

ROUTING SECURITY

BGP INCIDENTS, MITIGATION TECHNIQUES AND POLICY ACTIONS

OECD DIGITAL ECONOMY PAPERS

October 2022 **No. 330**

This paper was approved and declassified by written procedure by the Committee on Digital Economy Policy (CDEP) on 22 August 2022 and prepared for publication by the OECD Secretariat.

Note to Delegations:

This document is also available on O.N.E. under the reference code:

DSTI/CDEP/CISP/SDE(2021)4/FINAL

This document, as well as any data and any map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

The statistical data for Israel are supplied by and under the responsibility of the relevant Israeli authorities. The use of such data by the OECD is without prejudice to the status of the Golan Heights, East Jerusalem and Israeli settlements in the West Bank under the terms of international law.

@ OECD 2022

The use of this work, whether digital or print, is governed by the Terms and Conditions to be found at <http://www.oecd.org/termsandconditions>.

Foreword

This report on “Routing Security: BGP incidents, mitigation techniques and policy actions” was prepared jointly by the Working Party on Communication Infrastructure and Services Policy (WPCISP) and the Working Party on Security in the Digital Economy (WPSDE). It considers the challenges related to the digital security of the routing system, the solutions proposed to address some of these challenges and the role of policy makers to foster routing security. This report should be read in conjunction with the accompanying report on the “Security of the Domain Name System (DNS): an introduction for policy makers”.

This report was drafted by Lauren Crean with contributions by Verena Weber, Laurent Bernat and Ghislain de Salins from the OECD Secretariat and by WPCISP and WPSDE delegates. It was prepared under the supervision of Verena Weber and Laurent Bernat. This publication is a contribution to IOR 1.3.1.3.3, “How to analyse and enhance the security of communication networks” of the 2021-2022 Programme of Work and Budget of the CDEP.

The Secretariat wishes to thank external expert, Sara Alamin, as well as several experts from the technical community and industry, for their contributions to the report. These include: Amy Alvarez, AT&T; Stewart Bamford, Principle Architect, Lumen Technologies, Inc.; Einar Bohlin; Chris Boyer, AT&T; Randy Bush, Research Fellow, Internet Initiative Japan, Member of Technical Staff, Arrcus Inc.; Kathryn Condello, Senior Director National Security, Lumen Technologies, Inc.; David Conrad, ICANN; Alberto Dainotti, Georgia Tech; Alain Durand, ICANN; Patrik Fältström, Netnod; Laurent Ferrali, ICANN; Marco Hogewoning; Geoff Huston, APNIC; Anne- Rachel Inne; Olaf Kolkman, Internet Society; Alexander Lyamin, Qrator Labs; Doug Madory, Director of Internet Analysis, Kentik; Jason Olson, AT&T; Elena Plexida, ICANN; Andrei Robachevsky, Internet Society; Nicola Rustignoli, ETH Zürich; Chelsea J. Smethurst, Microsoft; Job Snijders, Principal Engineer at Fastly and OpenBSD developer; Mark Svancarek, Microsoft; Cecilia Testart, Georgia Tech; Martin Thygesen, Cisco Systems Inc.; and Andree Toonk, MySocket.io.

Additionally, we thank the Swedish Post and Telecom Authority (PTS) and the Finnish Transport and Communications Agency Traficom National Cyber Security Centre (Traficom) for providing input on specific routing policy initiatives.

EXECUTIVE SUMMARY

Routing, the process by which data packets are directed across the Internet to their destinations, is central to the operation of the Internet. Networks exchange routing information about destinations and the paths by which they can be reached using the Border Gateway Protocol (BGP). When this system fails, the Internet is unable to forward packets between networks or facilitate communications between them. Given its fundamental importance, ensuring the security and availability of the routing system is critical.

This report aims to analyse three main questions:

- What is the scope and scale of routing incidents?
- Which security techniques have been proposed to address them and how effective are they?
- What is the role of policy makers in securing the routing system?

Internet routing is affected by accidental and intentional disruptions and security breaches. While the Internet is generally resilient, the vulnerability of its routing continues to grow more pressing. For example, Facebook's over five-hour global outage in October 2021 was caused in part by a failure of its routing, resulted in more than 1.2 trillion person-minutes of service unavailability, and has been cited as the largest communications outage in history (Madory, 2021^[1]). **Routing vulnerabilities have been understood for many years, but persist.** This is driven by several challenges, including the complexity of the issue, the interconnected nature of the global Internet requiring collective action to improve security overall, and a lack of incentives among some stakeholders to spend time, money, or resources to implement existing techniques or develop new ones.

Another aspect of the problem lies in understanding the scope and scale of routing incidents, for which **robust data over time is needed.** Only a few measurement efforts exist, all with varying degrees of robustness, availability and period of observation. Similarly, there are limited sources to track the *effectiveness* of security techniques to decrease the incidence of routing events. This data is required for policy makers to gauge how routing incidents are evolving over time and to measure the impact of policy initiatives.

The universe of available routing security techniques each address aspects of routing security but fail to provide a sufficiently comprehensive solution, even when implemented together. Routing security techniques consist principally of filtering incorrect or malicious routing information as it enters and leaves networks, using filters built with various methods. **However, currently, no single technique or combination of techniques will meet the various challenges facing routing security.** Nevertheless, while not a complete solution, network operators should implement the currently deployable techniques and good practices to ensure all available protection against routing incidents.

In light of the measurement challenges of tracking routing incidents, **governments have an important role to fund the collection and publication of time-series data to identify and analyse routing incidents. Funding can support existing, neutral multi-stakeholder initiatives and should be long-term and consistent. Data should be published at no cost, in an easy to read and process format.** Additionally, methods should be developed to **track the implementation and effectiveness of routing security techniques** in the same way. Without information on the scope of the problem and the effectiveness of implemented techniques, policy makers do not have the information necessary to follow an evidence-based approach to policy making.

Furthermore, **policy makers can promote the awareness and deployment of available routing security techniques** among industry by issuing tailored guidance and implementing relevant techniques in government-owned IP addresses and autonomous systems (ASes). This would not only serve to benefit the routing system more generally by putting in place good practices, it would also allow the government to be able to share its experience to help other networks with real-world implementation. **Policy makers could also facilitate information exchange on routing incidents between different stakeholders.**

4 | ROUTING SECURITY: BGP INCIDENTS, MITIGATION TECHNIQUES AND POLICY ACTIONS

This information sharing could take place in existing structures within national, regional or sectoral Computer Emergency Response Teams (CERTs) or Information Sharing and Analysis Centres (ISACs), regulatory bodies, or government agencies. Safeguards should be put in place to protect informants from liability or reprisal. Existing multi-stakeholder initiatives that encourage cross-border information sharing and exchange of good practices should also be recognised as key convening bodies that facilitate international cooperation.

Finally, governments can contribute to enhanced routing security by working with industry and technical experts to **define a common framework to improve routing security**. While many governments currently regard routing security as a subset of cybersecurity, a more specific focus is necessary to make a significant and lasting improvement on the security of the routing system. OECD members, including Finland, Japan, and Sweden, are exploring possible approaches from formalised partnerships, to regulatory monitoring of implemented techniques, to voluntary guidelines, to more defined secondary legislation.

Table of contents

Foreword	2
Routing Security: BGP incidents, mitigation techniques and policy actions	7
Introduction	7
Brief overview of the routing system	8
Routing security	10
Security challenges of the Border Gateway Protocol	10
Scope and scale of routing incidents	12
Current techniques to enhance routing security	20
The mechanisms of routing security	20
Current routing security techniques	23
Overarching challenges to improving routing security	33
Policy discussion	34
Policy actions to support stronger routing security	34
Concluding Remarks	40
Annex A. Technical description of techniques to improve routing security	41
RPKI	41
BGPsec	41
ASPA	42
SCION: Scalability, Control and Isolation on Next-Generation Networks	42
References	44
Tables	
Table 1. Additional examples of BGP events affecting availability	17
Figures	
Figure 1. Simplified diagram of Internet route selection	9
Figure 2. Diagram of Internet routing topology	10
Figure 3. BGP Events: Leaks and hijacks	14
Figure 4. Monthly BGP incidents	15
Figure 5. Mapping of current routing security techniques	20
Figure 6. Example validation flow of incoming BGP announcements, as implemented in the DE-CIX route server	22
Figure 7. Number of validated ROAs, by RIR	25
Figure 8. Rate of ROV filtering in OECD countries (September 2021, January 2022, June 2022)	26
Figure 9. RPKI-ROV History of Unique Prefix-Origin Pairs (IPv4)	27

6 | ROUTING SECURITY: BGP INCIDENTS, MITIGATION TECHNIQUES AND POLICY ACTIONS

Figure 10. IRR and RPKI adoption, OECD countries (June 2022)	28
Figure 11. Simplified route leak under BGPsec	30

Boxes

Box 1. BGP events with a possible impact on availability and confidentiality	18
Box 2. Autonomous System Provider Authorisation (ASPA)	31
Box 3. A Proposal for a new Internet Architecture: SCION, “Scalability, Control and Isolation on Next-Generation Networks”	32
Box 4. PTS Supervision on handling risks related to BGP	36

Routing Security: BGP incidents, mitigation techniques and policy actions

Introduction

The importance of the Internet has perhaps never been more apparent as countries have taken measures to avoid the spread of the COVID-19 pandemic. Digital technologies facilitated fundamental shifts towards teleworking, e-learning, e-commerce, and telemedicine seen in many OECD countries and around the world. As people live more of their lives online and the volume of data communicated over the Internet increases, there are more opportunities for digital interference. The need, therefore, to secure the Internet becomes imperative.

Internet protocols enable the transmission of data and control messages between networks¹. The routing system is one of the Internet's foundational underpinnings and is critical to its performance, availability and operation. The importance of this routing system was highlighted by the complete outage of Facebook's services for over five hours on 4 October 2021 (Madory, 2021^[1]). The event began with a mistake during a routine maintenance check that disconnected Facebook's data centres from the Internet (Janardhan, 2021^[2]). This ultimately caused the withdrawal of dozens of its route announcements and rendered its services unreachable (Madory, 2021^[1]). Facebook's application suite has an estimated 3.5 billion users worldwide, with many people relying on its services to conduct business and communicate (Madory, 2021^[1]). This resulted in more than 1.2 trillion person-minutes of service unavailability and has been cited as the largest incident in history, given its duration and the sheer volume of users across Facebook's application suite (Madory, 2021^[1]). Together with four earlier incidents in the year, Facebook accumulated close to twelve hours of unavailability in the first ten months of 2021². However, Facebook is not alone in reporting significant unavailability; CDN operator Cloudflare had around nine and a half hours of downtime in 2020³ and approximately ten hours of downtime in the first nine months of 2021⁴.

Computing availability is commonly expressed as a percentage: the industry standard benchmark for reliable service is "five nines," meaning 99.999% availability, measured over the course of a year (Cisco Certified Expert, 2021^[3]). There are 525 960 minutes in a year, and "five nines" of uptime means maintaining a service that is available at least 99.999%, or 525 955, of those minutes, or, conversely, unavailable five minutes or less, each year. Critical infrastructure is sometimes held to six nines. In this context, the "two nines" of availability that Facebook and Cloudflare are reporting (99.84% and 99.86% respectively) in 2021 is noteworthy.⁵ Cloudflare attributed a typical outage to "a router on our global backbone [that] announced bad routes and caused the network to not be available" (Cloudflare, 2020^[4]). Failures of this sort underscore the importance – and demonstrate the fragility – of one of the foundations of the Internet. The fragility of this important routing system has increased in profile and priority within the United States, for example, as shown by the Federal Communications Commission's (FCC) public inquiry on secure Internet routing, namely on the "vulnerabilities threatening the security and integrity of the Border Gateway Protocol (BGP), which is central to the Internet's global routing system" (Federal Register, 2022^[5]).

This report examines the current challenges and trends in the security⁶ of the routing system, from an economic and social policy point of view, rather than a purely technical perspective. A digital security incident is an event that can disrupt the availability, integrity or confidentiality (“AIC triad”) of the hardware, software, networks or data that support these activities. Digital security incidents can be unintentional (e.g. a human error) or intentional “attacks” (e.g. caused by malicious actors).

The goal of the report is to identify and highlight the digital security risk related to the routing system. An understanding and appreciation of this risk is essential for policy makers as they consider policies or measures aimed at improving the security of communication networks. The report endeavours to bring existing approaches and techniques to these problems to the attention of policy makers, so that support can be extended to the essential work that is being done in this area.

Brief overview of the routing system

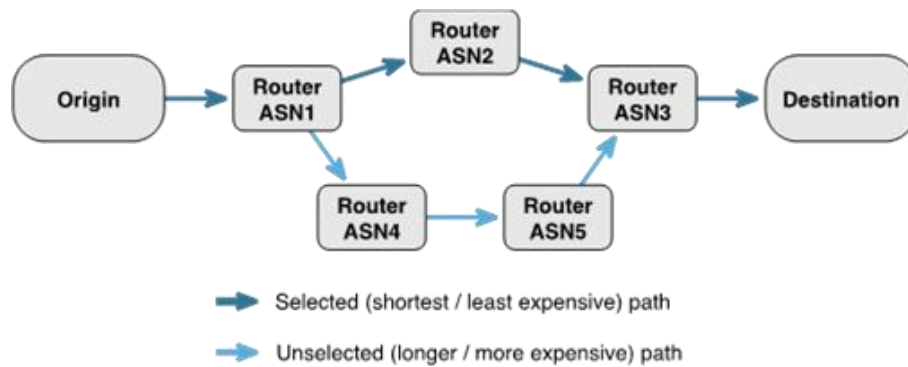
Routing is the process of selecting the optimal path for data packets to take across the Internet to reach their end destination. Simply put, it tells data packets how to get to their intended endpoint. Important to the functioning of the routing system, and used throughout this report, are Internet Protocol (IP) addresses, IP prefixes, Autonomous Systems (ASes) and Autonomous System Numbers (ASNs). IP addresses identify network endpoints, with each endpoint assigned a numeric value in the Internet Protocol (Postel, 1980^[6]). A contiguous group of IP addresses (individual network endpoints) sharing a common prefix (representing the network number) is referred to as an address prefix, or an “IP prefix” (Fuller and Li, 2006^[7]). This form of aggregation is fundamental to the scalability and stability of the Internet’s routing between Autonomous Systems. Since IP address prefixes cover a larger block of addresses that share a common network address prefix, the number of prefixes is significantly less than the number of individual addresses. This allows routers to be able to process and store this routing information, and eliminates the need to announce reachability to smaller blocks or even individual addresses, as such information is kept internal to the originating network.

Prefixes are in turn grouped into Autonomous Systems (ASes). Put simply, an Autonomous System is the network responsible for sending and receiving traffic on behalf of a group of IP prefixes and can independently define the policies that govern how its traffic is routed (OECD, 2013^[8])⁷. In order to exchange routing policy information with other network operators and control routing in their networks, network operators (e.g. Internet Service Providers (ISPs)) uniquely identify themselves using Autonomous System Numbers (ASNs) that, like IP addresses, are procured from Regional Internet Registries (RIRs).

Internet routing is a distributed, decentralised process. Each node in the network makes routing decisions independently, based upon its own understanding of the Internet’s topology. Each node calculates a “graph” of the Internet connection topology based upon routing data communicated to it by adjacent routers, and forwards packets toward their destination based upon its understanding of the shortest or least-expensive path between itself and the destination.

In the simplified example provided in Figure 1, the router associated with Autonomous System 1 (ASN1) makes a decision to forward packets toward the destination via ASN2, rather than ASN4, because the path through ASN2 consists of ASN2 – ASN3 – Destination, or two additional “hops.” This path is shorter than the path through ASN4, which consists of ASN4 – ASN5 – ASN3 – Destination, or three additional “hops.” These paths are called “AS paths” and are the fundamental building blocks of the communication and calculation of Internet routes.

Figure 1. Simplified diagram of Internet route selection



Source: Packet Clearing House.

Upon receiving a packet bound for the destination, the router associated with ASN2 similarly evaluates its options: it can forward the packet to ASN3, which has a “distance” of only one hop, or back to ASN1, which has two routes available, one at a distance of three hops (back via ASN2, which would constitute a problematic “routing loop”) or four hops (via ASNs 4, 5, and 3). ASN2’s choice is clear: forward onward to ASN3, the least-cost path. Upon receiving the packet, ASN3 similarly evaluates its options: it can pass the packet directly to its destination (zero additional hops) or back through ASN2 (two hops and a loop) or ASN5 (two hops and a loop).

Each router’s internal model of the Internet’s topology is built by listening to, and storing, routing announcements advertised to it by its neighbours, but each router makes its routing decisions independently and statelessly, each in the moment, without reference to any external direction or communication in that moment. Routing decisions are made principally based on AS path length, but typically include other rules as well, such as loop avoidance (by eliminating paths that contain one’s own AS) and financial evaluations of cost, which may take precedence over topological distance.

For the past thirty years, the Border Gateway Protocol (BGP) has been used to exchange routing information between networks (Autonomous Systems) in order to allow each route to construct its internal model of the Internet’s topology and make these “next-hop” decisions (Rekhter, 1991^[9])⁸.

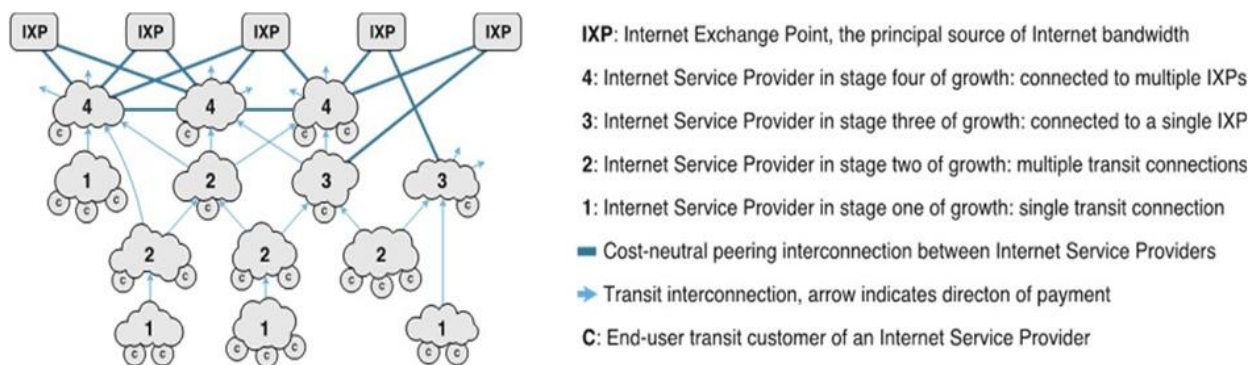
BGP routing information consists of lists of destination prefixes, each accompanied by an AS path through which that prefix can be reached. Each AS provides this routing information to its neighbouring ASes, detailing which prefixes can be reached through its network. To allow the Internet to reflect accurately the current inter-AS reachability and connectivity, ASes exchange this routing information continuously. This allows ASes to announce changes to their network (“announcements” or “withdrawals”) as they occur. This new information may cause routers of adjacent ASes to recalculate their selected (preferred or least-cost) paths, which in turn may cause them to advertise new paths to their neighbours. In this way, routing changes are propagated across the Internet. When a given AS’ selected path does not change as a result of the new routing information it has received, an AS will not pass these on to preserve bandwidth and resources.

Autonomous Systems interconnect in two ways, through peering and transit. These terms usually define a particular routing policy. Transit is usually a paid relationship and assumes that the transit provider accepts all prefixes announced by a customer (and that customer’s customers, if applicable) and provides the customer with a “global routing table”. Peering is usually a “settlement-free” relationship and assumes that only prefixes belonging to the peering networks (and their own customers) are exchanged. The path from any origin to any destination is commonly represented as a “ Λ ” lambda shape, where the origin and destination lie at the lower points, and a peering connection between two ASNs, generally at an Internet exchange point (IXP) or a common transit provider, lies at the top⁹. Internet exchange points are points of

production of Internet bandwidth, and the function of ISPs is to transport that bandwidth from its point of production to its point of consumption, the user's location, whether that be an office building in a fixed location, or a mobile device roaming on their person (Weller and Woodcock, 2013_[10]).

