NIST Special Publication 800
NIST SP 800-189r1 ipd

Border Gateway Protocol Security and Resilience

Initial Public Draft

Kotikalapudi Sriram
Doug Montgomery

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-189r1.ipd

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Special Publication 800
NIST SP 800-189r1 ipd

Border Gateway Protocol Security and Resilience

Initial Public Draft

Kotikalapudi Sriram
Doug Montgomery
Wireless Networking Division
Communications Technology Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-189r1.ipd

January 2025

U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Charles H. Romine, performing the non-exclusive functions and duties of the Under Secretary of Commerce for Standards and Technology and Director, National Institute of Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

Authority
This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

NIST Technical Series Policies
Copyright, Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

Publication History
Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added in final publication.]
Supersedes NIST Series XXX (Month Year) DOI [Will be added in final publication, which will supersede NIST SP 800-189.]

How to Cite this NIST Technical Series Publication:
Sriram K, Montgomery D (2025) Border Gateway Protocol Security and Resilience. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-189r1 ipd. https://doi.org/10.6028/NIST.SP.800-189r1.ipd

Author ORCID iDs
Kotikalapudi Sriram: 0000-0002-1585-0134
Doug Montgomery: 0000-0002-5364-9474

Public Comment Period
January 3, 2025 - February 25, 2025

Submit Comments
sp800-189@nist.gov

National Institute of Standards and Technology
Attn: Wireless Security Division, Communications Technology Laboratory
100 Bureau Drive (Mail Stop 6730) Gaithersburg, MD 20899-6730

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/sp/800/189/r1/ipd, including related content, potential updates, and document history.

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

This publication provides guidance on Internet routing security, preventing IP address spoofing, and certain aspects of DDoS detection and mitigation. It particularly focuses on Border Gateway Protocol, which is the routing protocol used to distribute and compute paths between the tens of thousands of autonomous networks that comprise the internet. Technologies recommended in this document for securing BGP routing include Resource Public Key Infrastructure, Route Origin Authorization, ROA-based route origin validation, and prefix filtering. Additionally, technologies recommended for mitigating DDoS attacks focus on preventing IP address spoofing using source address validation with access control lists and unicast Reverse Path Forwarding. Other technologies are also recommended as part of the overall routing security mechanisms, such as remotely triggered black hole filtering and flow specification.

Keywords

Autonomous System Provider Authorization (ASPA); Border Gateway Protocol (BGP) security; distributed denial-of-service (DDoS); Flowspec; Only to Customer (OTC); Resource Public Key Infrastructure (RPKI); ROA-based route origin validation (ROA-ROV); Route Origin Authorization (ROA); routing security and resilience.

Audience

This document gives technical guidance and recommendations for improving the security and resilience of Internet routing based on the Border Gateway Protocol. The primary audience includes Internet routing security engineers, information security officers, and managers of federal enterprise networks. The guidance also applies to the network services of hosting providers (e.g., cloud-based applications and service hosting) and Internet service providers (ISPs) when they are used to support federal IT systems. The guidance may also be useful for enterprise and transit network operators and equipment vendors in general.

The guidance and applicable security recommendations in this publication should be incorporated into the security plans and operational processes of federal enterprise networks. Likewise, applicable security recommendations should also be incorporated into federal contract requirements for Internet transit services and commercially-hosted application services (e.g., content distribution, remote storage, cloud services, email, domain name service).

Call for Patent Claims

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this NIST draft publication). Such guidance and/or requirements may be directly stated in this NIST Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this NIST draft publication and of any relevant unexpired U.S. or foreign patents.

NIST may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or

b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this NIST draft publication either:

   i. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or

   ii. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: sp800-189@nist.gov

Table of Contents

List of Tables

List of Figures

Acknowledgments

The authors thank William T. Polk, Scott Rose, Oliver Borchert, Susan Symington, William C. Barker, William Haag, Allen Tan, Isabel Van Wyk, and Jim Foti for their review and comments on this document or an earlier version. They are also grateful to the many reviewers who provided highly helpful comments on Public Drafts 1 and 2 during the publication of the original document in 2019.

Executive Summary

There have been numerous security and resilience incidents in recent years involving Border Gateway Protocol (BGP), including prefix hijacks, route leaks, and other forms of misrouting. These incidents include both malicious attacks and accidental misconfigurations that result in the denial of service (DoS), unwanted data traffic detours, and performance degradation [Madory]. Another form of abuse of Internet routing in the data plane is source Internet Protocol (IP) address spoofing, a technique often used in DoS attacks.

This document provides technical guidance and recommendations to improve the security and resilience of Internet routing based on BGP. It primarily focuses on the points of interconnection between enterprise networks or hosted service providers and the public Internet. These are commonly known as "stub" networks (i.e., those networks that only provide connectivity to their end systems) and transit networks (i.e., those networks that serve to interconnect and pass traffic between stub networks and other transit networks), and the points of interconnection between them are often referred to as the "Internet's edge." There is usually a contractual relationship between transit networks and the stub networks that they service, and the set of technical procedures and policies defined in that relationship is commonly called the "peering policy." Many of the recommendations in this document also apply to the points of interconnection between two transit networks, which may vary from those between stub and transit networks.

These recommendations can reduce the risk of accidental misconfigurations and malicious attacks on the Internet's BGP routing system and help prevent IP address spoofing and distributed DoS (DDoS) attacks. They primarily cover security and resilience technologies for routers that operate BGP (commonly called BGP routers) but also extend to other systems that support Internet routing security, such as Resource Public Key Infrastructure (RPKI) repositories.

The guidance in this publication should be incorporated into the security plans and operational processes of federal enterprise networks, and applicable recommendations should be incorporated into requirements for federal contracts for hosted application services and Internet transit services. This document also contributes to the ongoing broader efforts by the Federal Government to secure the foundational protocols of the Internet [NCSIP], particularly Internet routing [WH-ONCD][BITAG], with RPKI, Route Origin Authorization (ROA), ROA-based route origin validation (ROA-ROV), and prefix filtering. Additionally, the technologies recommended for mitigating DDoS attacks focus on the prevention of IP address spoofing using source address validation (SAV) with access control lists (ACLs) and unicast Reverse Path Forwarding (uRPF). Other technologies are also recommended as part of the overall security mechanisms, such as remotely triggered black hole (RTBH) filtering and flow specification (Flowspec).

1. Introduction

1.1. What This Guide Covers

This publication provides technical guidelines and recommendations for deploying protocols and technologies that improve Internet routing security, reduce the risk of accidental misconfigurations and malicious attacks in the routing control plane, and help detect and prevent IP address spoofing and resulting DDoS attacks. These recommendations primarily cover protocols and techniques to be used in BGP routers. However, they partly extend to other systems that support reachability on the Internet (e.g., RPKI repositories, DNS, and other open Internet services).

Technologies recommended in this document for securing interdomain routing control traffic include RPKI, ROA, ROA-ROV, and prefix filtering. Additionally, technologies recommended for mitigating DDoS attacks focus on the prevention of IP address spoofing using SAV with ACLs and uRPF. Other technologies, such as RTBH filtering and Flowspec, are also recommended as part of the overall security mechanisms.

This document addresses many of the same concerns regarding BGP vulnerabilities highlighted in [NCSIP][WH-ONCD][BITAG][FCC-NPR] but describes standards-based security mechanisms in greater technical depth and provides specific security recommendations.

1.2. What This Guide Does Not Cover

BGP origin validation relies on a global RPKI system (e.g., certificate authorities, publication repositories) as the source of trusted information about Internet address holders and their route origin authorization statements. Each RIR operates a trusted root certificate authority (CA) in the RPKI system and publishes a Certificate Practice Statement [RFC7382] that describes each implementation’s security and robustness properties. Each RPKI CA has integrity and authentication mechanisms for data creation, storage, and transmission. Nevertheless, compromise of the underlying servers and/or registry services is still a potential — if low probability — threat. Making security recommendations for mitigating against such threats is outside of this document's scope.

Additionally, while transport layer security is key to the integrity of messages that are communicated in BGP sessions, security recommendations for the underlying transport layer is also outside of this document’s scope.

DDoS attacks use spoofed IP addresses to exploit connectionless query-response services (e.g., DNS, Network Time Protocol [NTP], Simple Service Discovery Protocol [SSDP] servers) to “reflect” and amplify the impact on intended targets. While this document addresses SAV to detect and mitigate spoofed IP addresses, it does not address the security hardening of the servers that are exploited for reflection and amplification.

1.3. Document Structure

The rest of the document is presented in the following manner:

- **Section 2** describes routing control plane attacks, such as BGP prefix hijacking, autonomous system (AS) path modification, and route leaks.
- **Section 3** describes data plane attacks that involve source IP address spoofing and reflection amplification.
- **Section 4** describes solutions and makes security recommendations for BGP security.
- **Section 5** describes solutions and makes security recommendations for detecting and mitigating source IP address spoofing.

1.4. Conventions Used in This Guide

Throughout this guide, "**Security Recommendation**" denotes a recommendation that should be addressed in security plans, operational practices, and agreements for contracted services.

URLs and references are provided to guide readers to websites and online tools that are designed to aid administrators. This is not meant to endorse the website, or any product or service offered by the website publisher. All URLs were considered valid at the time of writing.

2. BGP Vulnerabilities

As initially designed and commonly deployed on the Internet, BGP lacked the security and resilience mechanisms to prevent malicious attacks and misconfigurations that can compromise Internet routing. BGP's original design lacked the capability to [RFC4272]:

- Validate the authority of remote networks to originate announcements to specific destinations,
- Verify the integrity and authenticity of messages exchanged between neighboring networks,
- Ensure the authenticity and integrity of information from remote networks, and
- Detect routing announcements that violate business policies between neighboring networks.

The lack of these capabilities often led to accidental misconfigurations that resulted in wide-scale impacts on Internet routing. As the Internet became essential to global commerce, critical infrastructure, and communications, malicious actors began purposefully exploiting these BGP vulnerabilities.

2.1. Unauthorized BGP Originations (Prefix Hijacks)

A BGP prefix hijack occurs when an autonomous system (AS) accidentally or maliciously originates a prefix that was not authorized or intended by the prefix owner. This is also known as false origination or announcement. In contrast, if an AS is authorized by the prefix owner to originate or announce a prefix, then such a route origination/announcement is legitimate. In Figure 1, the prefix 192.0.2.0/24 is legitimately originated by AS64500, but AS64510 falsely originates it. The path to the prefix via the false origin AS will be shorter for a subset of the ASs on the Internet, which will install the false route in their routing table or forwarding information base (FIB). That is, ASs for which AS64510 is closer (i.e., shorter AS path length) would choose the false announcement, and thus, data traffic from clients in those ASs destined for the network 192.0.2/24 will be misrouted to AS64510.

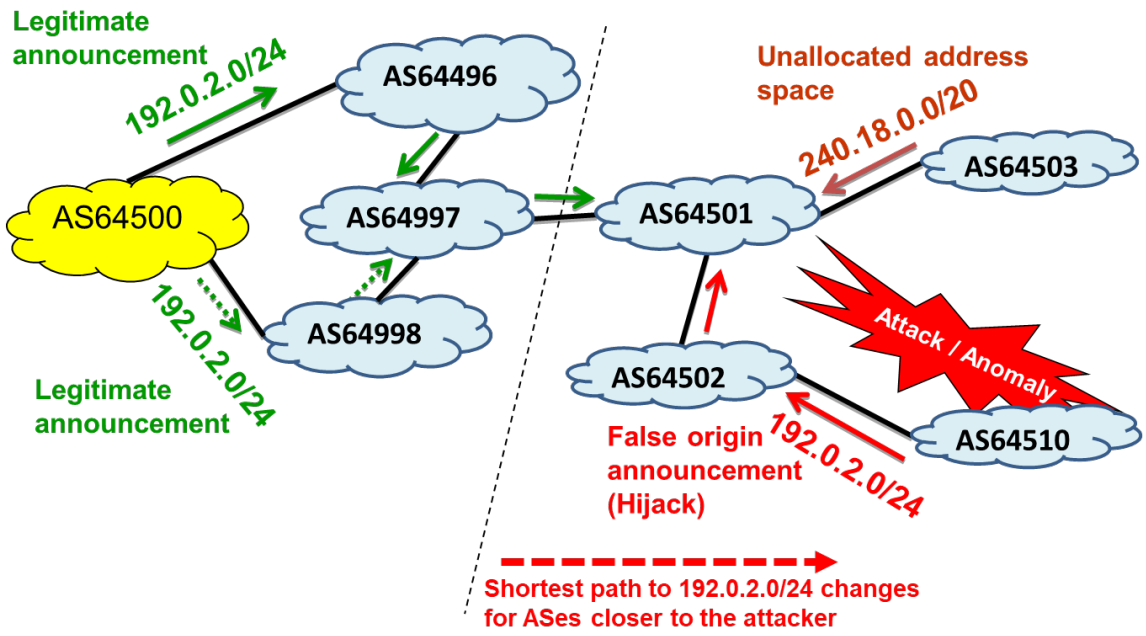

Adverse effects: denial-of-service, misrouting of traffic, unauthorized routing

Figure 1. Illustration of prefix hijacking and announcement of unallocated address space

The rules for IP route selection on the Internet always prefer the most specific (i.e., longest) matching entry in a router's FIB. When an offending AS falsely announces a more specific prefix than one announced by an authorized AS, the longer, unauthorized prefix will be widely accepted and used to route data.

Figure 1 also illustrates an example of unauthorized origination of unallocated (i.e., reserved) address space 240.18.0.0/20. Currently, 240.0.0.0/8 is reserved for future use [IANA-v4-r]. Similarly, an AS may falsely originate allocated but currently unused address space. This is referred to as "prefix squatting," where someone else's unused prefix is announced and used to send spam emails or for some other malicious purpose.

The unauthorized announcement of a prefix that is longer than the legitimate announcement is called a sub-prefix hijack. The consequences of such adverse actions can include DoS, eavesdropping, misdirection to imposter servers (e.g., to steal login credentials or inject malware), or the defeat of IP reputation systems to launch spam emails. Several commercial services and research projects that track and log anomalies in the global BGP routing system [BGPmon][ThousandEyes][BGPStream][ARTEMIS], and many of these sites provide detailed forensic analyses of observed attack scenarios.

2.2. Unauthorized BGP Update Modification (Path Hijacks)

BGP messages carry a sequence of AS numbers that indicates the "path" of interconnected networks over which data will flow. This "AS_PATH" [RFC4271] data is often used to implement routing policies that reflect the business agreements and peering policies negotiated between

networks. BGP is also vulnerable to unauthorized modification of the AS_PATH information that it conveys. For example, a malicious AS that receives a BGP update may illegitimately remove some of the preceding ASs in the AS_PATH attribute to make the overall path length seem shorter. When the update modified in this manner is propagated, the ASs upstream can be deceived into believing that the path to the advertised prefix via the adversary AS is shorter. By doing this, the adversary AS may seek to illegitimately increase its revenue from its customers or may be able to eavesdrop on traffic that would otherwise not transit through their AS.

Another example of maliciously modifying a BGP update is when an adversary AS replaces a prefix in a received update with a more specific sub-prefix and then forwards the update to neighbors. This attack is known as a Kapela-Pilosov attack [Kapela-Pilosov]. Only the prefix is replaced by a more specific prefix, but the AS path is not altered. In BGP path selection, a more specific prefix advertisement takes precedence over a less specific one. This means that ASs on the Internet would widely accept and use the adversary AS's advertisement for the more specific prefix. The exceptions are the ASs in the AS path from the adversary to the prefix. These exception ASs reject any advertisements that they may receive for the more specific prefix because they detect their own AS number in the AS path. This is called avoidance of loop detection and is a standard practice in BGP. Thus, the data path from the adversary AS to the prefix (i.e., the network in consideration) remains intact (i.e., unaffected by the malicious, more specific advertisement). The net result of this attack is that the adversary could force almost all traffic for the more specific prefix to be routed via their AS. Thus, they can eavesdrop on the data that was destined for the more specific prefix while channeling it back to the legitimate destination to avoid detection.

2.3. BGP Policy Violations (Route Leaks)

Previously, it was noted that the interconnections of networks on the Internet are dictated by contracted business relationships that express the policies and procedures for the exchange of routing and data traffic at each point of interconnection. Such peering policies often specify limits on what routing announcements will be accepted by each party. Often, these policies reflect the business relationship between networks.

Definitions of Peering Relations, Customer Cone: A transit provider typically provides service to connect its customer(s) to the global Internet. A customer AS or network may be single-homed to one transit provider or multi-homed to more than one transit providers. A stub customer AS has no customer ASes. A leaf customer is a stub customer that is single-homed to one transit provider and not connected to any other AS. Peering relationships considered in this document are provider-to-customer (P2C), customer-to-provider (C2P), and peer-to-peer (p2p). Here, "provider" refers to transit provider. The first two are transit relationships. A peer connected via a p2p link is known as a lateral peer (non-transit). A customer cone of AS A is defined as AS A plus all the ASes that can be reached from A following only P2C links [Luckie]. The term "customer cone prefixes" refers to the union of the prefixes originated by all networks in the customer cone of a specific AS. ASes that have a lateral peering (i.e., p2p) relationship typically announce their customer cone prefixes to each other and subsequently announce the lateral

peer's customer cone prefixes to their respective customers but not to other lateral peers or transit providers.

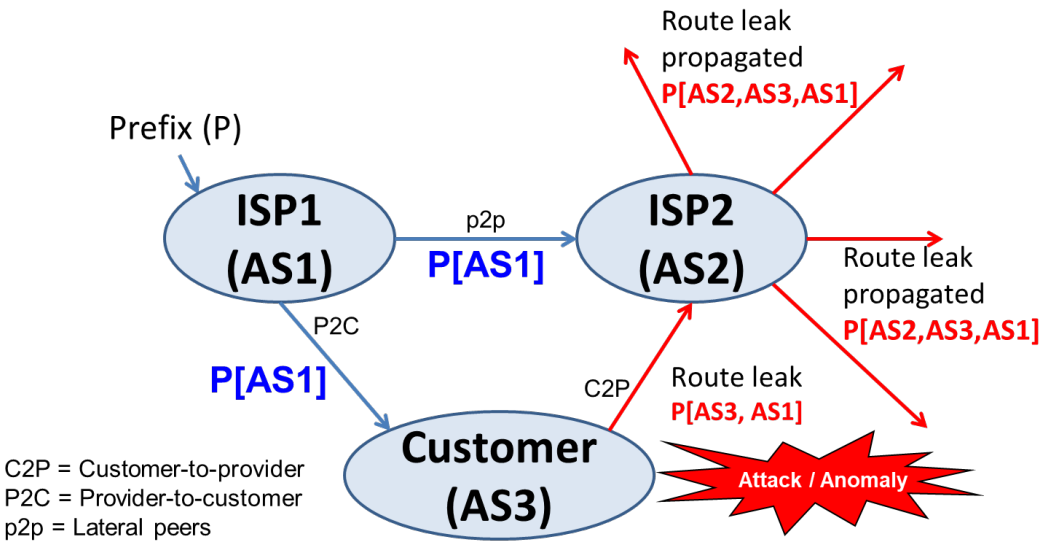

Figure 2. Illustration of the basic notion of a route leak

These relationships are significant because much of the operation of the global Internet is designed such that a stub or customer AS should never be used to route between two transit ASes. This policy ensures that stubs or customer ASes do not pass BGP routing information received from one transit provider to another. Figure 2 illustrates a common form of route leak that occurs when a multi-homed customer AS (such as AS3 in Figure 2 learns a prefix update from one transit provider (ISP1) and "leaks" the update to another transit provider (ISP2) in violation of intended routing policies. The second transit provider does not detect the leak and propagates the leaked update to its customers, lateral peers, and transit ISPs [RFC7908]. Some examples of recent route leak incidents include: 1) the MainOne (a Nigerian ISP) leak of Google prefixes, which caused an outage of Google services for over an hour in November 2018 [Naik]; (2) the Dodo-Telstra incident in March 2012, which caused an outage of Internet services nationwide in Australia [Huston2012]; and (3) the massive Telekom Malaysia route leaks, which Level3, in turn, accepted and propagated [Toonk-B].

More generally, as defined in [RFC7908], a route leak is the propagation of routing announcements beyond their intended scope. That is an AS's announcement 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. In the route leak depicted in Figure 2, the AS path violates the general routing policy that Internet paths should be "valley-free" [Rexford-Gao]. This term refers to the concept that once a BGP route is propagated "down" a provider-to-customer (P2C) peering path, it should never be propagated "up" a customer to the provider (C2P) peering path.

In [RFC7908], several types of route leaks are enumerated and described together with examples of recent incidents. The result of a route leak can include redirection of traffic through an unintended path, which may enable eavesdropping or malicious traffic analysis. When many routes are leaked simultaneously, the offending AS is often overwhelmed by the resulting unexpected data traffic and drops much of the traffic that it receives [Huston2012] [Toonk-A] [Naik] [Zmijewski]. This causes degradation and denial of service for the affected prefixes. Route leaks can be accidental or malicious but most often arise from accidental misconfigurations.

3. Other Internet Routing Related Vulnerabilities (IP Address Spoofing)

3.1. Spoofed Source Addresses

Distributed denial-of-service (DDoS) is a form of attack where malicious traffic is generated from distributed sources to achieve a high-volume denial of service attack and directed towards an intended victim (i.e., system or server) [Arbor] [Arbor2] [ISOC] [Huston2016] [Mirai1]. To conduct a direct DDoS attack, the attacker typically uses a few powerful computers or many compromised third-party devices (e.g., laptops, tablets, cell phones, Internet of Things (IoT) devices, etc.). The latter scenario is often implemented through botnets [Arbor] [Huston2016] [DOC-Botnet]. In many DDoS attacks, the IP source addresses in the attack messages are "spoofed" to avoid traceability [Arbor]. Some DDoS attacks are launched without using spoofed source addresses. For example, in the Mirai attacks [Mirai1] [Mirai2] [Winward] [TA16-288A], a huge number of compromised bots (IoT devices) sending the attack traffic used the normal source IP addresses of the IoT devices. Further, the source addresses could also belong to a hijacked prefix with the intention of deceiving source address validation (SAV) [BCP38] [BCP84] (see Section 5.1.7). If a hijacked prefix is being used, then the source addresses appearing in the DDoS attack packets are sometimes randomly selected from that prefix.

