, the attacker advertises a more specific prefix than the legitimate ones. Since BGP selects the longest prefix match, traffic to that subprefix is routed to the hijacker's AS. In *Interception Hijack (PRH-MIT)*, the attacker announces a prefix to attract traffic, forwards it to the legitimate destination after interception, and inspects the data in transit, effectively acting as a man-in-the-middle without visibly disrupting service. Finally, another type of hijacking can be identified, the so-called *Abusive Hijack (PRH-AB)*, where the attacker announces an IP prefix not currently advertised by its rightful owner This often occurs with an unassigned or unannounced address space, making detection more difficult since there is no conflicting legitimate announcement. As a result, this type of attack can remain unnoticed for extended periods. This subdivision aligns with established hijack classifications (e.g., exact-prefix, sub-prefix, and interception hijacks [50]).

There have been several significant hijacking incidents in the recent history of the Internet. One of the most notable occurred in 2022, when the Russian ISP RTComm began announcing a Twitter (now X) IP prefix, diverting a portion of the platform's traffic [54] [A1]. Similarly, in 2020, another Russian provider, Rostelecom (AS 12389), mistakenly announced numerous IP prefixes belonging to other networks, effectively hijacking traffic from multiple organizations worldwide [55] [A2].

2) AS_Path Manipulation: AS_Path Manipulation (ASM) refers to scenarios in which an attacker modifies the `AS_PATH` attribute of a BGP UPDATE message. One class, *AS_PATH*

poisoning (*ASM-PE*), involves the deliberate insertion of fake ASNs into the `AS_PATH` to influence routing decisions of other ASes. Upon receiving an UPDATE, an AS checks whether its ASN appears in the `AS_PATH`; if so, the route is discarded to avoid the loop. By "poisoning" the `AS_PATH`, an attacker can cause targeted ASes to drop specific routes, thereby rerouting traffic along an alternative, potentially malicious path. This mechanism can be used both offensively (to divert or block traffic) and defensively (to avoid certain unwanted paths). Although *ASM-PE* is sometimes used for legitimate traffic engineering, it can be exploited to intercept, block, or redirect Internet traffic [64]. As an example of legitimate use, suppose that AS64500 wants to send traffic to AS64540, but current routing forces it through AS64520 and AS64530, which is undesirable due to latency or business policy. To prevent AS64520 from using the undesirable route, AS64500 advertises the route to AS64540 with the modified `AS_PATH 64500 64520 64540` [A3]. By inserting AS64520 into the `AS_PATH`, AS64520 will detect its own ASN and discard the route to avoid a loop. As a result, traffic is routed through an alternative path that bypasses it.

A different technique is *Forged origin Injection* (*ASM-FI*). In this attack, an AS manipulates the `AS_PATH` to make it appear as though it is directly connected to a target AS, even when it is not. For instance, if AS64510 (a malicious AS) wanted to redirect traffic destined for AS64540, it could advertise an IP prefix belonging to AS64540, but manipulate the `AS_PATH` as `64510 64540` [A4]. From the perspective of other networks, this makes AS64510 appear directly connected to AS64540. Neighboring ASes may prefer this shorter path and send traffic to AS64510, enabling the attacker to intercept, edit, or drop the traffic before forwarding (or not) to AS64540.

Another class is the *Path Shortening attack* (*ASM-PS*), where an AS deliberately removes legitimate ASNs from the `AS_PATH`, to make its path appear shorter and more attractive. For example, if the original path is AS65001 AS65002 AS65003, the attacker at AS65001 could strip AS65002 from the path and announce it as AS65001 AS65003. This manipulation misleads other ASes into believing the path is shorter or more direct than it is, thus biasing route selection and attracting traffic to the attacker's network [A5].

Conversely, in a *path EXtension evasion (ASM-EX)* attack, an AS artificially inflates the `AS_PATH` length by inserting extra ASNs (sometimes fake or private ones [65]) to penalize a route or make it less likely to be selected. This may steer return traffic away from a particular AS or to avoid triggering specific routing policies (e.g., shortest-path preferences). For instance, AS64512 might prepend several fake ASNs to make its announcement appear less attractive, thus influencing the upstream AS to prefer a competitor's path [A6].

B. Attack Macro-Class #2: Routing Consistency Attacks

These attacks target the stability and integrity of BGP routing by destabilizing the network's routing tables or exploiting inconsistencies in routing information. They can cause frequent changes in the routing state, leading to instability, network congestion, or delayed traffic delivery [66]. Examples include attacks that exploit the volatility of routing paths or introduce false route advertisements that disrupt the expected path selection process, undermining the reliability of the entire routing system. In this category, we find attacks belonging to the micro-categories *State Volatility* and *Path Deaggregation*.

1) State Volatility: The stability of the BGP routing system is critical for efficient inter-domain traffic management. When UPDATE messages become unpredictable, they can trigger widespread routing disruptions. This instability is exploited in *State Volatility (VOL)* attacks, which induce persistent fluctuations in inter-AS routing behavior. A common example is BGP route flapping [52], where a BGP router rapidly and repeatedly advertises and withdraws the same routes. Such instability confuses neighboring routers and triggers frequent reconfiguration of routing tables across ASes receiving the unstable updates [67].

Such behavior can arise from a variety of causes, including technical issues (hardware or software instability), misconfiguration (incorrect or inconsistent policy settings), or intermittent network conditions where links repeatedly go up and down. This latter case exemplifies *Flooding Through Oscillation (VOL-FO)*, a scenario where physical or logical link instability triggers a wave of BGP updates, leading to unnecessary routing churn. In more severe cases, an attacker may intentionally induce route flapping as a form of *Disruptive Flapping (VOL-DF)*, using the instability itself as a DoS vector [68]. The consequences of such an attack or malfunction are significant; constant route changes force routers to recompute routes continuously, consuming CPU and memory resources, and potentially pushing devices to their processing limits, causing them to become congested or blocked.

As this routing churn propagates, it triggers a phenomenon known as *Amplified Churn Injection (VOL-AC)*, where a local instability leads to widespread updates across the global BGP system. This results in frequent changes to routing tables, increased control-plane traffic, and network instability.

Another critical consequence is *Delayed Convergence Attacks (VOL-DC)*, where the routing system takes significantly longer to reach a stable state due to persistent fluctuations. This extended convergence time can cause latency spikes, packet loss, and degraded service for end users. These phenomena, as well as the following Prefix Deaggregation, are well-known operational causes of instability; here we place them in the broader taxonomy to connect them to threats and mitigations.

2) Prefix Deaggregation: This class of attack involves a technique where an AS advertises more specific prefixes instead of aggregated ones. Although *Prefix Deaggregation (DEG)* can be a legitimate operational strategy, it becomes problematic when abused, either intentionally or negligently, to gain advantage or cause harm.

Malicious Deaggregation (DEG-MA) refers to the deliberate announcement of excessively specific prefixes to attract traffic or stress the routing system. In contrast, *Careless Deaggregation (DEG-CA)* arises from operational errors or lack of expertise, where overly granular prefixes are advertised without malicious intent, yet still contribute to routing instability.

Legitimate deaggregation may serve specific purposes, such as more precise traffic engineering, directing traffic to different exit points, or balancing load across infrastructure. However, when the practice results in the flooding of the Internet routing system with superfluous prefixes, especially without any operational justification, it falls under *DEG-MA*.

Consider the following example: an ISP assigned an IPv6 /29 prefix may initially advertise it as a single route in the global BGP routing table. To achieve finer traffic control, the same ISP may instead announce multiple more-specific /32 prefixes contained within it, up to eight. This scenario represents *Exploitative Deaggregation (DEG-EX)* attacks, where the use of more-specific prefixes serves a tactical, though potentially questionable, purpose [A7].

However, a malicious actor could escalate this approach by leveraging the fact that /48 is the smallest commonly accepted IPv6 prefix. Within a single /29, up to 524,288 distinct /48 prefixes can be generated and advertised. Doing so would constitute a *Fragmentation Flood Attack (DEG-FF)*, a large-scale expansion of routing announcements aimed at exhausting router memory and processing capacity. For comparison, a single actor alone could more than double the current global IPv6 BGP table size, which holds around 219,000 entries, potentially causing widespread service degradation or outages [A8].

C. Attack Macro-Class #3: Policy Violation Attacks

This category encompasses attacks that break the established BGP routing policies between different ASes. These attacks may involve the unauthorized propagation of routes between ASes that are not meant to be exchanged or violate local routing agreements. Policy violations can lead to unintended traffic leaks, misrouting, and even the exposure of sensitive network information. In this context, attackers may exploit misconfigurations or insufficient policy enforcement to compromise routing security and breach inter-AS agreements. In this category, we find attacks belonging to the micro-categories *Routing leaks* and *Policy Manipulation*.

1) Routing Leaks: This attack occurs when an AS propagates routing advertisements in violation of the intended routing policies of the sender, receiver, or any AS along the path [51]. Such violations, whether accidental or deliberate, result in a *Routing Leak (RLK)*, breaching core routing principles like valley-free routing.[4] *RLKs* can result from misconfigurations or intentional actions. They may cause suboptimal routing, increased latency, congestion, or even security vulnerabilities such as data interception or manipulation. In some cases, an AS may be forced to route traffic through a longer or costlier path. Since each AS may have transit (as a provider or a client) or peering relationships with other ASes, routing policies determine who can see and use certain routes.

On the one hand, a client AS must advertise all its networks to the provider, including those of its customers, but not those learned from peers and other providers. At the same time, it must not announce to peers or providers the routes learned via transit. Violation of these rules may result in an *Upward Leak* (*RLK-UL*). On the other hand, a peer AS must only announce its networks and those of its clients, i.e., without propagating routes received from other peers or providers, without causing a *Lateral Leak* (*RLK-LL*). A *downward leak* (*RLK-DL*) occurs when a provider advertises to a customer routes learned from another provider or peer. This behavior grants the customer unintended transit capabilities and breaks the valley-free principle by enabling the customer to act as a transit between two higher-tier ASes. Despite its theoretical relevance, there are currently no well-documented public incidents explicitly classified as *RLK-DL*. This does not imply that such leaks do not occur; rather, they may go unnoticed, unreported, or not identified under this specific category in public analyses. These RLK subtypes align with the route-leak definitions and classifications in RFC 7908 [51]

A more severe variant is the *Full Propagation leak* (*RLK-FP*), which happens when an AS propagates routes learned from one neighbor to all other neighbors, regardless of their role (provider, peer, or customer). This leads to large-scale dissemination of routes that should be selectively filtered based on policy. Documented incidents include the case of the Swiss colocation provider Safe Host, which in 2019 mistakenly advertised more than 70,000 routes to China Telecom, which subsequently propagated them across the global Internet (*RLK-FP*) [56]. This resulted in significant traffic redirection, affecting numerous networks [A9].

Another documented incident occurred in 2025, when an American technology company (AS 52863) began leaking routes from one of its providers (*RLK-UL*) [57]. The leak involved over 20,000 routes and lasted for approximately 15 minutes. Despite having a longer `AS_PATH`, the leaked routes were still preferred over more direct paths, highlighting that `AS_PATH` length is just one factor in BGP route selection [A10].

2) Policy Manipulation: Policy Manipulation (POL) exploits the flexibility and complexity of routing policies defined by ASes to influence inter-domain routing decisions. BGP is intentionally designed to allow each AS to apply its policies based on local preferences and business relationships with other operators (e.g., peers, providers, customers). However, this autonomy can be abused by malicious actors to manipulate network routes in their favor.

In practice, ASes define routing policies to prefer or reject specific paths based on criteria such as hop count, latency, cost, or commercial agreements. Although enforced locally, these policies have global effects on route propagation and selection. An attacker can exploit this by crafting announcements or influencing attributes to steer traffic away from optimal paths.

One such tactic is *AS_Path Length Manipulation (POL-PL)* [69], where the attacker advertises routes that are technically valid but constructed to make the `AS_PATH` appear shorter or more attractive, thus misleading others into preferring suboptimal routes. This can degrade performance or increase the costs for targeted networks.

Another vector is *Local Preference Manipulation (POL-LP)*. Local Preference is a BGP attribute used within an AS to

[4]Valley-free routing ensures that an AS does not forward traffic from one peer or provider to another peer or provider, preserving operational efficiency and reflecting business relationships rather than physical topology.

rank competing routes; higher values are preferred. An attacker may manipulate this attribute, directly or via misconfiguration in a partner AS, to ensure that traffic is directed through a chosen path, often through their network or those of colluding operators, even if better alternatives exist.

A further manipulation involves the MED attribute,[5] leading to *MED Abuse (POL-MD)*. By altering MED values, an attacker can influence how inbound traffic is routed, potentially causing congestion or traffic redirection.

Another form of manipulation is *Selective Propagation (POL-SP)*, where an AS deliberately announces specific prefixes or routes to only a subset of its neighbors while withholding them from others. This asymmetric propagation can be used to manipulate inbound or return traffic patterns, attract flows through specific links, or evade detection by keeping certain routes invisible to monitoring systems. Evidence of this behavior has been reported in the literature. For instance, Mittal et al. [58] analyze surgical interception attacks, where an attacker uses BGP communities to selectively propagate hijacked routes to carefully chosen upstream providers [A11]. By doing so, the attacker intercepts traffic from targeted regions or ASes without triggering widespread visibility or alarms, achieving stealthy and effective traffic redirection.

D. Attack Macro-Class #4: Session-Based Attacks

These attacks focus on destabilizing or manipulating the BGP session itself, which is the communication channel through which BGP messages are exchanged between routers. By exploiting session vulnerabilities, attackers can disrupt BGP's normal operation, cause session resets, or inject false updates. Such attacks can lead to service disruptions, loss of routing information, or unauthorized changes in routing behavior. Common methods include BGP session hijacking, session timeout attacks, or flooding BGP routers with excessive updates. In this category, we find attacks belonging to the micro-category *Attribute-Based Session Reset Attacks*.

1) Attribute-Based Session Reset: Attribute-Based Session Reset (ATR) attacks exploit vulnerabilities in the handling of attributes, particularly when malformed or improperly processed attributes cause disruptions in the routing system. This class relates to long-discussed UPDATE error-handling issues (RFC 7606 [53]) but is presented here as a dedicated attack family due to recent large-scale incidents [59].

A documented real-world incident illustrating this class occurred on May 20, 2025, underscoring the protocol's ongoing vulnerability to malformed attributes and the lack of consistent error-handling behavior across router vendors, a clear example of *Error Handling Policy Abuse (ATR-EH)* attacks [59] [A12]. The event began with a BGP UPDATE message containing a corrupted Prefix-SID attribute (Attribute Code 40), which was originated by an AS in the Asia-Pacific region. This message was improperly handled and subsequently propagated through the global routing system, exemplifying a *Malformed Optional Transitive Attribute Injection (ATR-MA)* attack scenario.

[5]It can be manipulated to influence which exit is chosen when there are multiple connection points between two ASes, affecting the direction of traffic.

While routers from vendors such as Cisco (IOS-XR) and Nokia (SR-OS) ignored the invalid attribute following RFC 7606, Juniper's JunOS implementation failed to do so, passing along the malformed message. This led to *Vendor-Specific Attribute Exploitation (ATR-VE)* attack, as Arista routers, upon receiving the malformed attribute, responded by resetting BGP sessions instead of discarding the update.

The situation was aggravated by route servers at IXPs, which relayed the malformed attribute without applying proper attribute filtering—again illustrating the risks of *ATR-MA* in shared interconnection environments. As a result, the global routing infrastructure experienced a sudden surge in BGP update volume (over 150,000 updates within 10 seconds), triggering session flaps and instability across multiple networks, including operators like Starlink, Disney, Zscaler, ByteDance, and others. This cascade of failures exemplifies how *ATR-EH*, stemming from inconsistent or lax error-handling policies, can escalate into widespread routing disruptions.

In response, Arista addressed the issue in EOS version 4.28.11 and subsequent releases by modifying their behavior to discard malformed Prefix-SID attributes instead of terminating sessions—mitigating the effects of *ATR-VE*. Nonetheless, the broader challenge of safeguarding against malformed or experimental attributes, especially in environments involving shared infrastructure like IXPs, remains unresolved—highlighting the ongoing risk posed by *ATR-MA*.

IV. Defense Approaches

So far, the R&I community has primarily concentrated on developing mechanisms to detect and counter attempts to corrupt, remove, or withhold routing information, while preserving the scaling properties of BGP. In the following, we report some of the well-known defense approaches to ensure BGP routing security. Please note that the proposed taxonomy (Fig. 3) provides a useful perspective for interpreting the defense approaches discussed in the following subsections. In particular, we distinguish between (i) structural or inter-domain policy attacks, whose mitigation mainly requires new validation frameworks, shared semantics, and coordinated adoption across multiple networks (e.g., PRH, RLK, POL, and part of ASM), and (ii) implementation or operational robustness attacks, whose mitigation mainly depends on software robustness, configuration hygiene, monitoring, and operational hardening (e.g., ATR, VOL, DEG, and part of ASM). This distinction helps explain how the different stakeholders tend to engage with routing security challenges. Academia often focuses first on the former, industry operators concentrate on the latter, while institutional actors become relevant once countermeasures reach sufficient maturity and the main challenge shifts from invention to coordinated adoption.

In the following, we examine the principal defense approaches from academic, industrial, and institutional perspectives. Table V summarizes these approaches across four operational dimensions: deployment and configuration cost, false positive risk, operational overhead, and sensitivity to partial adoption.

TABLE V
OPERATIONAL COMPARISON OF BGP DEFENSE APPROACHES

Defense	Deployment and configuration cost	False positive risk	Operational overhead	Partial-adoption sensitivity
Academic Perspective				
S-BGP	Hierarchical PKI + IPsec on all eBGP sessions; complexity scales with prefixes and peers.	Non-participating ASes indistinguishable from attackers.	Per-hop signing and verification overhead scales with AS_PATH length, increasing convergence time.	A single non-adopter breaks validation chain; universal deployment is required.
soBGP	Parallel certificate infrastructure; routers must maintain global security topology.	Stale certificates can cause filtering of legitimate routes.	Certificate processing and topology updates add overhead and latency to route installation.	Coverage degrades proportionally to non-participating ASes.
psBGP	Centralized PKI + decentralized PAL; burden scales with peering relationships.	Stale PALs can flag legitimate prefixes as unauthorized.	PAL consistency checks add modest per-advertisement overhead.	Local protection achievable; effectiveness scales with peer PAL accuracy.
Industrial Perspective				
IRR	No protocol changes; filter generation via external tools that scale with customer cone.	Inaccurate records cause legitimate routes filtering.	No runtime impact; overhead confined to offline filter construction.	Incomplete registration and filter enforcement each independently reduce effectiveness.
RPKI	No protocol changes; requires RIR contracts, ROA maintenance, and ROV on routers.	Stale ROAs cause filtering of legitimate routes.	Modest ROV overhead; periodic synchronization overhead; convergence time unaffected.	ROA coverage alone is insufficient; widespread ROV enforcement is required.
BGPsec	Per-hop signing with RPKI certificates; requires crypto-capable hardware; disables UPDATE packing.	Non-BGPsec ASes strip signature; legitimate and malicious routes indistinguishable.	Sustained CPU overhead scales with AS_PATH length. UPDATE packing prohibition slows convergence.	A single non-adopting AS renders end-to-end validation impossible.
ASPA	ASPA objects in RPKI; provider relationships updated per peering change; requires routing stack support.	Stale ASPA objects can cause legitimate routes to fail plausibility checks.	Modest per-UPDATE overhead. Does not alter convergence times.	Meaningful protection without universal adoption; key ASes limits route leak propagation.