The cost of Internet bandwidth is near zero at its point of production, while the cost of transportation is the product of speed and distance, typically contributing the vast majority of the ultimate cost (Weller and Woodcock, 2013_[10]). In the simplified Figure 2, for example, Internet users ("C") purchase transit service from ISPs ("1," "2," "3," or "4" depending upon their stage of growth), who either purchase transit from each other (light blue arrows) or peer with each other (heavy blue lines) at Internet exchange points. Money travels up this hierarchy following the direction of the arrows, until it reaches an IXP, where no value (and no money) passes¹⁰.

Figure 2. Diagram of Internet routing topology



Source: Packet Clearing House.

Referring back to the AS path of Figure 1, [Origin – AS1 – AS2 – AS3 – Destination], the financial relationships between ASes cannot be directly observed within routing data, but may be inferred or known through other means. For instance, if one knew that there existed a peering relationship between AS2 and AS3, one could infer that the Origin was a transit customer of AS1, that AS1 was a transit customer of AS2, and that the Destination was a transit customer of AS3 (Gao, 2001_[11]; Woodcock and Upadaya, 2005_[12])¹¹.

The decentralised nature of the routing system and its flexibility have been integral to its success and to the growth of the Internet by allowing it to scale through continuous sharing of information between networks (Weller and Woodcock, 2013_[10]). The routing system is essential to the Internet's resilience as it connects users from all over the world quickly and efficiently. However, the importance of the routing system to the overall functioning of the Internet, as well as its fragility, is often underestimated.

Routing security

Security challenges of the Border Gateway Protocol

Despite its strengths, BGP was not conceived with security in mind. When BGP was developed in 1989, the Internet was still the interest of a small homogenous group of experts. The security needs of information exchange at that time were much different from the requirements of today's networks. Trust is implicit in BGP, meaning that networks assume that the routing information provided by other networks is correct by default (Timberg, 2015_[13]). BGP operates in an authority-less mode where each routing element is trusted by all others to operate consistently and with integrity. This mutual trust was an integral and deliberate aspect of the protocol's design.

However, this mutual trust model makes BGP vulnerable from a security standpoint. Its general goal is to propagate information quickly and widely to reflect any network updates, or changes to the inter-network connections. However, it is ill-equipped to deal with hostile parties that deliberately alter or drop routing information in order to affect the forwarding decisions made by other ASes, or with possible mistakes and misconfigurations. This lack of means to determine the accuracy, authenticity and legitimacy of routing information communicated over BGP creates a risk for large-scale consequences across the global Internet caused by routing incidents, which can impact the availability, confidentiality and integrity of communication networks.

The security issues with BGP are widely recognised and discussed within the technical community and BGP's vulnerabilities have been catalogued comprehensively by the Internet Engineering Task Force IETF (Murphy, 2006^[14]).¹² Back in 1996, researchers were already considering BGP's vulnerabilities and proposing potential mitigating actions (Smith and Garcia-Luna-Aceves, 1996^[15]). While there have been other efforts to improve BGP security in the past, this report focuses on the techniques currently being implemented or that are at the forefront of discussions within the technical community¹³. In addition, it focuses on securing the information exchanged between BGP peers conducted to facilitate the routing of Internet traffic, although there are other aspects of BGP security.¹⁴

Routing incidents are often grouped into two categories: BGP leaks (also called "route leaks") and BGP hijacks (also called "BGP prefix hijacks" or "route hijacks").

- **BGP leak:** A BGP leak occurs when an AS announces a route it learned from another AS, which it was not supposed to pass on. An accidental route leak typically consists of a re-advertisement of routes learned via transit to another transit provider, or through peering, in violation of the routing policy of the involved ASes. The more official working definition from the IETF is "the propagation of routing announcement(s) beyond their intended scope. That is, an announcement from an Autonomous System (AS) of a learned BGP route to another AS is in violation of the intended policies of the receiver, the sender, and/or one of the ASes along the preceding AS path" (Sriram et al., 2016^[16]). The IETF outlines different types of leaks, some of which may cause disruptions in the functioning of the Internet through incorrect routing.
- **BGP hijack:** A BGP hijack occurs when one AS announces reachability to an IP prefix, or multiple IP prefixes, which is invalid (i.e., to which the AS does not have a valid route (or "reachability")). Researchers classify different types of BGP hijacks (Cho et al., 2019^[17]). Two commonly referenced types are origin hijacks and path hijacks. An origin hijack occurs when a hijacker announces a prefix that it does not have permission to announce, or announces a more specific prefix than the prefix announced by the authorised AS (Cho et al., 2019^[17]). A path hijack, also called an AS-path manipulation or forged AS path, is characterised by a hijacker announcing a forged AS path with its own ASN inserted into the path, but not as the originating AS. The hijacker sometimes deletes other ASes from the path ("pruning") or advertises a more-specific prefix before advertising it on to its neighbours (Cho et al., 2019^[17])¹⁵. Deleting the other valid ASes in the path results in a shorter path, while reducing the scope of the announced prefix makes it more specific; in both cases, this makes the hijacked route more preferable and more likely to attract traffic. If the forged path is chosen, a BGP hijack could effectively redirect traffic intended for the authorised AS through the attacking network, allowing this traffic to be inspected, discarded or altered in transit. When done with malicious intent, actors can use the hijacked prefixes for malicious purposes, or impersonate the victim ASes (Cho et al., 2019^[17]).

Routing incidents (leaks and hijacks) can be either intentional or accidental, although both types can affect the functioning and quality of service of the Internet. For instance, when done with malicious intent, actors can censor or intercept traffic through both BGP leaks and hijacks. As some aspects of router configuration are often performed manually, routing incidents can easily occur by mistake (e.g. typos), which lends fragility to the system. Even in automated cases, mistakes happen. Most BGP leaks are the result of

accidental misconfiguration (Sriram et al., 2016^[16]). However, it is important to note that many experts interviewed for this report considered that the majority of BGP hijacking incidents are accidental in nature as well. Not only is it difficult to determine the incidence of BGP events in general (both unintentional and intentional); it is especially challenging to classify a BGP event as intentional given the lack of concrete information surrounding these incidents. Nevertheless, efforts such as BGPStream and the Global Routing Intelligence Platform (GRIP) hosted at Georgia Tech aim to track BGP events in an effort to quantify the scope of the issue.

Scope and scale of routing incidents

Security risks from BGP vulnerabilities

BGP events are not isolated and occur much more frequently than the media reports. In 2018, the European cybersecurity agency (ENISA) asked 63 organisations in the European Union’s communication sector to assess the severity of the impact of BGP incidents. Organisations surveyed ranged from domestic ISPs to large international operators. Almost half (44%) of these organisations experienced incidents that had a “major impact” on their networks, i.e. long-lasting outages affecting many subscribers (ENISA, 2019^[18]). A third of respondents experienced incidents that had a medium impact on their network, meaning either that they were long-lasting but affected few subscribers, or that they had a short duration but affected many subscribers. Almost two-thirds of respondents strongly agreed that BGP hijacks are a serious issue that require an urgent solution.

ENISA distinguishes four types of security risk caused by incidents exploiting BGP vulnerabilities: altering Internet traffic content (breach of integrity), eavesdropping Internet traffic content (breach of confidentiality), performing Internet traffic and metadata analysis (breach of confidentiality), and creating connection outages (breach of availability) (ENISA, 2019^[18]). Other experts point to additional risks such as the potential to impersonate other sites by receiving rerouted traffic intended for the original site (breach of integrity) (Huston, 2021^[19]). Another highlighted risk is the possibility to obscure identity by misappropriating IP addresses, or by routing an IP address which has not been allocated, which could lend anonymity to launch an attack at the application layer (breach of integrity) (Huston, 2021^[19]). In addition, a BGP incident may impact more than one security dimension. For instance, a BGP path hijack could breach confidentiality, as traffic may be intercepted and read, integrity, as intercepted traffic may be modified, as well as availability, as the attacker may not reliably forward intercepted traffic to its intended destination.

The vulnerability of the routing system can also compromise the security of applications that run on top of it. As more people rely on digital tools and applications in their daily lives, the underlying foundation upon which these applications depend should not be overlooked. Attackers may leverage the routing system’s vulnerability to launch sophisticated cross-layer attacks on applications, whose goals vary based on the application’s functionality (Sun et al., 2021^[20]). Sun et al. outline three possible cross-layer application attacks: attackers identifying users attempting to communicate anonymously through the Tor Network, attackers intercepting traffic in the domain control verification process to allow adversaries to obtain a digital certificate for a victim domain, and attackers launching routing attacks on the bitcoin consensus process (Sun et al., 2021^[20]).¹⁶ There also have been instances of sophisticated attacks that leverage vulnerabilities of BGP and the DNS in conjunction, such as the 2018 BGP origin hijack of IP prefixes of Amazon Route 53, a DNS service offered by AWS (see BGP events impacting integrity), as well as unintentional incidents that involve both BGP and DNS, such as the October 2021 Facebook outage mentioned previously.

General statistics

It is difficult to identify BGP routing events correctly and to quantify their impact. “Impact” in this report refers to the consequences arising from a BGP incident, which may or may not be felt by users. This


difficulty arises from various reasons. First, different entities have different vantage points of the Internet, making it impossible to obtain a single, globally, consistent “complete” view of the global routing table. Second, sophisticated attackers with knowledge of BGP can remain undetected by localising attacks such that they are not visible to the global routing system. Third, it is difficult to ascertain motive without auxiliary information or resorting to conjecture. Some efforts described later on in the report aim to establish a clear record of IP address holders’ intentions, such as records in the Internet Routing Registries (IRRs) and Route Origin Authorisations (ROAs). These records give clues as to whether a given routing incident was intentional or not.

More tools to monitor and track routing incidents and their impact would be needed in this regard, as would support for research initiatives currently undertaking this work. These include the Global Routing Intelligence Platform (GRIP), which focuses on hijacks, among others. There are also efforts to track the global routing system, with time series data made available to researchers for further analysis, such as RIPE NCC’s Routing Information Service (RIS), Packet Clearing House’s (PCH) daily routing snapshots and the University of Oregon’s Route Views project.¹⁷ In conjunction, encouraging ASes to establish baseline records on how a particular block of IP addresses should be routed or announced in BGP, such as through creating ROAs, would better assist researchers to accurately detect and classify routing incidents.

Experts further underscored that most BGP events are the result of unintentional misconfigurations and cautioned against seeing nefarious motives in every incident. In the same line of thought, focusing on large-scale events obscures both more focused malicious incidents of purposely-constrained impact as well as the many potential incidents that do not occur because of industry’s successes in implementing mitigation practices. Nevertheless, this section presents an overview of available information regarding BGP leaks and hijacks with the aim of better understanding the scope, scale and impact of routing incidents, taking due consideration of the challenges associated with undertaking this task.

In 2021, BGPStream, a free service provided by Cisco as part of its Crosswork Cloud, reported over 3 000 distinct BGP leaks or hijacks (Figure 3) (Cisco Crosswork Cloud, 2022^[21]). BGPStream uses an algorithm to analyse “hundreds of millions of BGP messages every day” and identify incidents occurring around the global Internet, through various route view collectors (Cisco Crosswork Cloud, 2022^[21]). While Figure 3 shows a gradual decline in overall numbers of BGP leaks and hijacks, the short time period of evaluation (2017 – 2021) does not allow for any conclusive statements regarding trends in BGP incidents. A longer time series would be needed in order to do so. Also, as will be seen below, different sources report vastly different results of routing incidents, making it difficult to make conclusions regarding their evolution.

Figure 3. BGP Events: Leaks and hijacks

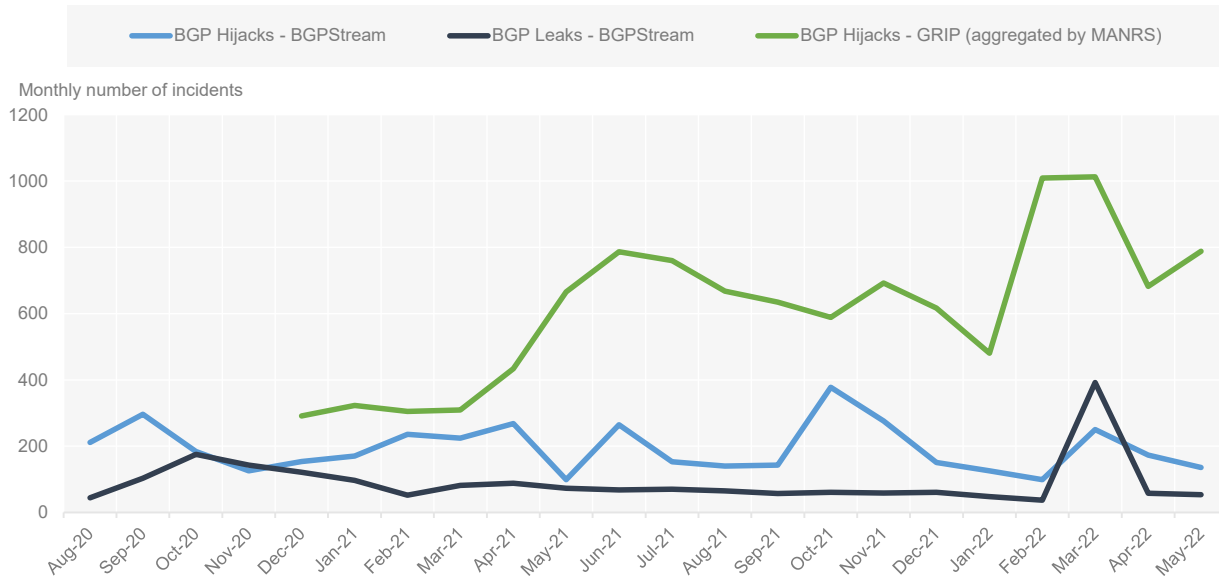


Note: Historical disaggregated data was unavailable for 2017 and 2018. BGPStream provides publicly available data on its website (bgpstream.crosswork.cisco.com) on BGP leaks and possible hijacks on a rolling 6-7 months view.

Source: BGPStream provided the underlying dataset used in all analyses (Cisco Crosswork Cloud, 2022^[21]). OECD elaboration using data collected from BGPStream (2021 data) (Cisco Crosswork Cloud, 2022^[21]). OECD elaboration using BGPStream data analysis from (Siddiqui, 2021^[22]) (for 2020 and 2019 data), (Robachevsky, 2019^[23]) (2018 data) and (Robachevsky, 2018^[24]) (2017 data).

Looking more granularly on a monthly basis, Figure 4 demonstrates that every month records several BGP incidents globally (Figure 4). There are two data sources shown: BGPStream and GRIP (Cisco Crosswork Cloud, 2022^[21]; Georgia Tech, 2022^[25]). GRIP is hosted at Georgia Tech and only looks at BGP hijacks (Georgia Tech, 2022^[25]). GRIP data is aggregated and included on the Mutually Agreed Norms for Routing Security (MANRS) Observatory (Robachevsky, 2021^[26]). As shown in Figure 4, GRIP data reports a higher incidence of hijacks than those reported by BGPStream. For instance, in May 2022, there were 136 possible BGP hijacks and 54 BGP leaks recorded on BGPStream, whereas GRIP reported 788 hijacks during the same period (Cisco Crosswork Cloud, 2022^[21]; Georgia Tech, 2022^[25]). Figure 4 shows an uptick in incidents, especially in March 2022, which coincides with the beginning of the Russian aggression in Ukraine. However further analysis would be needed to determine potential links, although this is outside the scope of the report.

Figure 4. Monthly BGP incidents



Note: GRIP data begins from December 2020 because that is the latest data available from the MANRS Observatory, which acts as an aggregator of this data to present on its website (<https://observatory.manrs.org/#/overview>). The BGPStream data for February 2022 may be underrepresented as the system appeared to be malfunctioning for a few days at the beginning of the month.

Source: OECD elaboration based on data from (Cisco Crosswork Cloud, 2022^[21]) and (MANRS, 2022^[27]). BGPStream is a free service provided by Cisco, as part of its Crosswork Cloud (Cisco Crosswork Cloud, 2022^[21]) and GRIP data (Georgia Tech, 2022^[25]) has been incorporated and aggregated by the MANRS Observatory (MANRS, 2022^[27]).

For means of comparison, Qrator Labs, a private company that provides a BGP monitoring tool called Qrator.Radar provides vastly different results to those presented above. For instance, from January to March 2021, Qrator Labs reported 7 822 532 leaks and 7 311 799 hijacks (Qrator Labs, 2021^[28]). Some of these events are classified as “global”, referring to their scale (e.g. the number of prefixes and ASes affected, propagation within global routing table). Of the close to 8 million leaks and over 7 million hijacks reported, nine leaks and two hijacks were classified as “global” over the same three-month period (Qrator Labs, 2021^[28]). The fourth quarter of 2021 reported a record high number of route leaks for Qrator Labs, with 19 852 504 leaks being recorded, four of which were classified as “global”, along with 6 076 144 hijacks, of which three were marked “global” (Qrator Labs, 2022^[29]). The wide discrepancy between the different data sources likely derives from the different methodologies and data collection methods between them, including choices regarding what constitutes a route leak or hijack. However, these are difficult to discern without detailed descriptions of the underlying methodology used.

The differences between these data sources underscore the difficulty in identifying and classifying BGP events and demonstrates that different methodological approaches lead to different results.¹⁸ While it is out of scope to discuss the merits and weaknesses of each approach in this report, more work and research should be conducted to come to agreed metrics, which may include refining existing techniques, in order to begin building trustworthy datasets and evaluate how the occurrence of leaks and hijacks evolves over time. In addition, the few sources of data, with limited detail published on their methodologies, make it difficult to determine the accuracy and potential pitfalls of such datasets. This will be further addressed in the policy discussion below.

Key routing incidents

Looking at overall numbers only fails to illustrate the impact and possible consequences from routing incidents. This section spotlights a few specific incidents that demonstrate possible risks, including

breaches of availability (e.g. connection outages, longer latencies), breaches of confidentiality (e.g. eavesdropping on traffic), and breaches of integrity (e.g. redirecting Internet traffic to a false site). However, these are not exhaustive and only aim to better illustrate the possible impacts arising from BGP events.

BGP events impacting availability

A commonly cited impact of a BGP incident refers to a risk of a connection and service outage, impacting the availability and quality of the service of the Internet connection. This could result in the unavailability of certain websites, longer latencies, or service degradation. One well cited example is the 2008 incident that resulted in the global outage of YouTube by Pakistan Telekom. Following orders from the government to censor YouTube, state-owned Pakistan Telecom began announcing a null or “dummy” route to YouTube’s address blocks in an effort to block local access to the site (Stone, 2008^[30]) The dummy route diverted traffic to another webpage, instead of YouTube, essentially discarding or “black-holing” local traffic to YouTube (Balakrishnan, 2009^[31]). However, the dummy route was inadvertently leaked to transit provider PCCW Global, which accepted the route in error and passed on these prefixes, or “propagated” it, to routers around the world (Balakrishnan, 2009^[31]).

This resulted in two available routes to YouTube, the dummy route and the legitimate route. Since Pakistan Telekom’s dummy route was more specific than the legitimate route, much of the worldwide traffic to YouTube chose the dummy route, due to BGP’s preference for more specific routes (Balakrishnan, 2009^[31]). This caused problems to access YouTube around the world, not only in Pakistan, which was far larger than the intended scope. This incident also serves as an example of how governments can be involved in and influence routing incidents. As seen here, state-owned operators, like Pakistan Telekom, can also perpetuate events.