3.2. Reflection Amplification Attacks

Source address spoofing is often combined with reflection and amplification from poorly administered open Internet servers (e.g., DNS, NTP) to significantly multiply the attack traffic volume [Azure] [TA14-017A] [ISOC]. Figure 3 illustrates an example of such attacks. The attacker sends query requests to high-performance Internet servers. The attacking systems employ source address spoofing, which inserts the IP address of the target (e.g., 203.0.113.1) as the source address in the requests. For Internet services that use the User Datagram Protocol (UDP) (e.g., DNS, NTP), the query and response are each contained in a single packet, and the exchange does not require the establishment of a two-way connection between the source and the server (unlike Transmission Control Protocol (TCP)). The responses from such open Internet servers are directed to the attack target since the target's IP address was forged as the source address field of the request messages. Often, the response from the server to the target address is much larger than the query itself, thus amplifying the effect of the DoS attack. Such reflection and amplification techniques can result in DDoS attacks with traffic volumes in the range of hundreds of Gbps [Azure] [Symantec] [ISTR-2015] [ISTR-2016] [ISTR-2017] [ISOC] [Verisign1] [Verisign2] [Bjarnason]. The attack volumes may still rise significantly if the Mirai-scale attacks are combined with reflection amplification attacks.

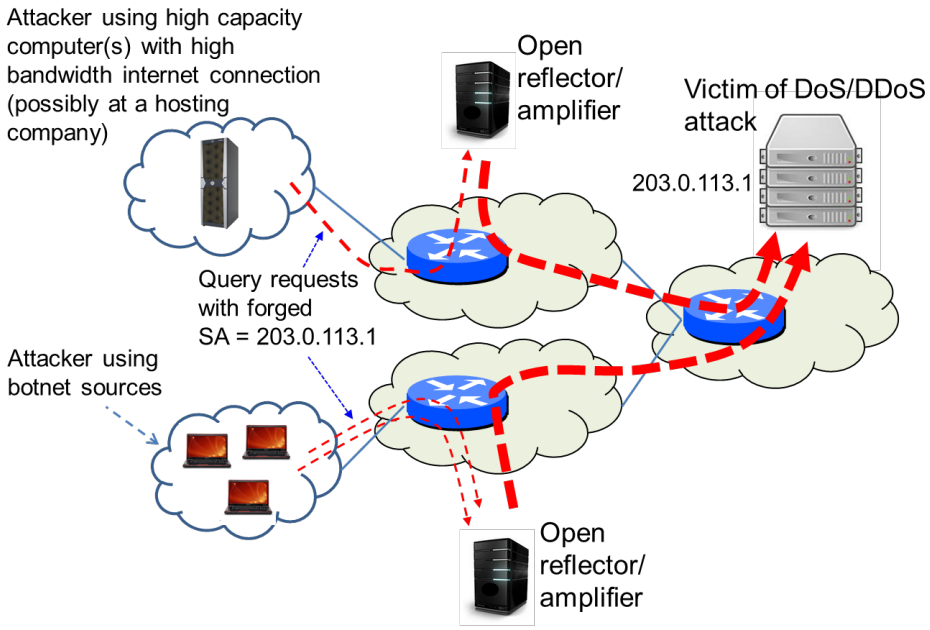

Figure 3. DDoS by IP source address spoofing and reflection and amplification

4. Improving BGP Security and Resilience — Solutions and Recommendations

BGP security vulnerabilities and mitigation techniques have been of interest within the networking community for several years (e.g., [IETF-SIDR] [RFC7454] [NANOG] [Murphy] [MANRS] [MANRS2] [ENISA] [Quilt] [NIST-RPKI] [CSRIC4-WG6] [CSRIC6-WG3] [RFC6811] [RFC8205] [NSA-BGP] [CSDE] [Chung] [Wishnick] [Yoo]). This section highlights key BGP security technologies that have emerged from such efforts and achieved some level of standardization or commercialization. Many of the solution technologies discussed here have been developed and standardized in the Internet Engineering Task Force (IETF) [IETF-SIDR] [IETF-SIDROPS] [IETF-IDR] [IETF-OPSEC] [IETF-GROW]. This document addresses many of the same concerns regarding BGP vulnerabilities and DDoS attacks as highlighted in other Government and industry initiatives [NCSIP] [WH-ONCD] [MANRS] [BITAG] [FCC-NPR] [OECD] [CableLabs] but goes into greater technical depth in describing standards-based and commercially available security mechanisms and providing specific security recommendations.

4.1. Registration of Route Objects in Internet Routing Registries

Declarative data about Internet resource allocations and routing policies have traditionally been available from regional Internet registries (RIRs) and Internet routing registries (IRRs). The RIR data are maintained regionally by ARIN in North America, RIPE in Europe, LACNIC in Latin America, APNIC in Asia-Pacific, and AfriNIC in Africa. The IRRs are maintained by the RIRs (RIPE NCC, APNIC, AfriNIC, and ARIN) as well as some major Internet service providers (ISPs). Additionally, Merit's Routing Assets Database (RADb) [Merit-RADb] and other similar entities provide a collective routing information base consisting of registered (at their site) as well as mirrored (from the IRRs) data. The route objects available in the IRRs provide routing information declared by network operators. Specifically, the route objects contain information regarding the origination of prefixes (i.e., the association between prefixes and the ASes that may originate them). Routing Policy Specification Language (RPSL) [RFC4012] [RFC7909] and the Shared Whois Project (SWIP) [SWIP] are two formats in which the data in RIRs/IRRs are presented. ARIN predominantly uses SWIP, but some use RPSL as well. LACNIC also uses SWIP. The rest of the RIRs and the ISPs' IRRs use only RPSL.

The completeness, correctness, freshness, and consistency of the data derived from these sources vary widely, and the data is not always reliable. However, there are efforts underway to make the data complete and reliable [RFC7909]. Network operators often obtain route object information from the IRRs and/or RADb, and they can make use of the data in the creation of prefix filters (see Sections 4.4 and 4.5) in their BGP routers.

It is worth noting that many of the RIRs run Internet routing registries (IRRs) that are integrated with regional Internet registry (RIR) allocation data that facilitate stronger authentication schemes. These are documented in [RFC2725].

468 **Table 1. Security recommendations related to IRR**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 1: All Internet number resources (e.g., address blocks and AS numbers) should be covered by an appropriate registration services agreement with an RIR, and all point-of-contact (POC) information should be up to date. The granularity of such registrations should reflect all sub-allocations to entities (e.g., enterprises with provider-based addresses, enterprises within the parent organization, branch offices) that operate their own network services (e.g., Internet access, email, DNS).	X	X
Security Recommendation 2: Route objects corresponding to the BGP routes originating from an AS should be registered and actively maintained in an appropriate RIR's IRR. Enterprises should ensure that appropriate IRR information exists for all IP address space used by them.	X	X

469 While efforts are encouraged to create complete and accurate IRR data in line with the current
470 operational reality, greater efforts should be devoted to creating route origin authorizations
471 (ROAs) (see Section 4.3) because RPKI provides a stronger authentication and validation
472 framework for network operators than IRR.

473 ### 4.2. Certification of Resources in Resource Public Key Infrastructure

474 Resource Public Key Infrastructure (RPKI) is a standards-based approach for providing
475 cryptographically secured registries of Internet number resources, and routing policy [RFC6480]
476 [RFC9582] [NANOG] [Murphy]. The IPv4/IPv6 address and AS number resource allocations
477 follow a hierarchy. The Internet Assigned Numbers Authority (IANA) allocates resources to the
478 regional Internet registries (RIRs) (e.g., ARIN, RIPE, etc.), and the RIRs suballocate resources to
479 ISPs and enterprises. The ISPs may further suballocate to other ISPs and enterprises. In some
480 regions, RIRs suballocate to local Internet registries (LIRs), which in turn suballocate to ISPs and
481 enterprises. RPKI is a global certificate authority (CA) and registry service offered by all regional
482 Internet registries (RIRs). The RPKI certification chain follows the same allocation hierarchy (see
483 Figure 4). Although RPKI certifications are illustrated only under ARIN in Figure 4, a similar
484 pattern is found in all other RIRs. Ideally, there should be a single root or trust anchor (TA) at
485 the top of the hierarchy, but currently, each of the five RIRs (AFRINIC, APNIC, ARIN, LACNIC, and
486 RIPE) maintains an independent TA for RPKI certification services in its respective region. Thus,
487 the global RPKI is currently operating with five TAs (see [ARIN1] [ARIN2] [RIPE1]). There are
488 various open-source Relying Party software tools available to perform RPKI validation [RIPE2]
489 [Routinator] [OctoRPKI] [FORT] [Phuntsho]. An analysis of the perceived legal barriers to the
490 adoption and use of RPKI services in the North American region is provided in [Wishnick] [Yoo].

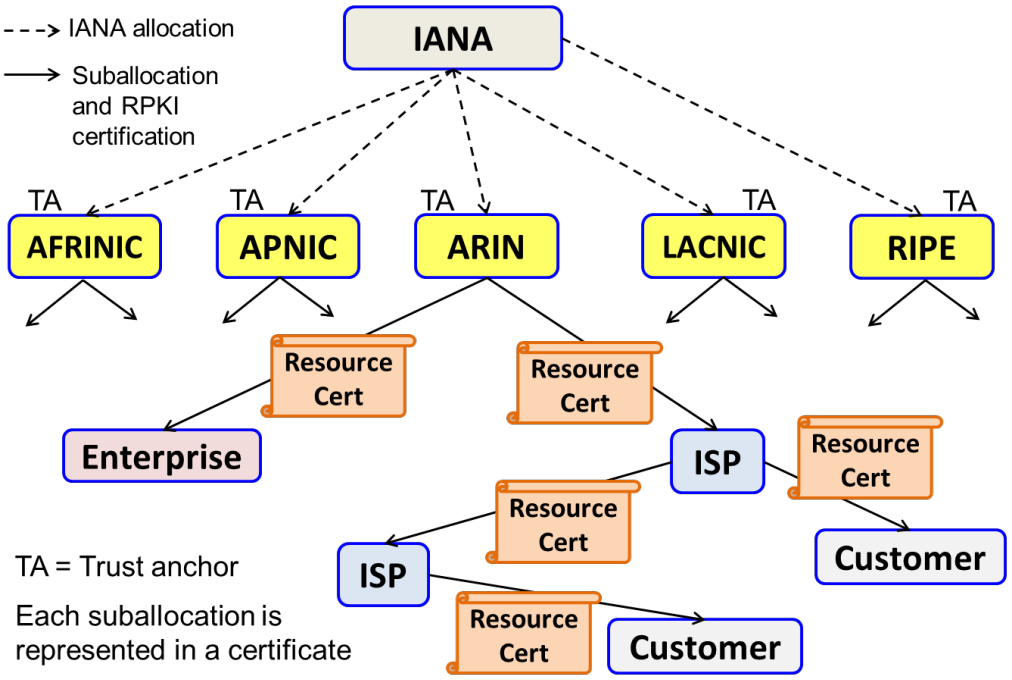

Figure 4. Illustration of resource allocation and certificate chain in RPKI

RPKI is based on the X.509 standard with RFC 3779 extensions that describe special certificate profiles for Internet number resources (prefixes and AS numbers) [RFC5280] [RFC6487] [RFC3779]. As shown in Figure 4, the RIRs issue resource certificates (i.e., certificate authority (CA) certificates) to ISPs and enterprises with registered number resource allocations and assignments. There are two models of resource certification: hosted and delegated [ARIN1] [RIPE1]. In the hosted model, the RIR keeps and manages keys and performs RPKI operations on their servers. In the delegated model, a resource holder (an ISP or enterprise) receives a CA certificate from their RIR, hosts their own certificate authority, and performs RPKI operations (e.g., signs route origin authorizations (see Section 4.3), issues subordinate resource certificates to their customers).

Table 2. Security recommendations related to resource certification

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 3: Internet number resource holders with IPv4/IPv6 prefixes and/or AS numbers (ASNs) should enroll those resources in the RPKI of the appropriate RIR so that RPKI certificates of those resources are issued.	X	X
Security Recommendation 4: Transit providers should provide a service where they facilitate creation, publication, and management of		X

Security Recommendation	Applicable to: Enter-prise	Applicable to: ISP
subordinate resource certificates for address space and/or ASNs suballocated to their customers. Note: Currently, RPKI services based on the hosted model and offered by RIRs are common. This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		
Security Recommendation 5: Legacy address space holders without an existing Registration Services Agreement with their RIR should establish an agreement and should enroll their number resources in the RPKI.	X	X

4.3. ROA-based Route Origin Validation (ROA-ROV)

This section describes route origin authorization (ROA) and ROA-based route origin validation (ROA-ROV) [RFC9582] [RFC6811] [RFC9319]. When reliable IRR data is available (see Section 4.1), ROA-ROV should be augmented with additional allowed {prefix, origin} pairs from the IRR data. There is also a proposal in the IETF for a new Signed Prefix List (SPL) object in RPKI and an ROV mechanism that combines ROA and SPL data [SPL-ROV]. Details of the SPL methodology [SPL-ROV] [SPL-profile] will be included in a future version of this document when the technology matures.

Once an address prefix owner obtains a CA certificate (Section 4.2), they can generate an end-entity (EE) certificate and use the private key associated with the EE certificate to digitally sign a route origin authorization (ROA) [RFC9582] [RFC6811] [RFC9319]. An ROA declares a specific AS as an authorized originator of BGP announcements for the prefix (see Figure 5). It specifies one or more prefixes (optionally a maxLength per prefix) and a single AS number. If a maxLength is specified for a prefix in the ROA, then a more-specific (i.e., longer) prefix (subsumed under the prefix) with a length not exceeding the maxLength is permitted to be originated from the specified AS. In the absence of an explicit maxLength for a prefix, the maxLength is equal to the length of the prefix itself. If the resource owner has a resource certificate listing multiple prefixes, they can create one ROA in which some or all those prefixes are listed. Alternatively, they can create one ROA per prefix.

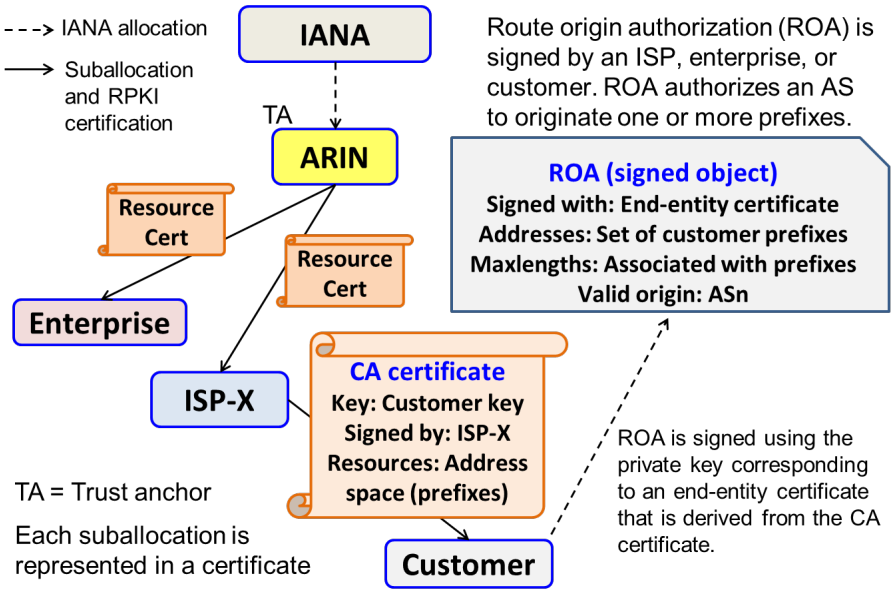

Figure 5. Creation of Route Origin Authorization (ROA) by prefix owner

ROAs can also be created (signed) by an ISP (transit provider) on behalf of its customer based on a service agreement, provided that the ISP suballocated the address space to the customer. The ISP can offer a service to its customers by creating and maintaining CA certificates for the customers' resources and ROAs for the customers' prefixes.

Once created, RPKI data is used throughout the Internet by relying parties (RPs). RPs, such as RPKI-validating servers, can access RPKI data from the repositories (see Figure 6) using either the rsync protocol [Rsync] [Rsync-RPKI] or the RPKI Repository Delta Protocol (RRDP) [RFC8182]. The RRDP protocol is often called "delta protocol" as shorthand. A BGP router typically accesses the required ROA data from one or more RPKI cache servers that are maintained by its AS. As shown in Figure 6, the RPKI-to-router protocol is used for communication between the RPKI cache server and the router [RFC8210] [RFC8210bis]. More details regarding secure routing architecture based on RPKI can be found in [RFC6480].

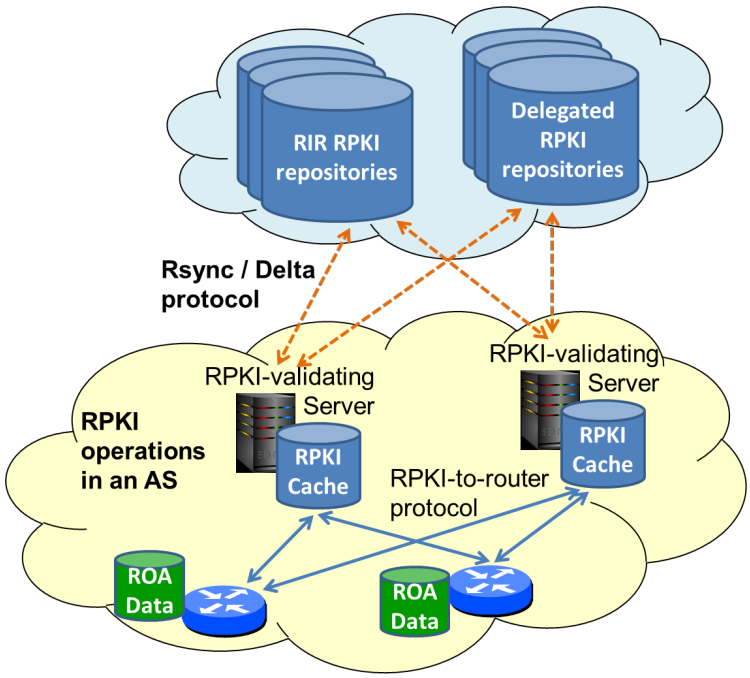

Figure 6. RPKI data retrieval, caching, and propagation to routers

A BGP router can use the ROA information retrieved from an RPKI cache server to mitigate the risk of prefix hijacks and some forms of route leaks in advertised routes. A BGP router would typically receive a validated list of {prefix, maxLength, origin AS} tuples (derived from valid ROAs) from one or more RPKI cache servers. This list may be called an allow-list. The router makes use of this list with the ROA-ROV process depicted in Figure 7 to determine the validation state of an advertised route [RFC6811]. A BGP route is deemed to have a "Valid" origin if the {prefix, origin AS} pair in the advertised route can be corroborated with the list (i.e., the pair is permissible in accordance with at least one ROA; see Figure 7 for the details). A route is considered "Invalid" if there is a mismatch with the list (i.e., AS number does not match, or the prefix length exceeds maxLength; see Figure 7 for additional details). Further, a route is deemed "NotFound" if the prefix announced is not covered by any prefix in the allow-list (i.e., there is no ROA that contains a prefix that equals or subsumes the announced prefix). When an AS_SET [RFC4271] is present in a BGP update, it is not possible to clearly determine the origin AS from the AS_PATH [RFC6811]. Thus, an update containing an AS_SET in its AS_PATH can never receive an assessment of "Valid" in the origin validation process (see Figure 7). The use of AS_SET (and AS_CONFED_SET) in BGP updates is prohibited [deprecate-as-set] (imminent IETF RFC). The ROA-based origin validation (ROA-ROV) may be supplemented by validation based on IRR data (see Section 4.1).

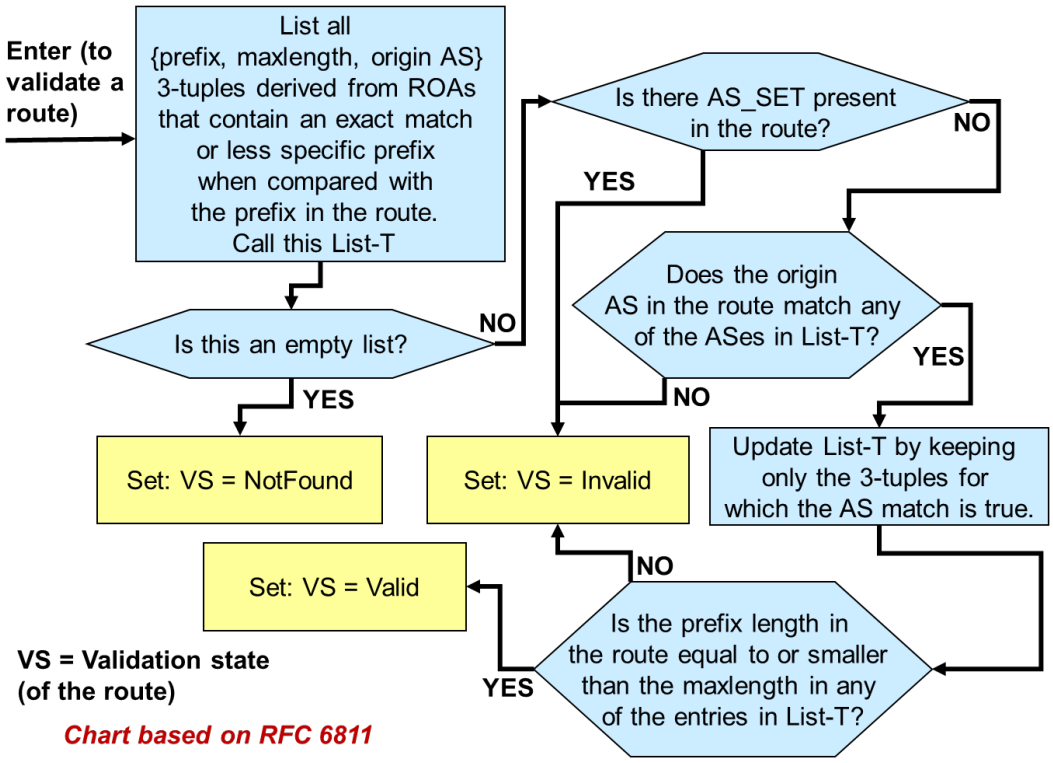

Figure 7. Algorithm for ROA-ROV (based on RFC 6811)

There are several implementations of ROA-ROV in commercial and open-source BGP router platforms [Juniper1] [Cisco1] [Patel] [Scudder] [NIST-SRx] [goBGP] [RTRlib]. Deployment guidance and configuration guidance for many of these implementations are available from several sources, including [NCCoE-sidr] [RIPE1] [MANRS]. Although ROA-ROV is already implemented in commercial BGP routers, the activation and ubiquitous use of RPKI and ROA-ROV in BGP routers require motivation and commitment on the part of network operators. Currently, 54% of unique {prefix, origin} pairs in routes propagated in the Internet are ROA-ROV Valid and about 0.4% are Invalid while the rest are NotFound [NIST-RPKI]. Network operators are turning on the ROA-ROV mechanism in their border routers, and some of them reject ROA-ROV Invalid routes (i.e., consider them ineligible for best path selection in BGP).