A. Academic Perspective

From the academic research perspective, defense approaches have historically focused on mitigating *PRH* and *ASM* attacks, both at the theoretical level and in the proof of concepts [1, § V-C2)i]. These approaches rely on external route monitoring databases/routing registries, on looking for multiple origin AS, or on continuously probing known transit ASes. In particular, the time period between 2000 and 2011 saw the emergence of three main proposals that serve as reference, that are, *Secure BGP (S-BGP)*, *Secure Origin BGP (soBGP)*, and *Pretty Secure BGP (psBGP)* – cf. [1, § V-C2)a,b,c].

S-BGP extends BGP to ensure the authenticity, integrity, and authorization of routing information by digitally signing the address and `AS_PATH` information contained in routing advertisements [70]. Specifically, to achieve these goals, it uses a Public Key Infrastructure (PKI) based on X.509 certificates that authenticate the ownership of ASes, routers, and IP prefixes through a trust hierarchy that starts with ICANN and includes RIRs. Advertised routes are accompanied by digital attestations: Route Attestations (RAs) certify the correctness of the AS route, while Address Attestations (AAs) confirm the legitimacy of the announced IP prefixes. In addition, S-BGP uses IPsec to protect communications between routers, ensuring the security of BGP messages in transit.

soBGP addresses some issues of S-BGP related to the processing overhead incurred when validating the chain of router attestations [71]. To ensure the authenticity of BGP information, soBGP relies on certificates and distributed trust mechanisms. Routers exchange digitally signed certificates containing information about their identities, AS relationships, and authorized IP prefix advertisements. This information is distributed across the network, enabling routers to build a global security topology used to verify advertised routes.

psBGP further enhances the previous two approaches by means of a centralized trust model to authenticate ASNs [72], whereby authentication is provided by certified authorities such as IANA and the RIRs. Each AS receives a public key certificate that binds its ASN to the corresponding public key, ensuring that only legitimate owners can use it. In fact, this reflects a centralized trust model for AS identity validation. In contrast, psBGP employs a decentralized model to verify the legitimacy of IP prefix origins. Each AS maintains a Prefix Assertion List (PAL) that includes the prefixes it owns and those of its peers. Consistency among these peer-maintained lists allows the network to detect unauthorized prefix advertisements without relying on a central authority.

However, these proposals did not lead to operational deployment. As observed by Testart [73], "*none of these proposals were fully implemented*", mainly due to "*the lack of agreement about what needs to be secure or validated*", as well as "*they directly impact the usual operation of network operators*", which further hinders industrial adoption.

B. Industrial Perspective

From the industry perspective, the main tools used to ensure BGP reliability against threats to the integrity of routing information are *Internet Routing Registries (IRRs)*, *Resource Public Key Infrastructure (RPKI)*, *BGP Security (BGPsec)*, and *Autonomous System Provider Authorization (ASPA)*.

IRRs are globally distributed databases that enable network operators to register routing policies for their networks [74]. These databases allow operators to specify the association

between the IP prefix and ASNs through *route:* objects, which define the ASN authorized to originate a given prefix. In addition, IRRs support *as-set:* objects, which define a group of ASNs that can be used to express routing policies and transit relationships. These objects help verify which ASes are legitimately allowed to advertise prefixes through a provider. Network operators can rely on IRR data to define and enforce routing policies, helping prevent incorrect or malicious announcements. IRRs are categorized into two types: *authoritative* registries, managed by RIRs, and *non-authoritative* registries, maintained by third-party organizations. The latter often apply less stringent verification procedures, which can undermine the reliability of the information they provide. For IRR-based filters to be effective, operators must keep their information up to date. Otherwise, inaccurate data may lead to the incorrect filtering of legitimate route announcements.

RPKI [75] is a framework that leverages X.509 digital certificates to verify the accuracy of routing announcements [1, § V-A]. Specifically, an AS that holds IP address blocks can create the so-called Route Origin Authorization (ROA) objects, which are digitally signed declarations specifying which ASNs are authorized to advertise those addresses via BGP. Each ROA includes not only the origin ASN and prefix, but also some other attributes, such as the `maxLength`, which allows network operators to specify a maximum prefix length for a given IP prefix. This attribute provides more granular control over which prefixes are considered valid, but it also introduces potential risks. Excessive use of `maxLength` may widen the attack surface, particularly in the context of *ASM-FI*. As a precaution, best practices recommend avoiding this attribute unless there is a clear operational need for it [76]. Using the ROA data, BGP routers can perform Route Origin Validation (ROV), a process that classifies incoming BGP route announcements into three categories: valid (prefix and origin ASN match a ROA), invalid (a ROA exists for the prefix, but the origin ASN or prefix length is inconsistent), or unknown (no relevant ROA exists). By enforcing ROV and filtering invalid routes, operators can mitigate both malicious and unintentional *PRH* or mistakes in RIRs delegation records [60], mentioned in Sec. II-B.

However, both tools have limitations. In fact,

- since route objects must be manually maintained by ASes' owners, IRRs may contain information that does not reflect the actual routing behavior. Furthermore, non-authoritative IRRs, which do not enforce strict validation procedures, can include malicious IRR records [77];
- to be effective, ROV adoption should ideally occur across *all ASes*, although this (at least, historically) has not been true [78]. Moreover, the design of the RPKI framework introduces certain vulnerabilities, such as downgrade attacks. In such attacks, the unavailability of RPKI can cause an AS to fall back to less secure routing policies that preceded RPKI implementation [79].

BGPsec is a security extension of BGP standardized in September 2017 as RFC 8205 [80], aimed at preserving the integrity of `AS_PATH` attribute in BGP UPDATE messages, ensuring it remains unmodified as it propagates through the Internet (*path validation*). BGPsec achieves this by requiring each AS to cryptographically sign the `AS_PATH` before forwarding the update to the next AS, thereby protecting each hop. Upon receiving a BGP UPDATE, a router validates the `AS_PATH` by verifying these signatures, using RPKI certificates. Although BGPsec has been implemented in major BGP routing daemons, such as BIRD [81], its adoption remains limited due to computational overhead and operational complexity [82]. Each router must validate and generate cryptographic signatures for every hop in the `AS_PATH`. Given the scale of the global routing table and an average `AS_PATH` length of four, startup may require roughly two million signature verifications. Furthermore, BGPsec enforces per-prefix integrity protection, preventing the packing of multiple prefixes into a single update, unlike conventional BGP. Deployment is also hindered by its reliance on end-to-end adoption: a single non-BGPsec-enabled AS in the `AS_PATH` strips the security information from the update, making partial deployment of limited benefit. Finally, even under full deployment, BGPsec lacks mechanisms to detect or mitigate *RLK* [83].

In addition, the Internet Engineering Task Force (IETF) has been developing standardized mechanisms that fall under the category of *path plausibility*, which is complementary to origin and path verification by validating whether a BGP path is consistent with known AS-level relationships. The main proposal is *ASPA*, a mechanism primarily targeting Policy Violation attacks, such as *RLK* and *ASM-FI*.

ASPA introduces cryptographically verifiable customer-provider relationships by listing the provider ASNs authorized to propagate routes for a given customer. Unlike BGPsec, ASPA reduces the risk of RLKs without requiring per-update cryptographic validation of BGP announcements. Recent analysis indicates that even partial deployment by strategically positioned ASes can significantly reduce the number of ASes affected by *RLKs* and *ASM* by up to 96% [84]. Nevertheless, adoption at the network edge provides limited benefit, and the mechanism requires operators to explicitly maintain provider relationship information. Even if, unlike AS-Cone, ASPA mitigates *ASM-FI* attacks by preventing an AS from modifying relationship objects belonging to another AS, ASPA does not eliminate all risks associated with *ASM*, particularly in scenarios involving transit provider routes [85]. In practice, ASPA has become one of the main mechanisms currently pursued by the operational community. The first IETF draft was published in 2018, and the standardization process is still ongoing [86], reflecting the extensive technical scrutiny typical of routing security mechanisms, as discussed in RFC 8963 [87]. Operational adoption has also depended on the availability of mature implementations in routing software. ASPA-based validation requires support from BGP daemons implementing the corresponding verification algorithms, which are more complex than traditional RPKI ROV. After significant engineering and testing efforts, support has now been integrated into widely used routing software such as OpenBGPD and BIRD, enabling broader deployment. As these technical prerequisites have been addressed, operators have started publishing ASPA objects and enforcing related policies. Since December 2025, operators have been able to publish ASPA objects in the RIR repositories [88], and 1,314 objects have

TABLE VI
OVERVIEW OF PUBLIC AND PRIVATE INITIATIVES ON BGP ROUTING SECURITY ASSURANCE

Institutions	Key Regulations/Reports	Focus / Highlights	Technical Measures / Recommendations
Publicly-led Initiatives			
OECD	OECD Report (2022) [61]	Routing security overview	General policy guidance
United States, NIST, NSA	US National Strategy (2024) [90], NIST SP 800-189 [91], NSA Technical Report (2018) [92]	Risk-based planning; incremental adoption of security technologies	Deployment guidelines for ROA and ROV; alignment with MANRS recommendations
European Commission, ENISA	Digital Networks Act [93], NIS2 Directive [94], CER Directive [95], EU Regulation 2024/2690 [96], ENISA Report (2019) [62]	Critical infrastructure security and resilience, EU-wide applicability, focus on IXPs and DNS	Technical, operational and organizational measures; best practices for DNS and routing security
Estonia, Italy, Netherlands	Cyber-conscious Estonia [97], Implementation plan [98], Action plan [99]	Necessity of protecting national digital infrastructure dealing with the Internet's core functionality	General policy guidance
Norway, Sweden	List of measures [100], Appendix 1: Action plan [101]	Improve the Internet routing security in the backbone network protection	General policy guidance
ANSSI	ANSSI BGP best practices [102]	Responsible BGP use by router administrators	Equipment-specific configuration best practices
UK NCSC	NCSC Technical Report (2021) [103]	Responsible BGP use by ISPs; need for updated best practices	AS_PATH length filtering, scalability (prefix limits, aggregation), IPv6 deaggregation
Australia	Australian Government Gateway Security Standard [104]	Strategic direction and minimum security standards	General policy guidance
Privately-led Initiatives			
MANRS	MANRS [105], MANRS for IXPs [106]	Best practices for network operators, IXPs, CPs and CDNs	Route filtering, improve coordination, secure peering and monitoring tools
EURO-IX	EURO-IX [107]	Collaboration among IXPs	Route server best practices and traffic analysis

been registered at the time of writing [89]. This trajectory indicates a technology transitioning from advanced validation toward early operational deployment.

Overall, industry efforts show a clear trend toward deployable and backward-compatible solutions, prioritizing incremental improvements over disruptive changes. Consequently, lightweight mechanisms have achieved wider adoption and practical impact, whereas more comprehensive security frameworks remain largely under-deployed due to operational and coordination constraints, leading to partial protection.

C. Institutional Perspective

As it emerges from Sec. III, most attacks against Internet routing stem from i) the decentralized nature of Internet routing, which does not assume the presence of a central entity checking and verifying the exchanged routing information, and ii) the untrustworthiness of ASes' owners. Over the years, several initiatives, both originating from public institutions as well as private organizations, have sought to collect best practices in the area of inter-AS routing security, and much has been done to promote a culture of collective responsibility for the security of the ecosystem itself. The main ones are reviewed in the following section and are summarized in Table VI. We now investigate the **publicly-led** initiatives.

1) International Bodies: In 2022, the Organization for Economic Cooperation and Development (OECD) released a report on routing security [61], which aimed to analyze i) the scope and scale of routing incidents, ii) the security techniques that have been proposed to address them, and iii) the role of policy makers in securing the routing system. As items i) and ii) have already been treated in Sec. III of this manuscript, respectively, it is worth focusing on item iii).

In this regard, governments are invited to fund publicly available data collection and measurement efforts to track BGP incidents over time, to, on the one hand, foster the establishment of formal information sharing mechanisms among industry stakeholders, and, on the other hand, enable tracking the effectiveness of the security measures that are progressively adopted. Moreover, nations are encouraged to promote awareness and deployment of best practices in Internet routing both in governmental ASes and in privately-owned ASes.

2) United States: In 2024, the United States considered a secure and open Internet as *critical* to its economic prosperity and national security [90]. While RPKI is considered the most important best practice (thus prioritizing defense against *PRH*), the challenges to its adoption already mentioned in Sec. IV-B lead to the key recommendation for all network operators to adopt a *risk-based planning* of their ASes. In fact, for each type of network and scope of operation, there are different risk management and cost models associated with adopting new security technologies. Being RPKI-based security measures one of such technologies, upon adopting them network operators and IXPs are invited to apply a risk-based approach that identifies, categorizes, and prioritizes the security of the systems, assets, and traffic they consider most valuable and critical to their operations. It is worth observing that the main principle underpinning [90], that is, the protection of federal systems, critical infrastructures, and digital supply chains, based on the pro-tempore national cyber strategy [108, cf. Strategic Objective 4.1], are *de facto* preserved in the current strategic document [109].

From the technical point of view, in 2018, the US National Security Agency (NSA) released a report addressing the cybersecurity of devices dedicated to running the BGP protocol (e.g., in terms of access control and logging) [92],

taking into consideration several widespread vendor models, to improve protection against *ATR*, *PRH*, and *ASM* attacks. Then, the National Institute of Standards and Technology (NIST) issued in 2019 a dedicated Special Publication (SP) 800-189 [91], which provides technical guidance and recommendations for technologies that facilitate resilient inter-domain traffic exchange (mainly, RPKI protecting from *PRH* attacks), as well as technical note 2060 [110], in 2021) containing a reference implementation for RPKI-based ROV and BGPsec path validation within a BGP router to improve protection against route manipulation attacks.

Additionally, the US has examined the underlying causes of limited RPKI adoption, identifying three primary factors: *insufficient awareness of routing security risks*, *limited resources to deploy new BGP security mechanisms*, and *administrative complexities associated with contractual arrangements with RIRs*. These factors contribute to a broader reluctance among network operators to prioritize routing security.

3) European Union: At continental level, a proposal for a Digital Networks Act (DNA) was released in January 2026 [93]. By being a regulation rather than a directive, the DNA aims at being directly applicable across EU Member States, updating the framework established by the European Electronic Communications Code (EECC). The final objective is to provide i) integration of European digital infrastructures, ii) requirements in terms of network resilience, iii) harmonised spectrum management, and iv) strengthened European governance of the sector. In particular, the DNA proposal considers security and resilience of critical digital infrastructures as essential for the EU's economic security, building on the solid existing legislative framework, notably the NIS2 Directive [94] and the Directive on the resilience of critical entities [95]. Both directives recognize the importance of taking appropriate and proportionate technical, operational, and organizational measures to manage the risks posed to the security of network and information systems, thus representing an important driver for securing important entities like Domain Name Systems (DNS) and IXPs, which are impacted by the present work. Specifically, the former directive explicitly defines an IXP as "[...] a network facility which enables the interconnection of more than two independent networks (ASes), primarily for the purpose of facilitating the exchange of Internet traffic, which provides interconnection only for ASes and which neither requires the internet traffic passing between any pair of participating ASes to pass through any third AS nor alters or otherwise interferes with such traffic" [94, Art. 6,subs. 18]. Moreover, the Commission Implementing Regulation (EU) 2024/2690 [96] laying down rules for the application of [94] mentions that "[...] network security measures in relation to [...] the application of best practices for DNS security, and for internet routing security and routing integrity entail specific challenges regarding the identification of best available standards and deployment techniques" [96, Subs. 8].

At the national level, various European countries have published their own cybersecurity strategies [111] that, at an action plan level, explicitly mentions the intention to improve the security of Internet routing. These countries include Estonia [97, cf. Annex 1], Italy [98, cf. Measure#17], Netherlands [99, cf. Aim 3,Internet Governance], Norway [100, cf. Measure#32], and Sweden [101, cf. Measure#3.2]. In particular, Estonia, Italy, and the Netherlands' strategies explicitly mention the necessity for collaboration between the respective governments and the national IXPs. On the other hand, Norway and Sweden's strategies relate the need to improve the Internet routing security to the broader objective of protecting the international backbone network infrastructure, which is largely composed of submarine optical fiber cables.

On the technical side, the French Cybersecurity Agency (ANSSI) released in 2014 an official document intended for router administrators, which describes the good configuration practices for the BGP [102] to improve protection mainly against *ATR* and *PRH* attacks. In 2019, the European Union Agency for Cybersecurity (ENISA) published a position paper on BGP routing security, which shortlists seven basic measures that are relatively easy to implement and relatively effective in mitigating PRH [62]. Notably, both of the documents acknowledge the collaboration with network operators and IXPs, testifying to the need to involve the private sector, as also recommended by OECD.

4) Other Countries: In 2021, the National Cyber Security Center (NCSC) of the United Kingdom released a technical report [103] outlining recommendations for the secure use of BGP by ISPs.[6] The NCSC identifies key areas to be addressed via best practices when designing, deploying, operating, and maintaining BGP peering relationships, including:

- *filtering*. The implementation of filters on various targets (e.g., the lengths of AS_PATH) ensures that only valid routes are accepted and advertised, thus preventing PRH;
- *scalability*. Adopting prefix-limits, route aggregation, and control-plane policing ensures that better traffic handling is achieved by limiting the ability of a single peer to send massive protocol updates to prevent processing of valid updates from other peers.
- *IPv6*. As the IPv6 address space is only partially allocated, as opposed to the IPv4 one, checks against IANA and RIR allocated address space in prefix-filters are recommended, with these updated on an iteration cycle of less than one month. Additionally, proper countermeasures against DEG attacks should be implemented

For each of these areas, where relevant, reference is made to other technical reports by international standardization bodies such as, e.g., NIST and IETF, dating back until 2016 – thus, possibly needing some further update at the time of writing.

Finally, it is worth citing the recent consultation paper [104] published by the Australian Government to engage with interested stakeholders on the so-called "Gateway Standard", which outlines the strategic direction and minimum security standards for the deployment of gateways and Security Service Edge (SSE) solutions across the Australian Government. In particular, §8.6 of [104] highlights that governmental entities' public IP addresses shall be signed and with valid ROAs.

[6]Notably, the NCSC report acknowledges the need to update its 2011 guidelines, noting the growing criticality of IP network stability and routing integrity over the past decade, a motivation aligned with this work.