Another example of the large-scale impact a BGP leak can have is Telekom Malaysia’s route leak in 2015, which caused significant decreases in Internet speed and quality around the world. Telekom Malaysia began announcing a large number of prefixes to network operator Level3 Communications,¹⁹ which mistakenly accepted and propagated these prefixes to its customers and peers (Toonk, 2015^[32]). In so doing, Telekom Malaysia inadvertently took responsibility for delivering these packets to the intended destinations, causing traffic to be routed via Level3 and Telekom Malaysia. Unfortunately, routing through Telekom Malaysia was often a longer route and Telekom Malaysia’s network was not equipped to handle the influx of traffic (Toonk, 2015^[32]). This resulted in service degradation and packet loss around the world for two hours before the issue was resolved. Despite the impact, analyses of Telekom Malaysia’s leak did not conclude nefarious motives. These are only a few examples of many; Table 1 and Box 1 summarise the details of a few other recent routing events affecting availability.

Table 1. Additional examples of BGP events affecting availability

Date of Incident	Event	Impact
29 July 2021	Telehouse, a Bulgarian provider, began erroneously announcing prefixes associated with Deutsche Telekom. This caused traffic intended for these networks to be redirected to Telehouse's network.	For around three hours, Deutsche Telekom customers were unable to access several services including Microsoft Office 365 and Cisco WebEx. This is likely due to traffic being redirected to Telehouse's network and then dropped, although the cause of why specifically Microsoft and Cisco services failed is unclear.
16 April 2021	Vodafone Idea began announcing over 30,000 prefixes, many more than it normally announces, including prefixes belonging to Google, Microsoft, Akamai, Cloudflare, Fastly, and others. Bharti Airtel, Vodafone Idea's upstream peer, passed on these route announcements.	During the incident, traffic was misdirected to Vodafone Idea, causing a sharp increase in inbound traffic to its network by up to thirteen times. Experts concluded that users trying to access IP addresses in the leaked routes as well as Vodafone Idea users likely experienced degraded service.
24 June 2019	DQE, an ISP in the United States, mistakenly announced over 20,000 more specific prefixes of popular services like Amazon, Cloudflare, and Facebook, to their customer, Allegheny Technologies. The routes were then leaked on to Verizon. Verizon propagated these routes to the global Internet, causing traffic to these popular sites to be redirected to Verizon, Allegheny Technologies, and DQE.	For about two hours, the route leak redirected traffic of several traffic-heavy sites to the networks of Verizon, Allegheny and DQE. The abrupt increase in traffic caused service degradation. For example, Cloudflare reported up to a 15% loss of its global traffic during the height of the incident.

Sources: In order of the events listed in the table (top to bottom): (Kleinz, 2021^[33]); (Siddiqui, 2021^[34]) and (Sharma, 2021^[35]); and (Strickx, 2019^[36]) and (Toonk, 2019^[37]).

Routing events that affect availability on a wide scale are among the easiest to detect and report, as they garner the attention of the press and the public. It is in part a testament to the resilience of the Internet that BGP events impacting availability on a large scale are relatively rare, considering the volume of traffic exchanged daily. Nevertheless, events that do not garner media attention may still impact the functioning and accessibility of the Internet. It is important to highlight that for every large-scale event noted above, there are many other incidents that are unreported.

BGP events impacting confidentiality

BGP incidents can also impact confidentiality. Malicious actors can potentially monitor or analyse traffic that is rerouted to pass under their observation. While encryption can curtail confidentiality breaches, such as by the use of protocols that employ Transport Layer Security (TLS), metadata still provides much information useful for surveillance purposes, for instance information about which servers have been accessed by which devices, at what time (ENISA, 2019^[18]). Researchers have also shown the ease with which malicious actors can obtain trusted digital certificates from Certificate Authorities (CAs) by feigning control of IP resources, which allows for impersonation and under some circumstances decryption of TLS traffic (Birge-Lee et al., 2018^[38]).

The extent of a third party's illicit observations, as well as the party's underlying intentions are often unknown, making it difficult to determine whether a breach of confidentiality has taken place. The examples presented here refer only to cases in which traffic surveillance could have occurred. One example of a possible breach of confidentiality is a 2017 incident in which an autonomous system (AS) out of the Russian Federation began announcing 80 prefixes normally announced by high-profile networks including Google, Apple, Facebook and Microsoft (Toonk, 2017^[39]). This may have been an intentional attack, first, because of the nature of the leaked prefixes, which were all well-known organisations with high traffic volumes. Second, the Russian origin AS had not been actively announcing prefixes prior to the event and some of the prefixes announced for the large players were more specific than those normally advertised by the affected companies (Goodin, 2017^[40]). It is unusual to see new prefixes appear in accidental misconfigurations and more specific prefixes are preferred in BGP route selection, suggesting that these additional prefixes were fabricated in an effort to encourage a higher rate of propagation and therefore a more effective attack (Goodin, 2017^[40]).

As other ASes began to propagate these prefixes onward through the Internet, traffic intended for the affected companies was routed through the Russian AS before reaching its final destination (Goodin, 2017^[40]). Given the volume of traffic handled by these sites, a large amount of data likely passed through the Russian AS' servers despite the incident's short duration. This raises questions as to why the traffic patterns changed, and what information may have been monitored during that timeframe. There have been other cases where traffic has been rerouted (Box 1).

Box 1. BGP events with a possible impact on availability and confidentiality

Swiss data centre route leaks result in traffic redirected to the People's Republic of China (hereafter "China")

On 6 June 2019, European mobile phone traffic was routed to China Telecom for over two hours. This was due to a route leak by Swiss data centre Safe Host, which mistakenly leaked over 70,000 routes to China Telecom. China Telecom then announced the routes onward to their peers, redirecting traffic intended for large European mobile providers (Swisscom, Bouygues Telecom, Numericable-SFR, KPN) through China Telecom's network.

Users of the affected mobile networks reported degraded service performance, including slow connectivity and unavailability of some services. Traffic analysis and eavesdropping would have been possible under the circumstances, but there is no specific information to suggest that such activities occurred in this case.

Route leak leads to traffic misdirection to China

Beginning in December 2015, an unusual routing misconfiguration caused Internet traffic destined for Verizon's Asia-Pacific (APAC) network (AS703) to traverse China Telecom's network (AS4134) in mainland China. Two mistakes occurred that ultimately inserted China Telecom into Verizon APAC's inbound traffic path:

1. SK Broadband leaked routes from one settlement-free peer (Verizon APAC) to another (China Telecom).
2. China Telecom leaked routes from one settlement-free peer (SK Broadband) to its peers and transit providers (Telia, Tata, GTT).

Despite attempts by independent researchers to alert Verizon of this issue, the phenomenon lasted over two years. For a period of ten days (within this two year timeframe), Verizon APAC sent routes from Verizon North America (AS701) to SK Broadband and China Telecom, resulting in some traffic addressed to Verizon's US destinations being temporarily routed through mainland China.

The incident's long duration is noteworthy. As the typical focus of BGP monitoring is to look for new and unexpected origins or transit providers, traffic misdirection that occurs at other parts of the AS path may be overlooked. Part of why this incident took so long to resolve may be because the problematic routing decisions occurred multiple AS hops from the origin.

Routing domestic traffic internationally results in a longer, indirect route, impacting quality of service of the connection. While there is also an associated risk of possible monitoring of traffic by the AS inserted in the path (China Telecom), it is difficult to attribute intent, as is often the case in these situations. Both route leaks by SK Broadband and China Telecom may have been the result of misconfigurations.

Note: This box was informed by written input and review by Doug Madory from Kentik.
Source: (Reporting and Analysis Centre for Information Assurance (MELANI), 2019^[41]).

BGP events impacting integrity

Another type of impact from a BGP event relates to a breach of integrity, in which traffic may be altered or directed to a spoofed site. Breaches of BGP's integrity sometimes affect the domain name system (DNS), which, as it relies on routing, is also vulnerable to BGP attacks. Some attackers have leveraged the vulnerabilities of both BGP and the DNS. For example, in May 2019, a Brazilian AS began announcing a more-specific prefix of Quad101's²⁰ prefix, which was propagated to the global Internet by Claro Brasil (Siddiqui, 2019_[42]). For a few minutes, traffic intended for Quad101 was rerouted to the Brazilian AS. Given that Quad101 is a public DNS resolver with users throughout the world, the event demonstrated the vulnerability of critical DNS infrastructure (Siddiqui, 2019_[42]).

In another instance, an AS (eNet) announced more specific prefixes of Amazon's Route 53 authoritative DNS service, which some of its peers passed on (Siddiqui, 2018_[43]). This caused some DNS queries for domains served by Amazon's Route 53 service to be redirected to a malicious DNS server (Siddiqui, 2018_[43]). The server only answered queries for the domain, "myetherwallet.com" and directed them to a false site, intending to steal login information and transfer cryptocurrency from members' wallets (Poinsignon, 2018_[44]). This also impacted availability, because in the areas where the more specific route propagated, only queries for one domain were answered ("myetherwallet.com") and all other queries from Amazon DNS customers were dropped for the duration of the incident, around two hours (Poinsignon, 2018_[44]). However, users browsing with HTTPS would have received a warning flag for the false site saying that the TLS certification was signed by an unknown authority (Poinsignon, 2018_[44]). This presumably limited its impact, as the user would have to choose to continue to the false site after the warning pop-up.

However, adversaries can sidestep this problem of TLS certification. As mentioned above, malicious actors can use BGP hijacking to subvert the verification process to obtain a trusted digital certificate from a CA, during which users are asked to demonstrate control over a domain (Birge-Lee et al., 2018_[38]). Researchers have duplicated the different ways malicious actors can use BGP hijacks to obtain trusted certificates by feigning control of network resources. With the trusted certificate, the malicious actor can impersonate a victim's servers, and under certain attack circumstances, decrypt transport layer security (TLS) traffic (Birge-Lee et al., 2018_[38]). During real-world trials, the researchers were able to obtain false certificates from leading CAs and "decrypt seemingly 'secure' HTTPS traffic within seconds" (Birge-Lee et al., 2018_[38]). These attacks can be quite stealthy and difficult to detect, especially if they do not disrupt traffic to the victim.²¹ Although fraudulent CA certificates are still rife today, the solution to this problem, DANE, has seen considerable uptake in some areas, notably server-to-server email (see the report, "Security of the Domain Name System (DNS): An introduction for policy makers", for further discussion).

Similar to the Amazon Route 53 example above, in 2014 Turk Telecom began announcing more specific IP addresses for the public resolvers of Google and OpenDNS in an effort to block certain websites (Toonk, 2014_[45]). Traffic and DNS queries intended for these public DNS recursive resolvers were instead directed to and answered by Turk Telecom's DNS resolver, which blocked access to certain domain names by returning incorrect IP addresses (Toonk, 2014_[45]). However, the fake servers answered other DNS queries normally, maintaining connectivity for the rest of the Internet.

Finally, malicious actors may hijack otherwise unused IP addresses for nefarious purposes. In one example, unused IP space owned by a Swiss cantonal administration was announced by an unauthorised AS for several months in 2015 for the purpose of sending out spam emails (GovCERT.ch, 2015_[46]). The spammer misappropriated and abused the identity and reputation of the Swiss cantonal government, which could have circumvented certain spam filtering mechanisms, which commonly take sender IP address reputation into account.

Current techniques to enhance routing security

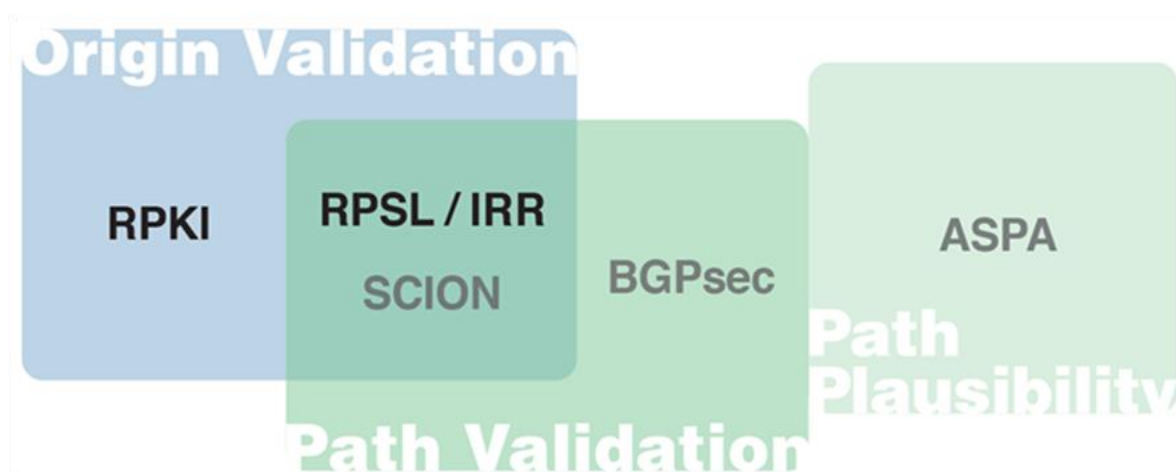
The mechanisms of routing security

BGP's vulnerabilities have prompted much discussion within the technical community over the past decades, the results of which have been comprehensively catalogued in RFC 4272 (Murphy, 2006^[14]). This section discusses the principle techniques being applied to routing security, from those that are quite mature, such as filtering and Routing Policy Specification Language (RPSL), to others that are in varying degrees of development, such as Resource Public Key Infrastructure (RPKI) and BGPsec.

Overall, each current routing security technique addresses limited aspects of the problem; even implemented all together, as many network operators and IXPs do (see Figure 6), they still do not fully solve the challenges facing routing security. Routing security consists principally of filtering incorrect or malicious routing information as it enters and leaves networks. Filters are currently built using a variety of methods and address different elements of routing security.

There are two main elements of routing security, origin validation and path validation. *Origin validation* is the process of verifying that a network that initiates the announcement of a route to a destination is authorised to do so (i.e. that the IP address space has not been hijacked by an unauthorised party). *Path validation* is the process of ensuring that no unauthorised network has diverted traffic from its authorised paths by announcing a false route (i.e. redirecting traffic through an unauthorised additional location). Related to path validation, but with more limited verification, *path plausibility* strives to determine the plausibility of a certain network being included in the AS path of a route announcement. Some techniques seek to address either origin validation or path validation/plausibility, while others seek to address both aspects ().

Figure 5. Mapping of current routing security techniques



Note: Filtering is not included in the diagram because it is a basic router function, which underlies and is used to implement all of the other techniques mentioned here. While other solutions may fit in the diagram, the focus is on those discussed in the report.

Source: OECD.

RPSL and Internet Routing Registries (IRRs) and RPKI, shown in dark bold font, are techniques that can be deployed today. RPSL and IRRs have existed in relatively unchanged form for 25 years and aim to address both path and origin validation. While there is broad adoption of IRR, many experts note the prevalence of inaccuracies in IRR routing records, limiting its effectiveness. RPKI is still under

development, but can be deployed now, and focuses on origin validation. While RPKI's rate of adoption appears to be picking up, there is still a way to go before reaching broad implementation in industry.

BGPsec has been proposed to extend RPKI techniques to perform path validation, while additional techniques, such as autonomous system provider authorisation (ASPA), have been proposed to perform a more limited evaluation of path "plausibility." However, these techniques are still under discussion and not ready for implementation. BGPsec, in particular, faces substantial challenges to deployment in the short to medium term, namely its inability to support incremental adoption and the added computational requirements placed on routers.

SCION is an academic networking effort, aiming for a "clean slate Internet" to address both path and origin validation. SCION has been deployed in full production at five ISPs at the time of writing. It can be seen as a testbed for highly secure routing infrastructure, however its ability to be incrementally deployed should be further investigated.

In addition to these algorithmic routing security techniques, network operators also manually apply routing policies that preclude many known destinations, ASes, or paths that should not appear in the routing table; these are referred to as "bogons." RFC 3871 describes a bogon as a packet that has a source address in an address block that either has not been allocated by the Internet Assigned Numbers Authority (IANA) or the RIRs, or that has been reserved for special or private use (see RFC 1918 and RFC 6890) (Jones, 2004^[47]). Since source IP addresses should not occur within these blocks, some network operators take a proactive filtering approach based on a bogon list, such as those maintained by Team Cymru (Team Cymru, 2022^[48]). These lists are dynamic and updated regularly to reflect changes (Team Cymru, 2022^[48])²². There are also some industry efforts to develop alternative methods to mitigate risks at a more local level, such as deploying monitoring tools such as BGPalerter and more bilaterally with "Peerlock" filters. A network operator may also filter out a certain address space based on previously observed incidents (e.g. malicious attacks originating from a certain address space). This is a more reactive and targeted approach.

Either by design or in practice, the techniques proposed by industry and the technical community solve only part of the challenges facing routing security. Currently, no single technique alone will meet the various routing security challenges. In addition, there are several challenges that impede the deployment of available techniques to improve routing security due to the vast and interconnected nature of the Internet and the low incentives to invest to implement new techniques, among others.

Route Filtering: the mechanism to apply routing security techniques

Route filtering is the process by which a router monitors incoming and outgoing routing information and discards information that does not adhere to its defined routing policies. It is a foundational function included in all BGP-speaking routers and it is the mechanism whereby all of the routing security techniques discussed here are applied, including RPSL/IRR and RPKI. The "best current practice" for BGP operations and security, defined in RFC 7454 of the IETF, includes recommended filters, including prefix, AS Path and Next-Hop filtering, and maximum prefix limits (Durand, Pepelnjak and Doering, 2015^[49]). RFC 7454 notes that "the main aspect of securing BGP resides in controlling the prefixes that are received and advertised on the BGP peerings", with filtering being the main mechanism to do so (Durand, Pepelnjak and Doering, 2015^[49]).

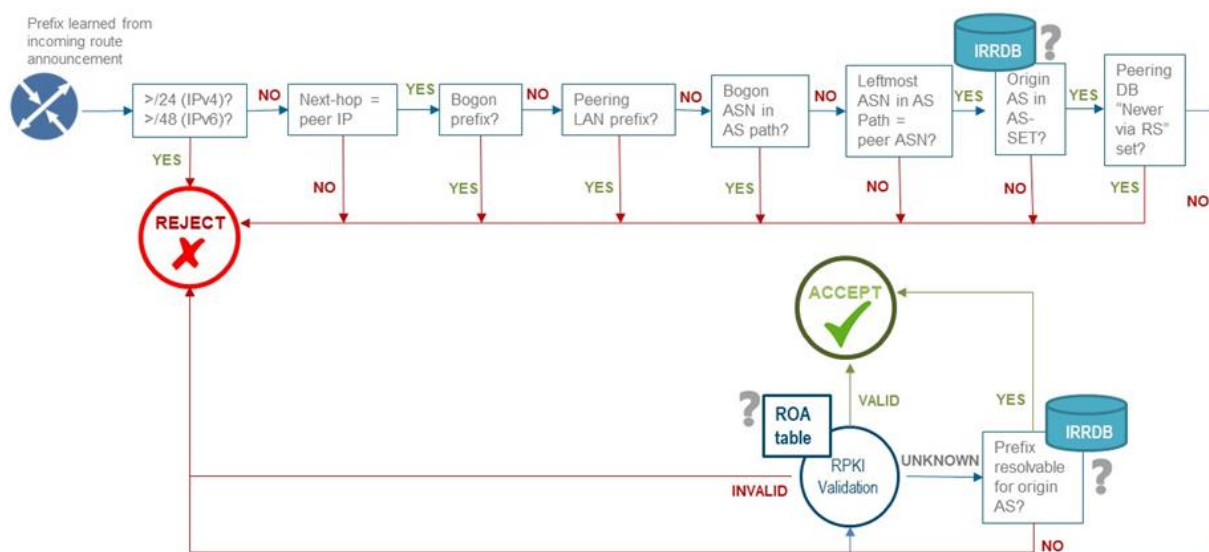
Some of the filters proposed in RFC 7454 make practical sense. These include rejecting prefixes that have not been allocated by IANA and/or the RIRs (bogons, as mentioned earlier) and prefixes that are more specific (longer) than commonly accepted levels of specificity²³ (Durand, Pepelnjak and Doering, 2015^[49]). Industry and the technical community have promoted best practices in route and packet filtering for more than two decades. In 2000, RFC 2827 (BCP 38) recommended the filtration of outbound packets in an effort to curtail denial of service (DoS) attacks perpetrated with spoofed IP source addresses (Ferguson and Senie, 2000^[50]). Some network operators also use BGP communities, which adds an attribute in BGP

announcements that gives additional information, as defined in RFC 1997 (Chandra, Traina and Li, 1996^[51]). This additional information can influence a router’s actions. When appropriately applied, the use of BGP communities can help network operators to avoid sending on the wrong routes, thereby limiting route leaks (among other uses) (Döring, 2018^[52]).²⁴ As RFC 1997 states, “a BGP speaker may use this attribute to control which routing information it accepts, prefers or distributes to other neighbours” (Chandra, Traina and Li, 1996^[51]).