Table 3. Security recommendations related to ROA

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 6: IP address space holders should register ROA(s) in the global RPKI for all prefixes that are announced or intended to be announced on the public Internet.	X	X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 7: Each transit provider (ISP) should provide a service where they facilitate creation, publication, and management of ROAs for prefixes suballocated to their customers. Note: This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		X
Security Recommendation 8: If a prefix that is announced (or intended to be announced) is multi-homed and originated from multiple ASes, then one ROA for each originating AS should be registered for the prefix (possibly in combination with other prefixes which are also originated from the same AS).	X	X
Security Recommendation 9: When an ISP or enterprise announces multiple prefixes that include less-specific and more-specific prefixes, they should ensure that the more-specific prefixes have published ROAs before creating ROAs for the subsuming less-specific prefixes.	X	X
Security Recommendation 10: A transit provider (ISP) should ensure that more specific prefixes announced by ASes within its customer cone have ROAs prior to the creation of its own ROAs for subsuming less-specific prefix(es).		X

AS0 is a special AS number that is not allocated to any autonomous system. AS0 is also not permitted in routes announced in BGP. An AS0 ROA is one which has an AS0 in it for the originating AS [RFC6483] [APNIC1]. An address resource owner can create an AS0 ROA for their prefix to declare the intention that the prefix or any more-specific prefix subsumed under it must not be announced until and unless a normal ROA simultaneously exists for the prefix or the more-specific prefix.

Table 4. Security recommendations related to ROA-ROV

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 11: An ISP or enterprise should have AS0 ROA coverage for any prefixes that are currently not announced or intended to be announced to the public Internet. However, this should be done cautiously only after ensuring that ROAs exist for more-specific	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
prefixes (if any) that are subsumed by the afore-mentioned prefixes and are announced or intended to be announced.		
Security Recommendation 12: A BGP router should be compliant with [deprecate-as-set] (imminent IETF RFC) which prohibits the use of AS_SET and AS_CONFED_SET in BGP Updates.	X	X
Security Recommendation 13: ISPs and enterprises that operate BGP routers should also operate one or more RPKI-validating caches that generate validated and distilled RPKI data for use by routers.	X	X
Security Recommendation 14: BGP routers used for inter-domain routing should implement ROA-based Route Origin Validation (ROA-ROV) [RFC6811].	X	X

Concerning Security Recommendation 14, ROA-ROV is implemented by most of major router vendors. The allow-list of {prefix, maxLength, origin ASN} 3-tuples is typically obtained and periodically refreshed by a router from a local RPKI cache server. As mentioned before, the RPKI-to-router protocol [RFC8210] [RFC8210bis] is used for this communication.

How ROA-ROV results are used in path selection is strictly a local policy decision for each network operator. Policy choices include:

- Tag-Only – ROA-ROV results are only used to tag/log data about BGP routes for diagnostic purposes.
- Prefer-Valid – Use local preference settings to give priority to valid routes. Note that this is only a tie-breaking preference among routes with the exact same prefix.
- Reject-Invalid – Use local policy to consider invalid routes as ineligible in the BGP decision process.

With the goal of not allowing Invalid routes to propagate in the Internet, the policy stated in the last bullet above is recommended.

Table 5. Security recommendations related to route selection policy

Security Recommendation	Applicable to	
	Enter-prise	ISP
Security Recommendation 15: In partial/incremental deployment state of the RPKI, the permissible {prefix, origin ASN} pairs for performing	X	X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
BGP origin validation should be generated by taking the union of such data obtained from ROAs, IRR data, and customer contracts.		
Security Recommendation 16: ROA-ROV results should be incorporated into local BGP policy decisions to select best paths. Note: How ROA-ROV results are used in path selection is strictly a local policy decision for each network operator. However, considering a route that is ROA-ROV Invalid to be ineligible for best path selection is recommended.	X	X

4.3.1. Forged-Origin Hijacks — How to Minimize Them

With ROA-based origin validation alone, it is possible to prevent accidental misoriginations. However, a purposeful malicious hijacker can forge the origin AS of any update by prepending the number of an AS found in an ROA for the target prefix onto their own unauthorized BGP announcement. For greater impact, in conjunction with forging the origin, the attacker may replace the prefix in the route with a more-specific prefix (subsumed under the announced prefix) that has a length not exceeding the maxLength in the ROA.

Security Recommendation 17 provides some degree of robustness against forged-origin attacks, and Security Recommendation 18 provides a greater degree of robustness[1] against the same.

Table 6. Security recommendations related to maxLength

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 17: The maxLength in a ROA should not exceed the length of the most specific prefix (subsumed under the prefix in consideration) that is originated or intended to be originated from the AS listed in the ROA.	X	X
Security Recommendation 18: If a prefix and select more-specific prefixes subsumed under it are announced or intended to be announced, then instead of specifying a maxLength, the prefix and the more-specific prefixes should be listed explicitly in the ROA.	X	X

[1] BGPsec [RFC8205] described in Section 4.7 is required for full protection against prefix and/or path modifications.

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Note: In general, the use of maxLength should be avoided unless all or nearly all more-specific prefixes up to a maxLength are announced (or intended to be announced) [RFC 9139].		

4.3.2. General Recommendations Related to RPKI and ROA-ROV

Some general security recommendations are provided below that pertain to sharing with neighbors about RAO-ROV deployment status, ensuring that resource certificates and ROAs are renewed before their expiry dates, and making use of BGP/RPKI monitoring tools/services.

Table 7. General recommendations related to RPKI and ROA-ROV

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 19: If ROA-ROV is deployed in the BGP routers of an entity, they should share that information with their BGP peers. ISPs and large enterprises should publish information about the types of peer interfaces (customers, lateral peers, etc.) on which ROA-ROV is deployed.	X	X
Security Recommendation 20: Resource holders should ensure all their resource certificates, ROAs, and other RPKI signed objects are up to date. Any such objects with an impending expiration date should be renewed well ahead of their expiry. Note: At ARIN, RPKI resource certs are set with a two-year lifespan, and they auto-renew after one year, resetting the two-year lifespan [ARIN2].	X	X
Security Recommendation 21: Internet number resource holders should employ BGP/RPKI monitoring tools/services to remain informed about changes in the RPKI system that may affect their BGP route originations.	X	X

4.4. Categories of Prefix Filters

BGP prefix filtering (also known as route filtering) is the most basic mechanism for protecting BGP routers from accidental or malicious disruption [RFC7454]. Prefix filtering differs from BGP origin validation in that only the prefixes expected in a peering (e.g., customer) relationship are accepted, and prefixes not expected—including bogons and unallocated—are rejected. Further, origin validation is not a part of traditional prefix filtering, but it is complementary. Filtering capabilities on both incoming prefixes (inbound prefix filtering) and outgoing prefixes (outbound prefix filtering) should be implemented. Route filters are typically specified using a syntax similar to that used for access control lists. One option is to list ranges of IP prefixes that are to be denied and then permit all others. Alternatively, ranges of permitted prefixes can be specified, and the rest denied. The choice of which approach to use depends on practical considerations determined by system administrators. Typically, BGP peers should have matching prefix filters (i.e., the outbound prefix filters of an AS should be matched by the inbound prefix filters of peers that it communicates with). For example, if AS 64496 filters its outgoing prefixes towards peer AS 64500 to permit only those in set *P*, then AS 64500 establishes incoming prefix filters to ensure that the prefixes it accepts from AS 64496 are only those in set *P*.

Different types of prefix filters are described in the rest of Section 4.4, and their applicability is described in the context of different peering relations in Section 4.5.

4.4.1. Unallocated Prefixes

The Internet Assigned Numbers Authority (IANA) allocates address space to RIRs. All the IPv4 address space (or prefixes), except for some reserved for future use, have been allocated by IANA [IANA-v4-r]. The RIRs have also nearly fully allocated their IPv4 address space [IANA-v4-r]. [2] The IPv6 address space is much larger than that of IPv4, and, understandably, the bulk of it is unallocated. Therefore, it is a good practice to accept only those IPv6 prefix advertisements that have been allocated by the IANA [IANA-v6-r]. Network operators should ensure that the IPv6 prefix filters are updated regularly (normally, within a few weeks after any change in allocation of IPv6 prefixes). In the absence of such regular updating processes, it is better not to configure filters based on allocated prefixes. Team Cymru provides a service for updating bogon prefix lists for IPv4 and IPv6 [Cymru-bogon].

[2] Some of the prefixes are designated for special use as discussed in Section 4.4.2.

637 **Table 8. Security recommendation related to filtering unallocated prefixes**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 22: IPv6 routes should be filtered to permit only allocated IPv6 prefixes. Network operators should update IPv6 prefix filters regularly to include any newly allocated prefixes [Cymru-bogon]. Note: If prefix resource owners regularly register AS0 ROAs (see Section 4.3) for allocated (but possibly currently unused) prefixes, then those ROAs could be a complementary source for the update of prefix filters.	X	X

638 If prefix resource owners regularly register AS0 ROAs (see Section 4.3) for allocated (but
639 possibly currently unused) prefixes, then those ROAs could be a complementary source for the
640 update of prefix filters.

641 **4.4.2. Special Purpose Prefixes**

642 IANA maintains registries for special-purpose IPv4 and IPv6 addresses [IANA-v4-sp] [IANA-v6-
643 sp]. These registries also include specification of the routing scope of the special-purpose
644 prefixes.

645 **Table 9. Security recommendation related to filtering special-purpose prefixes**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 23: Prefixes that are marked “False” in column “Global” [IANA-v4-sp] [IANA-v6-sp] are forbidden from routing in the global Internet and should be rejected if received from an external BGP (eBGP) peer.	X	X

646 **4.4.3. Single-Homed Prefixes**

647 An AS may originate one or multiple prefixes. In the inbound direction, the AS should (in most
648 cases) reject routes for the prefixes (subnets) it originates if received from any of its eBGP peers
649 (transit provider, customer, or lateral peer). In general, the data traffic destined for these
650 prefixes should stay local and should not be leaked over external peering. However, if the AS
651 operator is uncertain whether a prefix they originate is single-homed or multi-homed, then the

652 AS should accept the prefix advertisement from an eBGP peer (and assign a lower local
653 preference value) so that the desired redundancy is maintained.

654 **Table 10. Security recommendation related to filtering single-homed prefixes**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 24: For single-homed prefixes (subnets) that are originated by an AS, any routes for those prefixes received at that AS from eBGP peers should be rejected.	X	X

655 ### 4.4.4. Prefixes that Exceed a Specificity Limit

656 Normally, ISPs neither announce nor accept routes for prefixes that are more specific than a
657 certain level of specificity. For example, maximum acceptable prefix lengths are mentioned in
658 existing practices as /24 for IPv4 [RIPE-399] and /48 for IPv6 [RIPE-532]. The level of specificity
659 that is acceptable is decided by each AS operator and communicated with peers. In instances
660 when Flowspec (see Section 5.5) [RFC8955] [RFC8956] [RFC9117] [Ryburn] is used between
661 adjacent ASes for DDoS mitigation, the two ASes may mutually agree to accept longer prefix
662 lengths (e.g., a /32 for IPv4) but only for certain pre-agreed prefixes. That is, the announced
663 more-specific prefix must be contained within a pre-agreed prefix.

664 **Table 11. Security recommendation related to prefixes that exceed a specificity limit**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 25: It is recommended that an eBGP router should set a route specificity limit for each eBGP peer and reject prefixes that exceed the specificity limit on a per-peer basis. Note: The specificity limit may be the same for all peers (e.g., /24 for IPv4 and /48 for IPv6).	X	X

665 Some operators may choose to reject prefix announcements that are less-specific than /8 and
666 /11 for IPv4 and IPv6, respectively.

4.4.5. Default Route

A route for the prefix 0.0.0.0/0 is known as the default route in IPv4, and a route for ::/0 is known as the default route in IPv6. The default route is advertised or accepted only in specific customer-provider peering relations. For example, a transit provider and a customer that is a stub or leaf network may make this arrangement between them whereby the customer accepts the default route from the provider instead of the full routing table. In general, filtering the default route is recommended except in situations where a special peering agreement exists.

Table 12. Security recommendation related to default route

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 26: The default route (0.0.0.0/0 in IPv4 and ::/0 in IPv6) should be rejected unless there is an explicit peering agreement that permits accepting it.	X	X

4.4.6. IXP LAN Prefixes

Typically, there is a need for the clients at an Internet exchange point (IXP) to have knowledge of the IP prefix used for the IXP LAN which facilitates peering between the clients.

Table 13. Security recommendation related to filtering IXP LAN prefixes

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 27: An Internet exchange point (IXP) should announce—from its route server to all its member ASes—its LAN prefix or its entire prefix, which would be the same as or less specific than its LAN prefix. Each IXP member AS should, in turn, accept this prefix from the IXP and reject any more-specific prefixes (of the IXP announced prefix) from any of its eBGP peers.	X	X

Implementing Security Recommendation 24 will ensure reachability to the IXP LAN prefix for each of the IXP members. It will also ensure that the Path Maximum Transmission Unit Discovery (PMTUD) will work between the members even in the presence of unicast Reverse Path Forwarding (uRPF). This is because the “packet too big” Internet Control Message Protocol

683 (ICMP) messages sent by IXP members' routers may be sourced using an IP address from the
684 IXP LAN prefix. See [RFC7454] for more details on this topic.

685 4.5. Prefix Filtering for Peers of Different Types

686 The inbound and outbound prefix filtering recommendations vary based on the type of peering
687 relationship that exists between networks: lateral peer, transit provider, customer, or leaf
688 customer (see definitions in Section 2.3). The different types of filters that apply are from the
689 list described in Sections 4.4.1 through 4.4.6.

690 The security recommendations that follow apply to ISPs. They also apply to enterprises when
691 they have eBGP peering with neighbor ASes.

692 4.5.1. Prefix Filtering with Lateral Peer

693 **Table 14. Security recommendations for prefix filtering with lateral peer**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 28: Inbound prefix filtering facing lateral peer – The following prefix filters (disallowed prefixes) should be applied in the inbound direction: • Unallocated prefixes • Special-purpose prefixes • Prefixes that the AS originates • Prefixes that exceed a specificity limit • Default route • IXP LAN Prefixes	X	X
Security Recommendation 29: Outbound prefix filtering facing lateral peer – The allowed outbound prefixes are those that are originated by the AS in question and those originated by its downstream ASes (i.e., the ASes in its customer cone). The following prefix filters should be applied in the outbound direction: • Unallocated prefixes • Special-purpose prefixes • Prefixes that exceed a specificity limit • Default route • IXP LAN prefixes • Prefixes learned from AS's lateral peers	X	X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
• Prefixes learned from AS's transit providers		

4.5.2. Prefix Filtering with Transit Provider

Table 15. Security recommendations for prefix filtering with transit provider

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 30: Inbound prefix filtering facing transit provider – Case 1 (full routing table): In general, when the full routing table is required from the transit provider, the following prefix filters should be applied in the inbound direction: • Unallocated prefixes • Special-purpose prefixes • Prefixes that the AS originates • Prefixes that exceed a specificity limit • IXP LAN prefixes	X	X
Security Recommendation 31: Inbound prefix filtering facing transit provider – Case 2 (default route): If the border router is configured for only the default route, then only the default route should be accepted from the transit provider and nothing else.	X	X
Security Recommendation 32: Outbound prefix filtering facing transit provider: The same outbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1). Note: In conjunction with the outbound prefix filtering security recommendation, some policy rules may also be applied if a transit provider is not contracted (or chosen) to provide transit for some subset of outbound prefixes.	X	X

4.5.3. Prefix Filtering with Customer

Inbound prefix filtering: There are two scenarios that require consideration. **Scenario 1** is when there is full visibility of the customer and its cone of customers (if any) as well as knowledge of prefixes that are originated from such a customer and its cone. The knowledge of prefixes can

be based on direct customer knowledge, IRR data, and/or ROA data (if that data is known to be in a complete and well-maintained state for the customer in consideration and its customer cone). The prefixes thus known for the customer and its customer cone are listed in the configuration of the eBGP router in question. **Scenario 2** is when there is not a reliable knowledge of all prefixes originated from the customer and its cone of customers.

Table 16. Security recommendations for prefix filtering with customer

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 33: Inbound prefix filtering facing customer in Scenario 1 (see Section 4.5.3) – Only the prefixes that are known to be originated from the customer and its customer cone should be accepted, and all other route announcements should be rejected.		X
Security Recommendation 34: Inbound prefix filtering facing customer in Scenario 2 (see Section 4.5.3) – The same set of inbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1).		X
Security Recommendation 35: Outbound prefix filtering facing customer: The filters applied in this case would vary depending on whether the customer wants to receive only the default route or the full routing table. If it is the former, then only the default route should be announced and nothing else. In the latter case, the following outbound prefix filters should be applied: • Special-purpose prefixes • Prefixes that exceed a specificity limit Note: The default route may be added to the above filter list if the customer requires the full routing table but not the default route.		X

4.5.4. Prefix Filtering Performed in a Leaf Customer Network

A leaf customer network is one which is single-homed to a transit provider and has no lateral peers or customer ASes downstream.

709 **Table 17. Security recommendations for prefix filtering performed in a leaf customer network**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 36: Inbound prefix filtering for leaf customer facing transit provider – A leaf customer may request only the default route from its transit provider. In this case, only the default route should be accepted and nothing else. If the leaf customer requires the full routing table from the transit provider, then it should apply the following inbound prefix filters: • Unallocated prefixes • Special-purpose prefixes • Prefixes that the AS (i.e., leaf customer) originates • Prefixes that exceed a specificity limitDefault route	X	
Security Recommendation 37: Outbound prefix filtering for leaf customer facing transit provider – A leaf customer network should apply a very simple outbound policy of announcing only the prefixes it originates. However, it may additionally apply the same outbound prefix filters as those for a lateral peer (see Section 4.5.1) for extra caution.	X	

710 ## 4.6. Role of RPKI in Prefix Filtering

711 An ISP can retrieve (from RPKI registries) all available route origin authorizations (ROAs)
712 corresponding to autonomous systems (ASes) that are known to belong in their customer cone
713 (see definition in Section 2.3).[3] From the available ROAs, it is possible to determine the prefixes
714 that can be originated from the ASes in the customer cone. As the RPKI registries become
715 mature with increasing adoption, the prefix lists derived from ROAs will become useful for
716 prefix filtering. Even in the early stages of RPKI adoption, the prefix lists (from ROAs) can help
717 cross-check and/or augment the prefix filter lists that an ISP constructs by other means.

[3] The list of ASes in an AS's customer cone can be determined by forming the list of unique ASes in all BGP announcements received (i.e., currently in the Adj-RIB-ins [RFC4271]) on all customer interfaces at the AS under consideration (see additional details in Section 5.1.7 and [BAR-SAV]). This can be done in the network management system (off the router).

718 **Table 18. Security recommendation for use of ROA data in prefix filtering**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 38: The ROA data (available from RPKI registries) should be used to construct and/or augment prefix filter lists for customer interfaces. Note: This Security Recommendation is possibly more applicable to smaller ISPs that have accurate visibility of their customer cone. Larger ISPs tend not to have such visibility.		X

719 ## 4.7. AS Path Verification

720 As observed in Sections 4.3 and 4.3.1, ROA-ROV is necessary but, by itself, is insufficient for fully
721 securing the prefix and AS path in BGP announcements. BGP path verification is additionally
722 required to protect against prefix modifications and forged-origin attacks (see Section 4.3.1) as
723 well as other AS-path attacks such as path shortening and Kapela-Pilosov attacks (see Section
724 2.2). There is significant interest in the networking community to secure the AS path in BGP
725 updates so that a more comprehensive protection can be provided to BGP Updates [RFC8205]
726 [RFC8608] [RFC7353] [RFC8374] [ASPA-profile] [ASPA-verif].

727 BGPsec is one available technology and IETF standard [RFC8205] for AS path verification.
728 Autonomous System Provider Authorization (ASPA) is another technology and emerging IETF
729 standard for AS path verification [ASPA-profile] [ASPA-verif]. Both try to achieve AS path
730 security in BGP using cryptographical protections. BGPsec carries cryptographic signatures on
731 the wire in the Update messages and the signatures are processed on the routers. In contrast,
732 the cryptography is off-line or off the router in the ASPA technology. This difference makes
733 ASPA more suitable for deployment in the short term due to the reduced processing burden on
734 the routers when compared to BGPsec. BGPsec provides full cryptographic protection to the AS
735 path itself but does not protect against route leaks. On the other hand, ASPA together with
736 another technology called Only to Customer (OTC) [RFC9234] provides strong protection
737 against route leaks (accidental as well as malicious), while it provides protection against some
738 but not all forms of AS path manipulations. Open-source prototype implementations of BGPsec
739 are available [NIST-SRx] [Adalier2]. However, commercial vendor implementations of BGPsec,
740 ASPA-based AS path verification, and OTC are in the proof of concept (POC) stage and therefore
741 not readily available for broad deployment. This section briefly describes these technologies
742 and standards. The security recommendations for them are labeled as future planning (FP)
743 since their deployment is not viable until commercial router vendor implementations are
744 available.

4.7.1. BGPsec Protocol (Emerging/Future)

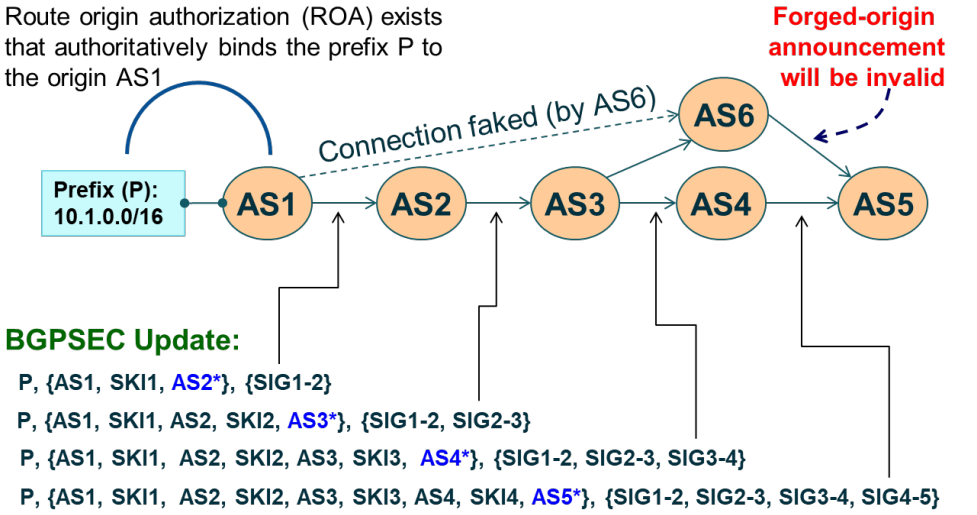

Figure 8. Basic principles of signing/verification of AS paths in BGP updates

The basic principles of BGPsec are illustrated in Figure 8 (see [RFC8205] for details). An ROA signed by the owner of the prefix 10.1.0.0/16 attests that AS1 is authorized to originate the prefix. Further, each network operator that has deployed BGPsec is given a resource certificate for their AS number, and the BGPsec routers within the AS are given router certificates and private keys for signing updates. The certificates for all BGPsec routers are retrieved by all participating ASes, and the public keys of all BGPsec routers are expected to be available at each BGPsec router. In Figure 8, AS1 uses its private key to generate its signature, SIG1-2, attesting that it sent a route for 10.1.0.0/16 to AS2. The target AS is included in the data that is under the signature. Likewise, AS2 signs the route to AS3 and so on. Each AS adds its signature as it propagates the update to its neighbors. The update includes the subject key identifier (SKI) for the public key of each AS in the path (i.e., the public key of the BGPsec router in the AS). AS5 receives an update with four signatures (one corresponding to each hop). If all signatures verify correctly at AS5, and the origin validation check also passes, then AS5 can be certain that the received update for 10.1.0.0/16 with AS path [AS1 (origin), AS2, AS3, AS4] is legitimate (i.e., not corrupted by prefix or path modifications along the way). For example, in Figure 8, AS6 would fail if it were to try to fake a connection to AS1 and announce a signed BGPsec update to AS5 (with a shorter path and a forged-origin AS1). This is because AS6 does not have an update signed to it directly from AS1.