We now move onto the **privately-led** initiatives.

5) MANRS: The most important and longest-running initiative on Internet routing security is Mutually Agreed Norms for Routing Security (MANRS) [105]. Launched in 2014, initially promoted by the Internet Society (ISOC), and since 2024, still under the umbrella of the Global Cyber Alliance (GCA), MANRS is a global initiative providing practical measures to reduce and mitigate major routing threats. Continuously updated to reflect the evolution of defense mechanisms, it serves as a reference framework for operators aiming to improve and maintain the security of their infrastructure. MANRS defines tailored programs for different roles in the ecosystem, including IXPs, CSPs, CDNs, vendors [106], and, most importantly, network operators. It currently includes over 1,260 participants, and its recommended practices are frequently referenced in governmental policy and strategy documents [90]. The basic idea is to i) raise awareness of routing security problems and encourage the implementation of actions that can solve them; ii) promote a culture of collective responsibility for security; iii) provide a framework for operators to better understand and address the problems.

MANRS identifies several key actions to address three classes of issues: i) incorrect reachability information (e.g., *PRH* and *ASM* attacks), ii) traffic with spoofed IP addresses (e.g., *ATR* attacks), and iii) lack of coordination among operators (e.g., *VOL* and *DEG*, *RLK* and *POL* attacks). In particular, the following actions are recommended for network operators:

1) *prevent the propagation of incorrect routing information.* Operators should ensure that only their prefixes and those of verified customers are announced, enforcing strict control over routing advertisements and validating their legitimate assignment.
2) *block traffic with forged source IP addresses.* Operators should implement mechanisms for source address validation within their networks and for their customers, enabling anti-spoofing filtering to prevent packets with invalid source IPs from entering or leaving the network.
3) *facilitate co-ordination and co-operation with other operators.* Contact information, such as Network Operations Center e-mail and telephone numbers, within IRRs and PeeringDB [112] must be kept up to date;
4) *document routing information.* Operators should register reachability information both within the IRRs (via *route:* and *route6:* objects) and customer cones (in the *as-set:* object), and create ROAs for prefixes.

Finally, it is worth remarking on the assessment done by the authors of [113], which deepens the relationship between MANRS recommendations and ISO/IEC 27001 security controls. Building on this, the 2022 study in [114] quantitatively evaluates the compliance of 870 MANRS members (849 ISPs and 21 CDNs) with key routing security standards using publicly available datasets. Their results show that over 95% of participants comply with the requirement to register their prefixes in RPKI or IRR databases, with around 58% of them having signed RPKI for their prefixes. In particular, small MANRS ASes are approximately 2.5 times more likely than non-members to originate only RPKI-valid prefixes. Moreover, more than 83% of MANRS ASes meet the filtering requirement by effectively preventing the propagation of invalid BGP advertisements. Specifically, 46% of large MANRS ASes propagate no invalid prefixes at all, compared to only 36% of similarly sized non-MANRS ASes. Furthermore, large MANRS ASes limit the propagation of invalid prefixes to at most 1.1% of their total advertisements. However, despite the improvement observed among MANRS participants, the broader adoption of MANRS recommendations across the Internet ecosystem remains limited. Several operational factors hinder wider compliance, including limited awareness and technical expertise among many networks, as well as the operational complexity associated with deploying mechanisms such as RPKI and ROV. In addition, the increasing use of IPv4 address leasing and transfers complicates the management and maintenance of ROAs, further slowing the adoption of best practices promoted by MANRS [115].

6) Euro-IX: Established in 2001, the European Internet Exchange Association (Euro-IX) [107] is an organization that brings together around 70 IXPs primarily from Europe but also from other parts of the world. Its primary goal is to strengthen the IXP community through collaboration, knowledge sharing, and joint development initiatives. Through collaborations with key Internet governance bodies such as RIPE NCC, ISOC, and ICANN, Euro-IX ensures that the interests and needs of IXPs are represented in broader policy discussions and the ongoing development of the global Internet.

Euro-IX collects and analyzes data on traffic patterns, infrastructure, and performance across IXPs to provide valuable information, support transparency, and help members benchmark their operations and plan for the future. Within the association, members engage in continuous exchange of technical expertise and best practices, improving the performance, reliability, and scalability of their infrastructure. In this context, it is worth observing that the ecosystem of IXPs exhibits tighter synergies with the developers of the major open-source BGP daemons (OpenBGPd [116] and Bird [81]) than with network equipment vendors, which are traditionally intertwined with ISPs.

Euro-IX members also participate in regular forums and meetings to discuss operational challenges, industry trends, and innovations to support mutual growth and learning. In cybersecurity, the Euro-IX community promotes the proper use of route servers and best practices for the route server validation chain, ensuring that only legitimate routing information is accepted and propagated at IXPs [117].

Overall, institutional initiatives play a key role in promoting best practices, coordination, and large-scale adoption of routing security mechanisms. However, their effectiveness remains constrained by voluntary compliance, heterogeneous deployment, and the absence of enforceable global guarantees.

D. Coverage and Residual Gaps Under Current Best practices

Although best practices such as IRR and RPKI provide effective origin validation, the taxonomy-driven analysis reveals that their coverage remains incomplete. Tab. VII captures this by mapping each attack class to i) the respective defense/best practice approaches, distinguishing between the

TABLE VII

TAXONOMY-DRIVEN SUMMARY OF CURRENT BEST-PRACTICE MITIGATIONS AND RESIDUAL GAPS. THE TABLE HIGHLIGHTS ATTACK CLASSES THAT REMAIN LARGELY UNMITIGATED UNDER COMMONLY DEPLOYED PRACTICES, MOTIVATING THE R&I DIRECTIONS DISCUSSED IN SEC. V

Class	Defense/Best practice approaches	Residual gap / largely unmitigated aspects
PRH	S-BGP, soBGP, psBGP (*academia*); IRR, RPKI (*industry*, *institutions*)	**Partial deployment** limit benefit; permissive `maxLength` expands the attack surface; PRH-AB remains feasible without maintained objects or filtering
ASM	S-BGP, soBGP (*academia*); BGPsec (*industry*, *institutions*)	**Operational complexity** and computational overhead; **Partial deployment** limits benefit; routing information correctness does **not guarantee forwarding-state alignment**
VOL	Operational hardening (*industry*); Control-plane policing (*institutions*)	Malicious injection or amplification can still overload routers and propagate globally; **Complexity** in distinguishing attacks from misconfigurations
DEG	Conservative ROA `maxLength` (*industry*); prefix-limits, filters for overly-specific prefixes (*industry*, *institutions*)	DEG-FF remains a DoS vector. Legitimate traffic engineering deaggregation complicates prevention. Mitigations are local using **auxiliary tools**
RLK	ASPA (*industry*), strict import or export policy filters (*institutions*)	**Partial deployment** and lack of coordination limit effectiveness. Some leak types may not be observable from public data.
POL	Policy hygiene (*industry*, *institutions*)	**Partially deployed** global validation for policy attributes. POL-SP is intrinsically stealthy and often relies on observability and response rather than prevention
ATR	Robust error-handling (*industry*); attribute filtering at route servers and vendor fixes (*institutions*)	Inconsistent vendor behavior and propagation through shared infrastructures can create instability. Prevention depends on software robustness and operational filtering leveraging **auxiliary tools**

preferred options of the three stakeholders and ii) the residual gaps.

Two patterns emerge. First, protection is uneven: while *PRH* is partially mitigated, classes such as *ASM* and *RLK* remain exposed due to the lack of widely deployed path-validation mechanisms, and even existing solutions (e.g., BGPsec, ASPA) are constrained by limited deployment and operational overhead. Second, *VOL, DEG, POL*, and *ATR* attacks fall outside cryptographic validation and are addressed through operational practices, such as filtering and coordinated incident response, which are inherently reactive and locally enforced.

Across all attack classes, four recurring limitations emerge: *partial deployment*, undermining global effectiveness; *operational complexity*, hindering adoption; the need for *auxiliary tools*, as complementary monitoring and response mechanisms; and the lack of *forwarding-state alignment guarantees*; all of which will be investigated in the next section.

V. STEPS AHEAD IN RESEARCH AND INNOVATION ON INTERNET ROUTING SECURITY

Collectively, the residual gaps highlighted in Tab. VII motivate further investigation of novel and/or innovative defenses to counter attacks against BGP routing. Thus, re-considering also the open questions raised in [1, §6] in light of the aforementioned observations, it is possible to identify the following four key questions.

1) **Routing Information versus Forwarding State.** Regarding the *alignment between routing information and forwarding state*, *is it possible to secure it such that any relying party can validate that the* `AS_PATH` *presented in* `BGP UPDATE` *not only matches the path taken by the prefix advertisement, but that the network's current forwarding state to reach the address prefix is aligned to this AS_PATH and this alignment can be validated?*
2) **Partial Deployment Problem.** Regarding the *overall deployment security*, *if a comprehensive security framework is proving to be elusive in terms of deployment considerations, then could a less comprehensive approach offer acceptable outcomes?*
3) **Improving Existing Routing Security Frameworks.** Regarding the *fundamental structural aspects of BGP routing security mechanisms*, *even assuming widespread adoption, is it possible to improve existing frameworks to better balance the trade-off between security guarantees and operational efficiency?*
4) **Auxiliary Tools to Mitigate BGP Weaknesses.** Regarding the *approaches to improve BGP routing security at large*, *is BGP too incomplete in terms of its information distribution properties to allow robust validation of the intended forwarding state? Does securing forwarding imply carrying additional information relating to the routing and forwarding state coupling in addition to routing that would be entirely impractical in a partial deployment scenario?*

In the remainder of this section, we review the solutions proposed in the literature for each of these questions. BGP attacks are referenced using the acronyms introduced in the taxonomy in Sec. III. Each solution is assigned a score based on the *interest metric*, defined in the following subsection.

A. Literature Inclusion Methodology

To ensure the reproducibility of the literature review, we outline the methodology used to identify and select relevant works. Publications were collected from IEEE Xplore, the ACM Digital Library, and Scopus, and complemented with RFCs and Internet-Drafts from the IETF Datatracker, as well as operational reports and technical documents from industry and institutional initiatives such as MANRS and NIST. Although the latter are not formally peer-reviewed, their inclusion reflects a deliberate methodological choice: BGP routing security is highly operational, in which relevant technical developments are often first disclosed through operator reports, and excluding them would produce an incomplete picture. The review covers work published between 2011 and March 2026, with the lower bound motivated by the baseline reference in [1]. Search queries were derived from the four R&I directions, combining terms such as "BGP security" and "partial deployment". A work was included if it primarily addressed at least one R&I direction and provided analytical, simulation-based, or experimental results, or if it constituted a relevant RFC, Internet-Draft, or recognized institutional report. Works not satisfying these criteria were excluded.

TABLE VIII
TRL-ALIGNMENT AND OBJECTIVE-RELEVANCE METRIC SCORES FOR DIFFERENT STAKEHOLDERS IN BGP SECURITY

Score	1	2	3	4
TRL-Alignment				
$\mathcal{T}_A$	High TRL: Mature, widely deployed standards with limited scope for further scientific advancement.	Mid-high TRL: Technologies close to full deployment, with mainly incremental research opportunities.	Low-mid TRL: Emerging technologies under active development, refinement, and evaluation.	Low TRL: Early-stage concepts with substantial scope for new theoretical or technical contributions.
$\mathcal{T}_I$	Low TRL: Conceptual or experimental technologies with unclear deployment pathways and limited implementation readiness.	Mid-low TRL: Developing of recently introduced technologies, with partial consideration of operational constraints.	Mid-high TRL: Technologies progressing toward standardization and industry-wide adoption.	High TRL: Mature technologies suitable for deployment with minimal integration complexity.
$\mathcal{T}_N$	Very high or very low TRL: Either widely deployed standards, or early-stage concepts not yet suitable for policy action.	Mid-low TRL: Developing technologies not yet sufficiently mature for regulatory or governance integration.	Mid TRL: Technologies approaching standardization and potentially informing policy or coordination mechanisms.	Mid-high TRL: Mature technologies with sufficient stability and consensus to support institutional adoption or policy enforcement.
Objective-Relevance				
$\mathcal{O}_A$	Limited research relevance: Action focuses on deployment or policy enforcement with minimal scientific contribution.	Moderate research relevance: Action incorporates some novel elements but mainly refines or applies existing concepts.	High research relevance: Action advances current knowledge through new methods, models, or empirical findings.	Very high research relevance: Action introduces new concepts or research directions with strong potential for scientific impact.
$\mathcal{O}_I$	Limited operational relevance: Action provides little benefit for deployment, cost efficiency, or reduction in real networks.	Moderate operational relevance: Action offers incremental improvements requiring significant adaptation or resources.	High operational relevance: Action supports deployment and improves network security or reliability with manageable effort.	Very high operational relevance: Action directly addresses industry needs, enabling near-term deployment with clear benefits.
$\mathcal{O}_N$	Limited policy relevance: Action has minimal implications for regulation, governance, or ecosystem stability.	Moderate policy relevance: Action may inform institutional strategies, but with limited immediate applicability to policies.	High policy relevance: Action contributes to policy development, coordination mechanisms, or regulatory guidance.	Very high policy relevance: Action directly supports standardization, regulatory enforcement, or large-scale Internet governance.

B. Interest Metric Methodology

To systematically evaluate the priorities of each considered stakeholder, i.e., academia, industry and institutions, for securing the BGP ecosystem, we define a numerical metric named **interest**, quantifying the degree to which a stakeholder is motivated to investigate a specific solution. We characterize the stakeholders through two components of the interest metrics, that are, the **TRL-Alignment** and the **Objective-Relevance**.

The TRL-Alignment, denoted as $\mathcal{T}_A, \mathcal{T}_I, \mathcal{T}_N$ for academia, industry, and institutions, respectively, evaluates the degree to which that TRL of a given approach is consistent with the stakeholder's preferred TRL range; it is defined on an ordinal scale from 1 to 4. A score of 1 indicates strong misalignment, meaning that the solution's TRL lies outside the maturity range typically associated with the stakeholder's activity domain. On the contrary, a score of 4 represents full alignment, indicating that the technology maturity is already suitable for stakeholder adoption without further changes.

On the other hand, the Objective-Relevance metric, denoted analogously $\mathcal{O}_A, \mathcal{O}_I, \mathcal{O}_N$ for each stakeholder, evaluates the degree to which a given action aligns with the stakeholder's primary objective. It is defined on an ordinal scale from 1 to 4, where a score of 1 indicates negligible alignment, meaning the action's purpose or impact is largely unrelated to the stakeholder's core objective; scores of 2 and 3 denote, respectively, that major or minor additional effort is required to adapt the initiative to the stakeholder's objective; and Instead, a score of 4 represents full alignment, indicating that the action directly supports the stakeholder's primary objective.

Table VIII provides a description of the criteria underlying each score assigned to the TRL-Alignment and Objective-Relevance metrics across the stakeholders. Scores assignment follows an expert-judgment process in which each author independently assigned scores for the stakeholder category they represent, guided by the definitions summarized in the aforementioned table.

With TRL-Alignment and Objective-Relevance defined, we introduce the **Interest** metric as $\mathcal{I}_A, \mathcal{I}_I, \mathcal{I}_N$ which respectively quantify the interest of the academia, industry and institutions of pursuing a specific initiative, as:

$$\mathcal{I}_i = \frac{\mathcal{T}_i \cdot \mathcal{O}_i}{\max(\mathcal{T}_i) \cdot \max(\mathcal{O}_i)}, \quad \forall i \in \{\mathrm{A}, \mathrm{I}, \mathrm{N}\} \tag{1}$$

In other words, the interest of academia, $\mathcal{I}_A$, toward a specific initiative is computed as the product of the TRL-Alignment and the Objective-Relevance associated with that initiative, normalized by the maximum attainable value. The same formulation applies to $\mathcal{I}_I$ and $\mathcal{I}_N$.

We adopt a multiplicative formulation, as an initiative is relevant only when aligned with both the objective and TRL dimensions. If either component is poorly aligned, the resulting interest decreases accordingly. For example, an initiative with aligned objectives but incompatible TRL would typically not be pursued.

C. Routing Information Versus Forwarding State

This first R&I direction explores the gap between control-plane routing information and data-plane forwarding behavior. Specifically, it investigates mechanisms to verify that BGP

TABLE IX
OVERVIEW OF APPROACHES TO ALIGN ROUTING INFORMATION WITH FORWARDING STATE VALIDATION

Ref.	Approach / Mechanism	Forwarding State Alignment	$\mathcal{O}_A$	$\mathcal{T}_A$	$\mathcal{I}_A$	$\mathcal{O}_N$	$\mathcal{T}_N$	$\mathcal{I}_N$	$\mathcal{O}_I$	$\mathcal{T}_I$	$\mathcal{I}_I$
Data-plane-informed BGP route selection											
[118]	BGP route validation via data-plane	Probing server within the Origin AS	3	4	**0.75**	1	2	0.12	1	2	0.12
[119]	Data-plane metrics in BGP route selection	Continuous performance monitoring	3	4	**0.75**	1	1	0.06	1	1	0.06
Path Enforcement Mechanisms											
[120]	Federated architecture of distributed PoPs	Traffic restricted to authorized paths	3	3	0.56	4	3	**0.75**	3	1	0.19
[121]	FC system for path verification	On-path ASes enforce forwarding; off-path filter anomalies	2	3	**0.38**	2	2	0.25	1	1	0.06
AS-Relationship-based route validation											
[122]	Bilateral filtering between neighbors	Propagation restricted to expected `AS_PATH` patterns	1	2	0.12	4	3	**0.75**	3	3	0.56
[123]	Cryptographically signed customer-cone declarations	Transit providers filter routes violating declared relationships	1	1	0.06	3	2	0.38	4	3	**0.75**

route announcements align with the actual forwarding paths used by packets, thereby improving the integrity and trustworthiness of end-to-end routing decisions.

Table IX summarizes the approaches discussed below, showing how data-plane-informed mechanisms remain largely academic, path enforcement solutions attract both academic and industrial attention, and AS-relationship-based validation is mainly driven by institutional and operational priorities.

Data-plane-informed BGP route selection. Early work in this direction explores integrating data-plane observations into routing decisions. The authors in [118] propose verifying that BGP-announced prefixes correspond to dataplane-reachable destinations by probing a designated server within the Origin AS. The binding between prefixes and validation endpoints is secured through an RPKI extension, and the mechanism is implemented in FRRouting with asynchronous validation to avoid routing instability. However, the guarantees remain limited, as the reachability of a single endpoint does not ensure the correctness of the full forwarding path. Moreover, the reliance on extra infrastructure and protocol extensions introduces deployment barriers, and the proposal remains at an early proof-of-concept stage.