Appropriately applied, even simple static bogon filters greatly reduce BGP’s vulnerabilities. Yet even this simplest and least controversial practice, which has neither cost nor risk, is not universally implemented. This underscores the difficulty of achieving widespread adoption of security norms in a heterogeneous network of global scale.

How filters are built and maintained varies: filters may be manually constructed or algorithmically generated. An operator at a keyboard may insert them into routers manually, or they may be inserted via automation. Often, both methods are used, creating potential conflicts between the filtering rules generated by different methods, or from inserting the rules into the router by different approaches. Filtering therefore can quickly become complex, and filter management errors are a common source of routing mistakes. Figure 6 illustrates a simplified representation of the routing security validation flow applied to the route server at the Frankfurt Internet Exchange, which includes numerous static bogon filters, filters generated using RPSL data from IRRs, filters generated using data from PeeringDB, and RPKI route origin validation (DE-CIX, 2021^[53]).

Figure 6. Example validation flow of incoming BGP announcements, as implemented in the DE-CIX route server



Note: This validation flow is based off the validation flow for DE-CIX Frankfurt route server; some aspects were omitted to simplify the graphic. Source: OECD elaboration based on detailed validation flow from DE-CIX route server and advice from Job Snijders; see (DE-CIX, 2021^[53]).

This complex layered approach to routing security, applying essentially all currently possible forms of protection, is necessary. However, it still does not provide sufficient protection against routing attacks and accidents due to a combination of uneven and insufficient deployment of good practices among ASes globally as well as the persistent vulnerability in the validation of the AS path. In addition, its complexity introduces fragility through a risk of routing mistakes and misconfigurations. Yet, this is the current best practice.

Current routing security techniques

Routing Policy Specification Language (RPSL) and the Internet Routing Registries (IRRs)

RPSL is a format to express the routing policies that ASes implement. Network operators can define what IP address blocks they are responsible for, who their customers are, who their transit providers are, to whom they delegate authority, and other useful information, using statements formatted in RPSL. Uniquely among the currently implementable techniques, RPSL data allows routing security implementers to validate both the *origin* of a route (its starting point), and the *path* of a route (the various valid and authorised midpoints through which traffic may flow to that destination). Of all the currently deployable techniques, it is the only one that has the *potential* to be a sufficient solution covering both aspects of routing security.

RPSL-formatted information is published in repositories called Internet Routing Registries (IRRs) (Internet Routing Registry, 2018^[54]). The present constellation of IRRs expanded gradually from an initial single registry, the Routing Arbiter Database (RADb), established in 1995. Today, each of the five RIRs maintains an IRR, along with the original RADb. Dozens of the largest network operators also maintain separate IRRs. The RIRs' IRRs can have an advantage over other third-party IRRs in terms of authorisation, authentication and validation of IRR objects. For instance, RIPE NCC's IRR only allows an address holder to create IRR objects over its own IP address space, which helps to establish legitimacy of the IRR object. This is not the case with other third-party IRRs, which have to rely on other means to authenticate and validate. However, the information published in any IRR (including the RIRs) must still be updated by IRR object holders, when required, to ensure that it remains current and accurate.

RPSL/IRR Adoption and Effectiveness

The well-defined and comprehensive language (RPSL), together with a developed ecosystem of tools, is appealing to many operators and has spurred participation. Globally, there is a high level of participation in the IRR with 88% of routing announcements registered in the IRR in January 2022²⁵ (MANRS, 2022^[27]). Nevertheless, RPSL is hamstrung by the deficiencies of the IRR system with which it has historically been paired. If RPSL were to form the basis of future routing security work, the IRR system would need to be jettisoned in favour of a modern end-to-end system of cryptographic signatures, rooted in RIR resource delegations, which would require new specifications from the IETF.

Internet Routing Registries seem to work best in defined contexts where they are carefully and continuously managed for consistency, coverage and accuracy (Huston, 2021^[55]). When they take on a broader scope or are not actively managed, their consistency and utility falls. However, metrics tracking the *effectiveness* of IRR in its current state are scarce. Available metrics focus on registration within the registries, not the accuracy of the data or their ultimate usefulness to network operators.

Challenges associated with RPSL and IRRs

RPSL gives network operators a common and sufficient language to communicate the policy information necessary to implement routing security and benefits from a broad and mature set of tools, both commercial and open-source, to manipulate RPSL-formatted data. However, the system has two weaknesses. First, the proliferation of only-partially-synchronised IRRs has led to the present situation, in which only a minority of the information published in IRRs is current and accurate. While it is easy to import RPSL data from IRRs and turn it into routing filters, those filters will only be as accurate as the data from which they were built. Unfortunately, the quality of the input data has been deteriorating over time.

Second, RPSL and the IRR system were designed prior to cryptographic security's commonplace adoption. As a result, like all early Internet protocols and systems, they depended upon trust and goodwill. Because routing security was recognised as critical at a very early juncture, attempts to secure the RPSL/IRR system began early on, before the cryptographic building blocks existed. Consequently, the

security model that was retrofit onto the RPSL/IRR system was based upon access-control, rather than end-to-end verification. The trustworthiness of RPSL data depends upon IRR operators prohibiting unauthorised parties from entering RPSL data into, deleting data from, or modifying data within, their IRR. This access control must be implemented independently by each IRR operator, and must be applied to each party who seeks to modify the content of the repository. Access control is a difficult task, and even more so in this system. Authenticating hundreds of thousands of individuals variably representing tens of thousands of different organisations, to dozens of IRRs, in constantly changing relationships as required under many-to-many access control has proven to be unsustainable.


Resource Public Key Infrastructure (RPKI)

RPKI is another method aimed at improving BGP security, first proposed in 2007 and published as an IETF RFC in 2012. RPKI establishes an architecture that allows an entity to define the IP addresses and/or ASNs of which it is the legitimate holder and to authorise one or more ASes to originate routes to a given set of the prefixes it holds (Lepinski and Kent, 2012^[56]). These authorisations are cryptographically verifiable (Lepinski and Kent, 2012^[56]). Current applications of the RPKI architecture focus on origin validation²⁶ for which two pieces are critical for its success: the creation of Route Origin Authorisations (ROAs) and the validation of ROAs through Route Origin Validation (ROV) and subsequent filtering on any entries deemed to be invalid. For a further technical description of ROA creation and ROV filtering, please see Annex A.

The RPKI architecture brings a few advantages. Most importantly, it offers a way to provide cryptographic origin validation of a route announcement. The technique can be implemented today, as evidenced by various organisations that have announced their active deployments of RPKI (ROA + ROV) in their networks. Its structure allows for incremental deployment and provides direct benefits to the ASes implementing it, as it allows them to protect their own prefixes (ROAs) and filter invalid routes from others (ROV). It also benefits its neighbours, which only receive valid routes. Certain network effects exist in the architecture (e.g. the more entities that sign ROAs and conduct ROV filtering, the more benefit adopting the technique will bring). In addition, the RPKI infrastructure puts in place certain safeguards to ensure the quality of ROAs. The architecture only allows an entity to digitally sign a route origin authorisation (ROA) for the prefixes it holds (Levy, 2018^[57]). Further, each ROA has a finite life, avoiding long-lived and outdated ROAs.²⁷ By comparison, the IRR does not have similar safeguards to protect against invalid or outdated data (Levy, 2018^[57]).

RPKI Adoption and Effectiveness

The two steps necessary to the implementation of RPKI are the creation of ROAs by prefix holders, and the rate that network operators validate routes through ROV. The usefulness of ROV filtering depends on the number of prefix holders that have created ROAs, and the accuracy of the ROAs themselves. Assuming that the information published in the ROAs is maintained and up-to-date, a sufficient coverage of the IP address space by ROAs is a prerequisite for operators to gain benefits from ROV filtering. Regarding the creation of ROAs, Figure 7 shows a steady increase in the number of validated ROAs since 2014, picking up speed in 2019. There are other sources tracking the creation of ROAs, including NLnet Labs and APNIC, which both show this information on a country basis (APNIC, 2022^[58]; NLNet Labs, 2022^[59]).


Figure 7. Number of validated ROAs, by RIR


Source: OECD elaboration based on RIPE NCC statistics; (RIPE NCC, 2022^[60]).

However, registering ROAs provides no inherent benefit if ASes do not filter routes to drop RPKI invalid results. To estimate the prevalence of networks filtering invalids (ROV), APNIC Labs measures the rate of propagation for an invalid route (Huston, 2021^[61]). A network that filters invalids would discard the invalid route, thereby limiting its propagation and preventing its users from reaching destinations that are announced only by invalid routes. A higher percentage for a given country means that more Internet users in that country are covered by networks that are ROV-filtering (Huston, 2021^[61]). Figure 8 shows that the OECD average rate was roughly 27% as of September 2021, rising to 28% in January 2022, before falling to around 24% in June 2022. There are also important differences between members, although what is perhaps most clear from the data is the degree of change over the three time periods under study, including in the ranking of countries (APNIC Labs, 2022^[62]). These fluctuations may suggest that network operators are still getting familiar with ROV filtering and may turn off filtering to fix errors (e.g., mistakenly filtering valid traffic). In addition, a small country's rate may be particularly impacted if one large provider starts or stops ROV filtering for a period of time. Globally, around 10-20% of users are covered by networks that are dropping invalid routes, with the rate fluctuating between that range from 2021-2022 and peaking at 38% in December 2020 (APNIC Labs, 2022^[63]).²⁸

There are other crowdsourced efforts that allow users to test whether their Internet providers are conducting RPKI filtering, such as Cloudflare's "Is BGP safe yet?" and RIPE Lab's RPKI test (Cloudflare, 2021^[64]; RIPE Labs, 2019^[65]). However, these are ad-hoc efforts without a systematic approach to track the actual rate of ROV filtering (e.g. one user performs a test for one ISP at a given point in time). Nevertheless, having additional sources to track ROV filtering would provide a more comprehensive view of the actual rate of filtering.

Figure 8. Rate of ROV filtering in OECD countries (September 2021, January 2022, June 2022)



Note: Figures shown are from three 7-day timespans (September 2021: 30 August – 5 September 2021, January 2022: 24 January – 30 January 2022 and June 2022: 21 June – 27 June 2022). Further discussion of methodology can be found at (Huston, 2021^[61]).
 Source: OECD elaboration based on data from APNIC Labs; (APNIC Labs, 2022^[62]).


Over the past few years, several large players have begun to drop RPKI invalids (ROV filtering), such as AT&T, Comcast and Hurricane Electric, as well as IXPs such as DE-CIX, Seattle IX, and YYCIX (Borkenhagen, 2019^[66]; Livingood, 2021^[67]; Hurricane Electric, n.d.^[68]; DE-CIX, 2021^[69]; Seattle Internet Exchange (SIX), n.d.^[70]; YYCIX, n.d.^[71]).²⁹ As large ISPs and network operators adopt the policy of discarding routes that are invalid according to published ROA information (e.g. RPKI invalids), the potential benefit of publishing ROAs increases. In one academic study, Testart et al. demonstrate that as more players filter ROV invalid responses, registering prefixes in RPKI provides direct benefits by curbing the propagation of an illicit announcement in actual cases, like origin hijacks (Testart et al., 2020^[72]). As more operators begin to enforce RPKI, this benefit is likely to improve and may incentivise prefix holders to sign and maintain ROAs in a beneficial cycle.

Considering the adoption of both ROA creation (Figure 7) and ROV filtering (Figure 8), the most pertinent question for policy makers concerns the impact the current state of RPKI implementation has had to decrease BGP incidents. One possible method to evaluate RPKI’s effectiveness to decrease incidents is to consider the percentages of “valid” responses returned during route origin validation, compared to total routes being announced. In theory, “valid” responses would indicate the proportion of prefix-origin pairs that would benefit from origin validation and therefore be protected from origin hijacks or misoriginations. However, this has a large caveat in that it depends on ASes actually filtering invalid responses.

The RPKI Monitor operated by the US National Institute of Standards and Technology (NIST) tracks the results from ROV validation of unique prefix-origin pairs on a global basis. In May 2022, 61.9% of unique prefix-origin pairs in terms of IPv4 (/24s) address space resulted in an “unknown” response in ROV, compared to 36.51% “valid” responses and 1.59% “invalid” (NIST, 2022^[73])³⁰. The percentage of “valid” responses indicates that just over one-third of overall prefix-origin pairs would benefit from origin validation and therefore would be protected from origin hijacks or misoriginations if its peers were actively filtering invalid responses. Further, Figure 9 shows that the percentage of “Valid” responses being returned for prefix-origin pairs has steadily increased, especially since 2019, coinciding with the increase in ROA creation (NIST, 2022^[73]). However, NIST notes that it “does not attempt to measure the extent of

deployment of RPKI-ROV in operational networks (that is, which networks are filtering BGP based upon RPKI data)” (NIST, 2022^[74]). MANRS Observatory provides a similar view as the numbers reported by NIST on a global scale, reporting 61.7% “unknown” results, 37.8% “valid”, and 0.5% “invalid” for May 2022 (MANRS, 2022^[27]).³¹

Figure 9. RPKI-ROV History of Unique Prefix-Origin Pairs (IPv4)




Note: The graph shows the percentages of valid, invalid, and not found prefix-origin pairs of all prefix-origin pairs extracted from BGP table “dumps”, which NIST then validates against RPKI databases (e.g., ROV filtering), according to the methodology described here at https://rpki-monitor.antd.nist.gov/Methodology#ROV_History.

Source: OECD elaboration based on data from NIST RPKI Monitor (v2.0), (NIST, 2022^[73]).

However, these results as a measure of RPKI’s effectiveness are incomplete, as they do not take into account the rate of ASes conducting ROV nor provide evidence on the extent to which implementation decreases the likelihood of BGP incidents (e.g. origin hijacks)³². This is not an easy task given the difficulty in measuring effect; existing research and measurement efforts should be leveraged and expanded upon to better evaluate the impact of implementing RPKI (ROA creation + ROV filtering) on routing incidents. This would greatly improve understanding in this area.

Comparing RPKI versus IRR among OECD countries, information aggregated by MANRS Observatory shows that registration in the IRR is on average higher than the evidence of ROAs in RPKI (Figure 10) (MANRS, 2022^[27]). This is evidenced by comparing the percentages of RPKI “unknown” to the percentage registered in IRRs, as an RPKI validator returns an unknown response if there is no matching ROA found. The OECD average as of June 2022 was 51.5% “valid”, 48.2% “unknown”, and 0.3% “invalid”, showing an increase in “valid” responses up from 45.7% as of August 2021 (MANRS, 2022^[27]). 94% of OECD countries had routes registered in an IRR as a route object (MANRS, 2022^[27]). However, the existing information on the IRR figures only looks at coverage, not at the accuracy of IRR records, which is an identified challenge of the IRRs, as noted above.

Figure 10. IRR and RPKI adoption, OECD countries (June 2022)



Note: “IRR, % registered” refers to whether a network’s routing announcements have been registered in the Internet Routing Registries (IRRs) as a percentage of prefixes from selected networks; RPKI, % valid/ % invalid/ % unknown refers to the numbers of valid, invalid, or unknown responses returned by a validator in the resource public key infrastructure (RPKI) as a percentage of prefixes for selected networks.
 Source: OECD elaboration based on data on OECD countries from Mutually Agreed Norms for Routing Security (MANRS) (2021), *MANRS Observatory: Overview*, <https://observatory.manrs.org/#/overview>.

Challenges associated with RPKI

While creating ROAs is a relatively straightforward process, implementing ROV filtering on a network is a more complex process, though a necessary step to receiving security benefits from the technique. This requires planning, time and resources. Given the potential impacts to the network (i.e. dropping announcements that should not be dropped), network operators often need to work with customers to ensure their ROAs are correct and carefully test ROV filtering prior to full deployment. Deploying ROV filtering may also demand upgrades to network equipment and/or software, training for staff, and stress testing to ensure proper implementation in the network.

To illustrate resources needed, in one recent real-world example, Lumen Technologies, an important network provider, implemented RPKI on the Lumen network (AS3356) in the first quarter of 2021, following years of considering the technology and planning its deployment (Pfaff, 2021^[75]). Before implementing RPKI fully, Lumen needed to undertake significant hardware and software upgrades, develop its routing policy to handle RPKI invalid results, implement RPKI validation servers including evaluating, testing and deploying validation software, and train staff. All of these actions required time and resources to ready Lumen’s system to turn on ROV filtering. While preparing for its RPKI rollout, Lumen contacted over 200 unique customers regarding over 2 600 routes that were suspected of being invalid. This outreach prompted many of its customers to validate and correct those ROAs found to be invalid, preventing them from being dropped once its RPKI solution was deployed. While the overall project took around 18 months to complete, the actual implementation itself was completed within 24 hours, without any issues due to Lumen’s careful approach. Because Lumen is one of the most deeply peered ASes on the Internet, this

validation process also enabled many customers to reclaim Internet Protocol (IP) ranges that were being hijacked.

Some also argue that the RPKI architecture has certain vulnerabilities, the largest of which being that it does not protect the AS path or provide path validation.³³ As the whole architecture is based upon trusting route announcements with valid ROAs, one possible attack vector is on the architecture itself. For instance, if RPKI were made to return an invalid response to a *valid route* announcement, that route would be discarded by any AS conducting ROV filtering. In addition, the possibility to impersonate an AS exists, which would not be normally detected through the RPKI ROV.³⁴ Another potential vulnerability arises from the integral role the five RIRs play in the RPKI architecture, as trust anchors and publishers of ROAs. This makes their continued proper functioning important to the whole infrastructure. The concentration of these activities in the five RIRs may also present an additional risk in case of an accidental or malicious act (for instance, if an outage of the RPKI ROA repository occurred or if an attacker leveraged a security breach to assume control over the repository) (Durand, 2020_[76]). Indeed, there have been reported outages of the RPKI repositories maintained by RIPE-NCC, ARIN, APNIC, and AFRINIC, although outages are unlikely to cause operational issues unless they last longer than a number of hours (Durand, 2020_[76]).³⁵ The risk under a security breach at one RIR is exacerbated because each RIR RPKI repository covers the entire Internet, not just the address space and ASs under their remit. Thus, a security breach in one could affect ROAs related to address space under the remit of the other RIRs (Durand, 2020_[76]).³⁶ However, the five RIRs acting as overlapping trust anchors can also provide a safety net if one RIR were to be compromised for a substantial length of time (Durand, 2020_[76]).

In addition, the supporting infrastructure of the RPKI (e.g. RPKI software packages, validator software) is still developing, with a limited number of software options. Among the limited options available, there have been bugs and security vulnerabilities reported. However, as the RPKI ecosystem becomes more developed, the software will likely improve as well, as is often the case with a developing software ecosystem. Finally, while there are safeguards in place for ROA quality, the possibility of human errors remains; a ROA only attests that the owner of the prefix signed it, not that the information in the ROA is necessarily correct (Durand, 2020_[76]). Nevertheless, many in industry see the benefits of the RPKI architecture for origin validation, even if it is only one piece of the puzzle for routing security.