The ECDSA-P256 algorithm is currently recommended for signing BGPsec updates between ASes that peer with each other [RFC8608]. Updates will have a larger size due to the addition of a 64-byte ECDSA P-256 signature for each hop. Also, the route processors in BGPsec routers will be required to perform additional processing due to signing and verification of path signatures.

The performance characterization of BGPsec quantifying routing information base (RIB) size and routing convergence time has been reported in [Sriram1]. High performance implementations of the cryptographic operations (ECC signing and verifications) associated with BGPsec update processing are available [Adalier1] [Adalier2] [NIST-SRx]. Optimization algorithms for BGPsec update processing are proposed and analyzed in [Sriram2]. BGPsec design choices and a summary of discussions leading to design decisions are presented in [RFC8374].

To reduce upgrade costs and encourage faster deployment, a leaf or stub AS is allowed to trust its upstream AS and negotiate to receive unsigned updates while it sends signed updates to the upstream AS [RFC8205].

The comprehensive set of standards for BGPsec are documented in [RFC8205] [RFC8206] [RFC8207] [RFC8608] [RFC8209] [RFC8210] [RFC8210bis]. For now, the security recommendation below concerning BGPsec is labeled as future planning (FP) since its deployment is not viable until router vendor implementations are available.

Table 19. Security recommendations (future) related to BGPsec

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation FP1: ASes should implement in their border routers the BGPsec-based AS path signing and verification procedures to protect AS paths in BGP Updates from path manipulations [RFC8205].	X	X

4.7.2. ASPA-based AS Path Verification (Emerging/Future)

The essential principles of the Autonomous System Provider Authorization (ASPA) object and ASPA-based AS path verification are described here. The details are available in [ASPA-profile] [ASPA-verif] [aspa-nanog89]. ASPA is a digitally signed object that is registered in an RPKI repository by a customer AS (CAS) to attest its set of provider ASes [ASPA-profile]. If an AS has no providers and is also not a route server (RS) client of a non-transparent IXP RS AS, it registers an AS0 ASPA, i.e., only AS 0 is included in the set of provider ASes (SPAS) field.

ASPA-based AS path verification is described in [ASPA-verif] [aspa-nanog89]. The AS path received by a receiving/verifying AS is represented as {AS(N), AS(N-1), …., AS(2), AS(1)}, where only the unique ASes are shown, N is the AS path length, AS(1) is the origin AS, and AS(N) is most recently added AS (Figure 10). Available ASPAs are cryptographically validated (X.509 validation) and from the validated ASPAs, the set provider ASes (SPAS) corresponding to each signing AS are obtained. An ASPA authorization check function for a pair of ASes {AS(i), AS(j)} as defined below (Figure 9) is used to verify the AS path.

auth(AS(i), AS(j)) = { **P** if AS(i) attests AS(j) is a provider; **nP** if AS(i) attests AS(j) is not a provider; **nA** if AS(i) does not have an ASPA

P: Provider
nP: not Provider
nA: no Attestation

Figure 9. ASPA authorization check function for a pair of ASes {AS(i), AS(j)}

With the help of Figure 10, the principle of detection of an Invalid (route leak) AS path can be explained for the case when the Update is received from a provider (i.e., in the downstream direction). The AS path is Invalid if there exist hops {AS(I), AS(I+1)} and {AS(J), AS(J-1)} with J > = I+2 such that both auth(AS(I), AS(I+1)) and auth(AS(J), AS(J-1)) are nP (see Figure 10). In this case, the AS path has a valley and hence it is a route leak (Section 2.3).

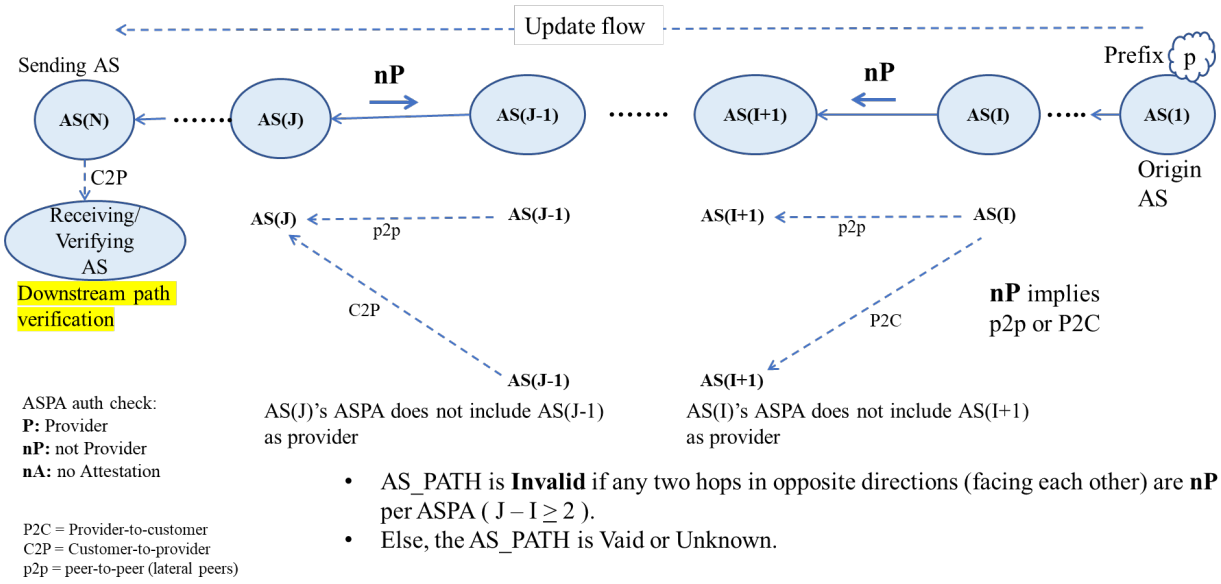

- AS_PATH is **Invalid** if any two hops in opposite directions (facing each other) are **nP** per ASPA (J – I ≥ 2).
- Else, the AS_PATH is Vaid or Unknown.

Figure 10. Basic principles of detection of Invalid AS path (route leak) using ASPA for downstream paths

With the help of Figure 11, the principle of detection of a Valid (i.e., not route leak) AS path can be explained for the case when the Update is received from a provider. If available ASPAs can establish that there are the Up-ramp and Down-ramp as illustrated in Figure 11 and there is no hop or just one hop (lateral peers) at the top between the apexes (AS(K) and AS(L)) of the two ramps, then the Update is Valid. If the Update in consideration was evaluated neither Valid nor Invalid per the described procedures, then it will be evaluated as Unknown (i.e., the ASPA data is insufficient due to partial deployment and the path validity cannot be ascertained).

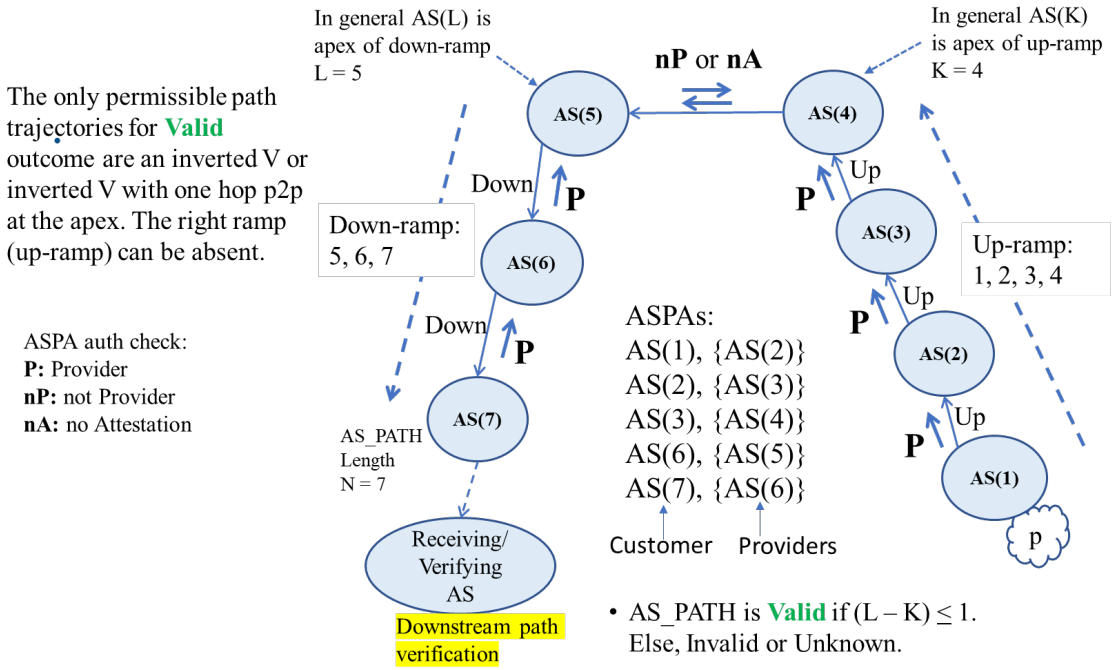

Figure 11. Basic principles of detection of Valid AS path (i.e., no route leak) using ASPA for downstream paths

If the Update is received from a customer or lateral peer (i.e., in the upstream direction), then the existence of even one hop for which auth$\{AS(I), AS(I+1)\}$ = nP is sufficient to evaluate the Update as Invalid. The Update will be evaluated as Valid only if each hop in the AS path is P, i.e., each hop is a C2P starting from $AS(1)$ to $AS(N)$. If the Update in consideration was evaluated neither Valid nor Invalid per the described procedure, then it will be evaluated as Unknown.

The algorithms for AS path verification using ASPA are described with additional considerations and details in [ASPA-verif]; slides with illustrations and video presentation are available in [aspa-nanog89]. For now, the security recommendations below concerning ASPA are labeled as future planning (FP) since its deployment is not viable until router vendor implementations are available.

Table 20. Security recommendations (future) related to ASPA

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation FP2: An AS owner should register its Autonomous System Provider Authorization (ASPA) object(s) per specification in [ASPA-prefix].	X	X
Security Recommendation FP3: Transit providers should provide a service where they facilitate creation, publication, and management of ASPAs for their customer ASes.		X

Security Recommendation	Applicable to	
	Enter-prise	ISP
Note: This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		
Security Recommendation FP4: ASes should deploy ASPA-based AS path verification and route leak mitigation procedures in their border routers per specification in [ASPA-verif].	X	X
Security Recommendation FP5: An AS owner doing ASPA should periodically check their own ASPA object(s) for correctness and completeness. They should also ensure that the same are renewed well before their expiry dates.	X	X
Security Recommendation FP6: An AS owner doing ASPA should periodically monitor all the ASPAs in the RPKI repositories to check if their AS number is incorrectly included as a provider in an ASPA (cryptographically valid), and if so, they should report it to the responsible party (or parties) so that the ASPA can be rectified.	X	X
Security Recommendation FP7: An AS owner doing ASPA should periodically monitor the ASPAs in the RPKI repositories to check if their AS number is incorrectly not included as a provider in the ASPA (cryptographically valid) of a customer AS, and if so, they should report it to the customer AS owner so that the ASPA can be rectified.	X	X

4.7.3. BGP Roles and OTC Attribute Solution for Route Leaks (Future)

A route leak solution technology using BGP Roles and the Only to Customer (OTC) Attribute has been standardized by the IETF (see [RFC9234]). This RFC specifies five BGP Roles: Provider, Customer, Route Server (RS), and RS Client, and Peer. Here Peer means the same as lateral peer. These Roles are initially locally configured for BGP peering sessions at an AS and are exchanged in the BGP OPEN messages using the BGP Role capability during a BGP session setup. The exchange of BGP Roles enables the cross-checking of the same between two neighbor ASes for the BGP session in consideration. If the exchanged BGP Roles indicate a mismatch, it means that the two neighbors are not in agreement about their BGP Roles, and they abstain from establishing the BGP session. That is, in this case, the BGP connection request is rejected using the Role Mismatch Notification [RFC9234]. If the exchanged BGP Roles match, the ASes proceed to establish the BGP session.

[RFC9234] also specifies a new Only to Customer (OTC) Attribute. The BGP Role value for the local AS and the OTC Attribute in BGP Update messages are used in the route leak prevention and detection procedures (Section 5 of RFC 9234). OTC contains the AS number (ASN) of the AS that attached it to the Update. The principle of OTC is that this Attribute is attached (if not

already present) by a compliant AS whenever an Update is advertised to a Customer, RS Client, or Peer. Subsequently, the Update with OTC can propagate to a customer or RS Client, but it must not be propagated to a Provider, RS, or Peer. If an Update with OTC is received from a Customer or RS Client, the routes conveyed in the Update are considered leaks and hence ineligible for path selection. If an Update with OTC is received from a Peer, the routes conveyed in the Update are considered leak and ineligible for path selection if the AS number (ASN) value in the OTC does not match the ASN of the Peer. If a route is received from a Provider, a Peer, or an RS and the OTC Attribute is not present, then it must be added (at ingress) with a value equal to the AS number of the remote AS (i.e., the neighbor AS that is sending the Update).

The OTC Attribute also helps to prevent the local AS from generating a route leak. This is because the presence of an OTC Attribute indicates to the egress router that the route was learned from a Provider, a Peer, or an RS, and it can be advertised only to the Customers.

There is at least one open-source implementation of RFC 9234 available [OpenBSD] and it has been deployed at some IXP RS ASes. Commercial implementations of RFC 9234 by major router vendors are still awaited. For now, the security recommendations concerning BGP Roles and OTC [RFC9234] are labeled as future planning (FP) since their deployment is not viable until router vendor implementations are available.

Table 21. Security recommendations (future) related to BGP Roles and OTC Attribute

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation FP8: ASes should implement in their border routers the procedures with BGP Roles as specified in [RFC9234].	X	X
Security Recommendation FP9: ASes should implement in their border routers the procedures with the OTC Attribute for route leak detection and mitigation as specified in [RFC9234].	X	X

4.8. Route Leak Solution Using BGP Community Tagging

Section 2.3 described the route leaks problem space and noted that in RFC 7908 [RFC7908], the various types of route leaks are enumerated. Section 2.3 also defined some basic terms used in discussions of route leaks. Route leak solutions fall into two categories: intra-AS and inter-AS (across AS hops). Many operators currently use an intra-AS solution, which is done by tagging BGP updates from ingress to egress (within the AS) using a BGP community [NANOG-list]. The BGP community used is non-transitive because it does not propagate in eBGP (between ASes). Each BGP update is tagged on ingress to indicate that it was received in eBGP from a customer, lateral peer, or transit provider. Further, a route that originated within the AS is tagged to indicate the same. At the egress point, the sending router applies an egress policy that makes

use of the tagging. Routes that are received from a customer are allowed on the egress to be forwarded to any type of peer (e.g., customer, lateral peer, or transit provider). However, routes received from a lateral peer or transit provider are forwarded only to customers (i.e., they are not allowed to be forwarded to a lateral peer or transit provider). These ingress and egress policies are central to route leak prevention within an AS (intra-AS).

Table 22. Security recommendations related to community tagging for intra-AS route leak prevention

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 39: An AS operator should have an ingress policy to tag routes internally (locally within the AS) to communicate from ingress to egress regarding the type of peer (customer, lateral peer, or transit provider) from which the route was received.	X	X
Security Recommendation 40: An AS operator should have an egress policy to utilize the tagged information (in Security Recommendation 37) to prevent route leaks when routes are forwarded on the egress. The AS should not forward routes received from a transit provider to another transit provider or a lateral peer. Also, the AS should not forward routes received from a lateral peer to another lateral peer or a transit provider.	X	X

The above intra-AS solution for the prevention of route leaks can also be implemented using a BGP attribute (instead of BGP community) – see description of the OTC Attribute in see Section 4.7.3 and [RFC9234]. The advantage of the attribute-based solution is that it can be made available in commercial routers as an RFC-standard feature, which in turn minimizes manual network operator actions. Note that the OTC Attribute based solution [RFC9234] (Section 4.7.3) is intra-AS as well as inter-AS solution for route leaks.

4.9. Checking AS Path for Disallowed AS Numbers

The AS path in an update received in eBGP is checked to make sure that there is no AS loop [RFC4271]. This is done by checking that the AS number of the local system does not appear in the received AS path. The AS path is also checked to ensure that AS numbers meant for special purposes [IANA-ASN-sp] are not present. Note that the special purpose ASN 23456 is allocated for AS_TRANS [RFC6793] and can be present in an AS_PATH in conjunction with an AS4_PATH [RFC6793] in the update.

891 **Table 23. Security recommendation related to checking AS path for disallowed AS numbers**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 41: The AS path in an update received in eBGP should be checked to ensure that the local AS number is not present. The AS path should also be checked to ensure that AS numbers meant for special purposes [IANA-ASN-sp] are not present. In case of a violation, the update should be rejected. Note: The special purpose ASN 23456 is allocated for AS_TRANS [RFC6793] and is allowed to be present in an AS_PATH in conjunction with an AS4_PATH [RFC6793] in the update.	X	X

892 ## 4.10. Generalized TTL Security Mechanism (GTSM)

893 Time to Live (TTL) is an 8-bit field in each IP packet and is decremented by one on each hop. The
894 Generalized TTL Security Mechanism (GTSM) [RFC5082] makes use of the TTL to provide an
895 additional security mechanism for BGP messages. Typically, a BGP session runs between
896 adjacent BGP routers, meaning BGP messages come from one hop away. Across such a BGP
897 session, the sending router sets TTL to 255 on each BGP message, and the receiving router
898 expects the incoming TTL to be 255 and rejects any BGP messages that have incoming TTL <
899 255. The expected TTL value in GTSM can be applied on a per-peer basis for each BGP session.
900 In rare instances, if a BGP session with a specific peer is known to run over n hops, then the
901 expected TTL for that session can be adjusted to a suitable value (255-n+1 in this case) in
902 accordance with the number of hops. Thus, GTSM helps detect and reject spoofed BGP
903 messages that may come from an attacker. Additional details regarding the operation of GTSM
904 can be found in [RFC5082].

905 **Table 24. Security recommendation related to GTSM**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 42: The Generalized TTL Security Mechanism (GTSM) [RFC5082] should be applied on a per-peer basis to provide protection against spoofed BGP messages.	X	X

4.11. Default External BGP Route Propagation Behavior without Policies

RFC 8212 emphasizes how critically important it is to explicitly configure import and export polices in eBGP. The following default behaviors are specified in [RFC8212]:

- Routes contained in an Adj-RIB-In associated with an eBGP peer SHALL NOT be considered eligible in the Decision Process if no explicit Import Policy has been applied.
- Routes SHALL NOT be added to an Adj-RIB-Out associated with an eBGP peer if no explicit Export Policy has been applied.

Once significant progress is made with implementation and operational experience with RFC 8212 recommendations, making those part of the security recommendations in this document (in a future revision) will be considered.

5. Source Address Validation and DDoS Mitigation[4]

There are various existing techniques and recommendations for deterrence against DDoS attacks with spoofed addresses [BCP38] [BCP84] [NABCOP] [CSRIC4-WG5]. Source address validation (SAV) of Internet Protocol (IP) packets is an effective anti-spoofing technique [BCP38] [BCP84]. BGP Flow Specification (Flowspec) [RFC8955] [RFC8956] [RFC9117] can also be used for DDoS mitigation. Employing a combination of these preventive techniques in enterprise and ISP border routers, hosted-service (Cloud) provider networks, broadband and wireless access networks, and data centers provides the necessary protections against DDoS attacks. The Spoofer project [Spoofer] [Luckie2] assesses and reports on the deployment of SAV in multiple dimensions: across time, autonomous systems, countries, and by IP version.

5.1. Source Address Validation Techniques

Source address validation (SAV) is performed in network edge devices, such as border routers, cable modem termination systems (CMTS) [RFC4036], digital subscriber line access multiplexers (DSLAM), and packet data network gateways (PGW) in mobile networks [PGW]. Ingress/egress access control lists (ACLs) and unicast Reverse Path Forwarding (uRPF) are techniques employed for implementing SAV [BCP38] [BCP84] [ISOC] [RFC6092; REC-5, REC-6] [Cisco3] [Juniper3]. Ingress SAV applies to incoming (received) packets, and egress SAV applies to outgoing (transmitted) packets.

Definitions of terms used in this section such as transit provider, lateral peer, peering relationship (C2P, p2p), and customer cone were provided in Section 2.3. In addition, the Reverse Path Forwarding list (RPF list) is defined as the list of permissible source-address prefixes for incoming data packets on a given interface.

5.1.1. SAV Using Access Control Lists

Ingress/egress access control lists (ACLs) are maintained with a list of acceptable (or alternatively, unacceptable) prefixes for the source addresses in the incoming/outgoing IP packets. Any packet with a source address that does not match the filter is dropped. The ACLs for the ingress/egress filters need to be maintained to keep them up to date. Hence, this method may be operationally difficult or infeasible in dynamic environments, such as when a customer network is multi-homed, has address space allocations from multiple ISPs, or dynamically varies its BGP announcements (i.e., routing) for traffic engineering purposes.

Typically, the egress ACLs in access aggregation devices (e.g., CMTS, DSLAM, PGW) permit source addresses only from the address spaces (prefixes) that are associated with the interface on which the customer network is connected. Ingress ACLs are typically deployed on border routers and drop ingress packets when the source address is spoofed (i.e., belongs to obviously

[4] Parts of the material in this section related to the review of existing SAV/uRPF technology read like corresponding parts in [RFC8704] since the authors worked on both documents and found it prudent to use the same or similar review material in both places. The IETF general rule is that original authors retain copyright. See https://trustee.ietf.org/reproduction-rfcs-faq.html .

disallowed prefix blocks—prefixes marked “False” in column “Global” [IANA-v4-sp] [IANA-v6-sp], the enterprise’s own prefixes, or the ISP’s internal-use only prefixes).