Extending this idea beyond binary reachability, the authors in [119] explore the integration of quantitative data-plane measurements into the BGP decision process. The authors introduce a Route Quality Measurement (RQM) framework that continuously evaluates metrics such as latency and encodes them in a dedicated attribute to guide route selection toward higher-quality paths. Stability is improved through statistical ranking mechanisms that mitigate metric oscillations. Implemented in the goBGP stack using Virtual Routing and Forwarding (VRF) isolation for measurement traffic, the prototype demonstrates the feasibility of incorporating data-plane awareness into route selection and convergence behavior. Nevertheless, measurement overhead, scalability concerns, and interoperability with existing inter-domain policies introduce deployment challenges, and the approach has so far been validated only through laboratory experimentation.

Path Enforcement Mechanisms. While previous approaches improve the trustworthiness of routing decisions, some aim to enforce forwarding behavior more directly. The authors in [120] adopt an architectural approach based on the Secure Backbone AS (SBAS), which virtualizes a federation of geographically distributed PoPs into a single logical AS connected to the Internet via BGP. Within SBAS, routing exchanges are authenticated, and forwarding is protected, while customer prefixes are validated at ingress through RPKI. Traffic destined to protected prefixes is steered through SBAS PoPs and carried over the secure backbone, enforcing path compliance through architectural isolation rather than protocol-level verification. Although prototype experiments in controlled environments indicate improved resilience to routing attacks, the approach requires a cooperating backbone infrastructure and introduces additional architectural complexity.

A different direction is explored in [121], where the authors propose FC-BGP, a secure inter-domain routing system that provides verifiable forwarding assurance without requiring universal adoption across ASes. The approach relies on Forwarding Commitments (FCs), cryptographic constructs through which ASes commit to forwarding traffic over specific two-hop path segments, enabling authentication of partial paths and supporting incremental deployment. FCs are attached to BGP advertisements during route propagation, allowing downstream validation of path authenticity, while complementary data-plane mechanisms enforce forwarding compliance through coordinated filtering by both on-path and off-path ASes. Evaluation with real BGP traces shows that many path segments remain stable across updates, enabling FC-BGP to reduce verification overhead by $\approx$ 36% compared to BGPsec, while even limited deployment can filter a significant fraction of malicious traffic. However, the approach introduces additional cryptographic state management and coordination requirements among participating ASes, and its validation currently remains limited to prototype-based experimentation.

AS-Relationship-based route validation. An operational mitigation is proposed in [122], where the authors introduce Peerlock, a mechanism based on bilateral filtering agreements between neighboring ASes to prevent the propagation of *RLKs*.

TABLE X
OVERVIEW OF APPROACHES TO OVERCOME THE PARTIAL DEPLOYMENT PROBLEM

Ref.	Approach / Mechanism	Effectiveness in partial deployments	$\mathcal{O}_A$	$\mathcal{T}_A$	$\mathcal{I}_A$	$\mathcal{O}_N$	$\mathcal{T}_N$	$\mathcal{I}_N$	$\mathcal{O}_I$	$\mathcal{T}_I$	$\mathcal{I}_I$
Enhancements of ROV											
[124]	Enhanced ROV with improved route selection and proactive signaling	Designed for partial deployment; lightweight variant simplifies adoption	3	2	0.38	3	2	0.38	4	2	**0.5**
[125]	Decentralized validation combining online and offline checks	Extends protection beyond RPKI-ROV coverage	4	4	**1.0**	1	2	0.12	2	1	0.12
[126]	SDX-based ROV mechanism at IXPs	Operates with partial SDN/ROV adoption	4	3	**0.75**	3	2	0.38	2	1	0.12
Lightweight Path Validation Mechanisms											
[127]	BGPsec extension with Protected-OTC, Up attribute, and ProConID checks	Effective even with ∼10% deployment	3	3	**0.56**	3	2	0.38	3	1	0.19
[128]	Lightweight RPKI extension validating final AS hops	Regional or ISP-level deployment provides local protection	4	3	**0.75**	2	2	0.25	3	1	0.19
Policy-Based Solutions											
[129]	Blockchain-based bi-hierarchical ledger for BGP announcement validation	Cooperating AS groups; effectiveness depends on honest-majority consensus	4	4	**1.0**	3	1	0.19	3	1	0.19
[13]	Analysis of routing policies under partial path-validation deployment	Effectiveness depends strongly on operator policy choices	1	2	0.12	2	4	0.5	4	4	**1.0**
[130]	OTC attribute and DO community for route propagation control	With ASPA, prevents ∼98% of *RLKs* with top 5% adoption	3	1	0.19	3	2	0.38	4	3	**0.75**

A protecting AS discards routes that appear to traverse a protected AS unless they originate from that AS or from its authorized upstream providers. The approach is particularly suited to Tier-1 networks, where stable interconnection relationships facilitate the required coordination. The authors also propose Peerlock-lite, a simplified variant that removes the need for explicit coordination by exploiting the heuristic that transit providers should not receive routes containing Tier-1 AS numbers in the AS-PATH from customers, enabling lightweight deployment through local filtering rules. Measurements in [131] show partial deployment of Peerlock among Tier-1 ASes and broader adoption of Peerlock-lite among smaller providers, while simulations suggest that wider deployment across large ISPs could significantly mitigate Tier-1 *RLKs*. Despite its practicality and low deployment barrier, the mechanism mainly protects against *RLK* attacks and depends on accurate relationship knowledge and operator coordination, limiting its effectiveness outside well-connected core networks. The approach nevertheless reflects a relatively mature operational practice, with deployment in production networks despite the absence of universal adoption.

Within the same direction, the IETF has been developing *AS-Cone*, a mechanism primarily target Policy Violation attacks, such as *RLK*. AS-Cone defines the customer cone of an AS using digitally signed objects published through the RPKI infrastructure [123]. Similar to the *as-set* objects in IRR databases (Sec. IV), these records explicitly list the transit customers of an AS, enabling transit and peering providers to construct filtering policies that block route advertisements inconsistent with the declared customer cone, thereby helping to detect *RLKs*. However, because AS-Cone objects are self-declared, a malicious AS may fabricate customer–provider relationships, preventing the mechanism from mitigating *ASM-FI*. The proposal remained limited to a set of IETF drafts and never progressed to RFC status, reflecting an early experimental stage with no known operational deployment.

Like BGPsec, the AS-relationship-based validation mechanisms verify the integrity of the `AS_PATH`, by ensuring that the reported path is consistent with authorized routing relationships. While this limits the propagation of invalid routes, it does not guarantee that data traffic will actually follow the path indicated in the `AS_PATH` [132].

D. Partial Deployment Problem

The second R&I direction investigates the feasibility of security solutions that remain effective under partial adoption. This is particularly relevant in real-world scenarios, where only a subset of ASes may be willing or able to deploy a common security framework. In such cases, it becomes essential to design mechanisms that provide tangible security benefits without requiring universal adoption. Table X provides a summary of the following approaches, indicating that ROV enhancements are predominantly explored in academic contexts, lightweight path validation techniques receive interest from multiple stakeholders, and policy-oriented solutions are chiefly driven by operational and industry needs.

Enhancements of ROV. As discussed in Sec. IV-C, both public and private organizations increasingly rely on RPKI-based solutions to strengthen BGP routing security. However, the effectiveness of such solutions depends on the degree of RPKI adoption by the ASes. The NIST RPKI Monitor [133] provides a state-of-the-art platform for tracking the evolution of ROA coverage in global Internet routing. According to the latest statistics reported by the monitor, it can be seen that:

- since 2020, the trend of the amount of unique IPv4 address prefix/origin AS pairs has significantly increased with respect to the historical trend – cf. Fig. 4a;
- as for IPv6, the positive trend has always (since the start of the historical series, in 2019) been linear – cf. Fig. 4b;
- among the 25 ASes with the highest number of valid pairs (IPv4), there are well-known content distribution networks and cloud providers – cf. Fig. 5.

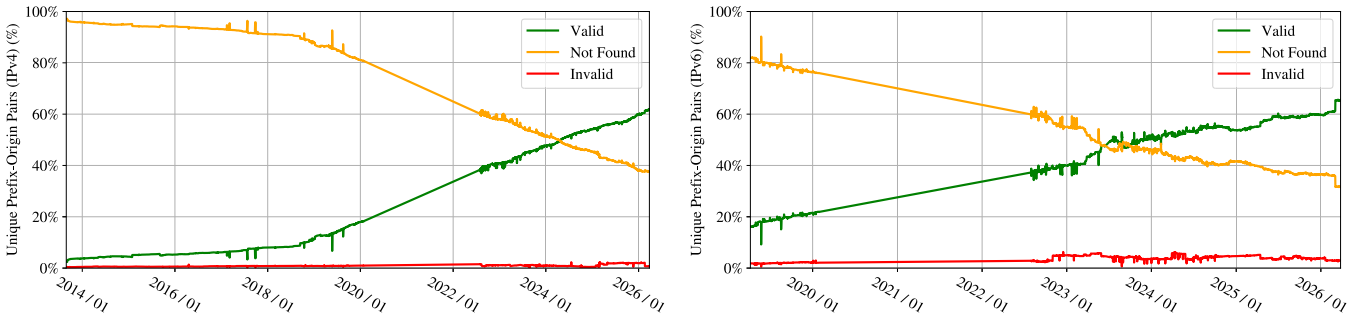

(a) Trend of unique IPv4 address prefix/origin AS pairs.

(b) Trend of unique IPv6 address prefix/origin AS pairs.

Fig. 4. Trend of unique IPv4 and IPv6 address prefix/origin AS pairs [133].

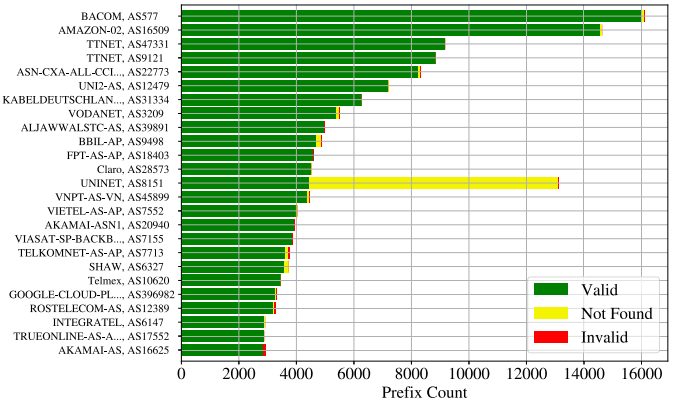

Fig. 5. Top-25 ASes by valid prefix/origin pairs as of March 2026 [133].

From these figures, the RPKI framework seems to have finally become the benchmark to prevent *PRH* [134], especially considering that the ASes of large CPs, which deliver most of the Internet traffic, are all protected by ROV. However, it is important to understand that these ASes represent the minority of the total number of ASes that, in turn, may not have embraced RPKI-based BGP routing security.

In addition to ROA coverage, the deployment of ROV plays a crucial role. Even with complete ROA coverage, *PRH* cannot be effectively mitigated unless ASes actively validate and filter RPKI-Invalid announcements. The longitudinal study in [135], covering more than 20 months and over 28k ASes, reports that 36.2% do not implement ROV, while only 12.3% achieve full protection. Some ASes may still benefit indirectly if their upstream providers enforce ROV, but heterogeneous deployment conditions, including selective filtering policies and multi-homing, often lead to uneven protection across routing paths. As a result, *partial RPKI deployment* remains the most common operational scenario.

Several research efforts, therefore, aim to enhance the effectiveness of ROV under partial deployment conditions. The authors in [124] propose ROV++, an extension of standard ROV by addressing interception scenarios caused by the coexistence of a valid covering prefix and an invalid sub-prefix from the same neighbor. Due to longest-prefix forwarding, routers may still select the hijacked route. ROV++ modifies the route selection process to select alternative safe paths, or locally blackhole announcements, or the generation of preventive sub-prefix announcements toward the legitimate origin. While these mechanisms improve interception resistance, they introduce trade-offs such as increased disconnections or incompatibility with path authentication schemes. The proposal has been evaluated through large-scale simulations and includes a deployable software-compatible variant (ROV++ Lite), but no operational deployment has been reported.

A different direction is explored in [125], where the authors propose Comp-RPKI, a complementary decentralized validation mechanism based on bilateral trust relationships. Neighboring ASes exchange public keys during offline business negotiations, allowing routers to authenticate routing updates when no valid ROA is available. In this way, each hop is authenticated either through RPKI or bilateral offline cryptographic validation, reducing reliance on centralized trust anchors. However, Comp-RPKI provides only local authentication and relies on the exchange and secure management of keys, which may introduce operational overhead. Moreover, the proposal remains largely conceptual, supported mainly by analytical discussion rather than large-scale experimentation.

Another direction focuses on improving validation capabilities at IXPs. The authors in [126] introduce SD-BROV, a Software-Defined eXchange (SDX)-based ROV mechanism aimed at enabling centralized validation within the exchange fabric. Incoming updates are checked against a local RPKI cache before datapath establishment, allowing IXPs to filter invalid announcements even when participant networks do not deploy ROV locally. Experimental evaluation shows that SD-BROV can detect and block invalid announcements before datapath installation. However, protection is limited to the SDX domain, and scalability concerns arise due to the processing overhead associated with validating large routing tables. Current evaluations are limited to emulated environments, and operational deployment has not yet been reported.

Lightweight Path Validation Mechanisms. Beyond origin validation, several proposals aim to secure the integrity of the `AS_PATH` while remaining deployable under partial adoption. Previous work has shown that the effectiveness of path validation approaches, such as e.g., S-BGP, soBGP, and BGPsec, not only on the level of deployment but also on the routing policies adopted by network operators. The authors in [13], based on a survey of 100 operators, show that even

partial deployment can provide strong protection when routing policies prioritize security (e.g., preferring secure routes). However, if network operators are likely to apply policies that privilege performance over security, the effectiveness of the protection can decrease strongly – this is reasonable, since path validation aims at preventing path-shortening attacks and performance-based routing prefers (possibly-bogus) short insecure routes over longer secure routes.

The authors in [127], propose BGP-iSec, an enhancement of BGPsec. The key idea is the reintroduction of mandatory transitive signatures: each adopting AS signs the announcement, and missing signatures from known adopters lead to rejection. This removes the "deployment island" limitation of BGPsec and enables partial path validation across non-adopting segments. To mitigate *RLKs*, BGP-iSec integrates three mechanisms: (i) Protected-OTC, a signed variant of the Only-to-Customer (OTC) attribute [136]; (ii) a hash-based UP attribute to restrict unauthorized upstream propagation; and (iii) ProConID, an RPKI object encoding provider-cone information to detect policy-inconsistent paths. Simulation results show significant reductions in interception rates, with performance at low adoption levels exceeding that of BGPsec near universal deployment. However, these benefits come with increased cryptographic overhead, additional transitive attributes, and new RPKI objects, while ProConID requires accurate provider-cone information. Overall, BGP-iSec remains a research proposal supported by simulation-based evaluation, with no reported operational adoption.

To explore deployable yet effective path-validation mechanisms, the authors in [128] introduce *path-end validation*, a lightweight extension to RPKI that authenticates only the last AS hop of the AS-PATH. The mechanism relies on offline-signed records in which an origin AS declares its authorized neighbors. Adopting ASes use these records to filter advertisements where the penultimate AS is not explicitly authorized, preventing *ASM-FI* without requiring modifications to BGP routers. The same records can also support validation of longer path suffixes when multiple adjacent ASes publish records, and a simple "non-transit" flag is proposed to mitigate certain *RLKs* originating from stub ASes. Simulation results indicate that even limited adoption by major ISPs significantly reduces attack success rates, approaching the protection achieved by fully deployed BGPsec. However, the mechanism does not guarantee full path integrity and remains vulnerable to collusion and certain policy-consistent manipulations. Although an open-source prototype is available, the approach has not been demonstrated in operational deployments.

Policy-Based Solutions. Finally, some approaches attempt to mitigate routing attacks through policy enforcement or alternative architectures rather than cryptographic path validation.

The authors in [129] propose RouteChain, a distributed architecture designed to mitigate *PRH* by recording route announcements as transactions in a bi-hierarchical ledger. ASes are grouped graphically into subgroups maintaining local ledgers, while a global chain records validated prefix ownership. Experiments on a 250-node testbed report millisecond-level confirmation latency. Security relies on majority agreement within and across subgroups: *PRH-CH* are treated as double-spending conflicts, while *PRH-IH* are detected via cross-checking with the global ledger. Although consensus can be achieved faster than typical hijack durations, the approach requires substantial coordination, continuous blockchain maintenance, and partial changes to the inter-domain trust model. Moreover, subgroup-based information sharing may limit routing policy autonomy, and evaluation remains confined to a controlled testbed with no reported operational adoption.

Instead, authors in [130] explore the use of existing BGP primitives, namely the Only to Customer (OTC) attribute and the Down Only (DO) community, to prevent *RLKs* through policy-constrained propagation. Both mechanisms rely on BGP Roles and encode Gao-Rexford-compliant forwarding semantics via explicit marking of UPDATE messages. By enforcing ingress and egress rules, ASes can detect and suppress multi-hop *RLKs* without modifying the BGP decision process. OTC is defined as an optional transitive path attribute [136], whereas DO relies on BGP Large Communities [137]. OTC requires explicit software support, but is less likely to be stripped in transit, while DO lowers the upgrade barrier but is more susceptible to removal. Simulations show that selective deployment among highly connected Tier-1 and Tier-2 ASes can mitigate over 98% of *RLKs*. OTC is already supported in several production BGP implementations, and early operational adoption has been observed, although deployment remains partial and the mechanism offers no cryptographic protection against malicious removal of markings.

E. Improving Existing Routing Security Frameworks

The third R&I direction investigates the feasibility of improving the state-of-the-art techniques mentioned in Sec. IV without impacting the efficiency of the routing operations. It is worth noting that many of the security mechanisms historically considered impractical due to the high computational cost on routing equipment are now increasingly viable. Thanks to significant improvements in the processing capabilities of modern routers – and the growing adoption of such hardware by network operators – cryptographic and validation-heavy solutions such as BGPsec are no longer inherently unfeasible. As also discussed in [150], performance limitations are no longer the primary obstacle to deployment: attention must now shift to operational complexity, incentive structures. Table XI provides a summary of the following approaches, indicating that path validation mechanisms are primarily explored in academic research, enhancements to RPKI-ROV attract strong institutional and industry interest, and decentralized or clean-slate solutions are largely driven by academic efforts with selective uptake from operational stakeholders.

Path Validation Mechanisms. Several works extend the path-validation paradigm originally introduced by S-BGP [70] (Sec. IV) to reduce the overhead of cryptographic verification.