BGPsec

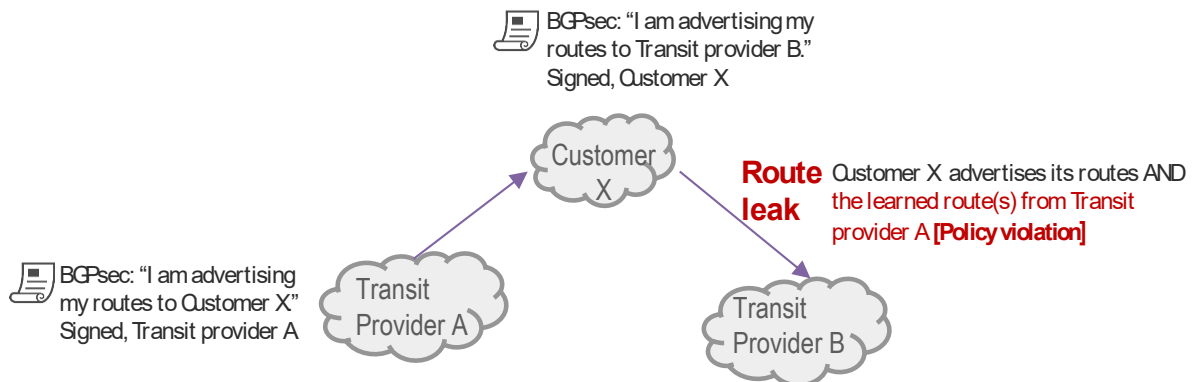
The RPKI architecture outlined above for route origin validation is only one aspect of routing security; it does not validate the AS path (e.g. path validation). Therefore, a determined actor could still manipulate the AS path and perpetrate a route hijack. BGPsec has been developed to address this aspect of BGP security, building upon the RPKI architecture, and indeed was conceived to be employed in tandem with RPKI (ROA creation and ROV filtering) (Lepinski and Sriram, 2017_[77]). Namely, a valid BGPsec “update” message confirms that every AS listed in the message has authorised the route to be advertised to the next AS in the path, by cryptographically signing the information it has added to it (Lepinski and Sriram, 2017_[77]). The RPKI architecture is the foundation for the cryptographic digital signatures. For further technical description of BGPsec, please see Annex A. The IETF has published an RFC describing the BGPsec Protocol (RFC 8205), which is labelled as a “proposed standard”. However, while there is a published RFC for the technique, BGPsec is not ready for widespread deployment, due to some of the challenges outlined below. Therefore, its adoption and effectiveness cannot be accurately measured.

Challenges associated with BGPsec

BGPsec has some limitations at the architectural level, the first of which being that it does not offer origin validation, as noted above. Secondly, route leaks can still occur even under perfect BGPsec deployment. Consider the case where customer X (e.g. a regional ISP) leaks a route it learned from transit provider A to transit provider B: even if both transit provider A and customer X both employed BGPsec, the customer

could still mistakenly leak routes. BGPsec affirms that an AS (e.g. transit provider A) announced the route to the next AS (e.g. customer X) – this makes sense, as a transit provider would normally announce routes to its customers in order to fulfil its role. The mistake occurs when customer X re-advertises the learned route(s) (from transit provider A) to transit provider B, which is a violation of policy between transit provider A and customer X (White, 2015^[78]). However, from BGPsec's perspective, customer X *did* legitimately receive the route (it was signed by transit provider A), so BGPsec would not automatically flag an error (Figure 11).

Figure 11. Simplified route leak under BGPsec



Note: This is a simplified example in order to aid comprehension.
Source: OECD.

In addition to its structural limitations, practical deployment of BGPsec raises challenges. Implementing BGPsec will require using a new BGP path attribute and negotiating a new BGPsec capability between peers (Lepinski and Sriram, 2017^[77]). This can be a complex undertaking for network operators. In addition, BGPsec requires routers to perform cryptographic functions, including signing outgoing BGP messages and validating those they receive. For comparison, ROV outsources these cryptographic functions to RPKI validators. The “update” messages themselves will also likely be larger, given the inclusion of digital signatures, and more numerous, as one “update” message is required for each prefix, for each AS to which it is sending the update message (Lepinski and Sriram, 2017^[77]). This together places a heavier burden on the router and whether today's routers are equipped to handle this increased load has been met with scepticism (Huston, 2021^[55]).

Perhaps an even more important challenge is that BGPsec requires every router in the path to deploy BGPsec to reap the security benefits of the protocol. If there is one non-BGPsec router in the chain, the path is converted back to a normal AS path update, losing any potential benefits of BGPsec in the process (Lepinski and Sriram, 2017^[77]). While perhaps incremental deployment may be possible, for instance, by focusing on deploying BGPsec in groups of ASes, BGPsec only works if each router in the path supports BGPsec. Given the scale of today's Internet and the number of routers that would have to implement BGPsec, this is a large challenge to BGPsec's adoption. An added challenge is that the number of possible connections and paths is exponential to the number of ASes, as ASes may exchange routing information with many other ASes through peering.

These challenges have caused some experts to label BGPsec as an unviable option. However, others maintain that while progress is slow, future implementation may still be possible. For example, future developments in technology may allow routers to handle cryptographic functions at the hardware level (although the challenge to incremental deployment remains). Nevertheless, most experts expect BGPsec deployment to be limited in the short to medium term. However, having a viable solution for origin validation without one for path validation leaves a gaping hole in the overarching security of the routing system. In

response, the technical community is considering other ways to provide a degree of protection to the AS path. One such effort being discussed at the IETF is the autonomous system provider authorisation (ASPA) (Box 2).

Box 2. Autonomous System Provider Authorisation (ASPA)

ASPA is an alternative response to path validation, providing a way to detect route leaks and hijacks through “path plausibility”. At a high level, the solution proposes a new digitally signed object, an ASPA, building off existing RPKI architecture. A customer AS would sign an ASPA to define and authorise a set of Provider ASes to propagate its route announcements. The set of Provider ASes would include a provider’s upstream provider(s) and peers. For further technical description of ASPA, please see Annex A.

ASPA addresses some of BGPsec’s shortcomings to implementation at scale, namely by allowing for incremental adoption and security benefits for participants even at an early stage of adoption. It also leverages existing RPKI infrastructure and does not change the BGP protocol (e.g., by introducing a new attribute). However, it does introduce a new signed object (ASPA) and requires a database where ASPA data could be queried (similar to RPKI, where databases store ROAs). ASPA can detect malicious attacks and mistakes from customers and peers, but only detects mistakes from upstream providers and route servers, as noted in its recent IETF draft.³⁷ Like RPKI, its benefits would increase as more ASes create ASPAs.

ASPA is currently being discussed within IETF as an “Internet-Draft”, with the latest draft published on 11 July 2022. Internet-Drafts are revised and improved within the IETF, and may be published as a “Request for Comments” document, more commonly referred to as an “RFC”. BGPsec, by comparison, is a published RFC (RFC 8205) and is labelled as a “proposed standard”.

Source: (Azimov et al., 2022^[79]); (IETF, n.d.^[80]); (Lepinski and Sriram, 2017^[77])

Industry solutions and initiatives

Industry efforts

The solutions highlighted above concentrate on structural solutions to improve routing security discussed by technical bodies, such as the IETF. However, this process is slow and widespread deployment of solutions at the global level takes time. Given the slow progress of these solutions, industry and other stakeholders have considered other ways to mitigate risks and address insecurity. As one example, some operators have begun to deploy BGP network monitoring tools that alert network anomalies or flag suspicious events in their network. Tools include both open source applications, as well as both in-house and third party solutions.

BGPalerter is one such open-source tool gaining traction among industry, with users such as Cloudflare, Fastly, the Latin America and Caribbean Network Information Centre (LACNIC), and the Seattle Internet Exchange, among others (BGPalerter (GitHub), 2021^[81]). NTT developed BGPalerter as an open source tool to allow other networks to use the solution to monitor BGP announcements and receive alerts in case of abnormalities on their networks³⁸ (Candela, 2020^[82]). Others include ARTEMIS,³⁹ developed with funding from the European Research Council, the National Science Foundation, and the Department of Homeland Security, among others, and Radar by Qrator, a third-party BGP monitoring service (ARTEMIS, 2020^[83]; Qrator, 2021^[84]).

Since 2017, Microsoft has operated an in-house solution called the Route Anomaly Detection and Remediation (RADAR) system that aims to identify BGP hijacks and leaks and protect against them in “near real-time” (Microsoft, 2021^[85]). Microsoft also publishes ROAs and conducts ROV filtering, and a first step in its route hijack analyser checks for a route’s validity (e.g. ROV) (Microsoft, 2021^[85]).

Another solution that has been adopted within the industry is to use “Peerlock” or “Peerlock-lite” filters. Put simply, Peerlock is a bilateral agreement between peering partners to designate authorised upstream networks. One ISP would ask its peering partners to define which networks they authorise as upstreams, if any, and apply those preferences into their routing policy (NTT, 2016^[86]). In 2016, NTT began deploying Peerlock as a way to prevent route leaks and noted a “vast improvement for networks that agreed to be a Protected ASN” (NTT, 2016^[86])⁴⁰. Academic simulations estimate that a wider deployment of Peerlock-lite filters at about 600 large ISPs (non-Tier 1) would mitigate up to 80% of simulated Tier 1 route leaks (McDaniel, Smith and Schuchard, 2021^[87]). Both Peerlock and Peerlock-lite have a relatively low technical complexity and a low burden to deploy, compared to other possible solutions.

Industry and the technical community have also come together on routing security through multistakeholder partnerships. One such example is the Mutually Agreed Norms for Routing Security (MANRS), a voluntary initiative backed by the Internet Society that aims to encourage stakeholders to adopt key actions to address common routing threats to improve the stability of the Internet. MANRS proposes a set of specific recommendations for network operators, IXPs, and content delivery providers (CDNs) and cloud providers, respectively (MANRS, 2021^[88]). In September 2021, it launched a program for routing equipment vendors to encourage the inclusion of security features in their products that will enable customers to route Internet traffic securely by default and implement actions recommended in the other MANRS programs (ISOC, 2021^[89]). MANRS has approximately 700 participants worldwide (Wan, 2021^[90]).⁴¹ The MANRS Observatory also aggregates various data sources in an easy to understand dashboard (MANRS, 2022^[27]).

Academic research and development

In addition to the industry solutions mentioned above, researchers in academia have also been engaged to develop ways to address common challenges facing routing security. One such effort is an innovative solution proposed by researchers from the ETH Zürich, a new Internet architecture that does not rely upon BGP, called “SCION” (Box 3). While it is unclear if such innovative efforts will come to widespread deployment, such research is important to cultivate new ideas and approaches to routing security.

Box 3. A Proposal for a new Internet Architecture: SCION, “Scalability, Control and Isolation on Next-Generation Networks”

SCION, or “Scalability, Control and Isolation on Next-Generation Networks” is a network architecture developed by researchers from the ETH Zürich that proposes a new way to route traffic on the Internet, without using BGP. It aims to solve many of the existing challenges to routing security, reliability and performance, by isolating routing failures by grouping ASes into separate domains called “isolation domains” (ISDs), giving more control over the routes packets take and ensuring explicit trust for end-to-end communication. For further technical description of the SCION architecture, please see Annex A.

While promising, the architecture is at the beginning stages of deployment and implementation, with formal standardisation being currently pursued. SCION can be deployed on existing commodity hardware at the network edges where traffic is exchanged between ASes, while reusing existing intra-domain network infrastructure, both of which make it relatively easy to deploy. While SCION connections can interoperate with IP/BGP, the full benefits of the architecture can only be realised

when there is a full end-to-end SCION connection, which would require substantial take-up of the technology by ASes around the world.

At the current stage, SCION is well suited to use cases that require multiple independent parties to reliably exchange sensitive information, such as in the healthcare, energy, finance, or governmental sectors. As one key example, the Swiss National Bank and SIX, which operates the infrastructure for the Swiss financial centre, announced the Secure Swiss Finance Network (SSFN) in July 2021. The SSFN is a new communication network based on SCION architecture to support the Swiss financial sector and the Swiss Interbank Clearing (SIC) system. In addition, the architecture already has been implemented at several Swiss operators, which now offer market-ready SCION products. In Korea, IoTcube is a SCION partner and provides SCION service in the country, with the first SCION node installed in the LGU+ network. Several research and educational institutes from the United States, Europe and Asia also are experimenting with the new architecture in the SCIONLab global research testbed.

Note: The information included in this box has benefited from input and review by Nicola Rustignoli from the Network Security Group at ETH Zürich.

Source: (SCION, 2021^[91]; SCIONLab, 2020^[92]; Swiss National Bank, 2021^[93]; LEE, 2021^[94]).

Overarching challenges to improving routing security

The insecurity of the routing system is not a new problem. Several solutions have emerged over the years, with continual efforts going on in the technical community to refine, promote, or develop new solutions. The persistence of efforts to improve the security of the routing system point to not only the continuance of the problem itself, but also to the fundamental challenges in implementing lasting solutions. While some challenges may be specific to certain techniques, others are more general.

One such challenge relates to the vast and interconnected nature of the global Internet and the diversity of actors participating in it. At its core, the Internet is a “network of networks”, all interconnecting with one another and exchanging traffic. Due to its very foundation, therefore, the actions of one AS alone will not safeguard itself from routing incidents, because it can be impacted by the actions of another AS. Conversely, one AS’ actions could have long-ranging impacts, for the same reason. This requires collective action and a sufficient threshold of ASes around the global Internet must adopt security solutions before security of the overall routing system substantially improves.

The interconnectedness of the Internet can also lead to misaligned incentives. As outlined above, the routing decisions of one network can affect the security of many ASes with whom the network connects. However, the impact of a given AS’ routing decisions may not be felt by the perpetrating network. This takes away some incentive for implementing good routing hygiene, because one AS’ actions may impact other ASes’ routing security more than on its own. Likewise, the benefit felt from implementing routing techniques often depends on how many other ASes have also implemented it. This is especially true for BGPsec, for example, which requires every AS in the path to implement it to reap security benefits. However, it can also be applied to RPKI, as ASes will benefit more from the solution if adoption increases (even if the implementing party does gain some benefit from the beginning).

A subsequent challenge is that some actors do not see sufficient incentive to spend time, money, or resources to implement existing solutions or develop new ones. Implementing any solution takes time and resources to roll out and businesses must be convinced of the benefits of doing so. Unfortunately, security improvements may be difficult to monetise or market to customers as a differentiating factor, meaning any investment therein is a sunk cost. Another aspect that factors into a business’ decision is the potential risk

to traffic exchange. For a network operator whose main concern is maintaining connectivity for its customers, turning on a solution that may drop a customer's traffic by mistake poses a substantial risk⁴².

Finally, the insecurity of the routing system is a large and multi-faceted problem. BGP incidents can take the form of leaks and hijacks and can be intentional or accidental, and some of them can be very hard to detect. Issues can occur at the origin or over the path. Given this complexity, some of the solutions proposed consider aspects of the problem, but that means that deploying one solution will not solve all issues facing routing security. Adopting a layered approach to routing security that implements different techniques can also add complexity for a network operator. This could inadvertently increase the risk of mistakes and misconfigurations in applying the various rules, which could then lead to routing incidents (and may impact availability).

Policy discussion

Policy makers can take a series of actions to improve routing security. These actions include funding the collection and publication of data on routing incidents and the effectiveness of different techniques, promoting awareness and deployment of good practices and security techniques, facilitating information sharing through formalised feedback loops, and defining a common framework with industry to improve routing security.

As routing security is a subset of digital security, routing topics may fall under broader legislation on digital security or communications, either explicitly or indirectly. In some countries, regulatory agencies may have the possibility of establishing secondary legislation on a more technical basis.

Routing security is a technical topic and industry players and the technical community hold substantial expertise gained from practical experience managing complex network routing. However, they are similarly bound by business constraints, which may impede investments to improve routing security. Therefore, the challenge facing policy makers is to strike the right balance between giving stakeholders the freedom to decide how best to protect networks from routing insecurity and ensuring that adequate measures are indeed being implemented.

While some policy actions can help to support the overall development of routing security, policy makers should be wary of introducing measures that place undue regulatory burden or centralise control of the routing system. The unintended consequences of any potential policy action should be duly considered to ensure the continued functioning of the Internet. The policy actions presented below include examples of how policy makers around the OECD have engaged to promote a stronger, more secure routing system. There are a range of approaches that countries can take to support this goal, which can be adapted to national circumstances and policy objectives. However, all policy makers should recognise that securing the routing system is foundational to ensure the digital security of communication networks.

Policy actions to support stronger routing security

Promote measurement efforts

Evidence-based policy making, a core tenet of OECD work, requires robust data that policy makers can use to inform their policy approach. Without information on the scope of routing incidents, policy makers and regulators cannot develop informed policies to improve routing security. Robust time-series data on the scope, scale and impact of routing incidents is currently lacking and a more systematic approach to tracking the scope and scale of routing incidents, therefore, is needed to form a solid basis of understanding of the topic.

The challenges to tracking and collecting such data are by no means small and have been duly noted above. In addition, only a handful of sources currently exist, of which some have been discontinued even

during the months this report was written⁴³ and some, if not most, lack long-term funding. These initiatives are also often driven by a small group of engaged researchers and engineers that strive to enhance routing security, but lack a more formal and consistent funding framework that would allow them to continue to invest their time and resources. This uncertainty makes collecting data and maintaining longevity a difficult endeavour to ensure time-series data.

Governments thus can fund the collection and publication of time-series data on identifying and analysing routing incidents by supporting private (multi-stakeholder) initiatives. Funding should be long-term and consistent, and awarded to neutral third parties. Funding requirements should further stipulate that results be published at no cost, in an easy to read and process format.

Besides allocating funding to research institutes, academics and technical experts in this area, an additional possible avenue could be for a national statistics bureau, regulator or other relevant government agency to begin measuring and collecting such data, in order to ensure more long-term provision for this information. However, this does not imply any obligation of network operators to report incidents to the government, as many methods do not rely on reporting to detect routing incidents. For instance, the main incident reporting tools used above, the GRIP platform and BGPstream, do not rely on incident reporting from operators but rather use algorithms that flag suspicious events or changes observed in the global routing table as potential leaks or hijacks.

Some initiatives are already underway. For instance, the Japanese Ministry of Internal Affairs and Communications (MIC) funded research efforts from 2006-2009 to develop a tool to detect possible route hijacking events, which culminated in a detection tool called “Keiro Bugyo” (or “route magistrate”) (MIC, n.d.^[95]). ICT-ISAC, an industry association, operates Keiro Bugyo and together with the Japan Network Information Center (JPNIC) notifies registrants in the JPIRR, a routing registry, of possible route hijacks (JPNIC, 2019^[96]). It would be helpful to make aggregate totals of these notifications publicly available, along with the tool’s methodology, to help policy makers, researchers and industry assess the long-term evolution of routing hijacks⁴⁴. This would provide an interesting perspective, even noting its limitations⁴⁵. This effort could also be refined and extended to provide a baseline of routing incidents moving forward.

In addition to tracking the incidence of routing events, it is also important to track the implementation of routing security techniques, and their impact to mitigate routing incidents. For example, to monitor RPKI adoption, metrics which track ROA creation from the five RIRs⁴⁶ over time and ROV filtering by country over time are helpful (RIPE NCC, 2022^[97]; APNIC Labs, 2022^[62]). In addition, NIST’s RPKI Monitor tracks the global implementation of RPKI by providing time series data on several useful metrics (NIST, 2022^[73]). It is an encouraging sign that these metrics have time series data, which hopefully will continue with dedicated funding for these data collection efforts. Making this information publicly available may also provide a further benefit by motivating private sector actors (e.g. network operators) to adopt good practices and available techniques to uphold routing security.

However, the above metrics do not yet track the extent to which existing techniques decrease the likelihood of BGP incidents. As a technique is deployed more broadly, being able to measure its effectiveness to reduce routing incidents provides critical insights that can guide future action by both the technical community and policy makers. For policy makers, seeing the impact of implementing certain techniques will inform their policy approach. For industry stakeholders, seeing the effect of their efforts to improve routing security can justify the investment taken to implement them, or spur further research and development efforts in areas where routing vulnerabilities remain.