5.1.2. SAV Using Strict Unicast Reverse Path Forwarding

Terminology: In the figures (scenarios) in this section and the subsequent sections, the following terminology is used: "fails" means drops packets with legitimate source addresses; "works (but not desirable)" means passes all packets with legitimate source addresses but is oblivious to directionality; "works best" means passes all packets with legitimate source addresses with no (or minimal) compromise of directionality. Further, the notation Pi [ASn ASm ...] denotes a BGP update with prefix Pi and an AS_PATH as shown in the square brackets.

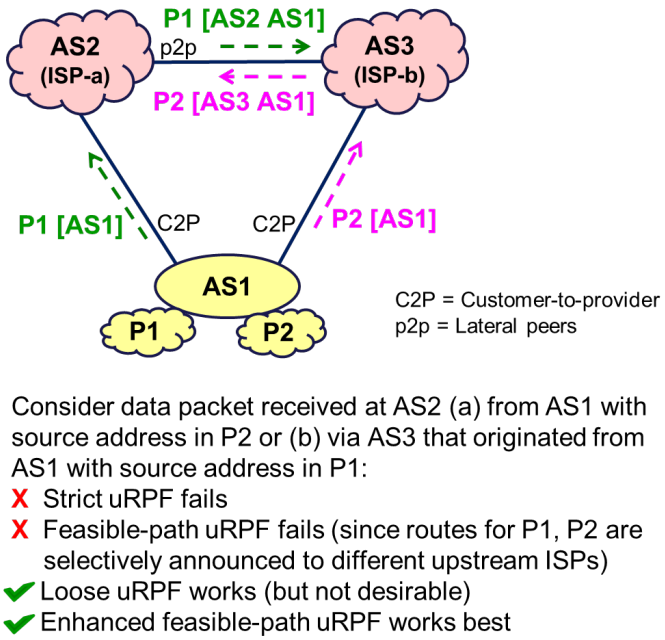

Figure 12. Scenario 1 for illustration of efficacy of uRPF schemes

In the strict unicast Reverse Path Forwarding (uRPF) method, an ingress packet on an interface at the border router is accepted only if the forwarding information base (FIB) contains a prefix that encompasses the source address and packet forwarding for that prefix points to the interface in consideration. In other words, the selected best path for routing to that source address (if it were used as a destination address) should point to the interface under consideration. This method has limitations when a network or autonomous system is multi-homed, routes are not symmetrically announced to all transit providers, and there is asymmetric routing of data packets. As an example, asymmetric routing occurs (see Figure 12, Scenario 1) when a customer AS announces one prefix (P1) to one transit provider (ISP-a) and a different prefix (P2) to another transit provider (ISP-b) but routes data packets with source addresses in the second prefix (P2) to the first transit provider (ISP-a) or vice versa. Then data packets with a source address in prefix P2 that are received at AS2 directly from AS1 will be

dropped. Further, data packets with a source address in prefix P1 that originate from AS1 and traverse via AS3 to AS2 will also be dropped at AS2.

5.1.3. SAV Using Feasible-Path Unicast Reverse Path Forwarding

The feasible-path uRPF helps partially overcome the problem identified with the strict uRPF in the multi-homing case. The feasible-path uRPF is similar to the strict uRPF, but in addition to inserting the best-path prefix, additional prefixes from alternative announced routes (on the interface under consideration) are also included in the RPF list (see definition at the top of Section 5.1). This method relies on either (a) announcements for the same prefixes (albeit some may be prepended to affect lower preference) propagating to all transit providers performing feasible-path uRPF checks or (b) announcement of an aggregate less-specific prefix to all transit providers while announcing more-specific prefixes (covered by the less-specific prefix) to different transit providers as needed for traffic engineering. As an example, in the multi-homing scenario (see Figure 13, Scenario 2), if the customer AS announces routes for both prefixes (P1, P2) to both transit providers (with suitable prepends if needed for traffic engineering), then the feasible-path uRPF method works. The feasible-path uRPF only works in this scenario if customer routes are preferred at AS2 and AS3 over a shorter non-customer route.

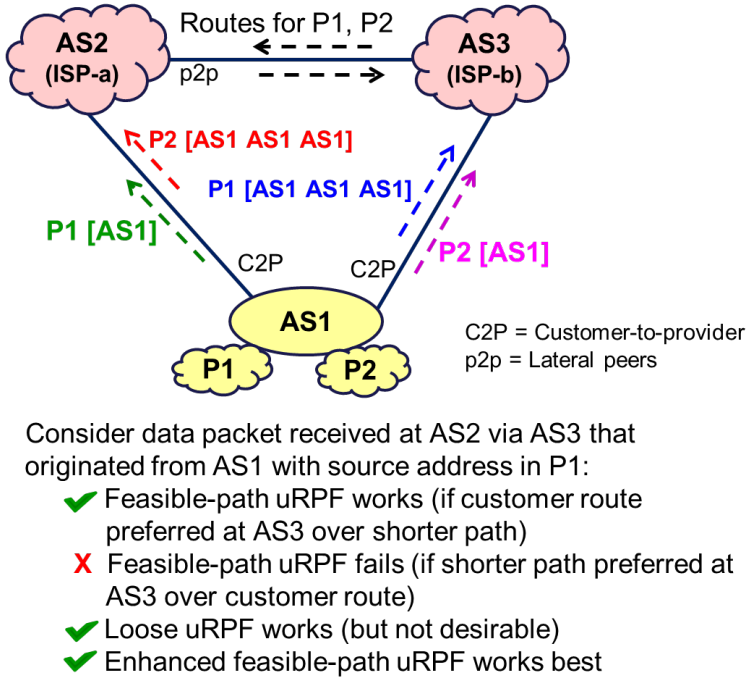

Figure 13. Scenario 2 for illustration of efficacy of uRPF schemes

However, the feasible-path uRPF method has limitations as well. One form of limitation naturally occurs when the recommendation of propagating the same prefixes (or combined address space) to all routers is not heeded. Another form of limitation can be described as follows: in Scenario 2 (illustrated in Figure 13), it is possible that the second transit provider AS3 (ISP-b) does not propagate the prepended route (i.e., P1 [AS1 AS1 AS1]) to the first transit

provider AS2 (ISP-a). This is because ISP-b's decision policy permits giving priority to a shorter route to prefix P1 via ISP-a over a longer route learned directly from the customer (AS1). In such a scenario, AS3 (ISP-b) would not send any route announcement for prefix P1 to AS2 (ISP-a). Then, a data packet originated from AS1 with a source address in prefix P1 that traverses via AS3 (ISP-b) will be dropped at AS2 (ISP-a).

5.1.4. SAV Using Loose Unicast Reverse Path Forwarding

In the loose unicast Reverse Path Forwarding (uRPF) method, an ingress packet at the border router is accepted only if the FIB has one or more prefixes that encompasses the source address. That is, a packet is dropped if no route exists in the FIB for the source address. Loose uRPF sacrifices directionality. This method is not very effective for preventing address spoofing. It only drops packets if the spoofed address is non-routable (e.g., belongs to obviously disallowed prefix blocks—prefixes marked “False” in column “Global” [IANA-v4-sp] [IANA-v6-sp], unallocated, or allocated but currently not routed). It may be noted that the method would seem more useful for IPv6 than IPv4.

5.1.5. SAV Using VRF Table

Virtual routing and forwarding (VRF) technology [RFC4364] [Juniper5] allows a router to maintain multiple routing table instances separate from the global routing information base (RIB). External BGP (eBGP) peering sessions send specific routes to be stored in a dedicated VRF table. The uRPF process queries the VRF table (instead of the FIB) for source address validation. A VRF table can be dedicated per eBGP peer and used for uRPF for only that peer, resulting in a strict mode operation. For implementing loose uRPF on an interface, the corresponding VRF table would be global (i.e., contains the same routes as in the FIB).

5.1.6. SAV Using Enhanced Feasible-Path uRPF (Emerging/Future)

The enhanced feasible-path uRPF (EFP-uRPF) method [RFC8704] provides a significant improvement in effectiveness and deployability over the feasible-path uRPF. This section briefly describes the technology and standards effort but does not make a security recommendation concerning the use of EFP-uRPF currently.

EFP-uRPF adds greater flexibility and accuracy to uRPF operations than the existing uRPF methods discussed in Sections 5.1.2 through 5.1.5. The basic principle of the EFP-uRPF method for enhancing efficacy in multi-homing and asymmetric routing scenarios is as follows. Looking at Figure 14, if a route for prefix P1 is received on customer interface X and has origin AS1, and routes for P2 and P3 are received on other peering interfaces Y and Z but have the same origin AS1, then allow the flexibility that data packets with a source address in any of these three prefixes (P1, P2, P3) may be legitimately received on customer interface X. Thus, based on the common origin AS principle, the prefix list for allowable source addresses in data packets (i.e., the RPF list) is expanded to include all three prefixes (P1, P2, P3) for customer interface X.

Further, the same principle is applied for determining the prefix list for allowable source addresses for each customer interface and possibly lateral peer interfaces.

As shown in Scenarios 1 and 2 (Figure 12 and Figure 13), the EFP-uRPF provides comparable or better performance than other uRPF methods for those scenarios. Scenario 3 (Figure 14) further illustrates that the EFP-uRPF method works best even in much more complex asymmetric routing scenarios. In Scenario 3 (Figure 14), the focus is on AS4 receiving data packets with a source address in {P1, P2, P3}. If the EFP-uRPF method (as described above) is used at AS4, then {P1, P2, P3} would be included in the RPF lists corresponding to the customer interfaces facing AS2 and AS3. Further, if EFP-uRPF is also applied at AS4 towards peer AS5, then {P1, P2, P3} would be included in the RPF list corresponding to the peer interface facing AS5. Thus, the operator (at AS4) can be assured that their SAV would work effectively, and none of the data packets originated from AS1 (and received via neighbors AS2, AS3, or AS5) with source addresses in {P1, P2, P3} would be denied due to the SAV. Thus, the EFP-uRPF method aims to eliminate or significantly reduce false positives regarding invalid detection in SAV compared to other uRPF methods. The details concerning EFP-uRPF can be found in [RFC8704]. Since it is still a work in progress, no security recommendations involving EFP-uRPF are offered here.

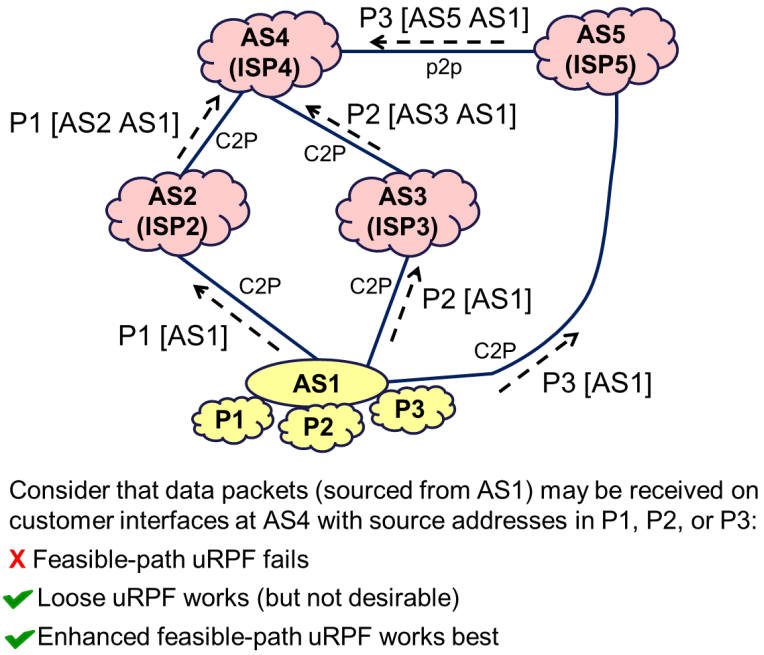

Figure 14. Scenario 3 for illustration of efficacy of uRPF schemes

5.1.7. SAV Using BAR-SAV (Emerging/Future)

BAR-SAV stands for SAV using BGP Updates, ASPA, and ROA. The BAR-SAV technique [BAR-SAV] [BAR-SAV-IETF121] is currently work in progress in the IETF and is an enhancement over the EFP-uRPF (Section 5.1.6). First, BAR-SAV improves on EFP-uRPF by making more efficient use of the BGP Update data. As illustrated in Figure 15, when an Update is received on a customer

interface with AS3, the BAR-SAV algorithms considers all ASes present in the AS_PATH (i.e., AS3 and AS1) to be within the customer cone of AS3. In the example in Figure 15, when AS4 learns on a different interface (i.e., the AS5 interface) that prefixes Q1 and P3 are originated by AS1 and AS3, respectively, it includes those prefixes also (along with P1) in the SAV filter allow-list towards AS3. The same principle is used also while designing a SAV filter for a lateral peer interface.

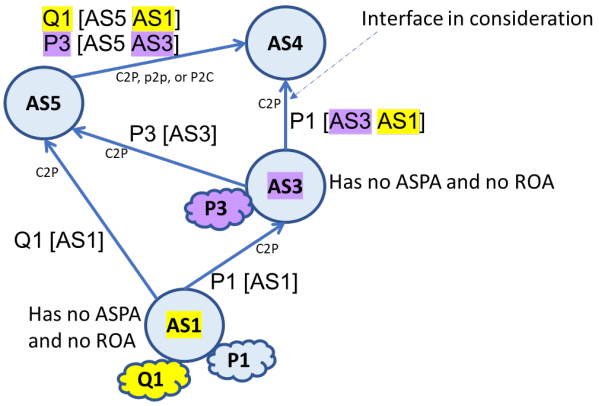

Figure 15. Refinement in BAR-SAV (over EFT-uRPF) for better utilization of BGP Update data

Second, BAR-SAV additionally improves on EFP-uRPF by making use of ASPA and ROA data pertaining to the customer cone (CC) in consideration (CC of AS3 in Figure 16). As illustrated in Figure 16, BAR-SAV makes complementary use of BGP, ASPA, and ROA data to find all ASes and prefixes in the CC of AS3. If an AS or prefix belonging in the CC is invisible in BGP Update data (possibly due to NO_EXPORT), BAR-SAV first finds the AS with help of ASPA data and then finds the prefixes associated with the AS with the help of ROA data. BAR-SAV has an efficient algorithm to first find the ASes at each level of hierarchy in the CC by recursively working its way from top to bottom. Here BGP and ASPA data are utilized. Once the list of ASes in the CC are found, the complete list of prefixes originated by those ASes or belonging to them are found from BGP and ROA data. (Note: ROAs registered with AS 0 as the origin AS are not used in the BAR-SAV procedures because such a ROA is used only for preventing squatting of allocated but unused prefixes.) Additional details of the BAR-SAV procedures can be found in [BAR-SAV]; slides with illustrations and video presentation are available in [BAR-SAV-IETF121].

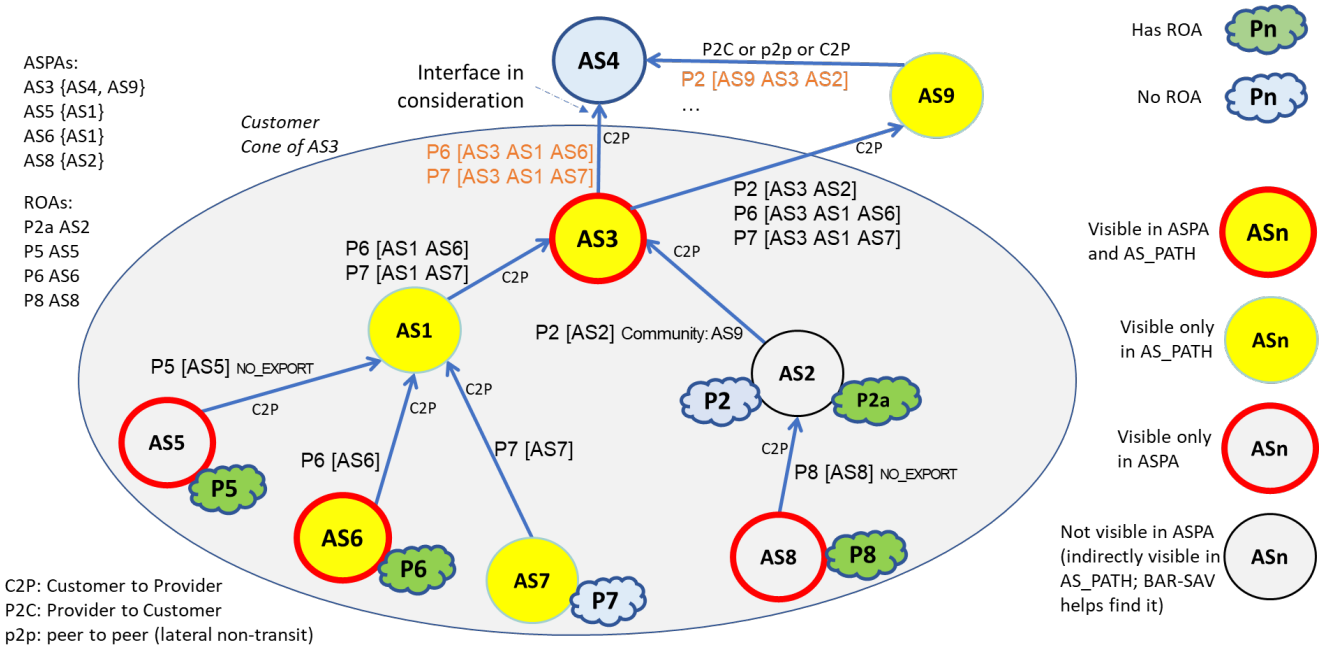

Figure 16. Efficient use of BGP Update, ASPA, and ROA data in BAR-SAV for discovery of source address prefixes

5.1.8. More Effective Mitigation with Combination of Origin Validation and SAV

With the combination of ROA-ROV (see Section 4.3) and the SAV (uRPF) techniques discussed above, a stronger defense against address spoofing and DDoS is made possible. A determined DDoS attacker can subvert any of the uRPF methods by performing prefix hijacking followed by source address spoofing as illustrated in Figure 17. In the scenario in Figure 17, the attacker first compromises routers (or perhaps owns some of them) at AS98 and AS99, and then falsely announces a less-specific prefix (e.g., 10.1.0.0/21) encompassing the target's prefix (e.g., 10.1.0.0/22). It is assumed that there is currently no legitimate announcement of the less-specific prefix (10.1.0.0/21). The feasible-path uRPF (FP-uRPF) filters at AS5 and AS6 are effectively deceived, and the attacker possibly stays under the radar because the hijacked prefix is a less-specific prefix. The attacker would then be able to successfully perform address spoofing and DDoS with reflection amplification. To protect against this type of multipronged attack, the combination of ROA-ROV (to prevent the hijacking) and FP-uRPF or EFP-uRPF (to prevent the address spoofing) should be employed. For this to work, the owners of the prefixes (10.1.0.0/22 and 10.1.0.0/21) should create ROAs, and all ASes (especially, AS5 and AS6) in Figure 17 should perform ROA-ROV in addition to employing SAV using the FP-uRPF/EFP-uRPF method.

1098

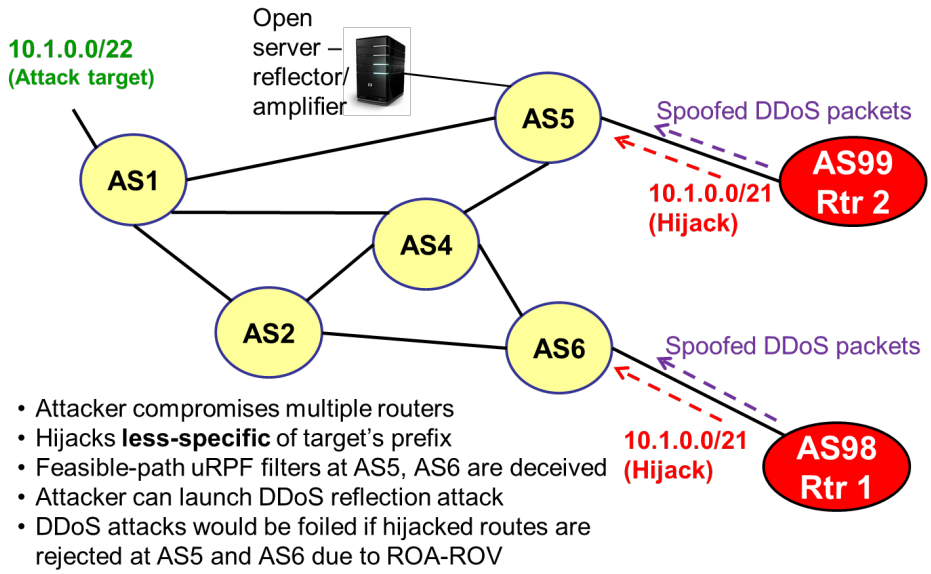

1099 **Figure 17. Illustration of how origin validation complements SAV**

1100 ## 5.2. SAV Recommendations for Various Types of Networks

1101 Three types of network scenarios are considered here, and SAV security recommendations are
1102 provided for each scenario. The network types are: 1) networks that have customers with
1103 directly connected allocated address space, such as broadband and wireless service providers;
1104 2) enterprise networks; and 3) Internet service providers (ISPs).

1105 When a government agency or enterprise procures the services of a hosted service provider or
1106 transit ISP, the security recommendations listed here should be considered for inclusion in the
1107 service contracts as appropriate.

1108 ### 5.2.1. Customer with Directly Connected Allocated Address Space: Broadband and Wireless
1109 Service Providers

1110 SAV with ACLs is relatively easy when a network served by an ISP's edge device (e.g., border
1111 router, CMTS, DSLAM, PGW) is directly connected and using an IP address space that is
1112 suballocated by the ISP. Hence, SAV using the ACL method should always be used in such cases.
1113 For the egress packets (i.e., packets transiting via the edge device onto the Internet), the source
1114 address must be within the allocated space. As an example, the Data Over Cable Service
1115 Interface Specification 3.1 (DOCSIS 3.1) standard for CMTS already incorporates this security
1116 check [DOCSIS] [Comcast] [RFC4036].

1117 **Table 25. Security recommendation related to SAV for directly connected customer**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 43: BGP routers that have single-homed directly connected customers, CMTS (or equivalent) in broadband access networks, and PGW (or equivalent) in mobile networks should implement SAV using ACLs (Section 5.1.1). The BGP routers in this context may alternatively use the strict uRPF method (Section 5.1.2).		X

1118 ### 5.2.2. Enterprise Border Routers

1119 The SAV security recommendations for enterprise border routers vary based on the
1120 egress/ingress nature of the data packets. Included here are recommendations concerning the
1121 routing control plane (BGP updates) as well.