In this direction, the authors in [138] propose Credible BGP (C-BGP), introducing a control layer composed of selected trusted ASes. Unlike S-BGP, which requires verification of every signature in the AS-PATH, C-BGP limits verification to the segment between the last trusted AS and the receiving AS,

TABLE XI
OVERVIEW OF APPROACHES TO IMPROVE EXISTING SECURITY MECHANISMS

Ref.	Approach / Mechanism	Security Guarantee	$\mathcal{O}_A$	$\mathcal{T}_A$	$\mathcal{I}_A$	$\mathcal{O}_N$	$\mathcal{T}_N$	$\mathcal{I}_N$	$\mathcal{O}_I$	$\mathcal{T}_I$	$\mathcal{I}_I$
Path Validation Mechanisms											
[138]	Trusted-AS control layer for selective validation	Near immunity to *PRH* with ∼60% trusted ASes	3	4	**0.75**	1	2	0.12	3	1	0.19
[139]	Fast Secure BGP using Critical Segment Attestations	Path authenticity; mitigation of *ASM*, especially *ASM-PS*	3	4	**0.75**	1	2	0.12	3	2	0.38
[140]	Hierarchical Origin AS and `AS_PATH` validation	Protect against *PRH* and *ASM*	4	3	**0.75**	2	1	0.12	3	2	0.38
[141] [142]	Path validation using aggregate signatures to compress AS-level attestations	Security guarantees equivalent to BGPsec	3	3	**0.56**	2	2	0.25	3	2	0.38
Enhancements to RPKI											
[143]	Distributed RP validation service	Consensus-based TLS validation	3	3	**0.56**	3	1	0.19	1	1	0.06
[144]	Signed Prefix List (SPL) objects for prefix authorization	Mitigates *ASM* and *PRH* via cryptographic validation	2	2	0.25	4	4	**1.0**	3	4	0.75
[145]	Route Path Authorization (RPA) objects for path policy validation	Cryptographic AS_PATH verification; mitigates *ASM* and route leaks	2	2	0.25	4	3	**0.75**	3	3	0.56
Decentralized Trust and Alternative Infrastructures											
[146]	Blockchain-based decentralized routing infrastructure	Consensus-based validation prevents *PRH*	4	4	**1.0**	3	1	0.19	1	1	0.06
[147]	Blockchain framework for IP, ASN, and BGP message authentication	Cryptographic ownership verification via multi-signatures	4	4	**1.0**	3	1	0.19	1	1	0.06
[148]	Decentralized prefix ownership certification in BGP attributes	Alternative to RPKI for prefix ownership validation	3	4	**0.75**	4	2	0.5	1	2	0.12
[149]	Clean-slate architecture with Isolation Domains	Path control and authenticated forwarding state	3	2	0.38	4	4	**1.0**	4	3	0.75

while prefix ownership is checked only if it has not already been validated by a trusted node. Analytical modeling and simulations show that when roughly 60% of ASes belong to the trusted layer, the network becomes nearly immune to *PRH* while reducing signature verifications by about 31% compared with full S-BGP deployment. However, the approach critically relies on the correct behavior of trusted ASes and introduces a hierarchical trust model without defining mechanisms for trust establishment, revocation, or governance, shifting to a trust management problem. Consequently, the proposal remains largely at a conceptual stage, with validation limited to analytical modeling and simulation, highlighting the trade-off between computational efficiency and trust centralization.

A different approach is proposed in [139], where authors introduce Fast Secure BGP (FS-BGP), a cryptographic mechanism to provide AS-PATH authentication with lower computational overhead than S-BGP. FS-BGP introduces Critical Segment Attestations (CSAs), where each AS signs only adjacent AS triplets rather than the full path, enabling path validation through chained segment signatures. The design relies on the Neighbor Based Import and Export (NBIE) observation that routing policies typically depend on adjacent AS relationships, thereby reducing the number of required cryptographic operations. Experiments show orders-of-magnitude reductions in cryptographic overhead compared to S-BGP while preserving comparable path authentication guarantees. However, the approach still requires a PKI infrastructure similar to S-BGP and relies on the NBIE assumption. Moreover, the segment-based scheme enables AS-path manipulation by concatenating valid segments. To mitigate this issue, the authors introduce Suppressed Path Padding (SPP), which artificially increases the length of suppressed paths but alters routing behavior and adds operational complexity. Overall, the proposal has been validated through trace-driven experiments and analytical evaluation, but lacks deployment evidence in operational networks.

Taking a different route, the authors in [140] propose the Hierarchical Origin and Path (HOP) verification method, which aims to validate both the Origin AS and the AS-PATH without modifying the BGP protocol itself. The architecture introduces a hierarchy of verification servers, each responsible for a group of ASes, which maintain routing topology and prefix ownership information in an Inter-AS Routing Table. Upon receiving an UPDATE, an AS queries the corresponding server to verify the origin and path information, and the result is summarized through a Route Credibility Score (RCS). The design avoids cryptographic attestations and instead relies on a server-based verification infrastructure, allowing lightweight operation and incremental deployment alongside legacy BGP. Routes involving non-participating ASes are not rejected but labeled with an *UNKNOWN* credibility state to preserve interoperability. However, the approach depends on centralized verification servers and on the continuous maintenance of accurate data. Moreover, the evaluation focuses mainly on analytical comparisons of message overhead, without Internet-scale experimentation, leaving open questions about scalability and resilience to server failures or compromise.

Other works address the practical feasibility of BGPsec. The work in [141] proposes APVAS, a path validation protocol that reduces BGPsec memory overhead by leveraging aggregate signatures. As in BGPsec, each AS signs routing information along the AS-PATH, but individual signatures are compressed into a single aggregate signature supporting sequential

aggregation, which preserves the ordered chain of AS attestations and interactive aggregation across routing updates. A prototype implemented in BIRD shows substantial reductions in routing attribute memory consumption compared to BGPsec. However, aggregation is limited to a single Secure-PATH, requiring routers to maintain separate signatures per path, so memory usage still grows linearly with the number of routes. To address this limitation, its advanced version APVAS+ [142], enables aggregation across multiple routing paths through a new cryptographic construction, Optimistic Bimodal Aggregate Signatures (OBAS). The scheme introduces partial signatures that support partial route advertisement while preserving compatibility with standard BGP operations. A prototype evaluated on Internet-scale datasets of roughly 800,000 routes shows further reductions in router memory requirements relative to APVAS and other BGPsec-based solutions. Nevertheless, the protocol still relies on pairing-based cryptographic operations that impose non-negligible computational overhead, and the evaluation is limited to controlled experimental environments rather than operational networks, leaving open questions regarding deployment feasibility, routing convergence impact, and compatibility with existing inter-domain routing policies.

Enhancements to RPKI. Beyond path validation, several works focus on improving the capabilities of RPKI ecosystem.

The authors in [143] propose ByzRP, a Byzantine-secure distributed Relying Party (RP) system to improve the robustness of ROV. Instead of relying on a single validator, ByzRP deploys multiple RP nodes that independently retrieve and validate ROAs from public repositories. The nodes form a permissioned network secured via mutual TLS and exchange their validated dataset to derive a consensus output through a threshold voting mechanism tolerant to Byzantine failures. The resulting set is distributed to routers via the RTR protocol, enabling ROV without requiring each network to maintain its own RP instance. The architecture also includes monitoring and skiplisting mechanisms to detect faulty or malicious repositories. While this approach improves availability, correctness, and resilience while remaining compatible with RPKI, it assumes coordinated deployment and majority honesty among participants, raising operational and trust concerns. In addition, the evaluation is limited to controlled experimental deployments rather than Internet-scale operational environments.

In contrast, the authors in [144] propose the Signed Prefix List (SPL), a new RPKI object intended to complement traditional ROV. Unlike ROAs, which authorize an AS to originate a prefix up to a maximum prefix length, SPLs allow an origin AS to publish a cryptographically signed list explicitly enumerating all prefixes it intends to announce. This enables an extended validation procedure, SPL-ROV, where routers verify routes against both ROA and SPL data and treat a route as invalid if either validation fails. By requiring the explicit prefix enumeration, SPLs provide finer control over origin authorization and reduce risks associated with permissive ROA configurations, particularly those involving large maximum prefix lengths. However, the approach introduces additional operational complexity, as operators must maintain accurate prefix lists and update SPL objects whenever routing policies change. Deployment also depends on support from routers and validators to integrate SPL-ROV into existing RPKI validation workflows, which currently limits operational adoption.

A complementary extension is proposed in [145] where authors introduce Route Path Authorizations (RPA), a new RPKI signed object designed to enable cryptographic verification of BGP propagation paths. RPA objects describe the authorized routing paths by specifying permitted upstream and downstream AS relationships for given prefixes. Routers can then compare the received AS-PATH with the authorized path segments encoded in RPA objects, classifying routes into states such as Valid, Weakly Valid, Invalid, or Unknown depending on their consistency with the declared propagation policies. By encoding routing intent directly in the RPKI, RPA aims to mitigate both *ASM* and *RLKs*. However, the approach introduces additional complexity in RPKI object management, as operators must explicitly describe and maintain routing ath policies. Enumerating valid path segments may become challenging in dynamic routing environments, raising scalability concerns for large-scale deployment.

Decentralized Trust and Alternative Infrastructures. A different line of research explores more fundamental changes to the trust infrastructure underlying inter-domain routing.

The authors in [146] propose ROAchain, a distributed infrastructure that decentralizes the management of ROAs. Each AS maintains a synchronized ROA repository stored on a blockchain ledger, where ROA registration, updates, and revocations are recorded as transactions. To improve scalability and security, the design combines sharding-based parallel verification, collective signatures, and a credence-based leader election mechanism, while a Management Authentication List (MAL) restricts participation to authorized nodes to limit Sybil attacks. Despite these design improvements, the approach introduces a dedicated blockchain infrastructure that requires the deployment of ROAchain nodes at each AS, potentially creating significant adoption barriers. The system also relies on assumptions about network synchrony and honest participants that may not fully reflect Internet conditions. Moreover, the evaluation is conducted in a controlled cloud environment with a limited number of nodes, leaving open questions regarding scalability and integration with the existing RPKI ecosystem.

Another proposal is presented in [147], where the authors explore the use of a distributed ledger to manage Internet resources such as IP prefixes, ASNs, DNS domains, and BGP advertisements. Resource allocations and authorizations are encoded as blockchain transactions, enabling verification of ownership and routing permissions through an immutable transaction history without relying on a centralized PKI root of trust. The framework also supports multi-signature transactions for joint authorization and extends validation beyond origin authentication by recording BGP advertisements as blockchain entries referencing upstream and downstream ASes. A translation module placed in front of BGP speakers converts BGP UPDATE messages into blockchain transactions and vice versa, allowing integration without modifying the BGP protocol. However, recording Internet-scale routing dynamics on a blockchain raises major scalability concerns, as the transaction throughput required to capture BGP

TABLE XII
OVERVIEW OF AUXILIARY TOOLS FOR ENHANCING BGP ROUTING SECURITY WITHOUT MODIFYING CORE PROTOCOLS

Ref.	Approach / Mechanism	Security Focus	$\mathcal{O}_A$	$\mathcal{T}_A$	$\mathcal{I}_A$	$\mathcal{O}_N$	$\mathcal{T}_N$	$\mathcal{I}_N$	$\mathcal{O}_I$	$\mathcal{T}_I$	$\mathcal{I}_I$
Data-Driven Attack Detection											
[154]	Random Forest on AS-triplet features	*RLK* detection via `AS_PATH` anomaly	2	2	0.25	4	3	**0.75**	2	3	0.38
[155]	Federated learning with blockchain over anonymized AS triples	Privacy-preserving *RLK* detection	3	4	**0.75**	2	1	0.12	1	1	0.06
[156]	Random Forest using new AS links and topology features	Detection of *ASM-FI*	3	3	**0.56**	3	2	0.38	2	1	0.12
[157], [158]	Origin-AS change detection	*PRH* detection	3	1	0.19	4	3	0.75	4	4	**1.0**
[159]	Analysis of NO EXPORT misuse to evade BGP monitoring	Detection of stealthy *PRH-IH* hidden from monitors	4	3	**0.75**	3	2	0.38	1	1	0.06
Operational Tools											
[160]	Resilience-based Tor relay selection with real-time monitoring	*PRH* detection affecting Tor anonymity	3	3	**0.56**	3	1	0.19	2	1	0.12
[161]	Cross-validation of BGP updates with local configuration	Detection of *PRH-CH/IH* and *ASM*	3	2	0.38	4	3	**0.75**	4	2	0.5
[162]	Configuration validation using IRR, BGP data, and emulation	MANRS compliance verification	1	2	0.12	4	4	**1.0**	4	4	**1.0**
[163]	Global routing database with enhanced SDN route reflectors	*PRH* and *ASM* detection	3	2	0.38	4	3	**0.75**	4	2	0.5
Risk Assessment											
[164], [165]	Risk-level metric from IXP RS RIB	*PRH* detection and AS Risk Level	4	3	0.75	2	3	0.38	4	4	**1.0**
[166]	Path deviation and credibility scoring	*ASM* detection	4	3	**0.75**	3	3	0.56	4	3	**0.75**

updates would exceed the capacity of current blockchain systems.

Furthermore, the authors in [148] propose DISCO (Decentralized Infrastructure for Securing and Certifying Origins), a system to automate the certification of IP prefix ownership, without relying on hierarchical PKI. DISCO verifies the ownership of prefixes by observing long-term routing behavior: a software agent within an AS attaches a public-key ownership indicator to BGP announcements via an optional transitive attribute, while independent registrars monitor BGP updates from multiple vantage points to verify that the prefix is consistently advertised with the same indicator during a certification interval. Once a sufficient agreement is reached, the prefix-to-key association is certified and can be used to generate ROA-like records. This design enables automated certification without coordination with upstream providers and requires only configuration changes rather than protocol modifications. Nevertheless, DISCO certifies operational possession rather than legal ownership, relies on long-term observations that may delay certification during ongoing attacks, and cannot verify prefixes that are not actively advertised in BGP. The proposal has been demonstrated through experimental evaluation but lacks validation in operational deployments.

Finally, the authors in [149] propose SCION, a clean-slate Internet architecture designed to replace BGP and address challenges in routing security, reliability, and performance. SCION introduces Isolation Domains (ISDs), which group ASes into administrative regions sharing a common routing infrastructure. This structure limits the propagation of misconfigurations or malicious activity beyond the affected ISD. Each ISD operates under a shared trust model with locally defined trust anchors, while maintaining connectivity with other ISDs for global communication. A key feature of the architecture is path-aware networking, which allows ASes to select and control packet forwarding paths. This capability improves transparency and resilience by avoiding untrusted routes, supporting multipath communication, and enabling rapid failover. However, several limitations may hinder adoption. The architecture requires new infrastructure components and operational practices beyond those of the current Internet, increasing deployment and management complexity [151]. Path availability may vary dynamically, reducing performance predictability, and multipath transmission does not always improve latency [152]. Early deployments have also revealed challenges such as interoperability between implementations and the need for automated certificate management [153]. As with other clean-slate proposals, adoption is constrained by strong network effects, since large-scale benefits emerge only with broad participation [153]. Nevertheless, the architecture has progressed beyond conceptual designs and has been validated through operational testbeds and pilot deployments.

F. Auxiliary Tools to Mitigate BGP Weaknesses

The fourth and last R&I direction investigates mechanisms externally to BGP that do not alter the protocol itself or any security frameworks described in Sec. IV. Instead, they leverage auxiliary data and real-time monitoring systems to provide early detection and risk assessment capabilities. Table XII provides a summary of the following approaches, showing that data-driven detection techniques are largely explored in academia, operational tools are primarily supported by institutional and industry actors, and risk assessment methods receive strong interest from both academic and industrial communities.

Data-Driven Attack Detection. A first class focuses on detecting routing anomalies by analyzing the BGP data stream.

The authors in [154] propose RoLL, a system designed to detect and locate *RLKs* by analyzing AS triplets extracted

from the AS-PATH. Each triplet is characterized through five features obtained from external datasets: distance from Tier-1 ASes, AS degree, IPv4 address space, geographic location, and AS type. The system architecture includes a preprocessor that extracts triplets from AS-PATHs, a feature extractor that retrieves the associated attributes, and a detection module that applies a Random Forest classifier to identify suspicious patterns. Experimental evaluation reports approximately 91% detection accuracy and leak localization latency below 10 ms. However, the approach relies on multiple external datasets describing AS topology and attributes, which may affect feature reliability. In addition, the model is trained on a relatively limited set of known *RLKs* events, potentially limiting generalization to previously unseen routing behaviors. The proposal has been validated through controlled experiments but has not yet been evaluated in operational Internet deployments.

The authors in [155] propose FL-RLD, a privacy-preserving framework for detecting *RLKs* based on a blockchain-enabled federated learning architecture. Instead of sharing sensitive business relationships, routing data is represented as anonymized AS triples labeled as malicious or regular according to local routing policies. Each AS trains a local model using its own data and shares only model updates through a blockchain network, where updates are aggregated into a global model via federated averaging. The blockchain layer provides integrity and traceability of updates while avoiding the need for a centralized aggregation server. To address the scarcity of labeled *RLK* events, FL-RLD employs a self-labeling mechanism derived from local routing behavior, particularly the valley-free rule. Detection is performed using an LSTM-based model. Experimental results show that the federated approach achieves performance close to centralized learning while preserving the confidentiality of routing policies. Despite these advantages, the integration of both federated learning and blockchain infrastructures introduces significant operational complexity, requiring ASes to deploy dedicated nodes and participate in consensus operations.

The authors in [156] presents DFOH, a system for detecting *ASM-FI* through the analysis of public BGP data. The approach identifies newly observed AS links in routing paths that are absent from historical AS-level topologies, which may signal malicious path manipulation. To distinguish benign changes from attacks, DFOH extracts features related to network topology, peering characteristics, routing policies, and link bidirectionality and processes them using a supervised Random Forest classifier. Experimental results show that DFOH detects approximately 90.9% of *ASM-FI* within about five minutes while generating only a limited number of daily alerts. Despite these promising results, the approach relies on the appearance of previously unseen AS links, which an attacker could evade by reusing existing topology links. The inference process also depends on auxiliary datasets such as IRR records, PeeringDB information, and inferred AS relationships, which may contain inaccuracies. Additionally, operational misconfigurations, such as incorrect AS-path prepending, can introduce previously unseen AS adjacencies that resemble malicious topology changes, potentially leading to false positives requiring manual validation by operators.

A major network operator developed a *PRH* detection system to identify unauthorized announcements in near real time [157] by monitoring changes of origin AS. The architecture includes a prefix-origin change detection module that analyzes live and historical BGP updates, and a detection module that correlates multiple datasets, such as ROAs, IRR records, bogon lists, and inferred AS relationships, to estimate the likelihood of a *PRH* via a scoring function. Detected incidents are then aggregated and distributed through an alerting component. Comparable functionality is offered by the authors in [158], which combines public BGP feeds with anonymized private data to improve routing visibility and support near real-time anomaly detection. These platforms typically provide ASN and prefix monitoring, AS-PATH tracking, multi-vantage-point reachability analysis, historical inspection, and correlation with traffic telemetry. However, their effectiveness depends on the accuracy and completeness of external datasets, while scoring-based classification mechanisms often require continuous calibration and manual verification to distinguish malicious events from legitimate routing changes.