Another way to better understand the potential domestic impact of routing events can also take the form of observing national stakeholders’ actions to mitigate known routing vulnerabilities to evaluate risk. In Sweden, the communication regulator, the Swedish Post and Telecom Authority (PTS), has a national mandate to regulate and supervise providers of publicly available communication networks and services. Since 2014 and 2015 respectively, PTS has secondary legislation in force, in the form of more detailed provisions regarding the technical and organisational security measures for operational reliability

(i.e. measures to ensure available communication networks and services) and for confidential communications services (Hersaeus, 2021^[98]). Under this overarching mandate, PTS conducted an exercise to monitor and assess key domestic stakeholders' awareness of the vulnerabilities related to their use of the BGP protocol and the effectiveness of existing security measures to mitigate such vulnerabilities (Box 4).

Box 4. PTS Supervision on handling risks related to BGP

In October 2020, PTS initiated a supervision to assess providers' awareness of known vulnerabilities related to the use of BGP protocol and the security measures put in place in response to the identified risks and risk analysis of the security of communication services. PTS monitored the providers' security measures that contribute to improved availability, reliability and confidentiality in communication services, which ultimately contributes to more secure routing on the Internet. The largest domestic Internet service providers were put under supervision (Telia Company, Tele2 Sweden, Telenor Sweden and Hi3G Access) as well as one Internet Exchange Point provider, Netnod Internet Exchange. PTS has supervised security measures within the following areas: monitoring and detection capabilities and implementation of security measures to mitigate consequences of possible BGP incidents, for instance, using and providing accurate data in acknowledged IRRs, implementing appropriate filters, and deploying RPKI. Security recommendations from MANRS and ENISA's "7 Steps to shore up BGP" have been used as background material. This is the first supervision where PTS has monitored providers' security measures in this area.

Rationale for the exercise

The basis for carrying out this supervision was PTS's acknowledgement of BGP's fundamental importance to the functioning of communication services and its awareness of BGP's inherent vulnerabilities and the violations possible over BGP. PTS also recognises the possibility of severe impacts if BGP incidents are realised, considering especially the risks to data integrity and to the availability of confidential services. Moreover, PTS takes into account available data on routing incidents from MANRS, which show that BGP incidents occur daily on the Internet. From PTS' standpoint, it is important that communication providers take measures to secure BGP (i.e. external BGP) in order to contribute to a more secure Internet.

Results of the supervision and next steps

Over the course of the supervision, PTS has observed that Swedish providers take appropriate security measures concerning identified vulnerabilities and risks related to BGP. Overall, the results have been positive: PTS has found that the providers work in a systematic and continuous way, according to best practice, and that their security measures are based on appropriate risk analysis.

In the future, PTS might initiate further supervisory activities regarding risk analysis and the security measures taken to secure BGP.

Note: This box was informed by an informational interview with Erika Hersaeus, Senior Advisor in Cyber Security at PTS, conducted in September 2021 and subsequent written material provided thereafter.

Source: (Hersaeus, 2021^[98])

Promote awareness and deployment of routing security techniques

While network operators are the key stakeholders to ensure wide-scale adoption of routing security techniques, if domestic ISPs are not willing or able to implement the routing security technologies, governments can play a role to increase awareness and spur the implementation of existing tools. First, governments can request that the ISPs that transit their traffic implement routing good practices, including filtering, and adopt techniques such as RPKI (e.g., ROA creation and ROV filtering). Second, governments can also apply relevant good practice and security techniques in government-owned IP addresses and ASes. In the Netherlands, for example, the Dutch government added RPKI, both the creation of ROAs and ROV filtering, to its list of mandatory standards under the “comply or explain” regime that governs all government ICT projects (Forum Standaardisatie, 2019_[99]). In practice, this requires the Dutch government, including the central government, municipalities and provinces, to adopt RPKI for public IP addresses of government IT systems (Forum Standaardisatie, 2019_[99]). This serves to promote the awareness and adoption of RPKI in the broader community and creates incentives for the market to develop the supporting technical ecosystem for RPKI.

Third, governments can act as a resource to encourage other network operators and key stakeholders to accelerate the adoption of secure techniques. For example, the National Cybersecurity Center of Excellence (NCCoE) at NIST issued guidance focusing on the practical implementation of ROV filtering based on experience from NCCoE’s implementation in a testbed setting (Haag et al., 2019_[100]). Tailored practical guidance can act as a tool for governments to promote awareness and facilitate implementation of techniques by stakeholders. In certain cases where private sector deployment is lagging, governments could consider working more directly with key stakeholders, for instance with large ISPs who provide transit services, to implement routing techniques and best practice or extend incentives to motivate deployment.

Facilitate information sharing

Clear mechanisms to share information on routing incidents between different stakeholders, both domestically and on an international basis, could be further defined and supported. Frameworks to share such information make it easier for network operators to identify and solve root causes and learn from experiences to become better equipped to handle or mitigate similar situations in the future.

Information sharing could take place in existing structures within national, regional or sectoral Computer Emergency Response Teams (CERTs) or Information Sharing and Analysis Centres (ISACs), governmental agencies or regulatory bodies. For example, CERTs could play a useful role to facilitate information exchange and promote mitigation efforts, as they have a broad perspective and may be able to identify trends regarding routing incidents occurring locally. Safeguards to protect industrial interests from regulatory recrimination could be introduced to encourage participation and to protect informant individuals and companies from liability or reprisal. The goal of such information sharing is to promote awareness among industry, to learn more about common routing incidents that occur and their impact, and to determine appropriate mitigation actions to prevent their reoccurrence.

As one national example, Article 28 of Japan’s Telecommunications Business Law requires operators to report serious accidents to MIC (Government of Japan, 1979 (as amended)_[101]). Furthermore, MIC convenes “telecommunication accident verification meetings” following large communication events, including routing incidents. One such meeting was held following a 2017 route leak involving Google and several Japanese operators. This led to a revision of MIC’s “Standards for Safety and Reliability of Information and Telecommunications Networks” to include provisions related to routing and proposed further information sharing between operators during such incidents to determine cause and examine appropriate countermeasures (MIC, 2017_[102]). In September 2021, MIC issued a report considering possible revisions of the legal framework for reporting, which included a review of the communication accident reporting and verification system (MIC, 2021_[103]). The review proposes to improve the reporting system by clarifying the incidents that should be reported and simplifying administrative procedures to

reduce reporting burden on operators. The review also considered whether to establish incentives to encourage reporting (MIC, 2021_[103]). Based on the review's outcomes, MIC will begin adapting the reporting system and will consider possible incentives to encourage reporting by operators (MIC, 2021_[103]).

Along with more formalised approaches, multi-stakeholder initiatives to facilitate international information sharing are also important. Given that the Internet and therefore, the routing system, crosses borders, international cooperation is particularly important to facilitate cross-border information exchange on routing security. At the multi-stakeholder level, Network Operator Groups (NOGs) provide an arena to bring together various stakeholders and exchange on a variety of operational networking topics, including routing. These groups are organised at the national and regional level, are open to all interested parties (regardless of geography), and encourage information exchange, training and collaborative problem solving. Indeed, there is an active and vibrant history of informal cross-border cooperation at the network engineer level to understand and respond quickly to routing incidents observed on their networks (e.g. NOG email Listservs⁴⁷ provide one avenue to poll a highly-skilled technical community to jointly solve problems). Governments can support such multi-stakeholder initiatives by recognising their importance as key convening bodies, taking note of relevant outputs from national and regional forums (e.g. NOGs) and potentially sending relevant staff to participate and follow discussions.

Define a common framework with industry to improve routing security

Finally, governments can contribute to enhanced routing security by working with industry and technical experts to define a common framework to improve it. This framework would establish targeted actions for stakeholders to take to improve routing security according to a set time frame and define actions to promote awareness of the framework and gain participation (e.g. establish incentives or consider more binding requirements). While many governments currently regard routing security as a subset of digital security, more government attention is necessary to improve routing security within the broader digital security policy agenda.

Industry and the technical community are key stakeholders in routing security and have a vested interest in upholding a high level of security in the routing system and the Internet more broadly. Regulators and policy makers may not be aware of all implications or consequences that could arise from implementing certain technical solutions or the potential cost of implementation, which could be material (e.g. costs related to time and resources spent), or immaterial (e.g. risks to accessibility of service). Policy makers should also keep in mind that the routing system is global and the actions of one stakeholder cannot “fix” the problem. However, inappropriate actions of one stakeholder can negatively impact a much broader range of actors, including across borders. Rather, governments should aim to encourage local actors to promote global security through their own individual actions. Therefore, any policy guidance, initiative or incentive should recognise the role of the different stakeholders and the attack vectors that apply given their role in the routing system.

There are several possible approaches to develop this common framework. Currently, countries are exploring different approaches to routing security, ranging from flexible ad-hoc meetings and workshops to more formalised partnerships, to legal frameworks granting power to regulate on specific issues related to digital security.

In Brazil, for example, collaboration between policy makers and industry is a foundational element of Brazil's Internet Steering Committee (CGI.br). Its legislative mandate is to develop strategic guidance on various topics related to the Internet, including on network security (Government of Brazil, 2003_[104]). CGI.br has a multi-stakeholder structure and is composed of members from the technical community, governmental agencies, industry sectors and civil society (Government of Brazil, 2003_[104]). In 2017, CGI.br, in conjunction with the Brazilian Network Information Center (NIC.br) launched the “Program for a Safer Internet” to support the technical community on several issues, including to “reduce prefix hijacking, route leak, and source IP spoofing” (NIC.br and CGI.br, 2019_[105]). Under the program, NIC.br has worked

with domestic stakeholders on awareness building and training, developed educational materials and best practice, and engaged with industry to encourage adherence to MANRS recommendations, among other actions (NIC.br and CGI.br, 2019_[105]). In the United States, the NCCoE within NIST is a public-private partnership that convenes stakeholders from government, industry and academia to address cybersecurity issues, with one project dedicated to secure inter-domain routing (NCCoE, n.d._[106]; NCCoE, n.d._[107]).

Moving to more formalised guidelines, Japan developed a set of voluntary guidelines broadly aimed at providing stable and reliable communication, called the "Standards for Safety and Reliability of Information and Telecommunications Networks", as mentioned above (Ministry of Posts and Telecommunications, 1987 (as amended)_[108]).⁴⁸ Some of the measures included in the standards concern routing security, following consultations with industry regarding the aforementioned 2017 routing incident involving Google. The voluntary measures include putting in place filtering best practices and ensuring adequate checks of router configuration to avoid human error. The standards further recommend that operators obtain as much information about a given BGP event as possible and share information among themselves (Ministry of Posts and Telecommunications, 1987 (as amended)_[108]).

Other countries have a more legal and binding approach, whereby there may be broad legislation or security guidelines in place, with the option of issuing secondary measures as needed. Switzerland, for example, does not have specific regulations related to routing security, but rather has general guidelines for communication service providers that aim to establish a minimum level of security to ensure the reliability and availability of communication infrastructure and services (Federal Office of Communications (OFCOM), 2009_[109]). The Swiss Federal Office of Communications (OFCOM) "may issue the technical and administrative security regulations and declare internationally harmonised standards concerning security and availability of telecommunications infrastructures and services to be binding" (Swiss Federal Council, 2007 [Status as of 1 July 2022]_[110]).⁴⁹

The Electronic Communications Law No. 5 809 of the Republic of Türkiye gives legal power to the national communication regulator, the Information and Communication Technologies Authority (BTK), to undertake tasks related to digital security (Republic of Türkiye, 2008_[111]). Under this mandate, BTK may impose obligations on network operators to ensure network security against unauthorised access (see Art. 12(2) (j)) (Republic of Türkiye, 2008_[111]). The legislation on network and information security of the communication sector introduces general obligations on operators, including measures applied to certain operators to protect network elements against cyberattacks, including routers. Under this framework, a Cyber Incident Response Team was formed within BTK to promote coordination and facilitate information sharing within the industry (BTK, 2017_[112]).

Finland has formalised some commonly accepted best practice into secondary legislation on "Information security in telecommunications operations", which details requirements for the protection of ISP interfaces, including some that outline measures to uphold basic security of the BGP (Finnish Transport and Communications Agency (Traficom), 2015_[113]). In particular, there are clauses related to the prevention of IP traffic in interconnection and customer interfaces, which both refer to common practices for filtering route announcements (Finnish Transport and Communications Agency (Traficom), 2015_[113]).⁵⁰

As the above examples show, OECD countries have taken different approaches related to routing security, ranging from collaborative efforts and voluntary measures, to more formalised legal obligations and requirements. Regardless of the specific approach, defining a common framework on routing security can structure action by defining the actions different stakeholders should take to improve routing security collectively.

Concluding Remarks

Given the routing system's importance to the Internet's overall functioning, ensuring its security is crucial. Border Gateway Protocol (BGP), which networks use to exchange routing information, is integral to the functioning of the routing system but fraught with vulnerabilities. These security issues have been well understood by the technical community for many years, but the problem of routing security persists. The overarching question is why, which the report aimed to address by examining, first, the scope and scale of routing incidents, and second, the techniques that have been proposed and their effectiveness at addressing the challenges of routing security. Following this evaluation, the report considered the role of policy makers to enhance security in the routing system and proposed possible policy tools.

There have been many reported routing incidents (BGP leaks and hijacks) that affected availability, confidentiality and integrity of communications, demonstrating the routing system's insecurity and fragility. However, when attempting to observe the trend of routing incidents over time, there is a lack of publicly available time series data (e.g. metrics tracked over time) to evaluate the total number of routing incidents globally. This hampers the ability of policy makers and other stakeholders to understand the evolution of routing incidents and their impact.

While several techniques to address routing's security challenges have been proposed by industry and the technical community, no single technique alone meets routing's various challenges, either by its design or in practice. Furthermore, there are limited sources to measure how the deployment of security techniques decreases the incidence of routing events. In addition, there are several challenges that impede the deployment of available techniques to improve routing.

In light of the current industrial and technical landscape of routing security, there are several policy tools that governments can consider to support the overall improvement of routing security of the Internet. First, governments could consider funding publicly available data collection and measurement efforts to track BGP incidents over time. It is important for funding of data collection efforts to be long-term, consistent and awarded to third-party actors. As deployment of some techniques increases (e.g. RPKI), further funding of efforts to track their effectiveness to decrease the incidence of routing events would help to build a more comprehensive evidence base. As a complementary exercise, policy makers could consider establishing more formal mechanisms to allow industry stakeholders to share information about routing incidents.

Furthermore, governments could promote awareness and deployment of routing good practices and current techniques to improve routing security and implement relevant techniques in government-owned IP addresses and ASes. Finally, policy makers could work with industry and technical experts to define a common framework to improve routing security. More specific government attention could be helpful to encourage industry to take action to improve routing security and more proactive policies can help governments ensure that stakeholders take appropriate action to promote routing security.

Annex A. Technical description of techniques to improve routing security

RPKI

There are two pieces critical to the success of RPKI for origin validation: the creation of Route Origin Authorisations (ROAs) and the validation of ROAs through Route Origin Validation (ROV) and subsequent filtering on any entries deemed to be invalid.

First, prefix holders must create Route Origin Authorisations (ROAs) for their prefixes. ROAs authorise an AS to originate a prefix (called the “origin AS”) and may also include maximum prefix length of the IP prefix assets to be announced. The entity signing the ROA must sign with the private key corresponding to the public key included in the end-entity (EE) certificate, which validates that the signing entity holds the prefix specified in the ROA (Lepinski and Kent, 2012_[56]). The main function of an EE certificate is to verify that the ROA relates to resources defined in the certificate, such that there is one EE certificate per signed object (e.g., ROA) (Lepinski and Kent, 2012_[56]). If an EE certificate is revoked or becomes invalid, the corresponding ROA would no longer be valid due to the EE certificate’s role to validate that ROA (Lepinski and Kent, 2012_[56]). The revoked EE certificate means that the corresponding ROA is no longer valid and would then return an “unknown” response from the RPKI validator.

Second, ASes must configure their routers to validate the BGP-advertised prefixes against the list of ROAs via a local RPKI validator, through a process called Route Origin Validation (ROV). Based on the result from the validator (either valid, invalid, or unknown), the BGP router can be configured to give preference to valid ROAs (authorised origins) and/or to drop routes that are reported to be invalid through the ROV process. A “valid” response means the route announcement matches the corresponding ROA for the relevant prefix and origin ASN, and the announced prefix length does not exceed the maximum prefix length, if specified in the ROA (Durand, 2020_[76]). In an “invalid” response, the BGP announcement does not match one or more of the attributes in the ROA for that prefix; either the origin ASN is different and/or the prefix length is longer than the max length specified (Durand, 2020_[76]). An “unknown” response means that the validator was unable to find a ROA for that prefix.

BGPsec

BGPsec and RPKI are envisioned to go hand-in-hand; RFC 7454 presents them as the two components of secure inter-domain routing, with RPKI handling origin validation and BGPsec addressing path validation (Durand, Pepelnjak and Doering, 2015_[49]). Indeed, RFC 8205 on the BGPsec protocol clearly states that “BGPsec relies on the Resource Public Key Infrastructure (RPKI) certificates that attest to the allocation of AS number and IP address resources” (Lepinski and Sriram, 2017_[77]). Similar to RPKI, the entity signing the BGPsec “Update” message must sign with the private key corresponding to the public key included in its end-entity (EE) certificate. Since routers are issuing the BGPsec Update message in BGPsec, the RPKI end-entity certificates are issued to routers within an AS and are called “router certificates” (Reynolds, Turner and Kent, 2017_[114]). This validates that the signing entity (router) holding the associated private key is authorised to issue secure route announcements for the AS(es) identified in the router certificate (Reynolds, Turner and Kent, 2017_[114]). However, a router can only send an “Update” message to one AS

at a time (unlike in normal BGP), as the BGPsec “Update” message includes the AS to which the message is being sent as a critical piece to ensure the security of the path (Lepinski and Sriram, 2017^[77]). An “Update” message can also only advertise a single prefix at a time, so a separate message must be sent for each specific prefix (Lepinski and Sriram, 2017^[77]).

A router receiving the “Update” message can trust that each AS that propagated the announcement intended to send it to the subsequent AS in the path and that the AS path reflects the order in which the announcement was propagated (Huston, 2021^[55]). This means that each “hop” in the path is protected and that the route was sent to its intended recipients. BGPsec defends against path hijacks by including the AS a router is sending the update to, making it difficult for a malicious AS to inject itself falsely in the AS path, as it would have had to receive the “Update” message legitimately from a peer first. Another aspect of BGPsec’s defence is protecting each hop with a digital signature to prevent a malicious actor from changing the path falsely, for instance by shortening it (e.g. deleting ASes in the path). A shortened fake path would not pass BGPsec validation checks.

ASPA

ASPA is an alternative response to path validation, providing “path plausibility”, and a way to detect route leaks and hijacks. At a high level, the solution proposes a new digitally signed object, an ASPA, building off existing RPKI architecture. A customer AS would sign an ASPA to define and authorise a set of Provider ASes to propagate its route announcements. The set of Provider ASes would include a provider’s upstream provider(s) and peers.

In the validation process, the first step is to validate that the provider AS (of the route announcement) was authorised to propagate a given customer AS’ routes. Validation would occur by querying a local cache of signed (“cryptographically valid”) ASPAs, retrieving any ASPAs with the customer AS in question, and checking whether the provider AS is included in the customer AS’ list of Provider ASes which it authorised to propagate its route announcements (Azimov et al., 2022^[79]). The next step is to verify the AS Path, by outlining processes to verify routes received from a customer (i.e. upstream paths), a provider (e.g. downstream paths), and from a route server (Azimov et al., 2022^[79]). If the outcome results in an invalid response, then the route should be rejected.