1122 **Table 26. Security recommendations related to SAV for enterprise border routers**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 44: An enterprise border router that is multi-homed should always announce all its address space to each of its upstream transit providers to enable more effective SAV. This can be done in one of two ways: 1) announce an aggregate less-specific prefix to all transit providers and more-specific prefixes (covered by the less-specific prefix) to different transit providers as needed for traffic engineering, or 2) announce the same prefixes to each transit provider (albeit with suitable prepending for traffic engineering).	X	
Security Recommendation 45: This is the exception case when the enterprise border router does not adhere to Security Recommendation 41 and instead selectively announces some prefixes to one upstream transit ISP and other prefixes to another upstream transit ISP. In this case, the enterprise should route data (by appropriate internal routing) such that the source addresses in the data packets towards each upstream transit ISP belong in the prefix or prefixes announced to that ISP.	X	
Security Recommendation 46: On the ingress side (i.e., for data packets received from the transit ISP), enterprise border routers should deploy loose uRPF (Section 5.1.4) and/or ACLs (Section 5.1.1) to drop	X	

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the enterprise's own prefixes).		
Security Recommendation 47: An enterprise should allow on the egress side (i.e., for data packets sent to the transit ISP) only those packets with source addresses that belong in their own prefixes.	X	

5.2.3. Internet Service Providers

The SAV security recommendations for ISPs vary based on the ingress/egress of packets as well as the relationship with the peer (e.g., customer, lateral peer, transit provider).

Table 27. Security recommendations related to SAV for ISPs

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 48: On customer-facing interfaces, smaller ISPs should perform SAV on ingress packets by deploying the feasible-path uRPF (see Section 5.1.3). They should avoid using strict or loose uRPF as they are not effective, especially in the case of multi-homed customers. It is recognized that larger ISPs may use loose uRPF on customer interfaces.		X
Security Recommendation 49: For feasible-path uRPF to work appropriately, a smaller ISP (especially one that is near the Internet edge) should propagate all its announced address space to each of its upstream transit providers. This can be done in one of two ways: 1) announce an aggregate less-specific prefix to all transit providers and announce more-specific prefixes (covered by the less-specific prefix) to different transit providers as needed for traffic engineering, or 2) announce the same prefixes to each transit provider (albeit with suitable prepending for traffic engineering).		X
Security Recommendation 50: ISPs should prefer customer routes over other (i.e., transit provider or lateral peer) routes. (This is also normal ISP policy in most cases.)		X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Note: Following this recommendation facilitates a basis for adhering to Security Recommendation 48. It is also one of the stability conditions on BGP policy for ensuring stable convergence of routing information [Gao-Rexford].		
Security Recommendation 51: On interfaces with lateral (i.e., non-transit) peers, smaller ISPs (near the edge of the Internet) should perform SAV on ingress packets by deploying the feasible-path uRPF (see Section 5.1.3). They should avoid using strict or loose uRPF as they are not very effective for SAV on the lateral peer interfaces. It is recognized that larger ISPs may use loose uRPF on the interfaces with lateral peers.		X
Security Recommendation 52: On interfaces with transit providers, ISPs should perform SAV on ingress packets by deploying loose uRPF (see Section 5.1.4) and/or ACLs (see Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the ISP's internal-use only prefixes).		X
Security Recommendation 53: On the egress side towards customers, lateral (i.e., non-transit) peers, and transit providers, the ISP's border routers should deploy ACLs (see Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the ISP's internal-use only prefixes).		X

5.3. BGP Flow Specification (Flowspec)

Destination-based remotely triggered black-holing (D/RTBH) [RFC3882] [RFC7999] and source-based remotely triggered black-holing (S/RTBH) [RFC5635] (the latter in conjunction with uRPF) have been used as techniques for DDoS mitigation. However, with the standardization and vendor support of Flowspec [RFC8955] [RFC8956] [RFC9117] [Ryburn] [Cisco4] [Juniper4], the basic principles of D/RTBH and S/RTBH are significantly enhanced and can be operationally deployed in a fine-grained, dynamic, and efficient way. Operational experience with Flowspec for DDoS mitigation has been reported in [Levy] [Compton] [Hinze]. It may be noted that an updated version of Flowspec referred to as Flow Specification v2 (FSv2) is work in progress in the IETF [FSv2-ip-basic].

In D/RTBH, a BGP message is sent to trigger the provider edge (PE) routers (within the victim's AS or its transit provider AS) to block ingress traffic to the specified IP address where the affected server resides. In S/RTBH, a BGP message is sent to trigger the provider edge (PE)

routers (within the victim's AS or its transit provider AS) to block ingress traffic from the specified IP address that is the source address employed by the attacker. In S/RTBH, loose uRPF is used to filter traffic from the specified source address.

In the BGP Flowspec mechanism, flow specification NLRIs are defined and used to convey (intra-domain and inter-domain) traffic Flow Specifications for IPv4/IPv6 unicast and IPv4/IPv6 BGP/MPLS VPN services [RFC8955] [RFC8956]. The Flow Specification pertains to rate limiting or filtering IPv4/IPv6 protocol data packets. As an example, this mechanism can be used by a downstream AS (customer) to request an upstream AS (ISP) to perform inbound filtering in their edge routers on unwanted (suspected DoS) traffic. SAFI values 133 and 134 are assigned, respectively, to "Dissemination of Flow Specification rules" and "L3VPN Dissemination of Flow Specification rules" [RFC8955] [RFC8956]. Table 28 shows the Flow Spec Component Types for IPv4 that are defined in [RFC8955]. The same or similar names of these components apply to IPv6 also [RFC8956].

Table 28. BGP Flowspec component types

Type 1	Destination Prefix
Type 2	Source Prefix
Type 3	IP Protocol
Type 4	Source or Destination Port
Type 5	Destination Port
Type 6	Source Port
Type 7	ICMP Type
Type 8	ICMP Code
Type 9	TCP flags
Type 10	Packet length
Type 11	DSCP
Type 12	Fragment Encoding

In Table 29 below shows selected Traffic Filtering Action Extended Communities (EC) including the tuple {EC value, action, encoding}. Table 8 in [RFC8955] provides the full list.

Table 29. Extended community values defined in Flowspec to specify various types of actions

EC Value	Extended Community	Encoding
0x8006	traffic-rate-bytes (set to 0 to drop all traffic)	2-octet as#, 4-octet float
0x800c	traffic-rate-packets (set to 0 to drop all traffic)	2-octet as#, 4-octet float
0x8007	traffic-action	bitmask
0x8008	route-target redirect AS-2octet	2-octet AS, 4-octet value
0x8009	traffic-marking	DSCP value

In the table above, VRF stands for "virtual routing and forwarding," and DSCP stands for "differentiated services code point".

1159 **Table 30. Security recommendations related to RTBH and Flow Specification**

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 54: Edge routers should be equipped to perform destination-based remotely triggered black hole (D/RTBH) filtering and source-based remotely triggered black hole (S/RTBH) filtering.	X	X
Security Recommendation 55: Edge routers should be equipped to make use of BGP flow specification (Flowspec) to facilitate DDoS mitigation (in coordination between upstream and downstream autonomous systems).	X	X
Security Recommendation 56: Edge routers in an AS providing RTBH filtering should have an ingress policy towards RTBH customers to accept routes more specific than /24 in IPv4 and /48 in IPv6. Additionally, the edge routers should accept a more specific route (in case of D/RTBH) only if it is subsumed by a less-specific route that the customer is authorized to announce as standard policy (i.e., the less-specific route has a registered IRR entry and/or a ROA). Further, the edge routers should not drop RTBH-related more-specific route advertisements from customers even though BGP origin validation may mark them as “Invalid”.		X
Security Recommendation 57: A customer AS should make sure that the routes announced for RTBH filtering have NO_EXPORT, NO_ADVERTISE, or similar communities.	X	X
Security Recommendation 58: An ISP providing an RTBH filtering service to customers must have an egress policy that denies routes that have community tagging meant for triggering RTBH filtering at the local AS. This is an additional safeguard in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails.		X
Security Recommendation 59: An ISP providing an RTBH filtering service to customers must have an egress policy that denies prefixes that are longer than expected. This provides added safety in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails.		X

1160

6. General: Outsourced Services, Supporting Standards, Open Source, and Measurements

In this section, some security recommendations are mentioned that are of a general nature.

Table 31. Some general security recommendations

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation 60: Enterprises should require their Internet transit providers to adhere to the relevant security recommendations (from this document) by including them in service contracts.	X	
Security Recommendation 61: Enterprises that outsource applications/services (e.g., Email, DNS, cloud hosted systems, etc.) should require their outsource service providers to adhere to the relevant security recommendations (from this document) by including them in service contracts.	X	
Security Recommendation 62: Government agencies, ISPs, and enterprises should support standards development and open-source implementation efforts related to standards-based routing security technologies.	X	X
Security Recommendation 63: To the extent possible, ISPs and enterprises should facilitate collection of routing data by trusted organizations engaged in or supporting R&D efforts related to routing robustness and security monitoring.	X	X

References

[Adalier1] M. Adalier, K. Sriram, O. Borchert, K. Lee, and D. Montgomery, "High Performance BGP Security: Algorithms and Architectures", North American Network Operators Group (NANOG 69), Washington D.C, February 2017. Available at https://archive.nanog.org/sites/default/files/1_Sriram_High_Performance_Bgp_v1.pdf (slides). Available at https://www.youtube.com/watch?v=Yp03po5WJP0 (video)

[Adalier2] M. Adalier, "Efficient and Secure Elliptic Curve Cryptography Implementation of Curve P-256," NIST Workshop on ECC Standards, June 2015. Available at http://csrc.nist.gov/groups/ST/ecc-workshop-2015/papers/session6-adalier-mehmet.pdf

[APNIC1] G. Michaelson, "MyAPNIC RPKI service now supports AS0 ROA creation," APNIC technical note online, November 2018. Available at https://blog.apnic.net/2018/11/09/myapnic-rpki-service-now-supports-as0-roa-creation/

[Arbor] "DDoS Threat Intelligence Report", Available at https://www.netscout.com/threatreport

[Arbor2] "NETSCOUT Arbor's 14th Annual Worldwide Infrastructure Security Report" (2019). Available at https://www.netscout.com/sites/default/files/2019-03/SECR_005_EN-1901%E2%80%93WISR.pdf

[ARIN1] "Using RPKI at ARIN to certify resources," ARIN online. Available at https://www.arin.net/resources/rpki/using_rpki.html#hosted

[ARIN2] "Resource Public Key Infrastructure (RPKI) FAQs & Best Practices" ARIN online. Available at https://www.arin.net/resources/manage/rpki/faq/#what-is-the-lifespan-of-an-rpki-resource-certificate

[ARTEMIS] Automatic and Real-Time dEtection and Mitigation (ARTEMIS). Available at http://www.inspire.edu.gr/artemis/

[aspa-nanog89] K. Sriram, "ASPA-based BGP AS_PATH Verification and Route Leaks Solution," Presented at NANOG 89, San Diego, USA, October 2023. Available at https://storage.googleapis.com/site-media-prod/meetings/NANOG89/4809/20231017_Sriram_Aspa-Based_Bgp_As_Path_v1.pdf (slides). Available at https://www.youtube.com/watch?v=GdVnZGd7jMo (video)

[ASPA-profile] A. Azimov, E. Uksov, R. Bush, J. Snijders, R. Housley, and B. Maddison, "A Profile for Autonomous System Provider Authorization," Internet Engineering Task Force, Internet Draft, June 2024. Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-profile/

[ASPA-verif] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects," Internet Engineering Task Force, Internet

Draft, July 2024. Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification

[Azure] "Anatomy of a DDoS amplification attack", blog by Azure Network Security Team, May 2022. Available at https://www.microsoft.com/en-us/security/blog/2022/05/23/anatomy-of-ddos-amplification-attacks/?msockid=141a63b6f426691714d376e9f5696878

[BAR-SAV] K. Sriram, I. Lubashev, and D. Montgomery, "Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)," Internet Engineering Task Force (IETF) Internet Draft, July 2024. Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-bar-sav

[BAR-SAV-IETF121] K. Sriram, I. Lubashev, and D. Montgomery, "Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)," Proceedings of the IETF 121, November 2021. Available at https://datatracker.ietf.org/meeting/121/materials/slides-121-savnet-update-on-the-bar-sav-draft-00 (slides). Available at https://youtu.be/VoN-DdoXF0U?si=iVsSfhu3lpbbEx4z&t=3642 (video)

[BCP38] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," BCP 38 (RFC 2827), May 2000. Available at https://tools.ietf.org/html/bcp38

[BCP84] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," BCP 84 (RFC 3704), March 2004. Available at https://tools.ietf.org/html/bcp84

[BGPmon] BGPmon. Available at https://bgpmon.net/

[BGPStream] BGPStream. Available at https://bgpstream.caida.org/

[BITAG] Security of the Internet's Routing Infrastructure, Broadband Internet Technical Advisory Group, November 2, 2022. Available at https://www.bitag.org/Routing_Security.php

[Bjarnason] S. Bjarnason, "Withstanding the Infinite: DDoS Defense in the Terabit Era," Presentation at NANOG-74, October 2018. Available at https://pc.nanog.org/static/published/meetings/NANOG74/1789/20181001_Bjarnason_Withstanding_The_Infinite__v1.pdf

[Botnet-Roadmap] "A Road Map Toward Resilience Against Botnets," Joint US DOC/DHS report, November 2018. Available at https://www.commerce.gov/sites/default/files/2018-11/Botnet%20Road%20Map%20112918%20for%20posting_0.pdf

[CableLabs] "Cybersecurity Framework Profile for Internet Routing," CableLabs Security, October 2024. Available at https://www.cablelabs.com/specifications/CL-GL-RS-Profile

[CC-DDoS-Resp] "Cybersecurity Framework DDoS Threat Mitigation Profile," Cybersecurity Coalition. Available at https://www.cybersecuritycoalition.org/frameworks/ddos-profile

[Chung] T. Chung, et al., "RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins," Proceedings of the Internet

Measurement Conference, Pages 406-419, October 2019. Available at https://dl.acm.org/citation.cfm?id=3355596

[CISA-DDoS-Resp] "Understanding and Responding to Distributed Denial-of-Service Attacks." Available at http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-origin-as.pdf https://www.cisa.gov/sites/default/files/2024-03/understanding-and-responding-to-distributed-denial-of-service-attacks_508c.pdf

[Cisco1] "BGP—Origin AS Validation." Available at http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-origin-as.pdf

[Cisco3] "Unicast reverse path forwarding enhancements for the Internet service provider—Internet service provider network edge," Cisco WP. Available: http://www.cisco.com/c/dam/en_us/about/security/intelligence/urpf.pdf

[Cisco4] "Routing Configuration Guide for Cisco ASR 9000 Series Routers, IOS XR Release 7.8.x – Chapter: Implementing BGP Flowspec," Cisco Configuration Guides, November 2022. Available at https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r7-8/routing/configuration/guide/b-routing-cg-asr9000-78x/implementing-bgp-flowspec.html

[Cloudflare-RPKI] Cloudflare's RPKI Monitor. Available at https://rpki.cloudflare.com/

[Comcast] "Comcast network management (prevent network spoofing)", November 2023. Available at https://www.xfinity.com/networkmanagement/oldarticles

[Compton] R. Compton, T. Bowlby, T. Harris, P. Lotia, "eBGP Flowspec Peering for DDoS Mitigation," NANOG 75, February 2019. Available at https://pc.nanog.org/static/published/meetings/NANOG75/1887/20190219_Compton_Ebgp_Flowspec_Peering_v1.pdf

[CSDE] "Cyber Crisis: Foundations of Multi-Stakeholder Coordination," Council for Secure Digital Economy (CSDE) report (2019). Available at https://securingdigitaleconomy.org/wp-content/uploads/2019/09/CSDE_CyberCrisis-Report_2019-FINAL.pdf

[CSRIC4-WG5] "Remediation of Server-Based DDoS Attacks," CSRIC IV Working Group 5 final report, September 2014. Available at https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG5_Remediation_of_Server-Based_DDoS_Attacks_Report_Final_(pdf)_V11.pdf

[CSRIC4-WG6] "Long-Term Core Internet Protocol Improvements," CSRIC IV Working Group 6 presentation, September 2014. Available at https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG6_Presentation_09242014.pdf

[CSRIC6-WG3] "Report on Best Practices and Recommendations to Mitigate Security Risks to Current IP-based Protocols," CSRIC VI Working Group 3 final

report, March 2019. Available at https://www.fcc.gov/files/csric6wg3finalreport030819pdf

[Cymru-bogon] "Bogon route server project: Bogons via BGP." Available at http://www.team-cymru.org/bogon-reference-bgp.html

[deprecate-as-set] W. Kumari, K. Sriram, J. Hass, L. Hannachi, "Deprecation of AS_SET and AS_CONFED_SET in BGP," IETF Internet Draft (imminent IETF RFC). Available at https://datatracker.ietf.org/doc/draft-ietf-idr-deprecate-as-set-confed-set/

[DOC-Botnet] U.S. Department of Commerce, U.S. Department of Homeland Security, "A Report to the President on Enhancing the Resilience of the Internet and Communications Ecosystem Against Botnets and Other Automated, Distributed Threats," May 22, 2018. Available at https://csrc.nist.gov/publications/detail/white-paper/2018/05/30/enhancing-resilience-against-botnets--report-to-the-president/final

[DOCSIS] "DOCSIS® 3.1 Technology", CableLabs. Available at https://www.cablelabs.com/technologies/docsis-3-1

[ENISA] "7 Steps to shore up the Border Gateway Protocol (BGP)", the EU Cybersecurity Agency, May 2019. Available: https://www.enisa.europa.eu/publications/7-steps-to-shore-up-bgp

[FCC-NPR] Reporting on Border Gateway Protocol Risk Mitigation Progress, PS Docket No. 24-146, Notice of Proposed Rulemaking, FCC 24-62, June 7, 2024. Available at https://docs.fcc.gov/public/attachments/FCC-24-62A1.pdf

[FISMA2002] Federal Information Security Management Act of 2002, Pub. L. 107-347 (Title III), 116 Stat. 2946. Available at http://www.gpo.gov/fdsys/pkg/PLAW-107publ347/pdf/PLAW-107publ347.pdf.

[FISMA2014] Federal Information Security Modernization Act of 2014, Pub. L. 113-283, 128 Stat. 3073. Available at http://www.gpo.gov/fdsys/pkg/PLAW-113publ283/pdf/PLAW-113publ283.pdf.

[FORT] FORT RPKI validator. Available at https://github.com/NICMx/FORT-validator

[FSv2-ip-basic] S. Hares, D. E. Eastlake, J. Dong, C. Yadlapalli, and S. Maduscke, "BGP Flow Specification Version 2 - for Basic IP," IETF Internet-Draft draft-ietf-idr-fsv2-ip-basic-02, October 2024. Available at https://datatracker.ietf.org/doc/draft-ietf-idr-fsv2-ip-basic/

[Gao-Rexford] Freedman, M., "Interdomain Routing Policy", Princeton University COS 461 Lecture Notes; Slides 25-27, Spring 2011. Available at http://www.cs.princeton.edu/courses/archive/spr11/cos461/docs/lec17-bgp-policy.ppt

[goBGP] Use of Resource Public Key Infrastructure (RPKI) server to do Origin AS Validation in goBGP. Available at https://github.com/osrg/gobgp/blob/master/docs/sources/rpki.md

[HelpNet] "DNS amplification attacks double in Q1 2018," Help Net Security blog, June 2018. Available at https://www.helpnetsecurity.com/2018/06/14/dns-amplification-attacks-q1-2018/

[Hinze] N. Hinze, M. Nawrocki, M. Jonker, A. Dainotti, T.C. Schmidt, M. Wählisch, "On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP," In: Proc. of ACM SIGCOMM. Poster Session, pp. 57--59, New York, NY, USA: ACM, August 2018. Available at http://www.caida.org/publications/papers/2018/potential_bgp_flowspec/potential_bgp_flowspec.pdf

[Huston2012] G. Huston, "Leaking Routes," Asia Pacific Network Information Centre (APNIC) Blog, March 2012. Available at http://labs.apnic.net/blabs/?p=139/

[Huston2016] G. Huston, "Taking a Closer Look at the Recent DDoS Attacks and What It Means for the DNS," CircleID Blog, October 2016. Available at http://www.circleid.com/posts/20161026_closer_look_at_recent_ddos_attacks_and_what_it_means_for_dns/

[IANA-ASN-sp] "Special-Purpose Autonomous System (AS) Numbers" IANA web page. Available at https://www.iana.org/assignments/iana-as-numbers-special-registry/iana-as-numbers-special-registry.xhtml

[IANA-v4-r] "IANA IPv4 Address Space Registry," IANA web page. Available at http://www.iana.org/assignments/ipv4-address-space

[IANA-v6-r] "Internet Protocol Version 6 Address Space," IANA web page. Available at http://www.iana.org/assignments/ipv6-address-space

[IANA-v4-sp] "IANA IPv4 Special-Purpose Address Registry," IANA web page. Available at https://www.iana.org/assignments/iana-ipv4-special-registry

[IANA-v6-sp] "IANA IPv6 Special-Purpose Address Registry," IANA web page. Available at http://www.iana.org/assignments/iana-ipv6-special-registry

[IETF-GROW] IETF Global Routing Operations (GROW) Working Group. Available at https://datatracker.ietf.org/wg/grow/documents/

[IETF-IDR] IETF Inter-Domain Routing (IDR) Working Group. Available at https://datatracker.ietf.org/wg/idr/documents/

[IETF-OPSEC] IETF Operational Security Capabilities for IP Network Infrastructure (OPSEC) Working Group. Available at https://datatracker.ietf.org/wg/opsec/documents/

[IETF-SIDR] IETF Secure Inter-Domain Routing (SIDR) Working Group. Available at https://datatracker.ietf.org/wg/sidr/documents/

[IETF-SIDROPS] IETF Secure Inter-Domain Routing Operations (SIDROPS) Working Group. Available at https://datatracker.ietf.org/wg/sidrops/documents/

[ISOC] P. Vixie (Ed.), “Addressing the challenge of IP spoofing,” ISOC report, September 2015. Available at https://www.internetsociety.org/wp-content/uploads/2017/08/ISOC-AntiSpoofing-20150909-en-2.pdf

[ISTR-2015] Internet Security Threat Report 2015, Volume 20, Symantec Corporation, Mountain View, CA, April 2015. Available at https://www.symantec.com/content/en/us/enterprise/other_resources/21347933_GA_RPT-internet-security-threat-report-volume-20-2015.pdf

[ISTR-2016] Internet Security Threat Report 2016, Volume 21, Symantec Corporation, Mountain View, CA, April 2016. Available at https://www.symantec.com/content/dam/symantec/docs/reports/istr-21-2016-en.pdf

[ISTR-2017] Internet Security Threat Report 2017, Volume 22, Symantec Corporation, Mountain View, CA, April 2017. Available at https://www.symantec.com/content/dam/symantec/docs/reports/istr-22-2017-en.pdf

[Juniper1] “Example: Configuring Origin Validation for BGP,” Juniper blog. Available at http://www.juniper.net/techpubs/en_US/junos12.2/topics/topic-map/bgp-origin-as-validation.html