While the aforementioned detection systems rely heavily on data from public BGP monitors, the work in [159] analyzes how attackers can exploit the behavior of the well-known NO EXPORT BGP community to prevent their malicious route announcements, a *PRH-IH*, from being propagated to BGP monitoring systems. They demonstrated that Tier 1 ASes commonly honor the NO EXPORT community in a way that hides routes from monitors. Internet topology simulations show that even compromising a few major networks can let attackers hijack traffic from up to 37% of AS-level paths.

Operational Tools. Beyond detection algorithms, several systems provide operational frameworks that enable network operators to monitor routing behavior and react to attacks.

The authors in [160] propose Counter-RAPTOR, a defense mechanism to protect the Tor anonymity network from *PRH*. A measurement study of the Tor ecosystem shows that several high-bandwidth relay ASes exhibit limited resilience to routing attacks due to their position in the Internet topology. Based on this observation, the authors design a guard relay selection algorithm that incorporates an AS resilience metric into the relay selection process. The algorithm combines relay bandwidth with the estimated probability that the relay's AS can resist malicious BGP advertisements. Experimental results indicate that this strategy can increase the probability that clients remain resilient to *PRH* by up to 36% on average. The system is complemented by a real-time BGP monitoring component that analyzes routing updates affecting Tor relays using origin checks and temporal analytics to identify suspicious announcements. However, the effectiveness of resilience-based relay selection depends on accurate inference of AS relationships and routing policies. Moreover, the evaluation is primarily conducted through simulations and controlled experiments, leaving open questions regarding the effectiveness of the approach under large-scale operational deployment.

A more comprehensive defense is presented in [161], where the authors propose ARTEMIS, a self-operated defense system

that enables ASes to detect and mitigate *PRH* within seconds using real-time BGP updates from public monitoring infrastructures and local routing information. The system cross-checks incoming routing advertisements against a locally maintained configuration specifying authorized prefixes, origin ASNs, and neighboring relationships. Detection follows a two-stage process: an initial alert triggered by suspicious events, followed by an optional refinement phase that gathers additional observations to improve accuracy and estimate the attack's impact. This design allows ARTEMIS to detect several attack types, including *PRH-CH*, *PRH-IH*, and *ASM*. Once an attack is detected, the system can automatically trigger mitigation actions, such as prefix deaggregation or outsourced MOAS announcements, enabling response times on the order of one minute. However, ARTEMIS requires operators to maintain accurate and updated local configuration data describing legitimate routing policies and neighbor relationships, which may introduce management overhead.

Another approach is the ROuting SEcurity Tool (ROSE-T), an open-source framework designed to automatically verify compliance with MANRS initiatives [162]. ROSE-T aims to facilitate MANRS adoption by enabling operators to assess whether their routing configurations satisfy two MANR actions focusing on route filtering and anti-spoofing. The verification collects routing policy information from IRRs, including (mp-)import and (mp-)export objects and AS relationships, together with BGP route dumps. These datasets are used to compare declared routing policies with observed announcements. The system then parses the vendor-specific router configuration and converts it into a vendor-agnostic representation. Finally, ROSE-T leverages the Kathará [167] network emulation framework to emulate the AS topology and configurations, enabling automated verification of compliance. However, the approach relies heavily on the accuracy and completeness of IRR records and configuration data, which are known to be inconsistently maintained across operators. The system has been demonstrated through prototype implementations and experimental validation [168].

The authors in [163] propose a framework to mitigate both *PRHs* and *RLKs*. The system constructs a global routing information database that aggregates datasets such as ROAs, IRR, and geographic routing data. This database supports large-scale monitoring and is integrated with an SDN-based control architecture that enhances BGP Router Reflectors to perform real-time route verification. Incoming routes are continuously compared against the trusted routing database to detect abnormal advertisements or policy violations. When suspicious routes are identified, the system can automatically enforce mitigation actions, such as filtering or blocking the routes at network boundaries. The architecture also includes monitoring and auditing functions to analyze route registration consistency and AS information, enabling minute-level interception of *RLKs* and prevention of *PRHs*. Although the system has been deployed within an operator's infrastructure and monitors tens of millions of routing sources, it relies on the construction and continuous maintenance of a centralized global routing information database integrating heterogeneous datasets such as IRR and RPKI records. In addition, the solution is coupled with the internal SDN-based routing control architecture of the deploying operator, which may limit its adoption.

Risk Assessment. If an AS owner, given its specific risk appetite, was to define a risk management policy as an acceptable trade-off between having all other ASes adopt RPKI to minimize both the possibility of human error and endogenous threats, and relying on trust between ASes owners, one direction is to investigate effective *risk-based route ranking* methods to support operators in the assessment of the exchanged routing information. For example, the authors of [164] propose a tool that assigns risk levels to ASes involved in *PRH* within an IXP. The system analyzes prefixes within the Routing Information Base (RIB) of the RS to identify suspicious announcements. For each case, two metrics are computed: the Source-Based Risk Level (RL_{SB}), which evaluates the risk associated with the origin AS, and the Path-Based Amplification Level (RL_{PB}), which estimates the amplification risk introduced by intermediate ASes along the path. A one-year analysis demonstrates the feasibility of the approach and its usefulness for operational monitoring. This framework is further extended in [165], where the RL metrics are operationalized into a service that IXPs can offer to their members to dynamically penalize high-risk ASes within the BGP decision process, limiting the propagation of prefix hijackings. However, the accuracy of the risk assessment depends on the reliability of external databases such as AS relationship and IRR records, which may contain incomplete or outdated information. Moreover, the weighting of risk metrics relies on the Analytic Hierarchy Process (AHP), introducing a degree of subjectivity that may require recalibration across different operational environments. The approach has been demonstrated through prototype analysis on real operational data but has not yet been widely deployed in production infrastructures.

Related to this idea, the work in [166] introduces B-Secure, a dynamic reputation system to detect *ASM* by analyzing anomalies in observed BGP paths. The system evaluates routing behavior using two metrics: path deviation and path credibility. The first measures how much an AS-PATH differs from historically observed routes between the same AS pair. In parallel, the latter evaluates the likelihood of observing a sequence of intermediate ASes based on historical routing data. Paths exhibiting both significant structural deviation and low statistical credibility are flagged as suspicious. Evaluation of 18 months identified 109 anomalous AS-PATHs, demonstrating the feasibility of reputation-based anomaly detection. Despite these results, the effectiveness of the approach depends heavily on the availability of extensive and representative historical routing data. Legitimate but infrequent events, such as new peering relationships or traffic engineering adjustments, may therefore appear anomalous and increase the risk of false positives. More broadly, the method assumes that historical routing behavior provides a stable baseline for normal operation, an assumption that may not always hold in the highly dynamic and evolving inter-domain routing ecosystem.

TABLE XIII
OUTCOMES FOR FUTURE DIRECTIONS OF R&I ON INTERNET ROUTING SECURITY

R&I direction	Academic focus	Industrial focus	Institutional focus
Routing information versus forwarding state.	Propose novel and potentially disruptive, and revolutionary security architectures.	Forecast precise traffic load and resource allocation; improve performance/cost ratio.	Prospective adoption of improved standard mechanisms ensuring a more secure network operation.
	$\mathcal{O}_A = 4,\ \mathcal{T}_A = 4,\ \mathcal{I}_A = \mathbf{1.0}$	$\mathcal{O}_I = 4,\ \mathcal{T}_I = 3,\ \mathcal{I}_I = \mathbf{0.75}$	$\mathcal{O}_N = 3,\ \mathcal{T}_N = 3,\ \mathcal{I}_N = \mathbf{0.5625}$
Partial deployment problem.	Analyze and propose security mechanisms resilient to partial deployment.	Investigate novel security frameworks under partial deployment status to assess their usability in the short-term.	Stimulate the adoption of appropriate BGP routing security solutions even in partial deployments.
	$\mathcal{O}_A = 3,\ \mathcal{T}_A = 4,\ \mathcal{I}_A = \mathbf{0.75}$	$\mathcal{O}_I = 2,\ \mathcal{T}_I = 2,\ \mathcal{I}_I = \mathbf{0.25}$	$\mathcal{O}_N = 4,\ \mathcal{T}_N = 3,\ \mathcal{I}_N = \mathbf{0.75}$
Improving existing routing security frameworks.	Investigate and propose small, incremental improvements to existing frameworks.	Integration of improvements to existing frameworks.	Prospective adoption of improved standard mechanisms ensuring a more secure network operation.
	$\mathcal{O}_A = 2,\ \mathcal{T}_A = 1,\ \mathcal{I}_A = \mathbf{0.125}$	$\mathcal{O}_I = 3,\ \mathcal{T}_I = 2,\ \mathcal{I}_I = \mathbf{0.375}$	$\mathcal{O}_N = 3,\ \mathcal{T}_N = 2,\ \mathcal{I}_N = \mathbf{0.375}$
Auxiliary tools to mitigate BGP weaknesses.	Develop detection tools separated from the BGP process; envision data features that may help better detect/classify attacks; propose methods to perform attribution.	Foster auxiliary tools adoption to monitor and respond to BGP routing security attacks, especially in real-time.	Proliferation of monitoring tools to assess the status of the network operations.
	$\mathcal{O}_A = 4,\ \mathcal{T}_A = 4,\ \mathcal{I}_A = \mathbf{1.0}$	$\mathcal{O}_I = 4,\ \mathcal{T}_I = 3,\ \mathcal{I}_I = \mathbf{0.75}$	$\mathcal{O}_N = 4,\ \mathcal{T}_N = 4,\ \mathcal{I}_N = \mathbf{1.0}$

VI. ENVISIONED TRENDS OF THE IDENTIFIED R&I DIRECTIONS

In this section, we outline actionable directions aligned with the four identified R&I directions on BGP routing security in Sec. V, showing how academia, industry, and institutions can contribute to improving BGP security under their specific domain of influence. The detailed description of each topic is presented throughout this section and summarized in Tab. XIII.

Routing information versus forwarding state. According to the contributions collected in Sec. V-F, this seems to be the most immature direction to explore, as valuable outputs from both the academic and industrial realms are available, but their practical adoption can be considered behind schedule. From an academic perspective, universities could propose novel, potentially disruptive security architectures that, despite having a low TRL, could pave the way for new security frameworks. Conversely, the industry, with its precise control over routing information and forwarding state, may forecast traffic load to perform a more precise resource allocation, thereby improving the performance-to-cost ratio.

Partial deployment problem. In this context, as can be seen in Sec. V-D, the research ecosystem could continue to propose security mechanisms that are resilient to a partial deployment scenario. On the other hand, the industry is focusing on ramping up RPKI deployments; thus, this direction could be of less interest. However, there is potential for collaboration between academia and industry, where academia proposes a new framework and a subset of the industry adopts it, allowing for real-world evaluation of its effectiveness. From the perspective of institutions, as anticipated in Sec. V-D, there exist real-world scenarios like, e.g., nationwide Public Administration networks, where only a subset of ASes would benefit from deploying a common BGP routing security framework. In these cases, institutions may be interested in stimulating the adoption of such solutions to secure the local, inter-AS level, without the requirement of global deployment.

Improving existing routing security frameworks. Unlike the previous direction, this is typically of greater interest to industry than academia, as the latter could only design incremental improvements to already existing frameworks – cf. Sec. V-E. Academia may be less interested in such incremental approaches, as they require addressing all the technicalities and constraints present in real-world production systems, tending instead to focus more on disruptive innovations.

Auxiliary tools to mitigate BGP weaknesses. Based on the contributions in Sec. V-F, both academia and industry could put effort in these directions, as the former could introduce novel solutions which are not tightly related to an existing framework, while the latter could directly benefit from such solutions in terms of detection and mitigation of cyber-threats. Moreover, in order to govern the complexity related to the global network architecture, institutions are likely to consider auxiliary tools as crucial because they allow for monitoring and assessing the current state of the cyber-threats against BGP routing security, other than allowing for smooth integration of them in cyber-risk assessment tools.

VII. CROSS-CUTTING INSIGHTS AND LESSONS LEARNED

Based on the broad landscape of BGP routing security surveyed throughout this work, we extract a set of cross-cutting insights that go beyond the individual analysis of each R&I direction, consolidating what has practically succeeded, what has repeatedly failed, which assumptions have proved unrealistic, and which deployment trade-offs consistently emerge across academia, industry, and institutions.

A. Lightweight and Incremental Mechanisms Achieve Broader Adoption

The history of BGP security clearly shows that mechanisms requiring minimal changes to existing infrastructure and operations achieve wider deployment. RPKI, IRR, and more recently ASPA [86], have gained traction precisely because they can be adopted incrementally, without requiring

simultaneous action from all ASes. In contrast, comprehensive frameworks such as S-BGP [70], soBGP [71], and psBGP [72], despite their stronger security guarantees, have remained undeployed due to their operational complexity and end-to-end adoption requirements. This pattern consistently emerges across all stakeholders and suggests that, in the inter-domain routing ecosystem, *good enough and deployable* systematically outperforms *optimal but impractical*.

B. Universal Adoption Is an Unrealistic Assumption

A recurring theme across the surveyed literature is the implicit or explicit assumption that security guarantees are only meaningful under full deployment. BGPsec [80] is the most prominent example: a single non-adopting AS in the path strips all security information, rendering partial deployment of limited benefit. This assumption has proven unrealistic in the highly heterogeneous and decentralized Internet ecosystem, where economic incentives, administrative constraints, and technical capabilities vary enormously across ASes. Research efforts that are explicitly designed for partial deployment, such as ROV++ [124], path-end validation [128], and approaches based on the OTC attribute [130], [136], represent a more pragmatic and impactful direction and should be considered the baseline design assumption for future proposals.

C. The Gap Between Control-Plane Security and Data-Plane Enforcement Remains Open

Despite significant progress in origin and path validation, a fundamental limitation persists: securing the routing information exchanged via BGP does not guarantee that actual packet forwarding follows the validated path. This gap between the control and the data plane is acknowledged across the surveyed literature but remains unaddressed in operational deployments. Proposals such as FC-BGP [121] and SBAS [120] attempt to bridge this gap, but they remain at early experimental stages. Closing this gap represents one of the most technically challenging and academically relevant open problems in Internet routing security.

D. Operational Complexity Is a First-Class Deployment Barrier

Across all surveyed mechanisms, operational complexity consistently emerges as a primary obstacle to adoption, often more so than computational overhead or cryptographic cost. Managing ROA objects, IRR records, ASPA provider relationships, and BGPsec certificates introduces non-trivial operational burdens, particularly for smaller ASes with limited technical staff. Official recommendations from MANRS [114] and, e.g., ENISA [62] have recognized this challenge and attempted to address it through best practice guidelines and tooling support. Future mechanisms should explicitly account for operational simplicity as a design objective, not an afterthought.

E. Incentive Misalignment Hinders Collective Security

BGP routing security is inherently a collective action problem: security benefits are shared across the Internet, while deployment costs are borne individually by each AS. This asymmetry explains why many well-designed mechanisms fail to achieve widespread adoption. Institutional frameworks and policy mandates, such as, e.g., the EU NIS2 Directive [94], represent an important lever to rebalance this calculus. However, their effectiveness depends on enforcement and compliance monitoring, which remain heterogeneous across jurisdictions.

F. Auxiliary Tools Fill the Gap Left by Incomplete Protocol-Level Security

Given the limitations of protocol-level security mechanisms, auxiliary tools for monitoring, anomaly detection, and risk assessment play a critical complementary role. Systems such as, e.g., ARTEMIS [161], ROSE-T [162], and IXP-based risk assessment tools [164], [165] provide operational value precisely because they do not require changes to BGP itself or coordination across multiple ASes. Their effectiveness, however, depends on the quality and completeness of external datasets such as IRR records, RPKI repositories, and AS relationship databases, which are known to be inconsistently maintained. Improving the accuracy and coverage of these datasets is therefore a prerequisite for the effectiveness of the broader BGP security ecosystem.

G. Stakeholder Priorities Shape the Evolution of BGP Security

A final cross-cutting observation is that the trajectory of BGP security is not determined solely by technical merit, but by the interplay of academic innovation, industrial pragmatism, and institutional governance. Academia tends to propose disruptive solutions at low TRLs, industry prioritizes deployable and backward-compatible mechanisms, and institutions focus on promoting adoption through policy and coordination. These roles are complementary but not always aligned: mechanisms that are academically rigorous may be operationally impractical, while institutionally mandated solutions may lag behind the state of the art. Bridging these gaps, as this survey has attempted to do, is essential for translating research advances into real-world improvements in Internet routing security.

VIII. Conclusion

In this survey, we provide a comprehensive, multi-stakeholder perspective on BGP routing security, integrating the viewpoints of academia, industry, and institutional entities. We first propose a structured taxonomy of BGP routing attacks, offering a framework to classify and prioritize threats according to their characteristics and target components of the protocol. In parallel, we review the current defense approaches proposed at academic and industrial level, together with policy frameworks and best practices released by various institutions, aimed at improving BGP security. Finally, we tackle open research questions identified in prior studies and survey the literature to identify novel research direction, which are mapped to the interests of the different stakeholders to we illustrate how they can contribute to improve BGP security under the

different domains of influence. The presented methodology offers a holistic understanding of the field, guiding future research, informing operational decisions, and supporting policy development to strengthen the security and resilience of the global Internet infrastructure.

ACKNOWLEDGMENT

The contributions provided by the authors to this manuscript are personal and do not necessarily reflect those of the National Cybersecurity Agency of Italy.

REFERENCES

[1] G. Huston, M. Rossi, and G. Armitage, "Securing BGP—A literature survey," *IEEE Commun. Surveys Tuts.*, vol. 13, no. 2, pp. 199–222, 2nd Quart., 2011.

[2] A. Chatzivasileiou, A. Kornilakis, K. Lionta, G. Nomikos, X. Dimitropoulos, and G. Smaragdakis, "How Russia's invasion of Ukraine impacted the internet peering of the conflicted countries," in *Proc. 8th Netw. Traffic Meas. Anal. Conf. (TMA)*, May 2024, pp. 1–10.

[3] R. Fontugne, K. Ermoshina, and E. Aben, "The internet in crimea: A case study on routing interregnum," in *Proc. IFIP Netw. Conf. (Netw.)*, Jun. 2020, pp. 809–814.

[4] W. Ji, L. Yuxin, C. Libo, and T. Jiuting, "Using the technology readiness levels to support technology management in the special funds for marine renewable energy," in *Proc. OCEANS*, Apr. 2016, pp. 1–5.

[5] A. H. Muosa and A. H. Ali, "Detecting bgp routing anomalies using machine learning: A review," in *Forthcoming Networks and Sustainability in the AIoT Era*. Cham, Switzerland: Springer, 2024.

[6] J. Raynor, T. Crnovrsanin, S. Di Bartolomeo, L. South, D. Saffo, and C. Dunne, "The state of the art in BGP visualization tools: A mapping of visualization techniques to cyberattack types," *IEEE Trans. Vis. Comput. Graph.*, vol. 29, no. 1, pp. 1059–1069, Jan. 2023.