SCION: Scalability, Control and Isolation on Next-Generation Networks

SCION is a network architecture developed by researchers from the ETH Zürich that proposes a new way to route traffic on the Internet, without using BGP. It aims to solve many of the existing challenges to routing security, reliability and performance, by isolating routing failures by grouping ASes into separate domains called “isolation domains” (ISDs), giving more control over the routes that packets take and ensuring explicit trust for end-to-end communication.

A key element of the architecture is the introduction of “isolation domains” (ISDs), which are groups of ASes that share a routing plane. This serves to limit the effect of any routing misconfiguration, failure or malicious behaviour. Members of an ISD share a mutual level of trust (i.e. they are within the same jurisdiction), agree on local roots of trust within the ISD and interconnect with other ISDs in order to provide global connectivity. In addition, the forwarding state is authenticated. These are all critical elements for secure end-to-end communication. Another key element is the introduction of path-aware communications, which gives users and ISPs a choice in how their packets are sent, from the start to the endpoint, and natively provides multipath communication. The ability to choose the end-to-end path negates the risk that traffic is routed via a circuitous route or through a suspicious actor without the sender’s knowledge (as seen in Box 1). Being aware of all available paths also allows for performance optimisation, by providing immediately available backup paths in case of a network failure, and resilience against DDoS attacks.

Importantly, all of the BGP incidents detailed in this paper would be negated, as BGP would not be used in a full SCION deployment. While attacks to end hosts may be possible, SCION offers additional mechanisms to authenticate and filter traffic and guarantees traffic for critical applications.

References

- Alaettinoglu, C. et al. (1999), *Routing Policy Specification Language (RPSL) [RFC 2622]*, [117]
<https://datatracker.ietf.org/doc/html/rfc2622>.
- APNIC (2022), *ROA data by country (%)*, <https://stats.labs.apnic.net/roas> (accessed on [58]
 1 February 2022).
- APNIC Labs (2022), *I-Rov filtering rate by country (%)*, <https://stats.labs.apnic.net/rpki> (accessed [62]
 on 28 June 2022).
- APNIC Labs (2022), *Use of RPKI Validation for World (XA)*, <https://stats.labs.apnic.net/rpki> [63]
 (accessed on January 2022).
- ARTEMIS (2020), *Home*, <https://bgpartemis.org/>. [83]
- Azimov, A. et al. (2022), *Verification of AS_PATH Using the Resource Certificate Public Key [79]
 Infrastructure and Autonomous System Provider Authorization draft-ietf-sidrops-aspa-
 verification-08*, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification>
 (accessed on 2 September 2022).
- Balakrishnan, H. (2009), *How Youtube was "Hijacked"*, [31]
<http://web.mit.edu/6.02/www/s2012/handouts/youtube-pt.pdf>.
- BGPalerter (GitHub) (2021), *BGPalerter*, <https://github.com/nttgin/BGPalerter#documentation>. [119]
- BGPalerter (GitHub) (2021), *Who is using BGPalerter*, [81]
<https://github.com/nttgin/BGPalerter/blob/main/docs/friends.md>.
- Birge-Lee, H. et al. (2018), *Bamboozling Certificate Authorities with BGP*, [38]
<https://www.usenix.org/conference/usenixsecurity18/presentation/birge-lee>.
- Borkenhagen, J. (2019), *AT&T/as7018 now drops invalid prefixes from peers*, [66]
<https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html>.
- BTK (2017), *Sectoral Cyber Incident Response Team [unofficial translation]*, [112]
<https://www.btk.gov.tr/sektorel-siber-olaylara-mudahale-ekibi> (accessed on 19 October 2021).
- Candela, M. (2020), *Easy BGP monitoring with BGPalerter*, [82]
https://labs.ripe.net/author/massimo_candela/easy-bgp-monitoring-with-bgpalerter/.
- Chandra, R., P. Traina and T. Li (1996), *BGP Communities Attribute [RFC 1997]*, [51]
<https://datatracker.ietf.org/doc/html/rfc1997> (accessed on 21 January 2022).

- Cho, S. et al. (2019), *BGP hijacking classification*, <https://doi.org/10.23919/TMA.2019.8784511>. [17]
- Cisco Certified Expert (2021), *Specifying Availability Requirements*, <https://www.ccexpert.us/network-design-2/specifying-availability-requirements.html>. [3]
- Cisco Crosswork Cloud (2022), *BGPStream*, <https://bgpstream.crosswork.cisco.com/>. [21]
- Cloudflare (2021), *Is BGP safe yet?*, <https://isbgpsafeyet.com/#faq> (accessed on 14 October 2021). [64]
- Cloudflare (2020), *Cloudflare System Status: Cloudflare Network and Resolver Issues*, https://www.cloudflarestatus.com/incidents/b888fyhbbyg8?_gl=1*yuurwf*_ga*MTM3MTYyMDQ4Ni4xNjM1MjUyNzly*_gid*MTU1Njg1NDc4LjE2MzUyNTI3MjI. (accessed on 22 October 2021). [4]
- Cloudflare (n.d.), *What is a protocol? Network protocol definition*, <https://www.cloudflare.com/learning/network-layer/what-is-a-protocol/>. [118]
- CODE BGP (2021), *Home*, <https://codebgp.com/>. [120]
- DE-CIX (2021), *Frankfurt route server guide*, <https://www.de-cix.net/en/locations/frankfurt/route-server-guide>. [53]
- DE-CIX (2021), *RPKI at the DE-CIX route servers*, <https://www.de-cix.net/en/resources/service-information/route-server-guides/rpki>. [69]
- Döring, G. (2018), *BGP communities 101*, <https://media.ccc.de/v/denog10-13-bgp-communities-101> (accessed on 21 January 2022). [52]
- Durand, A. (2020), *Resource Public Key Infrastructure (RPKI) Technical Analysis*, <https://www.icann.org/en/system/files/files/octo-014-02sep20-en.pdf> (accessed on 2021). [76]
- Durand, J., I. Pepelnjak and G. Doering (2015), *BGP Operations and Security*, <https://www.rfc-editor.org/info/rfc7454>. [49]
- ENISA (2019), *7 Steps to shore up BGP*, <https://www.enisa.europa.eu/publications/7-steps-to-shore-up-bgp>. [18]
- Federal Office of Communications (OFCOM) (2009), *Richtlinien zur Sicherheit und Verfügbarkeit von Fernmeldeinfrastrukturen und -diensten*, <https://www.bakom.admin.ch/bakom/de/home/telekommunikation/fernmeldedienstanbieter/richtlinien-zur-sicherheit-und-verfuegbarkeit-von-fernmeldeinfra.html> (accessed on 10 October 2021). [109]
- Federal Register (2022), "Secure Internet Routing", *National Archives: the daily journal of the United States Government*, <https://www.federalregister.gov/documents/2022/03/11/2022-05121/secure-internet-routing> (accessed on 24 March 2022). [5]
- Ferguson, P. and D. Senie (2000), *RFC 2827 (BCP 38): Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing*, <https://www.rfc-editor.org/info/rfc2827>. [50]
- Finnish Transport and Communications Agency (Traficom) (2015), *Regulation on information security in telecommunications operations [unofficial translation]*, https://www.finlex.fi/data/normit/44046/M67A_2015_EN.pdf (accessed on September 2021). [113]

- Forum Standaardisatie (2019), *RPKI*, <https://www.forumstandaardisatie.nl/open-standaarden/rpki> [99]
(accessed on 2 February 2022).
- Fuller, V. and T. Li (2006), *RFC 4632: Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*, IETF, <https://datatracker.ietf.org/doc/html/rfc4632> [7]
(accessed on 12 October 2021).
- Gao, L. (2001), *On inferring autonomous system relationships in the Internet*, pp. 733-745, [11]
<https://doi.org/10.1109/90.974527> (accessed on 19 October 2021).
- Georgia Tech (2022), *Global Routing Intelligence Platform (GRIP)*, [25]
<https://grip.inetintel.cc.gatech.edu> (accessed on 1 March 2022).
- Goodin, D. (2017), “Suspicious” event routes traffic for big-name sites through Russia, [40]
<https://arstechnica.com/information-technology/2017/12/suspicious-event-routes-traffic-for-big-name-sites-through-russia/>.
- GovCERT.ch (2015), *Cantonal IP space in Switzerland hijacked by spammers*, [46]
<https://www.govcert.ch/blog/cantonal-ip-space-in-switzerland-hijacked-by-spammers/>.
- Government of Brazil (2003), *Decree No. 4,829, of September 3, 2003*, [104]
<https://cgi.br/pagina/decretos/108/> (accessed on 11 October 2021).
- Government of Japan (1979 (as amended)), *Telecommunications Business Law, 1979 Law No. 86 [unofficial translation]*, [101]
<https://elaws.e-gov.go.jp/document?lawid=359AC0000000086>
(accessed on 19 October 2021).
- Haag, W. et al. (2019), *Protecting the Integrity of Internet Routing: L Border Gateway Protocol (BGP) Route Origin Validation*, [100]
<https://doi.org/10.6028/NIST.SP.1800-14> (accessed on 18 October 2021).
- Hawkinson, J. and T. Bates (1996), *RFC 1930 (Best Current Practice): Guidelines for creation, selection, and registration of an Autonomous System (AS)*, [115]
<https://datatracker.ietf.org/doc/html/rfc1930> (accessed on 2021).
- Hersaeus, E. (2021), *Interview with Erika Hersaeus, Senior Advisor in Cyber Security at PTS*. [98]
- Hurricane Electric (n.d.), *Hurricane Electric route filtering algorithm*, [68]
<https://routing.he.net/algorithm.html>.
- Huston, G. (2021), *A survey on securing inter-domain routing: Part 1*, [19]
<https://blog.apnic.net/2021/07/08/a-survey-on-securing-inter-domain-routing-part-1/>.
- Huston, G. (2021), *A survey on securing inter-domain routing: Part 2*, [55]
<https://blog.apnic.net/2021/07/09/a-survey-on-securing-inter-domain-routing-part-2/>.
- Huston, G. (2021), *Measuring ROAs and ROV*, [61]
<https://blog.apnic.net/2021/03/24/measuring-roas-and-rov/>.
- IETF (n.d.), *RFCs*, <https://www.ietf.org/id/draft-ietf-sidrops> (accessed on 8 October 2021). [80]
- Internet Routing Registry (2018), *Internet Routing Registry (IRR)*, <http://www.irr.net?index.html>. [54]

- ISOC (2021), *Arista, Cisco, Huawei, Juniper Networks, and Nokia Launch New MANRS Equipment Vendor Program to Improve Routing Security Worldwide*, <https://www.internetsociety.org/news/press-releases/2021/arista-cisco-huawei-juniper-and-nokia-launch-new-manrs-equipment-vendor-program-to-improve-routing-security-worldwide/> (accessed on 21 September 2021). [89]
- Janardhan, S. (2021), *More details about the October 4 outage*, <https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/> (accessed on 10 October 2021). [2]
- Jones, E. (2004), *Operational Security Requirements for Large Internet Service Providers (ISP) IP Network Infrastructure [RFC 3871]*, <https://datatracker.ietf.org/doc/html/rfc3871> (accessed on 19 January 2022). [47]
- JPNIC (2019), *Japan Network Information Center (JPNIC) route magistrate [unofficial translation]*, <https://www.nic.ad.jp/ja/ip/irr/jpnic-keirobugyou.html> (accessed on 18 October 2021). [96]
- Klein, T. (2021), *Wenn IP-Adressen entführt werden*, <https://www.spiegel.de/netzwelt/web/wenn-ip-adressen-entfuehrt-werden-a-c9ae6d81-671c-4387-a241-be12251eabf8>. [33]
- LEE, H. (2021), *Interview with Prof. Heejo LEE of Korea University on August 25th*. [94]
- Lepinski, M. and S. Kent (2012), *An Infrastructure to support secure Internet routing*, <https://tools.ietf.org/html/rfc6480>. [56]
- Lepinski, M. and K. Sriram (2017), *BGPsec Protocol Specification [RFC 8205]*, <https://datatracker.ietf.org/doc/html/rfc8205>. [77]
- Levy, M. (2018), *RPKI - The required cryptographic upgrade to BGP routing*, <https://blog.cloudflare.com/rpki/>. [57]
- Livingood, J. (2021), *Improved BGP routing security adds another important layer of protection to online networks (Comcast)*, <https://corporate.comcast.com/stories/improved-bgp-routing-security-adds-another-layer-of-protection-to-network>. [67]
- Madory, D. (2021), *Facebook's historic outage, explained*, <https://www.kentik.com/blog/facebooks-historic-outage-explained/> (accessed on 12 October 2021). [1]
- MANRS (2022), *MANRS Observatory: Overview*, <https://observatory.manrs.org/#/overview>. [27]
- MANRS (2021), *About MANRS*, <https://www.manrs.org/about/>. [88]
- MANRS (2021), *MANRS for Network Operators*, <https://www.manrs.org/isps/>. [121]
- MANRS (n.d.), *MANRS Observatory: About*, <https://observatory.manrs.org/#/about> (accessed on 14 October 2021). [116]
- McDaniel, T., J. Smith and M. Schuchard (2021), *Flexsealing BGP against route leaks: Peerlock active measurement and analysis*, <https://doi.org/10.14722/ndss.2021.23080>. [87]
- MIC (2021), *IP Network Facilities Subcommittee of the Department on Information and Communications Technology under the Information and Communications Council Fifth report*. [103]

- MIC (2017), *Verification report on large-scale Internet connection failure that occurred in August 2017 [unofficial translation]*, https://www.soumu.go.jp/main_content/000523153.pdf (accessed on 18 October 2021). [102]
- MIC (n.d.), *Research and development related to detection, recovery, and prevention of route hijacking [unofficial translation]*, https://www.soumu.go.jp/main_content/000394358.pdf (accessed on 18 October 2021). [95]
- Microsoft (2021), *How to trust your neighbors: Securing the edge network of a large cloud (submitted paper)*, SIGCOMM. [85]
- Ministry of Posts and Telecommunications (1987 (as amended)), *Information and Communication Network Safety and Reliability Standards [unofficial translation]*, https://www.soumu.go.jp/main_content/000755047.pdf (accessed on 18 October 2021). [108]
- Murphy, S. (2006), *BGP Security Vulnerabilities Analysis [RFC 4272]*, <https://datatracker.ietf.org/doc/html/rfc4272> (accessed on 17 October 2021). [14]
- NCCoE (n.d.), *About the Center*, <https://www.nccoe.nist.gov/about-the-center> (accessed on 19 October 2021). [106]
- NCCoE (n.d.), *Secure Inter-Domain Routing*, <https://www.nccoe.nist.gov/projects/building-blocks/secure-inter-domain-routing> (accessed on 19 October 2021). [107]
- NIC.br and CGI.br (2019), *Internet security in Brazil - MANRS: Program for a safer Internet [Powerpoint]*, <https://bcp.nic.br/i+seg/assets/pdfs/Program-Safer-Internet-Futuras-Tecnologias-SPO.pdf> (accessed on 11 October 2021). [105]
- NIST (2022), *NIST RPKI Monitor, version 2.0*, <https://rpki-monitor.antd.nist.gov/> (accessed on 20 September 2021). [73]
- NIST (2022), *NIST RPKI Monitor: NIST RPKI Monitor 2.0, Methodology and User's Guide*, <https://rpki-monitor.antd.nist.gov/Methodology> (accessed on 14 October 2021). [74]
- NLNet Labs (2022), *RPKI Maps (Coverage, Accuracy)*, <https://rpki-maps.nlnetlabs.nl/ui/world.html> (accessed on 1 February 2022). [59]
- NTT (2016), *NTT Peer Locking: Deployment of NTT "Peer Locking" route leak prevention mechanism*, http://instituut.net/~job/peerlock_manual.pdf. [86]
- OECD (2013), *OECD Communications Outlook 2013*, OECD Publishing, Paris, https://doi.org/10.1787/comms_outlook-2013-en. [8]
- Pfaff, R. (2021), *Lumen Enhances Routing Security With Resource Public Key Infrastructure (RPKI)*, <https://blog.lumen.com/lumen-enhances-routing-security-with-resource-public-key-infrastructure-rpki/> (accessed on 23 September 2021). [75]
- Poinsignon, L. (2018), *BGP leaks and cryptocurrencies*, <https://blog.cloudflare.com/bgp-leaks-and-crypto-currencies/>. [44]
- Postel, J. (1980), *RFC 760: DOD Standard - Internet Protocol*, <https://datatracker.ietf.org/doc/html/rfc760> (accessed on 12 October 2021). [6]
- Qrator (2021), *Radar by Qrator*, <https://radar.qrator.net/>. [84]

- Qrator Labs (2022), *Q4 2021 DDoS attacks and BGP incidents*, https://blog.qrator.net/en/q4-2021-ddos-attacks-and-bgp-incidents_153/ (accessed on 1 February 2022). [29]
- Qrator Labs (2021), *Q1 2021 DDoS attacks and BGP incidents*, https://blog.qrator.net/en/q1-2021-report_129/ (accessed on 15 October 2021). [28]
- Rekhter, Y. (1991), *Experience with the BGP Protocol*, <https://www.rfc-editor.org/rfc/rfc1266.txt>. [9]
- Reporting and Analysis Centre for Information Assurance (MELANI) (2019), *Information Assurance: Situation in Switzerland and internationally*, <https://www.ncsc.admin.ch/ncsc/en/home/dokumentation/berichte/lageberichte/halbjahresbericht-2019-1.html>. [41]
- Republic of Türkiye (2008), *Elektronik haberleşme kanunu [Electronic Communications Law]*, <https://www.mevzuat.gov.tr/mevzuat?MevzuatNo=5809&MevzuatTur=1&MevzuatTertip=5>. [111]
- Reynolds, M., S. Turner and S. Kent (2017), *A profile for BGPsec router certificates, certificate revocation lists, and certification requests (RFC 8209)*, <https://datatracker.ietf.org/doc/html/rfc8209> (accessed on 7 October 2021). [114]
- RIPE Labs (2019), *RPKI Test*, https://labs.ripe.net/author/nathalie_nathalie/rpki-test/ (accessed on 14 October 2021). [65]
- RIPE NCC (2022), *Resource Certification (RPKI) Statistics, Number of ROAs*, <http://certification-stats.ripe.net/?type=roa-v4u>. [97]
- RIPE NCC (2022), *RIR Trust Anchor Statistics*, <https://www.ripe.net/manage-ips-and-asns/resource-management/rpki/rir-trust-anchor-statistics>. [60]
- Robachevsky, A. (2021), *New Data Source, Feedback Loop Enhance MANRS Observatory*, <https://www.manrs.org/2021/07/new-data-source-feedback-loop-enhance-manrs-observatory/> (accessed on 16 August 2021). [26]
- Robachevsky, A. (2019), *Routing security - getting better, but no reason to rest!*, <https://www.manrs.org/2019/02/routing-security-getting-better-but-no-reason-to-rest/>. [23]
- Robachevsky, A. (2018), *14,000 Incidents: A 2017 Routing Security Year in Review*, <https://www.internetsociety.org/blog/2018/01/14000-incidents-2017-routing-security-year-review/>. [24]
- SCION (2021), *SCION: Scalability, Control, and Isolation on Next-Generation Networks*, <https://www.scion-architecture.net/>. [91]
- SCIONLab (2020), *Welcome to SCIONLab*, <https://www.scionlab.org/>. [92]
- Seattle Internet Exchange (SIX) (n.d.), *Route server details*, <https://www.seattleix.net/route-servers>. [70]
- Sharma, A. (2021), *Major BGP leak disrupts thousands of networks globally*, <https://www.bleepingcomputer.com/news/security/major-bgp-leak-disrupts-thousands-of-networks-globally/>. [35]
- Siddiqui, A. (2021), *A major BGP route leak by AS55410 - Vodafone Idea Ltd*, <https://blog.apnic.net/2021/04/26/a-major-bgp-route-leak-by-as55410/>. [34]

- Siddiqui, A. (2021), *APNIC Blog: BGP, RPKI, and MANRS: 2020 in review*, [22]
<https://blog.apnic.net/2021/02/05/bgp-rpki-and-manrs-2020-in-review/>.
- Siddiqui, A. (2019), *Public DNS in Taiwan the latest victim to BGP hijack*, [42]
<https://www.manrs.org/2019/05/public-dns-in-taiwan-the-latest-victim-to-bgp-hijack/>.
- Siddiqui, A. (2018), *What happened? The Amazon Route 53 BGP hijack to take over Ethereum cryptocurrency wallets*, [43]
<https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/>.
- Smith, B. and J. Garcia-Luna-Aceves (1996), *Securing the Border Gateway Routing Protocol*, [15]
<https://apps.dtic.mil/sti/pdfs/ADA461684.pdf>.
- Sriram, K. et al. (2016), *Problem Definition and Classification of BGP Route Leaks*, [16]
<https://tools.ietf.org/html/rfc7908>.
- Stone, B. (2008), "Pakistan cuts access to YouTube worldwide", *The New York Times*, [30]
<https://www.nytimes.com/2008/02/26/technology/26tube.html> (accessed on 19 January 2022).
- Strickx, T. (2019), *How Verizon and a BGP optimizer knocked large parts of the Internet offline today*, [36]
<https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-today/>.
- Sun, Y. et al. (2021), "Securing Internet Applications from Routing Attacks", *Communications of the Association for Computing Machinery (ACM)*, [20]
<https://cacm.acm.org/magazines/2021/6/252822-securing-internet-applications-from-routing-attacks/fulltext>.
- Swiss Federal Council (2007 [Status as of 1 July 2022]), *Ordinance on Telecommunication Services (OTS) [784.101.1]*, [110]
<https://www.fedlex.admin.ch/eli/cc/2007/166/en> (accessed on 25 March 2022).
- Swiss National Bank (2021), *SNB and SIX launch the communication network Secure Swiss Finance Network: Helping to strengthen cyber resilience in the Swiss financial sector*, [93]
https://www.snb.ch/en/mmr/reference/pre_20210715/source/pre_20210715.en.pdf.
- Team Cymru (2022), *The Bogon Reference*, [48]
<https://team-cymru.com/community-services/bogon-reference/> (accessed on 19 January 2022).
- Testart, C. et al. (2020), *To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today*, [72]
https://catalog.caida.org/details/paper/2020_filter_not_filter.
- Timberg, C. (2015), *The long life of a quick 'fix': Internet Protocol from 1989 leaves data vulnerable to hijackers*, *The Washington Post*, [13]
<https://www.washingtonpost.com/sf/business/2015/05/31/net-of-insecurity-part-2/>.
- Toonk, A. (2019), *Twitter, Tweet 24 June 2019*, [37]
<https://twitter.com/atoonk/status/1143139749915320321>.
- Toonk, A. (2017), *Popular destinations rerouted to Russia*, [39]
<https://www.bgpmon.net/popular-destinations-rerouted-to-russia/>.