[Juniper3] “Understanding Unicast RPF (Routers),” Juniper blog., July 2024. Available at https://www.juniper.net/documentation/us/en/software/junos/security-services/topics/topic-map/interfaces-configuring-unicast-rpf.html

[Juniper4] “Example: Enabling BGP to Carry Flow-Specification Routes,” Juniper TechLibrary. Available at https://www.juniper.net/documentation/en_US/junos12.3/topics/example/routing-bgp-flow-specification-routes.html

[Juniper5] “Creating Unique VPN Routes Using VRF Tables,” November 2023. Available at https://www.juniper.net/documentation/en_US/junos/topics/topic-map/l3-vpns-routes-vrf-tables.html#id-understanding-virtual-routing-and-forwarding-tables

[Kapela-Pilosov] A. Pilosov and T. Kapela, "Stealing the Internet: An Internet-Scale Man in the Middle Attack", 16th Defcon Conference, August 2008. Available at https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf

[Levy] N. Levy, D. Smith, and J. Schiel, “Bi-Lateral Security Management Framework (a.k.a. DDoS peering),” NANOG 71, October 2017. Available at https://pc.nanog.org/static/published/meetings/NANOG71/1447/20171003_Levy_Operationalizing_Isp_v2.pdf

[Luckie] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k. claffy, “AS Relationships, Customer Cones, and Validation,” Proceedings of the 2013 ACM Internet Measurement Conference (IMC), DOI

10.1145/2504730.2504735, October 2013. Available at http://www.caida.org/~amogh/papers/asrank-IMC13.pdf

[Luckie2] M. Luckie, R. Beverly, R. Koga, K. Keys, J. Kroll, and k. claffy, "Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet", in ACM Computer and Communications Security (CCS), Nov 2019. Available at https://dl.acm.org/citation.cfm?id=3354232

[Madory] D. Madory, "A Brief History of the Internet's Biggest BGP Incidents," Kentik blog, June 2023. Available: https://www.kentik.com/blog/a-brief-history-of-the-internets-biggest-bgp-incidents/

[MANRS] "Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide." Available at https://www.manrs.org/isps/guide/

[MANRS2] "State of Routing Security," MANRS Observatory. Available at https://observatory.manrs.org/#/overview

[Merit-RADb] "Merit RADb" (Merit Network Inc.). Available at http://www.radb.net

[Mirai1] "Mirai: what you need to know about the botnet behind recent major DDoS attacks," Symantec Security Response, October 27, 2016. Available at https://www.symantec.com/connect/blogs/mirai-what-you-need-know-about-botnet-behind-recent-major-ddos-attacks

[Mirai2] "Dyn Analysis Summary of Friday October 21 Attack," Dyn Company News, October 26, 2016. Available at https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/

[Murphy] S. Murphy, "RPKI Tutorial: Routing Security and RPKI", NANOG on the Road (NOTR), St. Louis, MO, November 2015. Available at https://www.nanog.org/sites/default/files/04-Murphy-StLouis.pdf

[NABCOP] "DDoS-DoS-attack-BCOP," North American BCOP. Available at http://nabcop.org/index.php/DDoS-DoS-attack-BCOP

[Naik] A. Naik, "Internet Vulnerability Takes Down Google," ThousandEyes report, November 2018. Available at https://blog.thousandeyes.com/internet-vulnerability-takes-down-google/

[NANOG] "Practical BGP Origin Validation using RPKI: Vendor Support, Signing and Validation Services, and Operational Experience," NANOG Track (multiple presentations) at NANOG 67, Chicago, IL, June 2016. Available at https://archive.nanog.org/meetings/nanog67/agenda

[NANOG-list] "Intra-AS messaging for route leak prevention," NANOG Email List - Discussion Thread, June 2016. Available at http://mailman.nanog.org/pipermail/nanog/2016-June/thread.html#86348

[NCCoE-sidr] W. Haag, D. Montgomery, W.C. Barker, A. Tan, "Protecting the Integrity of Internet Routing: Border Gateway Protocol (BGP) Route Origin Validation, Volume B," NIST Special Publication (SP) 1800-14B, August 2018. Available at

https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14b-draft.pdf

[NCSIP] National Cybersecurity Strategy Implementation Plan (2023) (Executive Office of the President). Available at https://www.whitehouse.gov/wp-content/uploads/2023/07/National-Cybersecurity-Strategy-Implementation-Plan-WH.gov_.pdf

[NIST-SP800-189] K. Sriram and D. Montgomery, "Resilient Interdomain Traffic Exchange: BGP Security and DDoS Mitigation," NIST Special Publication, NIST SP 800-189, December 2019. Available at https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-189.pdf

[NIST-CSF] Cybersecurity Framework, National Institute of Standards and Technology. Available at http://www.nist.gov/cyberframework/

[NIST-RIDR] "Robust Inter-Domain Routing," NIST RIDR project. Available at https://www.nist.gov/programs-projects/robust-inter-domain-routing

[NIST-RPKI] "RPKI Deployment Monitor," NIST's online monitor with Global and Regional views. Available at https://rpki-monitor.antd.nist.gov/

[NIST-SRx] NIST BGP Secure Routing Extension (BGP-SRx) Software Suite. Available at https://www.nist.gov/services-resources/software/bgp-secure-routing-extension-bgp-srx-software-suite

[NSA-BGP] "A guide to Border Gateway Protocol (BGP) Best Practices," NSA Technical Report, September 2018. Available at https://apps.nsa.gov/iaarchive/library/reports/a-guide-to-border-gateway-protocol-bgp-best-practices.cfm

[OctoRPKI] OctoRPKI: Cloudflare's RPKI Validator. Available at https://github.com/cloudflare/cfrpki#octorpki

[OECD-330] "Routing security: BGP incidents, mitigation techniques and policy actions," OECD Digital Economy Papers, No. 330, October 2022. Available at https://www.oecd.org/en/publications/routing-security_40be69c8-en.html

[OpenBSD] OpenBSD Project. Available at https://man.openbsd.org/bgpd.conf
https://github.com/cloudflare/cfrpki - octorpki

[Patel] K. Patel, "Cisco's Origin Validation Implementation," NANOG 67, June 2016. Available at https://www.nanog.org/sites/default/files/Patel.pdf

[PEO-13800] U.S. Presidential Executive Order 13800: Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure, May 2017. Available at https://www.whitehouse.gov/the-press-office/2017/05/11/presidential-executive-order-strengthening-cybersecurity-federal

[Phuntsho] T. Phuntsho, "How to install an RPKI validator," RIPE NCC blog. Available at https://labs.ripe.net/Members/tashi_phuntsho_3/how-to-install-an-rpki-validator

[PGW] "Cisco PGW Packet Data Network Gateway", Cisco online. Available: https://www.cisco.com/c/en/us/products/wireless/pgw-packet-data-network-gateway/index.html

[Quilt] "The Quilt security cookbook," published by the Quilt community. Available at https://www.nitrd.gov/nitrdgroups/images/d/db/Quilt-Network-Security-Cookbook-v7.pdf

[RFC2725] C. Villamizar, C. Alaettinoglu, D. Meyer, S. Murphy, "Routing Policy System Security," IETF RFC 2725, December 1999. Available at https://tools.ietf.org/html/rfc2725

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004. Available at https://www.rfc-editor.org/info/rfc3779

[RFC3882] D. Turk, "Configuring BGP to Block Denial-of-Service Attacks," IETF RFC 3882, September 2004. Available at https://tools.ietf.org/rfc/rfc3882.txt

[RFC4012] L. Blunk, J. Damas, F. Parent, and A. Robachevsky, "Routing Policy Specification Language next generation (RPSLng)," IETF RFC 4012, March 2005. Available at https://tools.ietf.org/html/rfc4012

[RFC4036] W. Sawyer, "Management Information Base for Data Over Cable Service Interface Specification (DOCSIS) Cable Modem Termination Systems for Subscriber Management", RFC 4036, DOI 10.17487/RFC4036, April 2005. Available at https://www.rfc-editor.org/info/rfc4036

[RFC4271] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF RFC 4271, January 2006. Available at https://tools.ietf.org/html/rfc4271

[RFC4272] Murphy S. L., "BGP Security Vulnerabilities Analysis", IETF RFC 4272, January 2006. https://doi.org/10.17487/RFC4272

[RFC4364] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006. Available at https://www.rfc-editor.org/info/rfc4364

[RFC5280] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, "Internet X.509 Public Key Infrastructure Certification and Certificate Revocation List (CRL) Profile," IETF RFC 5280, May 2008. Available at http://www.ietf.org/rfc/rfc5280.txt.

[RFC5635] W. Kumari and D. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)", RFC 5635, DOI 10.17487/RFC5635, August 2009. Available at https://tools.ietf.org/html/rfc5635

[RFC5802] V. Gill, J. Heasley, D. Meyer, P. Savola, Ed., C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)," IETF RFC 5082, October 2007. Available at https://tools.ietf.org/html/rfc5082

[RFC6092] J. Woodyatt, "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet

Service," IETF RFC 6092, January 2011. Available at https://tools.ietf.org/html/rfc6092

[RFC6480] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," RFC6480, February 2012. Available at https://tools.ietf.org/html/rfc6480

[RFC6481] G. Huston, R. Loomans, and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, February 2012. Available at https://tools.ietf.org/html/rfc6481

[RFC6483] G. Huston and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs) ", RFC 6483, February 2012. Available at https://tools.ietf.org/html/rfc6483

[RFC6487] G. Huston, G. Michaelson, and R. Loomans, "A Profile for X.509 PKIX Resource Certificates," RFC 6487, February 2012. Available at https://tools.ietf.org/html/rfc6487

[RFC6492] G. Huston, R. Loomans, B. Ellacott, and R. Austein, "A Protocol for Provisioning Resource Certificates," RFC 6492, February 2012. Available at https://tools.ietf.org/html/rfc6492

[RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012. Available: https://www.rfc-editor.org/info/rfc6793 .

[RFC6810] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol," RFC 6810, January 2013. Available at https://tools.ietf.org/html/rfc6810

[RFC6811] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," IETF RFC 6811, January 2013. Available at https://tools.ietf.org/pdf/rfc6811.pdf

[RFC7318] A. Newton and G. Huston, "Policy Qualifiers in Resource Public Key Infrastructure (RPKI) Certificates," RFC 7318, July 2014. Available at https://tools.ietf.org/html/rfc7318

[RFC7353] S. Bellovin, R. Bush, and D. Ward, "Security Requirements for BGP Path Validation," IETF RFC 7353, August 2014. Available at https://tools.ietf.org/html/rfc7353

[RFC7382] S. Kent, D. Kong, and K. Seo, "Template for a Certification Practice Statement (CPS) for the Resource PKI (RPKI)," IETF RFC 7382, April 2015. Available at https://tools.ietf.org/html/rfc7382

[RFC7454] J. Durand, I. Pepelnjak, and G. Doering, "BGP Operations and Security," IETF RFC 7454, February 2015. Available at https://tools.ietf.org/html/rfc7454

[RFC7908] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, June 2016. Available at https://tools.ietf.org/html/rfc7908

[RFC7909] R. Kisteleki and B. Haberman, “Securing Routing Policy Specification Language (RPSL) Objects with Resource Public Key Infrastructure (RPKI) Signatures,” IETF RFC 7909, June 2016. Available at https://tools.ietf.org/html/rfc7909

[RFC7935] G. Huston and G. Michaelson, “The Profile for Algorithms and Key Sizes for Use in the Resource Public Key Infrastructure,” IETF RFC 7935, August 2016. Available at https://tools.ietf.org/html/rfc7935

[RFC7999] T. King, et al., “BLACKHOLE Community,” IETF RFC 7999, October 2016. Available at https://tools.ietf.org/html/rfc7999

[RFC8182] T. Bruijnzeels, O. Muravskiy, B. Webre, and R. Austein, “RPKI Repository Delta Protocol (RRDP),” IETF RFC 8182, July 2017. Available at https://tools.ietf.org/html/rfc8182

[RFC8205] M. Lepinski (Ed.) and K. Sriram (Ed.), “BGPsec Protocol Specification,” IETF RFC 8205, September 2017. Available at https://tools.ietf.org/html/rfc8205

[RFC8206] George, W. and S. Murphy, “BGPsec Considerations for Autonomous System (AS) Migration”, RFC 8206, DOI 10.17487/RFC8206, September 2017, https://www.rfc-editor.org/info/rfc8206

[RFC8207] Bush, R., “BGPsec Operational Considerations”, BCP 211, RFC 8207, DOI 10.17487/RFC8207, September 2017, https://www.rfc-editor.org/info/rfc8207

[RFC8209] Reynolds, M., Turner, S., and S. Kent, “A Profile for BGPsec Router Certificates, Certificate Revocation Lists, and Certification Requests”, RFC 8209, DOI 10.17487/RFC8209, September 2017, https://www.rfc-editor.org/info/rfc8209 [RFC8210] R. Bush and R. Austein, “The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1,” IETF RFC 8210, September 2017. Available at https://tools.ietf.org/html/rfc8210

[RFC8210bis] R. Bush and R. Austein, “The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 2,” IETF Internet Draft. Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-8210bis/

[RFC8212] J. Mauch, J. Snijders, and G. Hankins, “Default External BGP (EBGP) Route Propagation Behavior without Policies”, IETF RFC 8212, DOI 10.17487/RFC8212, July 2017. Available at https://www.rfc-editor.org/info/rfc8212

[RFC8374] K. Sriram (Ed.), “BGPsec Design Choices and Summary of Supporting Discussions,” IETF RFC 8374, April 2018. Available at https://tools.ietf.org/html/rfc8374

[RFC8608] S. Turner and O. Borchert, “BGPsec Algorithms, Key Formats, & Signature Formats,” IETF RFC 8608, September 2017. Available at https://tools.ietf.org/html/rfc8608

[RFC8704] K. Sriram, D. Montgomery, and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding," IETF RFC 8704, Feb. 2020. Available at https://datatracker.ietf.org/doc/html/rfc8704

[RFC8955] C. Loibl, S. Hares, R. Raszuk, D. McPherson, and M. Bacher, "Dissemination of Flow Specification Rules," RFC 8955, December 2020. Available at https://www.rfc-editor.org/info/rfc8955

[RFC8956] C. Loibl, R. Raszuk, and S. Hares, "Dissemination of Flow Specification Rules for IPv6," RFC 8956, December 2020. Available at https://www.rfc-editor.org/info/rfc8956

[RFC9117] J. Uttaro, J. Alcaide, C. Filsfils, D. Smith, and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications, RFC 9117, August 2021. Available at https://www.rfc-editor.org/info/rfc9117

[RFC9234] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages," IETF RFC 9234, May 2022. Available at https://datatracker.ietf.org/doc/rfc9234/

[RFC9319] Y. Gilad, S. Goldberg, K. Sriram, J. Snijders, and B. Maddison, "The Use of maxLength in the Resource Public Key Infrastructure (RPKI)", IETF RFC 9319, October 2022. Available at https://tools.ietf.org/html/rfc9319

[RFC9582] J. Snijders, B. Maddison, M. Lepinski, D. Kong, and S. Kent, "A Profile for Route Origin Authorizations (ROAs)", RFC 9582, May 2024. Available at https://datatracker.ietf.org/doc/rfc9582/

[RIPE1] RIPE NCC Resource Certification: Using the RPKI System. Available at https://www.ripe.net/manage-ips-and-asns/resource-management/certification/using-the-rpki-system

[RIPE2] RIPE NCC RPKI Validator. Available at https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources

[RIPE3] "Router Configuration with JunOS and Cisco IOS," RIPE NCC blog. Available at https://www.ripe.net/manage-ips-and-asns/resource-management/certification/router-configuration

[RIPE-399] P. Smith, R. Evans, and M. Hughes, "RIPE-399 - RIPE Routing Working Group Recommendations on Route Aggregation", December 2006. Available at https://www.ripe.net/publications/docs/ripe-399

[RIPE-532] P. Smith and R. Evans, "RIPE-532 - RIPE Routing Working Group Recommendations on IPv6 Route Aggregation", November 2011. Available at https://www.ripe.net/publications/docs/ripe-532

[RouteLeak1] K. Sriram (Ed.) and A. Azimov (Ed.), "Methods for Detection and Mitigation of BGP Route Leaks", IETF Internet Draft. Available at https://datatracker.ietf.org/doc/draft-ietf-grow-route-leak-detection-mitigation/

[Routinator] Routinator: NLNetLabs' RPKI validator. Available at https://nlnetlabs.nl/projects/rpki/routinator/

[RPKI-software] RPKI Relying Party Software Projects. Available at https://rpki.readthedocs.io/en/latest/ops/tools.html

[Rsync] Wiki page on the Rsync protocol. Available at https://en.wikipedia.org/wiki/Rsync

[Rsync-RPKI] S. Kent and K. Sriram, "RPKI Rsync Download Delay Modeling," Presented at the IETF-86, IETF SIDR WG Meeting, March 2013. Available at https://www.ietf.org/proceedings/86/slides/slides-86-sidr-1.pdf

[RTRlib] "An open-source C implementation of the RPKI/Router Protocol client." Available at https://github.com/rtrlib and http://www.mi.fu-berlin.de/en/inf/groups/ilab/software/index.html

[Ryburn] J. Ryburn, "DDoS Mitigation using BGP Flowspec," NANOG 63, February 2015. Available at https://archive.nanog.org/sites/default/files/tuesday_general_ddos_ryburn_63.16.pdf

[Scudder] J. Scudder, "RPKI on Juniper Routers," NANOG 67, June 2016. Available at https://www.nanog.org/sites/default/files/Scudder.pdf

[SP800-53] Joint Task Force Transformation Initiative. (2013) Security and Privacy Controls for Federal Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-53r4. https://doi.org/10.6028/NIST.27TUSPU27T.800-53r4

[SP800-54] Kuhn DR, Sriram K, Montgomery D (2007) Border Gateway Protocol Security. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-54. https://doi.org/10.6028/NIST.27TUSPU27T.800-54

[SPL-profile] J. Snijders and G. Huston, "A profile for Signed Prefix Lists for Use in the Resource Public Key Infrastructure (RPKI)," IETF Internet Draft. Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-rpki-prefixlist/

[SPL-ROV] K. Sriram, J. Snijders, and D. Montgomery, "Signed Prefix List (SPL) Based Route Origin Verification and Operational Considerations," IETF Internet Draft, Available at https://datatracker.ietf.org/doc/draft-ietf-sidrops-spl-verification/

[Spoofer] CAIDA Spoofer Project: Assessment and reporting on the deployment of source address validation (SAV) best anti-spoofing practices. Available at https://www.caida.org/projects/spoofer/

[Sriram1] K. Sriram, D. Montgomery, and R. Bush, "RIB Size and CPU Workload Estimation for BGPSEC," Presentation at the IETF-91 Joint IDR/SIDR WG Meeting, November 2014. Available at http://www.ietf.org/proceedings/91/slides/slides-91-idr-17.pdf

[Sriram2] V.K. Sriram and D. Montgomery, "Design and analysis of optimization algorithms to minimize cryptographic processing in BGP security protocols," Computer Communications, volume 106, pages 75-85, July 2017. https://doi.org/10.1016/j.comcom.2017.03.007

[SWIP] S. Whipple, “The SWIP Template Tutorial,” ARIN VII, April 2001. Available at https://www.arin.net/vault/participate/meetings/reports/ARIN_VII/PDF/tutorials/swip_arin.pdf

[Symantec] C. Wueest, “Denial-of-service attacks – short but strong: DDoS amplification attacks continue to increase as attackers experiment with new protocols,” Symantec Blog, October 2014. Available at http://www.symantec.com/connect/blogs/denial-service-attacks-short-strong

[TA14-017A] “UDP-Based Amplification Attacks,” US-CERT alert TA14-017A, January 17, 2014. Available at https://www.us-cert.gov/ncas/alerts/TA14-017A

[TA16-288A] “Heightened DDoS Threat Posed by Mirai and Other Botnets,” US-CERT alert TA16-288A, November 30, 2016. Available at https://www.us-cert.gov/ncas/alerts/TA16-288A

[ThousandEyes] ThousandEyes: BGP Route Monitoring. Available at https://www.thousandeyes.com/solutions/bgp-and-route-monitoring

[Toonk-A] Toonk, A., "What caused the Google service interruption", BGPMON Blog, March 2015. Available at http://www.bgpmon.net/what-caused-the-google-service-interruption/

[Toonk-B] Toonk, A., "Massive route leak causes Internet slowdown", BGPMON Blog, June 2015. Available at http://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/

[Verisign1] “Verisign Releases Q4 2016 DDoS Trends Report: 167% Increase in Average Peak Attack from 2015 to 2016,” CircleID blog post, February 2017. Available at http://www.circleid.com/posts/20170214_verisign_releases_q4_2016_ddos_trends_report_167_increase/

[Verisign2] “Distributed Denial of Service Trends Report” by Verisign, Published quarterly. Available at http://www.verisign.com/en_US/security-services/ddos-protection/ddos-report/index.xhtml

[White] R. White, “Rethinking Path Validation,” NANOG 66, February 2016. Available at https://www.nanog.org/sites/default/files/White_Rethinking_Bgp_Path.pdf

[WH-ONCD] "Roadmap to Enhancing Internet Routing Security", The White House Office of the National Cybersecurity Director (ONCD), September 2024. Available at https://www.whitehouse.gov/wp-content/uploads/2024/09/Roadmap-to-Enhancing-Internet-Routing-Security.pdf

[Winward] R. Winward, “Mirai – Inside of an IoT Botnet,” NANOG 69, February 2017. Available at https://www.nanog.org/sites/default/files/1_Winward_Mirai_The_Rise.pdf

[Wishnick] D. Wishnick and C. Yoo, “Overcoming Legal Barriers to RPKI Adoption,” Presented at NANOG 74, October 2018. Available at https://pc.nanog.org/static/published/meetings//NANOG74/daily/day_2.html#talk_1767

[Yoo] C. Yoo and D. Wishnick, “Lowering Legal Barriers to RPKI Adoption,” University of Pennsylvania Law School publication, January 2019. Available at https://scholarship.law.upenn.edu/faculty_scholarship/2035/

[Zmijewski] E. Zmijewski, "Indonesia Hijacks the World", Dyn Research/Renesys Blog, April 2014. Available at http://research.dyn.com/2014/04/indonesia-hijacks-world

Appendix A. Consolidated List of Security Recommendations

Table 32 provides a consolidated list of the security recommendations from various sections throughout the document. If the "Enterprise" column is checked, it means that the security recommendation should be considered for implementation in enterprise and hosted service provider autonomous systems (ASes)—in some cases, action(s) to be performed by the AS operator, and in other cases, feature(s) that should be available in their BGP router(s). A similar statement applies for ISPs when the "ISP" column is checked. When an enterprise outsources services, then the feature/service corresponding to a security recommendation that applies to them would in turn apply to their hosted service provider. An enterprise should always consider (in their service contract) whether their transit ISP meets security recommendations that are checked in the ISP column. There is no column in Table 32 corresponding to an Internet exchange point (IXP), but the security recommendations for ISPs also often apply to IXPs with some variations depending on whether the IXP has transparent or non-transparent Route Server (RS) per specifications in related IETF RFCs (e.g., [ASPA-verif] [RFC8205]).