[7] N. H. Hammood, B. Al-Musawi, and A. H. Alhilali, "A survey of BGP anomaly detection using machine learning techniques," in *Proc. Int. Conf. Appl. Techn. Inf. Secur.*, 2022, pp. 109–120.

[8] B. A. Scott, M. N. Johnstone, and P. Szewczyk, "A survey of advanced border gateway protocol attack detection techniques," *Sensors*, vol. 24, no. 19, p. 6414, Oct. 2024.

[9] M. O. Nicholes and B. Mukherjee, "A survey of security techniques for the border gateway protocol (BGP)," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 1, pp. 52–65, 1st Quart., 2009.

[10] A. Mitseva, A. Panchenko, and T. Engel, "The state of affairs in BGP security: A survey of attacks and defenses," *Comput. Commun.*, vol. 124, pp. 45–60, Jun. 2018.

[11] B. Al-Musawi, P. Branch, and G. Armitage, "BGP anomaly detection techniques: A survey," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 1, pp. 377–396, 1st Quart., 2017.

[12] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," *Proc. IEEE*, vol. 98, no. 1, pp. 100–122, Jan. 2010.

[13] R. Lychev, S. Goldberg, and M. Schapira, "BGP security in partial deployment: Is the juice worth the squeeze?" in *Proc. ACM SIGCOMM Conf. SIGCOMM*, Aug. 2013, pp. 171–182.

[14] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proc. Conf. Appl., Technol., Archit., Protocols Comput. Commun.*, Aug. 2002, pp. 3–16.

[15] T. Alfroy, T. Holterbach, T. Krenc, K. C. Claffy, and C. Pelsser, "The next generation of BGP data collection platforms," in *Proc. ACM SIGCOMM Conf.*, Aug. 2024, pp. 794–812.

[16] P. Gill, M. Schapira, and S. Goldberg, "Let the market drive deployment: A strategy for transitioning to BGP security," in *Proc. ACM SIGCOMM Conf.*, Aug. 2011, pp. 14–25.

[17] R. Subramanian, "The growth of global internet censorship and circumvention: A survey," *Commun. IIMA*, vol. 11, pp. 69–90, Oct. 2011.

[18] S. Ginoza and H. Flanagan, *RFC Style Guide*, document RFC 7322, Sep. 2014.

[19] (Mar. 2002). *Directive 2002/21/EC of the European Parliament and of the Council of 7 March 2002 (Framework Directive)*. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2002/21/oj/eng

[20] (2026). *FTTH Council*. [Online]. Available: https://www.ftthcouncil.eu/

[21] (Mar. 2025). *Investment Tracker-March 2025 Update*. [Online]. Available: https://www.ftthcouncil.eu/resources/all-publications-and-assets/2357/investment-tracker-march-2025-update

[22] (2026). *Fiber Broadband Association*. [Online]. Available: https://fiberbroadband.org/

[23] (Mar. 2025). *The State of the North American Fiber Deployment*. [Online]. Available: https://fiberbroadband.org/resources/the-state-of-the-north-american-fiber-deployment-january-2025/

[24] (2026). *Submarine Cable Map*. [Online]. Available: https://www.submarinecablemap.com/

[25] A. Purvins, L. Sereno, M. Ardelean, C.-F. Covrig, T. Efthimiadis, and P. Minnebo, "Submarine power cable between Europe and North America: A techno-economic analysis," *J. Cleaner Prod.*, vol. 186, pp. 131–145, Jun. 2018.

[26] E. Dahlman, S. Parkvall, and J. Skold, *4G: LTE/LTE-Advanced for Mobile Broadband*. New York, NY, USA: Academic, 2013.

[27] E. Dahlman, S. Parkvall, and J. Skold, *5G/5G-Advanced: The New Generation Wireless Access Technology*. Amsterdam, The Netherlands: Elsevier, 2023.

[28] K. Alam et al., "A comprehensive tutorial and survey of O-RAN: Exploring slicing-aware architecture, deployment options, use cases, and challenges," *IEEE Commun. Surveys Tuts.*, vol. 28, pp. 1637–1678, 2026.

[29] G. Sebestyen, S. Fujikawa, N. Galassi, and A. Chuchra, *Low Earth Orbit Satellite Design*. Cham, Switzerland: Springer, 2018.

[30] (2023). *EUSPA Secure SATCOM -Market and User Technology Report*. [Online]. Available: https://www.euspa.europa.eu/sites/default/files/external/publications/euspasecuresatcomreport2023.pdf

[31] N. B. Ruparelia, *Cloud Computing*. Cambridge, MA, USA: MIT Press, 2023.

[32] (2026). *Cloud Infrastructure Map*. [Online]. Available: https://www.cloudinfrastructuremap.com/

[33] (Dec. 2023). *Cloud Computing-Statistics on the Use By Enterprises*. [Online]. Available: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Cloudcomputingstatisticsontheusebyenterprises

[34] (Jul. 2023). *Content Delivery Network Market*. [Online]. Available: https://www.marketsandmarkets.com/Market-Reports/content-delivery-networks-cdn-market-657.html

[35] J. A. Hawkinson and T. J. Bates, *Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)*, document RFC 1930, Mar. 1996.

[36] G. Huston. (2026). *BGP Routing Table Analysis Reports*. [Online]. Available: https://bgp.potaroo.net

[37] (2026). *Internet Number Resource Status Report*. [Online]. Available: https://www.nro.net/about/rirs/statistics/

[38] M. Goyal et al., "Improving convergence speed and scalability in OSPF: A survey," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 2, pp. 443–463, 2nd Quart., 2012.

[39] X. Wu, W. Yang, B. Cui, H. Yu, and R. Sui, "The security of RIP and IS-IS routing protocols: Research on routing attack experiment," in *Proc. 2nd Int. Conf. Mechatronics, IoT Ind. Informat. (ICMIII)*, Jun. 2024, pp. 202–209.

[40] Y. Dabone, T. F. Ouedraogo, and P. J. Kouraogo, "Impact of internet exchange points on ISPs speeds and latency," in *Proc. Int. Symp. Netw., Comput. Commun. (ISNCC)*, Jul. 2022, pp. 1–6.

[41] (2026). *Internet Exchange Map*. [Online]. Available: https://www.internetexchangemap.com/

[42] (2026). *Oneterabitclub*. [Online]. Available: https://github.com/tking/OneTeraBitClub

[43] N. Bargisen. (2023). *Speeding Up the Web: A Comprehensive Guide To Content Delivery Networks and Embedded Caching*. [Online]. Available: https://www.kentik.com/blog/speeding-up-the-web-comprehensive-guide-content-delivery-networks- embedded-caching/

[44] (2026). *Ix.Br*. [Online]. Available: https://ix.br/

[45] (2026). *DE-CIX*. [Online]. Available: https://www.de-cix.net/

[46] (2026). *PITCHILE*. [Online]. Available: https://www.pitchile.cl/wp/

[47] (2026). *AMSIX*. [Online]. Available: https://www.ams-ix.net/ams

[48] (2026). *LINX*. [Online]. Available: https://www.linx.net/

[49] T. Tofoni, F. Luciani, and A. Prado, *BGP From Theory to Practice*. L'Aquila, Italy: Reiss Romoli, Nov. 2023.

[50] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "BGP hijacking classification," in *Proc. Netw. Traffic Meas. Anal. Conf. (TMA)*, Jun. 2019, pp. 25–32.

[51] K. Sriram, D. Montgomery, D. R. McPherson, E. Osterweil, and B. Dickson, *Problem Definition and Classification of BGP Route Leaks*, document RFC 7908, Jun. 2016.

[52] C. Villamizar, R. Chandra, and D. R. Govindan, *BGP Route Flap Damping*, document RFC 2439, Nov. 1998.

[53] E. Chen, J. Scudder, P. Mohapatra, and K. Patel, *Revised Error Handling for BGP UPDATE Messages*, document RFC 7606, Aug. 2015.

[54] P. Spadaccino, S. Bruzzese, F. Cuomo, and F. Luciani, "Analysis and emulation of BGP hijacking events," in *Proc. NOMS - IEEE/IFIP Netw. Oper. Manage. Symp.*, May 2023, pp. 1–4.

[55] A. Siddiqui. (Apr. 2020). *Not Just Another BGP Hijack*. [Online]. Available: https://manrs.org/2020/04/not-just-another-bgp-hijack/

[56] D. Madory. (2026). *Large European Routing Leak Sends Traffic Through China Telecom*. [Online]. Available: https://manrs.org/2019/06/large-european-routing-leak-sends-traffic-through-china-telecom/

[57] (2025). *Beyond Their Intended Scope: BGP Goof-UPX*. [Online]. Available: https://www.kentik.com/blog/beyond-their-intended-scope-bgp-goof-upx/

[58] P. Mittal, W. Lee, H. Shulman, Y. Shavitt, and P. Winter, "SICO: Surgical interception attacks by manipulating BGP communities," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2019, pp. 1093–1110.

[59] (2026). *BGP Handling Bug Causes Widespread Internet Routing Instability*. [Online]. Available: https://blog.benjojo.co.uk/post/bgp-attr-40-junos-arista-session-reset-incident

[60] E. N. Nemmi, F. Sassi, M. La Morgia, C. Testart, A. Mei, and A. Dainotti, "The parallel lives of autonomous systems: ASN allocations vs. BGP," in *Proc. 21st ACM Internet Meas. Conf.*, Nov. 2021, pp. 593–611.

[61] *Routing Security: BGP Incidents, Mitigation Techniques and Policy Actions*, OECD, Paris, France, 2022, doi: 10.1787/40be69c8-en.

[62] A. Koukounas, E. Vytogianni, and M. Dekker. (May 2019). *7 Steps to Shore Up BGP*. [Online]. Available: https://www.enisa.europa.eu/publications/7-steps-to-shore-up-bgp

[63] E. Jaw, M. Müller, C. Hesselman, and L. J. M. Nieuwenhuis, "Reproducibility study and assessment of the evolution of serial BGP hijacking events," *IEEE Trans. Netw. Service Manage.*, vol. 23, pp. 3602–3621, 2026.

[64] Y. Zhang and M. Pourzandi, "Studying impacts of prefix interception attack by exploring BGP AS-PATH prepending," in *Proc. IEEE 32nd Int. Conf. Distrib. Comput. Syst.*, Jun. 2012, pp. 667–677.

[65] J. Mitchell, *Autonomous System (AS) Reservation for Private Use*, document RFC 6996, Jul. 2013.

[66] Z. Tsiatsikas, G. Karopoulos, and G. Kambourakis, "The effects of the Russo-Ukrainian war on network infrastructures through the lens of BGP," in *Proc. Comput. Secur. ESORICS Int. Workshops*, 2023, pp. 81–96.

[67] Y. Song, A. Venkataramani, and L. Gao, "Identifying and addressing reachability and policy attacks in 'secure' BGP," *IEEE/ACM Trans. Netw.*, vol. 24, no. 5, pp. 2969–2982, Oct. 2016.

[68] M. Jonker, A. Pras, A. Dainotti, and A. Sperotto, "A first joint look at DoS attacks and BGP blackholing in the wild," in *Proc. Internet Meas. Conf.*, Oct. 2018, pp. 457–463.

[69] P. Marcos, L. Prehn, L. Leal, A. Dainotti, A. Feldmann, and M. Barcellos, "AS-path prepending: There is no Rose without a thorn," in *Proc. ACM Internet Meas. Conf.*, Oct. 2020, pp. 506–520.

[70] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol (S-BGP)," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 4, pp. 582–592, Apr. 2000.

[71] R. White, "Securing BGP through secure origin BGP," *Internet Protocol J.*, vol. 6, pp. 47–53, Sep. 2003.

[72] P. C. V. Oorschot, T. Wan, and E. Kranakis, "On interdomain routing security and pretty secure BGP (psBGP)," *ACM Trans. Inf. Syst. Secur.*, vol. 10, no. 3, p. 11, Jul. 2007.

[73] C. Testart, "Reviewing a historical internet vulnerability: Why isn't BGP more secure and what can we do about it?" in *Proc. 46th Res. Conf. Commun., Inf. Internet Policy (TPRC)*, Aug. 2018, pp. 1–36, doi: 10.2139/ssrn.3141666.

[74] B. Kuerbis and M. Mueller, "Internet routing registries, data governance, and security," *J. Cyber Policy*, vol. 2, no. 1, pp. 64–81, Jan. 2017.

[75] R. Bush and R. Austein, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*, document RFC 6810, Jan. 2013.

[76] Y. Gilad, S. Goldberg, K. Sriram, J. Snijders, and B. Maddison, *The Use of MaxLength in the Resource Public Key Infrastructure (RPKI)*, document RFC 9319, Oct. 2022.

[77] B. Du et al., "IRRegularities in the internet routing registry," in *Proc. ACM Internet Meas. Conf.*, Oct. 2023, pp. 104–110.

[78] T. Chung et al., "RPKI is coming of age: A longitudinal study of RPKI deployment and invalid route origins," in *Proc. ACM Internet Meas. Conf. (IMC)*, New York, NY, USA, 2019, pp. 406–419.

[79] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Stalloris: RPKI downgrade attack," in *Proc. 31st USENIX Secur. Symp. (USENIX Secur.)*, 2022, pp. 4455–4471.

[80] J. Hawkinson and T. Bates, *Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)*, document RFC 8205, Sep. 2017.

[81] W. Almuhtadi, W. Fenwick, L. Henley-Vachon, and P. Mitchell, "Assessing network infrastructure-as-code security using open source software analysis techniques applied to BGP/BIRD," in *Proc. IEEE Int. Conf. Consum. Electron. (ICCE)*, Jan. 2022, pp. 1–6.

[82] N. Höger, C. Hesselman, M. Müller, and S. Kastanakis, "XBGPsec: Enhancing BGP security with transitive signatures and external validation," in *Proc. 21st Int. Conf. Netw. Service Manage. (CNSM)*, Oct. 2025, pp. 1–5.

[83] J. Oesterle, H. Kinkelin, and F. Rezabek, "Challenges with BGPSec," *Network*, vol. 5, 2021.

[84] N. Umeda, T. Kimura, and N. Yanai, "The juice is worth the squeeze: Analysis of autonomous system provider authorization in partial deployment," *IEEE Open J. Commun. Soc.*, vol. 4, pp. 269–306, 2023.

[85] N. Rodday, G. D. Rodosek, A. Pras, and R. Van Rijswijk- Deij, "Exploring the benefit of path plausibility algorithms in BGP," in *Proc. NOMS - IEEE Netw. Oper. Manage. Symp.*, May 2024, pp. 1–10.

[86] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS verification based on autonomous system provider authorization (ASPA) objects," Internet Eng. Task Force (IETF), IETF Admin. LLC, Wilmington, DE, USA, Tech. Rep., Jan. 2025.

[87] C. Huitema, *Evaluation of a Sample of RFCs Produced in* 2018, document RFC 8963, Jan. 2021.

[88] T. Bruijnzeels. (Dec. 2025). *ASPA in the RPKI Dashboard: A New Layer of Routing Security*. [Online]. Available: https://labs.ripe.net/author/timbruijnzeels/aspa-in-the-rpki-dashboard-a-new-layer-of-routing-security/

[89] (Dec. 2025). *Cloudflare Radar: RPKI Aspa Deployment*. [Online]. Available: https://radar.cloudflare.com/routing

[90] (Sep. 2024). *Roadmap to Enhancing Internet Routing Security*. [Online]. Available: https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/09/Roadmap-to-Enhancing-Internet-Routing-Security.pdf

[91] K. Sriram and D. Montgomery. (Feb. 2025). *Border Gateway Protocol Security and Resilience*. [Online]. Available: https://csrc.nist.gov/pubs/sp/800/189/r1/ipd

[92] (Sep. 2018). *A Guide Border Gateway Protocol (BGP) Best Practices*. [Online]. Available: https://www.nsa.gov/Press-Room/Digital-Media-Center/Document-Gallery/igphoto/2002158110/

[93] (Jan. 2026). *Proposal for a Regulation for the Digital Networks Act (DNA)*. [Online]. Available: https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-digital-networks-act-dna

[94] (Dec. 2022). *Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 (NIS 2 Directive)*. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2555/oj/eng

[95] (Dec. 2022). *Directive (EU) 2022/2557 of the European Parliament and of the Council of 14 December 2022 (CER Directive)*. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2557/oj/eng

[96] (Oct. 2024). *Commission Implementing Regulation (EU) 2024/2690 of 17 October 2024 Laying Down Rules for the Application of Directive (EU) 2022/2555*. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R2690

[97] (2024). *Cybersecurity Strategy 2024–2030–Cyber-Conscious Estonia*. [Online]. Available: https://www.enisa.europa.eu/sites/default/files/ncss-map/strategies/reports/EENCSS2024en.pdf

[98] (2022). *Implementation Plan-National Cybersecurity Strategy 2022-2026*. [Online]. Available: https://www.acn.gov.it/portale/documents/20119/531899/ACNENImplementationPlan.pdf/

[99] (2022). *Action Plan-Netherlands Cybersecurity Strategy 2022-2028*. [Online]. Available: https://www.enisa.europa.eu/sites/default/files/ncss-map/strategies/action-plans/NLACTIONPLAN2022en.pdf

[100] (2019). *List of Measures-National Cyber Security Strategy for Norway*. [Online]. Available: https://www.regjeringen.no/contentassets/c57a0733652f47688294934ffd93fc53/list-of-measures–national-cyber-security-strategy-for-norway.pdf

[101] (2025). *Action Plan National Cybersecurity Strategy 2025–2029*. [Online]. Available: https://www.enisa.europa.eu/sites/default/files/ncss-map/strategies/action-plans/SEACTIONPLAN2025se.pdf

[102] (Oct. 2014). *BGP Configuration Best Practices*. [Online]. Available: https://cyber.gouv.fr/en/publications/bgp-configuration-best-practices

[103] (Jan. 2021). *Technical Report: Responsible Use of the Border Gateway Protocol (BGP) for ISP Interworking*. [Online]. Available: https://www.ncsc.gov.uk/report/responsible-use-of-bgp-for-isp-interworking

[104] (2025). *Australian Government Gateway Security Standard-Consultation Paper*. [Online]. Available: https://www.homeaffairs.gov.au/cyber-security-subsite/files/australian-government-gateway-security-standard-consultation-paper.pdf

[105] (2026). *Mutually Agreed Norms for Routing Security (MANRS)*. [Online]. Available: https://manrs.org/

[106] (2026). *MANRS for IXPs*. [Online]. Available: https://manrs.org/ixps/

[107] (2026). *Euro-IX*. [Online]. Available: https://www.euro-ix.net/en/about-us/

[108] (Mar. 2023). *National Cybersecurity Strategy* 2023. [Online]. Available: https://bidenwhitehouse.archives.gov/oncd/national-cybersecurity-strategy/

[109] (Mar. 2026). *National Cybersecurity Strategy* 2026. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2026/03/President-Trumps-Cyber-Strategy-for-America.pdf

[110] O. Borchert, K. Lee, K. Sriram, D. Montgomery, P. Gleichmann, and M. Adalier, "BGP secure routing extension (BGP-SRX)," NIST, Gaithersburg, MD, USA, Tech. Rep. NIST Tech. Note 2060, Jul. 2021.