- Toonk, A. (2015), *Massive route leak causes Internet slowdown*, [32]
<https://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/>.
- Toonk, A. (2014), *Turkey hijacking IP addresses for popular global DNS providers*, [45]
<https://www.bgpmon.net/turkey-hijacking-ip-addresses-for-popular-global-dns-providers/>.
- Wan, A. (2021), *Why should you care about routing security? Introducing MANRS Primers*, [90]
<https://www.manrs.org/2021/07/why-should-you-care-about-routing-security-introducing-manrs-primers/>.
- Weller, D. and B. Woodcock (2013), “Internet Traffic Exchange: Market Developments and Policy Challenges”, *OECD Digital Economy Papers*, No. 207, OECD Publishing, Paris, [10]
<https://doi.org/10.1787/5k918gpt130q-en>.
- White, R. (2015), *BGPSEC: Leaks and Leaks*, https://packetpushers.net/bgpsec-leaks-leaks/?doing_wp_cron=1656663306.9865689277648925781250 (accessed on 1 July 2022). [78]
- Woodcock, B. and G. Upadhaya (2005), *AS-Path analysis to test claims of “Tier 1” status*, Packet Clearing House, <https://www.pch.net/resources/Papers/testing-tier1-status/testing-tier1-status.pdf> (accessed on 19 October 2021). [12]
- YYCIX (n.d.), *BGP sessions default configuration*, <https://yycix.ca/communities.html>. [71]

End Notes

¹ A protocol is an agreed set of rules to standardise data's format and its processing, which enables devices to talk to one another (Cloudflare, n.d.^[118]).

² OECD elaboration on 2021 incidents at Facebook based on the following sources: [19 March 2021] <https://downdetector.com/status/facebook/news/376207-problems-at-facebook/>, <https://www.reuters.com/article/us-facebook-outages-idUSKBN2BB232>; [8 April 2021] <https://downdetector.com/status/facebook/news/381098-problems-at-facebook/>, <https://www.theverge.com/2021/4/8/22374499/facebook-instagram-outage-down> ; [10 June 2021] <https://www.independent.co.uk/life-style/gadgets-and-tech/facebook-instagram-whatsapp-messenger-down-outage-reports-b1862990.html>, <https://downdetector.com/status/facebook/news/393018-problems-at-facebook/>; [3 July 2021] <https://www.reuters.com/business/media-telecom/facebook-instagram-down-thousands-users-downdetector-2021-07-03/>, <https://downdetector.com/status/facebook/news/397607-problems-at-facebook/>; [4 October 2021] (Madory, 2021^[11]). Please note that not every incident was necessarily a global service unavailability; for instance, some incidents may have impacted Facebook users in certain parts of the world, not everywhere.

³ OECD elaboration on 2020 incidents at Cloudflare based on the following sources: [1 April 2020] <https://www.thousandeyes.com/blog/rostelecom-route-hijack-highlights-bgp-security>; [17 July 2020] (Cloudflare, 2020^[4]); [30 August 2020] <https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage/>. Please note these incidents do not necessarily imply global service unavailability; some incidents may have impacted parts of Cloudflare's network (e.g., in some geographies, for certain clients).

⁴ OECD elaboration on 2021 incidents at Cloudflare based on the following sources: [17 March 2021] <https://mfmbpzc.share.thousandeyes.com/view/internet-insights/?roundId=1615993200&metric=interfaces&scenarioid=outageTraffic&filters=N4lgZglgNgLgpgJwM4gFyglYActQgYwxqgHsA7AERIFsMlyAVDAczQG0BdAGhG1wKKkyAeQCuMFnADK%2BEIjjsQAQQBiKgKIBhBuooB9HVIZ7hAJT0BJAHJ7NSHkoAywgOlhuvJGXYBGAMwBAKweMAgYYJD4jiSExOTsACx%2BAJx%2BAEw%2BAOxcgQAcAAzJuQkJIWERBGISzNky8oqqGtq6BupGJubWtvZOru4Avv1AA>; [19 March 2021] <https://www.cloudflarestatus.com/incidents/5bmpy4zj9nd9>; [3 May 2021] <https://www.thousandeyes.com/blog/internet-report-episode-37> ; [11 June 2021] <https://www.cloudflarestatus.com/incidents/0cvlzpvwg251>. Please note these incidents do not necessarily imply global service unavailability; some incidents may have impacted parts of Cloudflare's network (e.g., in some geographies, for certain clients).

⁵ Percentages calculated based on the number of minutes of reported service unavailability in 2021 for Facebook and Cloudflare, respectively, compared to minutes in the year up to October 2021 (therefore, ten out of twelve months in the year). For further details on sources, please see endnotes 2 and 4 above.

⁶ Often referred to as “cybersecurity” or “information security”, “digital security” at the OECD looks beyond purely technical security issues to the economic and social challenges they present. An important aspect of digital security is managing risk, a combination of the likelihood and severity of digital security incidents that may adversely impact the digital environment.

⁷ The more formal definition of an Autonomous System (AS) is defined by RFC 1930 as “a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy” (Hawkinson and Bates, 1996^[115]).

⁸ The Border Gateway Protocol (BGP) is the most commonly used Exterior Gateway Protocol (EGP), which refers to the exchange of routing information between ASes. Information exchanged between network elements within an AS is governed by Interior Gateway Protocols (IGPs). IGPs are not the subject of this report.

⁹ In the common case, the origin and destination do not share a common “upstream” ISP that they are both customers of, and in these cases, the two “uphill” chains of purchased transit are joined by a peering connection at the top. In the rare case that both origin and destination lie within the “transit cone” of transit customers of a single ASN, the bandwidth is produced within the ISP at the apex. Because the numbers needed to quantify the latter case are generally held as trade secrets, it’s difficult to say for certain how much of the Internet’s bandwidth is produced “on-net” by individual ISPs, but the portion is small and further declines over time as the Internet’s population becomes larger and its topology becomes more diverse.

¹⁰ There is no financial mechanism by which traffic may pass through more than one IXP, nor for a “valley” to exist in the topology. (That is, the topology must always be a Λ , and can never be an M.)

¹¹ This additional information is commonly annotated within an AS path in the form [Origin / AS1 / AS2 – AS3 \ Destination] in reference to their respective positions within the Λ topology, and such inferences are commonly based on knowledge of peering relationships at Internet exchange points.

¹² For instance, in the Internet Engineering Task Force’s (IETF) SIDR Operations Working Group.

¹³ Widespread countermeasures, in the form of the Routing Arbiter Database (RADB) and other RPSL-formatted Internet Routing Registries (IRRs), were standardised in 1999 through RFC 2622 (Alaettinoglu et al., 1999_[117]). There have been other proposals to improve BGP security in the past, including Secure BGP, Secure Origin BGP, Pretty Secure BGP, and Inter-domain Route Validation, among other efforts (Huston, 2021_[55]).

¹⁴ Other aspects of BGP security include ensuring the protection of the BGP speaker and of BGP sessions, as described in (Durand, Pepelnjak and Doering, 2015_[49]).

¹⁵ A malicious actor could also deliberately lengthen a path to make another path more desirable, but this technique is not as common.

¹⁶ The possible attacks detailed in (Sun et al., 2021_[20]) on Certificate Authorities during the domain control verification stage refer to the work of (Birge-Lee et al., 2018_[38]) that is mentioned in the BGP events impacting integrity section.

¹⁷ For more information regarding the Global Routing Intelligence Platform hosted at Georgia Tech, please see <https://grip.inetintel.cc.gatech.edu/method>. Please note that the GRIP platform was previously hosted at the Center for Applied Internet Data Analysis (CAIDA), based at the San Diego Supercomputer Center. For RIPE NCC’s RIS, see <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>, for PHC’s routing information, see https://www.pch.net/resources/Routing_Data/, and for the University of Oregon’s Route Views project, please see <http://www.routeviews.org/routeviews/>.

¹⁸ For further discussion of methodology, please see the following: <http://bgpstream.crosswork.cisco.com/about/> (BGPStream), <https://grip.inetintel.cc.gatech.edu/method> (GRIP), and https://blog.qrator.net/en/q1-2021-report_129/ (Qrator, see section “data sources and observation methodology”).

¹⁹ Level 3 was an ISP from the United States, became a part of CenturyLink and is now Lumen Technologies.

²⁰ Quad101 is a public DNS resolver operated by the Taiwan Network Information Center (TWNIC), which is also Chinese Taipei's country-code Top Level Domain (ccTLD) operator.

²¹ For example, if traffic is passed on to the victim AS, not “black-holed”.

²² Team Cymru maintains two lists, a “traditional bogon” and “fullbogon” list. Both lists include the reserved and special use IPv4 address space, as well as the IP addresses that have not been allocated to the RIRs by IANA (Team Cymru, 2022^[48]). The fullbogon list goes one step further and also includes the IP space that has been allocated to the RIRs but not yet assigned and covers both IPv4 and IPv6 address space (Team Cymru, 2022^[48]).

²³ For example, IPv4 prefixes longer than /24 or IPv6 prefixes longer than /48 have been documented by the RIPE community.

²⁴ As aptly described by Gert Döring in his presentation on “BGP communities 101”, one practical application of BGP communities would be for an ASN to tag the routes learned from its customers with a certain community value and establish that only routes with the customer tag be sent on to its upstream providers (Döring, 2018^[52]). If correctly applied, *only* the routes tagged as coming from an ASN's customer would be sent to their upstream providers, and therefore no routes learned from a peer would be leaked by mistake (Döring, 2018^[52]).

²⁵ This is down from 94% registered in January 2020 and 92% registered in January 2019, but is a slight increase from the 84% recorded in January 2021 (MANRS, 2022^[27]).

²⁶ However, as (Lepinski and Kent, 2012^[56]) note, “while the initial focus of this architecture is routing security applications, the PKI described in this document could be used to support other applications that make use of attestations of IP address or AS number resource holdings.”

²⁷ An EE certificate provides the private key with which to sign the ROA and validates that the signatory holds the prefix in question. These EE certificates expire, which requires entities to recreate their ROAs on a regular basis. More generally, the validity of the ROA is tied to the validity of its associated EE certificate (Lepinski and Kent, 2012^[56]). For instance, if an RIR revokes an allocation for a given IP address space that has an associated ROA, the EE certificate is automatically revoked. The revoked EE certificate means that the corresponding ROA is no longer valid and would then return an “unknown” response from the RPKI validator. Lapsing into an unknown and therefore unprotected state as a “fail safe” option seems to be a reasonable option to avoid inaccessibility of certain prefixes even if some ROAs expire.

²⁸ Metrics available from June 2020 (“Use of RPKI Validation for World”). See <https://stats.labs.apnic.net/rpki/XA>.

²⁹ The website <https://isbgpsafeyet.com/> maintains a list of public announcements regarding RPKI deployment (signing ROAs and dropping RPKI invalids through ROV) made by network operators around the world.

³⁰ As reported on 31 May 2022, according to the “RPKI-ROV analysis of unique Prefix-origin pairs (IPv4)” in the NIST RPKI Monitor, there were 367,263 “valid” unique prefix-origin pairs, 622,570 “not found”, and 15,955 “invalid” (NIST, 2022^[73]).

³¹ The MANRS Observatory aggregates several data sources, including various BGP collectors, Route views, RIPE RIS, RPKI, IRR, RIR whois, RIPEstat, and BGPStream among others, which is collected, aggregated and analysed, and used to populate its MANRS Observatory dashboard (MANRS, n.d.^[116]).

³² Even considering the global average of 15% of global Internet users being covered by an ISP conducting ROV filtering would not provide a completely accurate view of the benefit of implementing RPKI. The benefit of certain ISPs to conduct ROV filtering is greater than for others; for instance, a large ISP filtering invalids provides a benefit to their customers and peers (regardless of whether they do ROV filtering) by not passing on invalid route announcements.

³³ Exploiting this vulnerability allows an attacker to append the hijacked ASN to the announced path, thus inserting itself into the AS path by pretending to be a transit provider for the hijacked AS.

³⁴ Such a case may entail a party establishing a new connection, in which it represents itself as the hijacked ASN (i.e., impersonates the hijacked ASN) for the purposes of announcing at least one of the hijacked ASN's prefixes. Since RPKI would not trigger an error, a network operator would have to look at the transit provider and path to determine plausibility, but this requires knowledge of the hijacked ASN's transit provider and neighbouring ASes.

³⁵ ARIN experienced an outage in its RPKI repository in 2018 (see <https://www.arin.net/vault/announcements/2018/20181024.html>) and in 2020 when deploying a new version of its RPKI system (see <https://www.arin.net/announcements/20200826/>). APNIC reported a 20-minute outage in 2019 (see <https://www.apnic.net/about-apnic/service-updates/service-announcement-13-december-2019/>). AfriNIC had problems with its RPKI service in 2020 (see <https://lists.afrinic.net/pipermail/rpki-discuss/2020-March/000108.html>). RIPE NCC reported a few outages, including in February 2020 (see <https://www.ripe.net/ripe/mail/archives/routing-wg/2020-February/004015.html>) and in April 2020 (see <https://www.ripe.net/support/service-announcements/rsync-rpki-repository-downtime>).

³⁶ If a ROA already exists and an actor attacked the RIR that issued the prefix, he or she could revoke the pre-existing ROA and issue a false ROA in its place. Or, if the security breach occurred at another RIR that did not issue the prefix, the attacker could still issue a contradicting ROA, which could open the prefix up to vulnerabilities, such as route leaks via more specific announcements (Durand, 2020^[76]). If no ROA exists, a security breach at *any* of the RIRs would allow an attacker to issue a false ROA for a given prefix, which would label the real BGP route announcement with the true origin ASN as false, without altering anything in BGP (Durand, 2020^[76]).

³⁷ Other efforts in IETF are focused on improving the detection of accidental route leaks, such as “draft-ietf-idr-bgp-open-policy-16” (BGP roles) and “draft-ietf-grow-route-leak-detection-mitigation-00” (Route leak detection and mitigation).

³⁸ Users can set alerts for several cases, including for possible hijacks but also in case of anomalies in its own routing announcements (e.g., announcing RPKI invalid prefixes or ROAs expiring soon) (BGPalerter (GitHub), 2021^[119]).

³⁹ ARTEMIS is now being maintained by Code BGP, led by some of the original researchers, with plans to soon launch a third-party monitoring platform that builds on and expands ARTEMIS' capabilities (CODE BGP, 2021^[120]).

⁴⁰ The Peerlock-lite filter assumes that a non-Tier 1 network would not sell transit to a Tier 1 network, thereby it would reject any prefixes that it receives from customers with a Tier 1 AS in the AS path (McDaniel, Smith and Schuchard, 2021^[87]).

⁴¹ To join, each participant must meet certain requirements and agree to adhere to and continue to meet compulsory actions; MANRS also recommends several additional actions for stakeholders to undertake

(MANRS, 2021_[121]). The requirements differ across the four MANRS programs (for network operators, CDNs and cloud providers, IXPs and equipment vendors).

⁴² Taking the example of RPKI, once an AS decides to implement ROV filtering fully (e.g., by discarding invalid responses), any invalid route would be discarded. In a large ISP with many customer ASes, there is a chance of mistakenly dropping a customer's traffic when turning on ROV filtering (e.g., if they have an incorrect ROA, for example).

⁴³ For example, Oracle previously published publicly available reporting and analysis of BGP incidents through its Internet Intelligence platform, which was referenced in the first draft of this report; however, this functionality is no longer available on its website.

⁴⁴ JPNIC has reported on annual numbers in some presentations; for 2019 figures, see <https://www.nic.ad.jp/ja/materials/iw/2019/proceedings/s10/s10-watanabe-2.pdf>. For an overview of the trends from 2012-2018, see <https://www.nic.ad.jp/ja/materials/iw/2018/proceedings/s03/s3-kimura.pdf>. However, further explanation of the metrics would be helpful to ensure correct comprehensive of the high-level graphs shown.

⁴⁵ These limitations include, for instance, that the reporting focused only on JPIRR registrants, and that inaccuracies of IRR databases may trigger notifications.

⁴⁶ The five RIRs are AFRINIC, APNIC, ARIN, LACNIC and RIPE NCC.

⁴⁷ A Listserv is an online email discussion group.

⁴⁸ Further discussion of the Safety and Reliability Standards can be found at https://www.soumu.go.jp/menu_seisaku/ictseisaku/net_anzen/anshin/ and additional explanation of the measures included therein at https://www.soumu.go.jp/main_content/000694915.pdf [in Japanese only].

⁴⁹ Unofficial translation.

⁵⁰ For instance, related to interconnection interfaces, the regulation specifies that an operator should reject route announcements “belonging to the operator's own address blocks or to those provided by the telecommunications operator to one of its customers and that cannot be expected to be advertised by other telecommunications operators” (Finnish Transport and Communications Agency (Traficom), 2015_[113]). Similarly, regarding customer interfaces, the text directs operators to “filter any traffic from a customer interface to the communications network with a source address that is not assigned to the customer interface in question” (Finnish Transport and Communications Agency (Traficom), 2015_[113]).