Table 32. Consolidated list of the security recommendations

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
BGP Origin Validation (IRR, RPKI, ROA, ROV):		
Security Recommendation 1: All Internet number resources (e.g., address blocks and AS numbers) should be covered by an appropriate registration services agreement with an RIR, and all point-of-contact (POC) information should be up to date. The granularity of such registrations should reflect all sub-allocations to entities (e.g., enterprises with provider-based addresses, enterprises within the parent organization, branch offices) that operate their own network services (e.g., Internet access, email, DNS).	X	X
Security Recommendation 2: Route objects corresponding to the BGP routes originating from an AS should be registered and actively maintained in an appropriate RIR's IRR. Enterprises should ensure that appropriate IRR information exists for all IP address space used by them.	X	X
Security Recommendation 3: Internet number resource holders with IPv4/IPv6 prefixes and/or AS numbers (ASNs) should enroll those resources in the RPKI of the appropriate RIR so that RPKI certificates of those resources are issued.	X	X
Security Recommendation 4: Transit providers should provide a service where they facilitate creation, publication, and management of subordinate resource certificates for address space and/or ASNs suballocated to their customers.		X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Note: Currently, RPKI services based on the hosted model and offered by RIRs are common. This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		
Security Recommendation 5: Legacy address space holders without an existing Registration Services Agreement with their RIR should establish an agreement and should enroll their number resources in the RPKI.	X	X
Security Recommendation 6: IP address space holders should register ROA(s) in the global RPKI for all prefixes that are announced or intended to be announced on the public Internet.	X	X
Security Recommendation 7: Each transit provider (ISP) should provide a service where they facilitate creation, publication, and management of ROAs for prefixes suballocated to their customers. Note: This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		X
Security Recommendation 8: If a prefix that is announced (or intended to be announced) is multi-homed and originated from multiple ASes, then one ROA for each originating AS should be registered for the prefix (possibly in combination with other prefixes which are also originated from the same AS).	X	X
Security Recommendation 9: When an ISP or enterprise announces multiple prefixes that include less-specific and more-specific prefixes, they should ensure that the more-specific prefixes have published ROAs before creating ROAs for the subsuming less-specific prefixes.	X	X
Security Recommendation 10: A transit provider (ISP) should ensure that more specific prefixes announced by ASes within its customer cone have ROAs prior to the creation of its own ROAs for subsuming less-specific prefix(es).		X
Security Recommendation 11: An ISP or enterprise should have AS0 ROA coverage for any prefixes that are currently not announced or intended to be announced to the public Internet. However, this should be done cautiously only after ensuring that ROAs exist for more-specific prefixes (if any) that are subsumed by the afore-mentioned prefixes and are announced or intended to be announced.	X	X
Security Recommendation 12: A BGP router should be compliant with [deprecate-as-set] (imminent IETF RFC) which prohibits the use of AS_SET and AS_CONFED_SET in BGP Updates.	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
Security Recommendation 13: ISPs and enterprises that operate BGP routers should also operate one or more RPKI-validating caches that generate validated and distilled RPKI data for use by routers.	X	X
Security Recommendation 14: BGP routers used for inter-domain routing should implement ROA-based Route Origin Validation (ROA-ROV) [RFC6811].	X	X
Security Recommendation 15: In partial/incremental deployment state of the RPKI, the permissible {prefix, origin ASN} pairs for performing BGP origin validation should be generated by taking the union of such data obtained from ROAs, IRR data, and customer contracts.	X	X
Security Recommendation 16: ROA-ROV results should be incorporated into local BGP policy decisions to select best paths. Note: How ROA-ROV results are used in path selection is strictly a local policy decision for each network operator. However, considering a route that is ROA-ROV Invalid to be ineligible for best path selection is recommended.	X	X
Security Recommendation 17: The maxLength in a ROA should not exceed the length of the most specific prefix (subsumed under the prefix in consideration) that is originated or intended to be originated from the AS listed in the ROA.	X	X
Security Recommendation 18: If a prefix and select more-specific prefixes subsumed under it are announced or intended to be announced, then instead of specifying a maxLength, the prefix and the more-specific prefixes should be listed explicitly in the ROA. Note: In general, the use of maxLength should be avoided unless all or nearly all more-specific prefixes up to a maxLength are announced (or intended to be announced) [RFC 9139].	X	X
Security Recommendation 19: If ROA-ROV is deployed in the BGP routers of an entity, they should share that information with their BGP peers. ISPs and large enterprises should publish information about the types of peer interfaces (customers, lateral peers, etc.) on which ROA-ROV is deployed.	X	X
Security Recommendation 20: Resource holders should ensure all their resource certificates, ROAs, and other RPKI signed objects are up to date. Any such objects with an impending expiration date should be refreshed well ahead of their expiry. Note: At ARIN, RPKI resource certs are set with a two-year lifespan, and they auto-renew after one year, resetting the two-year lifespan [ARIN2].	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
Security Recommendation 21: Internet number resource holders should employ BGP/RPKI monitoring tools/services to remain informed about changes in the RPKI system that may affect their BGP route originations.	X	X
Prefix (Route) Filtering:		
Security Recommendation 22: IPv6 routes should be filtered to permit only allocated IPv6 prefixes. Network operators should update IPv6 prefix filters regularly to include any newly allocated prefixes [Cymru-bogon]. Note: If prefix resource owners regularly register AS0 ROAs (see Section 4.3) for allocated (but possibly currently unused) prefixes, then those ROAs could be a complementary source for the update of prefix filters.	X	X
Security Recommendation 23: Prefixes that are marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] are forbidden from routing in the global Internet and should be rejected if received from an external BGP (eBGP) peer.	X	X
Security Recommendation 24: For single-homed prefixes (subnets) that are originated by an AS, any routes for those prefixes received at that AS from eBGP peers should be rejected.	X	X
Security Recommendation 25: It is recommended that an eBGP router should set a route specificity limit for each eBGP peer and reject prefixes that exceed the specificity limit on a per-peer basis. Note: The specificity limit may be the same for all peers (e.g., /24 for IPv4 and /48 for IPv6).	X	X
Security Recommendation 26: The default route (0.0.0.0/0 in IPv4 and ::/0 in IPv6) should be rejected unless there is an explicit peering agreement that permits accepting it.	X	X
Security Recommendation 27: An Internet exchange point (IXP) should announce—from its route server to all its member ASes—its LAN prefix or its entire prefix, which would be the same as or less specific than its LAN prefix. Each IXP member AS should, in turn, accept this prefix from the IXP and reject any more-specific prefixes (of the IXP announced prefix) from any of its eBGP peers.	X	X
Security Recommendation 28: Inbound prefix filtering facing lateral peer – The following prefix filters (disallowed prefixes) should be applied in the inbound direction: • Unallocated prefixes	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
• Special-purpose prefixes • Prefixes that the AS originates • Prefixes that exceed a specificity limit • Default route • IXP LAN Prefixes		
Security Recommendation 29: Outbound prefix filtering facing lateral peer – The allowed outbound prefixes are those that are originated by the AS in question and those originated by its downstream ASes (i.e., the ASes in its customer cone). The following prefix filters should be applied in the outbound direction: • Unallocated prefixes • Special-purpose prefixes • Prefixes that exceed a specificity limit • Default route • IXP LAN prefixes • Prefixes learned from AS's lateral peers • Prefixes learned from AS's transit providers	X	X
Security Recommendation 30: Inbound prefix filtering facing transit provider – Case 1 (full routing table): In general, when the full routing table is required from the transit provider, the following prefix filters should be applied in the inbound direction: • Unallocated prefixes • Special-purpose prefixes • Prefixes that the AS originates • Prefixes that exceed a specificity limit • IXP LAN prefixes	X	X
Security Recommendation 31: Inbound prefix filtering facing transit provider – Case 2 (default route): If the border router is configured for only the default route, then only the default route should be accepted from the transit provider and nothing else.	X	X
Security Recommendation 32: Outbound prefix filtering facing transit provider: The same outbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1). Note: In conjunction with the outbound prefix filtering security recommendation, some policy rules may	X	X

Security Recommendation	Applicable to: Enter-prise	Applicable to: ISP
also be applied if a transit provider is not contracted (or chosen) to provide transit for some subset of outbound prefixes.		
Security Recommendation 33: Inbound prefix filtering facing customer in Scenario 1 (see Section 4.5.3) – Only the prefixes that are known to be originated from the customer and its customer cone should be accepted, and all other route announcements should be rejected.		X
Security Recommendation 34: Inbound prefix filtering facing customer in Scenario 2 (see Section 4.5.3) – The same set of inbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1).		X
Security Recommendation 35: Outbound prefix filtering facing customer: The filters applied in this case would vary depending on whether the customer wants to receive only the default route or the full routing table. If it is the former, then only the default route should be announced and nothing else. In the latter case, the following outbound prefix filters should be applied: • Special-purpose prefixes • Prefixes that exceed a specificity limit Note: The default route may be added to the above filter list if the customer requires the full routing table but not the default route.		X
Security Recommendation 36: Inbound prefix filtering for leaf customer facing transit provider – A leaf customer may request only the default route from its transit provider. In this case, only the default route should be accepted and nothing else. If the leaf customer requires the full routing table from the transit provider, then it should apply the following inbound prefix filters: • Unallocated prefixes • Special-purpose prefixes • Prefixes that the AS (i.e., leaf customer) originates • Prefixes that exceed a specificity limit • Default route	X	
Security Recommendation 37: Outbound prefix filtering for leaf customer facing transit provider – A leaf customer network should apply a very simple outbound policy of announcing only the prefixes it originates. However, it may additionally apply the same outbound prefix filters as those for a lateral peer (see Section 4.5.1) for extra caution.	X	

Security Recommendation	Applicable to	
	Enter-prise	ISP
Security Recommendation 38: The ROA data (available from RPKI registries) should be used to construct and/or augment prefix filter lists for customer interfaces. Note: This Security Recommendation is possibly more applicable to smaller ISPs that have accurate visibility of their customer cone. Larger ISPs tend not to have such visibility.		X
Route Leak Mitigation:		
Security Recommendation 39: An AS operator should have an ingress policy to tag routes internally (locally within the AS) to communicate from ingress to egress regarding the type of peer (customer, lateral peer, or transit provider) from which the route was received.	X	X
Security Recommendation 40: An AS operator should have an egress policy to utilize the tagged information (in Security Recommendation 37) to prevent route leaks when routes are forwarded on the egress. The AS should not forward routes received from a transit provider to another transit provider or a lateral peer. Also, the AS should not forward routes received from a lateral peer to another lateral peer or a transit provider.	X	X
Checking AS Path for Disallowed AS Numbers		
Security Recommendation 41: The AS path in an update received in eBGP should be checked to ensure that the local AS number is not present. The AS path should also be checked to ensure that AS numbers meant for special purposes [IANA-ASN-sp] are not present. In case of a violation, the update should be rejected. Note: The special purpose ASN 23456 is allocated for AS_TRANS [RFC6793] and is allowed to be present in an AS_PATH in conjunction with an AS4_PATH [RFC6793] in the update.	X	X
GTSM		
Security Recommendation 42: The Generalized TTL Security Mechanism (GTSM) [RFC5082] should be applied on a per-peer basis to provide protection against spoofed BGP messages.	X	X
Source Address Validation (Anti-spoofing):		
Security Recommendation 43: BGP routers that have single-homed directly connected customers, CMTS (or equivalent) in broadband access networks, and PGW (or equivalent) in mobile networks should implement		X

Security Recommendation	Applicable to	
	Enter-prise	ISP
SAV using ACLs (Section 5.1.1). The BGP routers in this context may alternatively use the strict uRPF method (Section 5.1.2).		
Security Recommendation 44: An enterprise border router that is multi-homed should always announce all its address space to each of its upstream transit providers to enable more effective SAV. This can be done in one of two ways: 1) announce an aggregate less-specific prefix to all transit providers and more-specific prefixes (covered by the less-specific prefix) to different transit providers as needed for traffic engineering, or 2) announce the same prefixes to each transit provider (albeit with suitable prepending for traffic engineering).	X	
Security Recommendation 45: This is the exception case when the enterprise border router does not adhere to Security Recommendation 41 and instead selectively announces some prefixes to one upstream transit ISP and other prefixes to another upstream transit ISP. In this case, the enterprise should route data (by appropriate internal routing) such that the source addresses in the data packets towards each upstream transit ISP belong in the prefix or prefixes announced to that ISP.	X	
Security Recommendation 46: On the ingress side (i.e., for data packets received from the transit ISP), enterprise border routers should deploy loose uRPF (Section 5.1.4) and/or ACLs (Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the enterprise's own prefixes).	X	
Security Recommendation 47: An enterprise should allow on the egress side (i.e., for data packets sent to the transit ISP) only those packets with source addresses that belong in their own prefixes.	X	
Security Recommendation 48: On customer-facing interfaces, smaller ISPs should perform SAV on ingress packets by deploying the feasible-path uRPF (see Section 5.1.3). They should avoid using strict or loose uRPF as they are not effective, especially in the case of multi-homed customers. It is recognized that larger ISPs may use loose uRPF on customer interfaces.		X
Security Recommendation 49: For feasible-path uRPF to work appropriately, a smaller ISP (especially one that is near the Internet edge) should propagate all its announced address space to each of its upstream transit providers. This can be done in one of two ways: 1) announce an aggregate less-specific prefix to all transit providers and announce more-specific prefixes (covered by the less-specific prefix) to different transit		X

Security Recommendation	Applicable to	
	Enter-prise	ISP
providers as needed for traffic engineering, or 2) announce the same prefixes to each transit provider (albeit with suitable prepending for traffic engineering).		
Security Recommendation 50: ISPs should prefer customer routes over other (i.e., transit provider or lateral peer) routes. (This is also normal ISP policy in most cases.) Note: Following this recommendation facilitates a basis for adhering to Security Recommendation 48. It is also one of the stability conditions on BGP policy for ensuring stable convergence of routing information [Gao-Rexford].		X
Security Recommendation 51: On interfaces with lateral (i.e., non-transit) peers, smaller ISPs (near the edge of the Internet) should perform SAV on ingress packets by deploying the feasible-path uRPF (see Section 5.1.3). They should avoid using strict or loose uRPF as they are not very effective for SAV on the lateral peer interfaces. It is recognized that larger ISPs may use loose uRPF on the interfaces with lateral peers.		X
Security Recommendation 52: On interfaces with transit providers, ISPs should perform SAV on ingress packets by deploying loose uRPF (see Section 5.1.4) and/or ACLs (see Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the ISP's internal-use only prefixes).		X
Security Recommendation 53: On the egress side towards customers, lateral (i.e., non-transit) peers, and transit providers, the ISP's border routers should deploy ACLs (see Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks—prefixes marked "False" in column "Global" [IANA-v4-sp] [IANA-v6-sp] and the ISP's internal-use only prefixes).		X
DDoS Mitigation (Remote Triggered Black Hole filtering, Flow specification):		
Security Recommendation 54: Edge routers should be equipped to perform destination-based remotely triggered black hole (D/RTBH) filtering and source-based remotely triggered black hole (S/RTBH) filtering.	X	X
Security Recommendation 55: Edge routers should be equipped to make use of BGP flow specification (Flowspec) to facilitate DDoS mitigation (in coordination between upstream and downstream autonomous systems).	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
Security Recommendation 56: Edge routers in an AS providing RTBH filtering should have an ingress policy towards RTBH customers to accept routes more specific than /24 in IPv4 and /48 in IPv6. Additionally, the edge routers should accept a more specific route (in case of D/RTBH) only if it is subsumed by a less-specific route that the customer is authorized to announce as standard policy (i.e., the less-specific route has a registered IRR entry and/or a ROA). Further, the edge routers should not drop RTBH-related more-specific route advertisements from customers even though BGP origin validation may mark them as "Invalid".		X
Security Recommendation 57: A customer AS should make sure that the routes announced for RTBH filtering have NO_EXPORT, NO_ADVERTISE, or similar communities.	X	X
Security Recommendation 58: An ISP providing an RTBH filtering service to customers must have an egress policy that denies routes that have community tagging meant for triggering RTBH filtering at the local AS. This is an additional safeguard in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails.		X
Security Recommendation 59: An ISP providing an RTBH filtering service to customers must have an egress policy that denies prefixes that are longer than expected. This provides added safety in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails.		X
General: Outsourced Services, Supporting Standards, Open Source, and Measurements		
Security Recommendation 60: Enterprises should require their Internet transit providers to adhere to the relevant security recommendations (from this document) by including them in service contracts.	X	
Security Recommendation 61: Enterprises that outsource applications/services (e.g., Email, DNS, cloud hosted systems, etc.) should require their outsource service providers to adhere to the relevant security recommendations (from this document) by including them in service contracts.	X	
Security Recommendation 62: Government agencies, ISPs, and enterprises should support standards development and open-source implementation efforts related to standards-based routing security technologies.	X	X
Security Recommendation 63: To the extent possible, ISPs and enterprises should facilitate collection of routing data by trusted	X	X

Security Recommendation	Applicable to	
	Enter-prise	ISP
organizations engaged in or supporting R&D efforts related to routing robustness and security monitoring.		
Emerging Technologies – Security Recommendations for Future Planning (FP) (Awaiting implementation in routers by commercial vendors)		
Security Recommendation FP1: ASes should implement in their border routers the BGPsec-based AS path signing and verification procedures to protect AS paths in BGP Updates from path manipulations [RFC8205].	X	X
Security Recommendation FP2: An AS owner should register its Autonomous System Provider Authorization (ASPA) object(s) per specification in [ASPA-prefix].	X	X
Security Recommendation FP3: Transit providers should provide a service where they facilitate creation, publication, and management of ASPAs for their customer ASes. Note: This security recommendation can be implemented in the hosted or delegated model based on service agreements with customers.		X
Security Recommendation FP4: ASes should deploy ASPA-based AS path verification and route leak mitigation procedures in their border routers per specification in [ASPA-verif].	X	X
Security Recommendation FP5: An AS owner doing ASPA should periodically check their own ASPA object(s) for correctness and completeness. They should also ensure that the same are renewed well before their expiry dates.	X	X
Security Recommendation FP6: An AS owner doing ASPA should periodically monitor all the ASPAs in the RPKI repositories to check if their AS number is incorrectly included as a provider in an ASPA (cryptographically valid), and if so, they should report it to the responsible party (or parties) so that the ASPA can be rectified.	X	X
Security Recommendation FP7: An AS owner doing ASPA should periodically monitor the ASPAs in the RPKI repositories to check if their AS number is incorrectly not included as a provider in the ASPA (cryptographically valid) of a customer AS, and if so, they should report it to the customer AS owner so that the ASPA can be rectified.	X	X
Security Recommendation FP8: ASes should implement in their border routers the procedures with BGP Roles as specified in [RFC9234].	X	X

	Applicable to	
Security Recommendation	**Enter-prise**	**ISP**
Security Recommendation FP9: ASes should implement in their border routers the procedures with the OTC Attribute for route leak detection and mitigation as specified in [RFC9234].	X	X

1786

Appendix B. List of Symbols, Abbreviations, and Acronyms

ACL
Access Control List

AfriNIC
African Network Information Center

APNIC
Asia-Pacific Network Information Centre

ARIN
American Registry for Internet Numbers

AS
Autonomous System

BGP
Broder Gateway Protocol

BGPsec
Broder Gateway Protocol with Security Extensions

DA
Destination Address

DDoS
Distributed Denial of Service

DHS
Department of Homeland Security

DNS
Domain Name System

DNSSEC
Domain Name System Security Extensions

DoS
Denial of Service

D/RTBH
Destination-Based Remotely Triggered Black-Holing

DSCP
Differentiated Services Code Point

eBGP
External BGP

EFP-uRPF
Enhanced Feasible Path Unicast Reverse Path Forwarding

FIB
Forwarding Information Base

FISMA
Federal Information Security Modernization Act

Flowspec
Flow Specification

FP-uRPF
Feasible Path Unicast Reverse Path Forwarding

GTSM
Generalized TTL Security Mechanism

IANA
Internet Assigned Numbers Authority

iBGP
Internal BGP

ICMP
Internet Control Message Protocol

IETF
Internet Engineering Task Force

IGP
Internal Gateway Protocol

IRR
Internet Routing Registry

ISP
Internet Service Provider

IXP
Internet Exchange Point

LACNIC
Latin America and Caribbean Network Information Centre

maxLength
Maximum allowed length of a prefix specified in RAO

NCCoE
National Cybersecurity Center of Excellence

NIST SP
NIST Special Publication

NLRI
Network Layer Routing Information (synonymous with prefix)

NTP
Network Time Protocol

RFC
Request for Comments (IETF standards document)

RFD
Route Flap Damping

RIB
Routing Information Base

RIPE
Réseaux IP Européens

RIR
Regional Internet Registry

RITE
Resilient Interdomain Traffic Exchange

RLP
Route Leak Protection

ROA
Route Origin Authorization

ROA-ROV
ROA-Based Route Origin Validation

RPKI
Resource Public Key Infrastructure

RPKI-to-router protocol
RPKI Cache to Router Protocol

RRDP
RPKI Repository Delta Protocol

RTBH
Remotely Triggered Black-Holing

SA
Source Address

SAV
Source Address Validation

SIDR
Secure Inter-Domain Routing

SIDR WG
Secure Inter-Domain Routing Working Group (in the IETF)

S/RTBH
Source-Based Remotely Triggered Black-Holing

SSDP
Simple Service Discovery Protocol

TCP
Transmission Control Protocol

TLS
Transport Layer Security

UDP
User Datagram Protocol

UPnP
Universal Plug and Play

uRPF
Unicast Reverse Path Forwarding

Appendix C. Change Log

In January 2025, the following changes were made to the document:

This document (NIST 800-189r1 ipd) contains changes that reflect (1) advances made in the IETF with standards (e.g., work that progressed from draft to RFC status and updates to existing RFCs), and (2) evolution of promising new technologies in the IETF that offer complementary and/or more effective solutions (e.g., ASPA, OTC, BAR-SAV). The latter are described (new Sections 4.7.2, 4.7.3, 5.1.7) but the security recommendations based on them are labeled FP (Future Planning) pending publication of the solutions as RFCs and availability of implementations.

Section 6 titled “General: Outsourced Services, Supporting Standards, Open Source, and Measurements” and the security recommendations included there are new.

A section titled “Monitoring UDP/TCP Ports with Vulnerable Applications and Employing Traffic Filtering” (Section 5.4 in the original publication [NIST-SP800-189]) has been removed because the techniques discussed in it were not related to BGP. This section can still be accessed in the original publication [NIST-SP800-189].