[111] (2026). *National Cybersecurity Strategies*. [Online]. Available: https://www.enisa.europa.eu/topics/national-cyber-security-strategies/ncss-map/national-cyber-security-strategies-interactive-map

[112] (2026). *PeeringDB*. [Online]. Available: https://www.peeringdb.com/

[113] J. M. Evang and I. Livadariu, "How large is the gap? Exploring MANRS and ISO27001 security management," in *Proc. Int. Conf. Softw., Telecommun. Comput. Netw. (SoftCOM)*, Sep. 2023, pp. 1–6.

[114] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and K. Claffy, "Mind your MANRS: Measuring the MANRS ecosystem," in *Proc. 22nd ACM Internet Meas. Conf.*, Oct. 2022, pp. 716–729.

[115] A. Robachevsky. (2026). *Reflecting on 'routing Security: Fostering Continuous Improvement' At Apricot2025*. [Online]. Available: https://manrs.org/2025/03/apricot2025-panel/

[116] N. Hart, C. Rotsos, V. Giotsas, N. Race, and D. Hutchison, "λBGP: Rethinking BGP programmability," in *Proc. NOMS - IEEE/IFIP Netw. Oper. Manage. Symp.*, Apr. 2020, pp. 1–9.

[117] S. Servillo, P. Spadaccino, S. Konstantaras, F. Luciani, and F. Cuomo, "Exploring the blind spot of internet exchange point route servers," in *Proc. NOMS IEEE Netw. Oper. Manage. Symp.*, 2026, pp. 1–5.

[118] T. Wirtgen and O. Bonaventure, "A first step towards checking BGP routes in the dataplane," in *Proc. ACM SIGCOMM Workshop Future Internet Routing Addressing*, Aug. 2022, pp. 50–57.

[119] E. Deprez, "Making BGP aware of the data plane," Master's thesis, École polytechnique de Louvain, Université catholique de Louvain, Louvain-la-Neuve, Belgique, 2024.

[120] H. Birge-Lee et al., "Creating a secure underlay for the internet," in *Proc. 31st USENIX Secur. Symp. (USENIX Secur.)*, Aug. 2022, pp. 2601–2618.

[121] X. Wang et al., "Secure inter-domain routing and forwarding via verifiable forwarding commitments," 2023, *arXiv:2309.13271*.

[122] (2026). *NTT Peer Locking*. [Online]. Available: https://instituut.net/~job/peerlockmanual.pdf

[123] J. Snijders, M. Stucchi, and M. Aelmans, "RPKI autonomous systems cones: A profile to define sets of autonomous systems numbers to facilitate BGP filtering," Internet Eng. Task Force (IETF), IETF Admin. LLC, Wilmington, DE, USA, Tech. Rep., Apr. 2020.

[124] R. Morillo, J. Furuness, C. Morris, J. Breslin, A. Herzberg, and B. Wang, "ROV++: Improved deployable defense against BGP hijacking," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2021, pp. 1–18.

[125] Y. Wu et al., "Comp-RPKI: A decentralized protocol for full route origin validation," in *Proc. 9th Int. Conf. Big Data Comput. Commun. (BigCom)*, Aug. 2023, pp. 301–308.

[126] P.-W. Tsai, A. C. Risdianto, M. H. Choi, S. K. Permal, and T. C. Ling, "SD-BROV: An enhanced BGP hijacking protection with route validation in software-defined eXchange," *Future Internet*, vol. 13, no. 7, p. 171, Jun. 2021.

[127] C. Morris, A. Herzberg, B. Wang, and S. Secondo, "BGP-iSec: Improved security of internet routing against post-ROV attacks," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2024, pp. 1–12.

[128] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "Jumpstarting BGP security with path-end validation," in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 342–355.

[129] M. Saad, A. Anwar, A. Ahmad, H. Alasmary, M. Yuksel, and D. Mohaisen, "RouteChain: Towards blockchain-based secure and efficient BGP routing," *Comput. Netw.*, vol. 217, Nov. 2022, Art. no. 109362.

[130] N. Höger, N. Rodday, and G. Dreo Rodosek, "Mitigating BGP route leaks with attributes and communities: A stopgap solution for path plausibility," *Int. J. Netw. Manage.*, vol. 35, no. 2, Mar. 2025.

[131] T. McDaniel, J. M. Smith, and M. Schuchard, "Flexsealing BGP against route leaks: Peerlock active measurement and analysis," 2020, *arXiv:2006.06576*.

[132] Q. Li, J. Liu, Y.-C. Hu, M. Xu, and J. Wu, "BGP with BGPsec: Attacks and countermeasures," *IEEE Netw.*, vol. 33, no. 4, pp. 194–200, Jul. 2019.

[133] (2026). *RPKI Monitor*. [Online]. Available: https://rpki-monitor.antd.nist.gov

[134] J. Snijders. (Jan. 2025). *Year to Year Growth of the Distributed RPKI Database*. [Online]. Available: https://manrs.org/2025/01/rpki-growth-2024/

[135] W. Li et al., "RoVista: Measuring and analyzing the route origin validation (ROV) in RPKI," in *Proc. ACM Internet Meas. Conf.*, Oct. 2023, pp. 73–88.

[136] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and K. Sriram, *Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages*, document RFC 9234, May 2022.

[137] K. Sriram and A. Azimov, "Methods for detection and mitigation of BGP route leaks," Internet Eng. Task Force (IETF), IETF Admin. LLC, Wilmington, DE, USA, Tech. Rep., Feb. 2025.

[138] J. Israr, Y. Gahi, M. Guennoun, and H. T. Mouftah, "Security analysis of C-BGP: A light alternative to S-BGP," in *Proc. IEEE Can. Conf. Electr. Comput. Eng. (CCECE)*, May 2016, pp. 1–6.

[139] Y. Xiang, X. Shi, J. Wu, Z. Wang, and X. Yin, "Sign what you really care about–secure BGP AS-paths efficiently," *Comput. Netw.*, vol. 57, no. 10, pp. 2250–2265, Jul. 2013.

[140] G. Sharma and L. Ragha, "Hierarchical origin and path verification for securing inter-domain routing protocol," in *Proc. 5th IEEE Int. Conf. Adv. Telecommun. Syst. Netw. (ANTS)*, Dec. 2011, pp. 1–6.

[141] O. Junjie, N. Yanai, T. Takemura, M. Okada, S. Okamura, and J. Paul Cruz, "APVAS: Reducing memory size of AS_PATH validation by using aggregate signatures," 2020, *arXiv:2008.13346*.

[142] T. Takemura, N. Yanai, N. Umeda, M. Okada, S. Okamura, and J. P. Cruz, "APVAS+: A practical extension of BGPsec with low memory requirement," in *Proc. ICC - IEEE Int. Conf. Commun.*, Jun. 2021, pp. 1–7.

[143] J. Frieß, D. Mirdita, H. Schulmann, and M. Waidner, "Byzantine-secure relying party for resilient RPKI," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Dec. 2024, pp. 49–63.

[144] K. Sriram, J. Snijders, and D. Montgomery, "Signed prefix list (SPL) based route origin verification and operational considerations," Internet Eng. Task Force (IETF), IETF Admin. LLC, Wilmington, DE, USA, Tech. Rep., Jun. 2025.

[145] K. Xu, S. Jiang, Y. B. Guo, and X. Wang, "BGP AS verification based on route path authorizations (RPA) objects," Internet Eng. Task Force (IETF), IETF Admin. LLC, Wilmington, DE, USA, Tech. Rep., Nov. 2025.

[146] G. He, W. Su, S. Gao, J. Yue, and S. K. Das, "ROAchain: Securing route origin authorization with blockchain for inter-domain routing," *IEEE Trans. Netw. Service Manage.*, vol. 18, no. 2, pp. 1690–1705, Jun. 2021.

[147] A. Hari and T. V. Lakshman, "The internet blockchain: A distributed, tamper-resistant transaction framework for the internet," in *Proc. 15th ACM Workshop Hot Topics Netw.*, Nov. 2016, pp. 204–210.

[148] Y. Gilad, T. Hlavacek, A. Herzberg, M. Schapira, and H. Shulman, "Perfect is the enemy of good: Setting realistic goals for BGP security," in *Proc. 17th ACM Workshop Hot Topics Netw.*, Nov. 2018, pp. 57–63.

[149] L. Chuat et al., *The Complete Guide to SCION-From Design Principles to Formal Verification*. Cham, Switzerland: Springer, 2022.

[150] (2026). *BGPsec-Could You Run It If You Wanted to?*. [Online]. Available: https://www.sidnlabs.nl/en/news-and-blogs/bgpsec-could-you-run-it-if-you-wanted-to

[151] L.-C. Schulz, R. Wehner, and D. Hausheer, "Cryptographic path validation for SCION in P4," in *Proc. 6th Eur. P4 Workshop*, Dec. 2023, pp. 17–23.

[152] L. Herschbach, D. Rossi, and S. Keshvadi, "Path dynamics in a deployed path-aware network: A measurement study of SCIONLab," 2025, *arXiv:2509.04695*.

[153] F. Wirz et al., "Scaling SCIERA: A journey through the deployment of a next-generation network," in *Proc. ACM SIGCOMM Conf.*, Sep. 2025, pp. 720–741.

[154] J. Li, J. Cao, Z. Meng, R. Xie, and M. Xu, "RoLL: Real-time and accurate route leak location with AS triplet features," in *Proc. IEEE Int. Conf. Commun.*, May 2023, pp. 5240–5246.

[155] M. Zeng, D. Li, P. Zhang, K. Xie, and X. Huang, "Federated route leak detection in inter-domain routing with privacy guarantee," *ACM Trans. Internet Technol.*, vol. 23, pp. 1–22, Feb. 2023.

[156] T. Holterbach, T. Alfroy, A. Phokeer, A. Dainotti, and C. Pelsser, "A system to detect forged-origin BGP hijacks," in *Proc. 21st USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, Apr. 2024, pp. 1751–1770.
[157] M. C. Z. Mingwei. (2023). *Cloudflare Radar's New BGP Origin Hijack Detection System*. [Online]. Available: https://blog.cloudflare.com/bgp-hijack-detection/
[158] A. Murty. (2026). *Introducing BGP Monitoring From Kentik*. [Online]. Available: https://www.kentik.com/blog/bgp-monitoring-from-kentik/
[159] H. Birge-Lee, M. Apostolaki, and J. Rexford, "Global BGP attacks that evade route monitoring," in *Passive and Active Measurement*. Cham, Switzerland: Springer, 2025.
[160] Y. Sun, A. Edmundson, N. Feamster, M. Chiang, and P. Mittal, "Counter-RAPTOR: Safeguarding tor against active routing attacks," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2017, pp. 977–992.
[161] P. Sermpezis et al., "ARTEMIS: Neutralizing BGP hijacking within a minute," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2471–2486, Dec. 2018.
[162] M. Scazzariello, A. Prado, and T. Caiazzi. (2026). *Verify MANRS Compliance Automatically With Rose-T*. [Online]. Available: https://manrs.org/2024/03/verify-manrs-compliance-automatically-with-rose-t/
[163] J. Huang, Z. Liu, Z. Zheng, X. Wei, M. Li, and M. Jia, "Research and development of intelligent protection capabilities against internet routing hijacking and leakage," in *Proc. Int. Conf. Artif. Intell., Inf. Process. Cloud Comput. (AIIPCC)*, Aug. 2022, pp. 50–54.
[164] S. Servillo, P. Spadaccino, F. Luciani, and F. Cuomo, "Estimating autonomous system risk levels by analyzing IXP route server RIB," *Comput. Commun.*, vol. 237, May 2025, Art. no. 108154.
[165] S. Servillo, A. P. Pompeo, P. Spadaccino, and F. Cuomo, "Risk level as a service: Enhancing BGP security through IXP-based risk assessment," in *Proc. IFIP Netw. Conf. (IFIP Netw.)*, May 2026, pp. 1–9.
[166] A. Prem Sankar, P. Poornachandran, A. Ashok, R. Manu, and P. Hrudya, "B-secure: A dynamic reputation system for identifying anomalous BGP paths," in *Proc. 5th Int. Conf. Frontiers Intell. Comput., Theory Appl.*, vol. 1, 2017, pp. 767–775.
[167] M. Scazzariello, L. Ariemma, and T. Caiazzi, "Kathará: A lightweight network emulation system," in *Proc. NOMS - IEEE/IFIP Netw. Oper. Manage. Symp.*, Apr. 2020, pp. 1–2.
[168] (Sep. 2024). *ROSE-T: Routing Security Tool*. [Online]. Available: https://github.com/OpenAccess-Italia/roset

Stefano Servillo (Graduate Student Member, IEEE) received the B.S. and M.S. degrees in engineering in computer science from the University of Rome "La Sapienza," Italy, in 2021 and 2023, respectively, where he is currently pursuing the Ph.D. degree with the Department of Information Engineering, Electronic and Telecommunications. He is collaborating with Namex, IXP of Rome, and AMS-IX, Amsterdam IXP. He has been a Visiting Ph.D. Student with the Design and Analysis of Communication Systems Group, University of Twente, Enschede. His research interests mainly focus on routing security, border gateway protocol, and internet exchange points.

Pietro Spadaccino received the Ph.D. degree from the University of Rome "La Sapienza," Italy, in 2024. He is currently a full-time Researcher (RTD-A) with the Department of Information Engineering, Electronic and Telecommunications. Since 2021, he has been an Assistant Professor in several courses for the master's degrees in cybersecurity, computer science engineering, AI and robotics, and data science, dealing with topics such as network configuration on Linux, routing, and network security. He was a Visiting Researcher with the Institute of Wireless Internet of Things, Northeastern University of Boston, and the Department of Engineering, University of Perugia. His research interests mainly focus on network protocols, security, edge computing, and the Internet of Things. He received the Best Paper Award in ALGOCLOUD 2025.

Marco Centenaro (Member, IEEE) received the Ph.D. degree in telecommunication engineering from the University of Padua, Italy, in 2018. His research focuses on advanced connectivity technologies (e.g., 5G and beyond and the IoT) and cybersecurity (in particular, network security and critical infrastructure resilience). With a solid experience across the academic community (University of Padua, the University of Aalborg, and Fondazione Bruno Kessler) and the telecommunications industry (Nokia Bell Labs and Athonet, a HPE acquisition), he contributed to this article as a Civil Servant with the National Cybersecurity Agency of Italy. As of the date of publication, he is an Associate Professor with the University of Padua. He was a recipient of the Ph.D. Award from Italian Association of Telecommunications and Information Technologies (Gruppo Telecomunicazioni e Tecnologie dell'Informazione, GTTI) in 2018, in recognition of the best Ph.D. theses defended at Italian University in the areas of communications technologies; and the Best Paper Award at IEEE PIMRC 2018.

Antonio Prado received the Ph.D. degree in computer science, with a research focus on internet routing security. His doctoral work focused on tools and methodologies for the automated verification of BGP configurations. He is the co-author of *BGP from Theory to Practice*, a book aimed at both technical and academic audiences, and the outreach volume *Parlami di TEch*, designed to make technology accessible to a broader readership. His latest book, *Crescere con l'Intelligenza Artificiale* (2026), helps adults understand how AI is reshaping children's everyday lives. He also holds a second-level Master's degree in Innovation and Digital Transformation in Public Administration, with a specialization in the use of Artificial Intelligence in public-sector processes and services, addressing both opportunities and ethical-organizational implications. Since 2022, he has been a member of the RIPE Programme Committee, contributing to the review and selection of scientific presentations for European community conferences. He has delivered seminars and lectures on Internet and Network Architectures for the Academia, and has taught DNS, Internet Governance, and BGP at the Namex School of Advanced Networking, helping to train the next generation of network professionals. At the time of writing, he is also an External Researcher at the University of Padova, Italy.

Massimiliano Rossi received the M.Sc. degree in automation and control engineering from the University of Rome "La Sapienza." He is currently with the Research Projects Division, National Cybersecurity Agency of Italy (ACN). He is a member of the IT Security Commission of the Order of Engineers of the Province of Rome. Previously, he was a Cyber Risk Manager and a member of the Acceptance Committee for the National Public Administration Network (SPC), CERT-PA, Agency for Digital Italy (AgID). He was also the National Contact Point for the Trans-European Services for Telematics between Administrations (TESTA) Project, a data communication network service operated by European Commission. He has over ten years' experience in sensor networks and geographical network security, which he gained at the National Institute of Geophysics and Volcanology (INGV).

Flavio Luciani was born in Rome in 1981. He received the degree in computer engineering from Roma Tre University in 2005. Since 2008, he has been part of the Namex Team, the internet exchange point in Rome, first as a member of the technical staff, and has been a Chief Technology Officer since 2020. He is involved in several initiatives within the internet community, and he collaborates with RIPE NCC and European Internet Exchange Association EURO-IX. Previously, he held a role on the Steering Committee of the Internet Society (ISOC) Initiative Mutually Agreed Norms for Routing Security (MANRS). Through workshops, courses, and in-depth articles, he promotes greater awareness of routing security issues.

Monica Scannapieco received the Ph.D. degree in computer science and engineering from the University of Rome "La Sapienza," Italy. She is currently the Head of the Research Projects Division, National Cybersecurity Agency of Italy (ACN). Formerly, she was with Italian National Institute of Statistics for more than 15 years, with leading roles in IT and R&I organizational structures, in charge of innovative projects on data quality, data integration, data security, data privacy, and big data management. She has collaborated with several organizations, including UNECE, World Bank, European Commission, and Italian National Presidency of the Council of Ministries.

Francesca Cuomo (Senior Member, IEEE) is a Full Professor and the Dean of the Faculty of Information Engineering, Informatics, and Statistics, University of Rome "La Sapienza," Italy. She has authored over 208 peer-reviewed papers published in prominent international journals and conferences. Her Google Scholar H-index is 35 with more than 5 000 citations. Her current research interests focus on low power wide area networks and the Internet of Things, 5G and open radio access networks, smart metering and vehicular networks, and multimedia networking. She has been the TPC Co-Chair of several editions of ACM PE-WASUN Workshop, the TPC Co-Chair of ICCCN 2016, the TPC Symposium Chair of IEEE WiMob 2017, the General Co-Chair of the First Workshop on Sustainable Networking Through Machine Learning and Internet of Things (SMILING) in conjunction with IEEE INFOCOM 2019, the Workshop Co-Chair of AmI 2019: European Conference on Ambient Intelligence 2019, and the Track Chair of IEEE WCNC 2025.

Open Access funding provided by 'Sapienza Universita di Roma' within the CRUI CARE Agreement