Draft NIST Special Publication 800-189

Secure Interdomain Traffic Exchange

BGP Robustness and DDoS Mitigation

Kotikalapudi Sriram
Doug Montgomery

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

C O M P U T E R S E C U R I T Y

National Institute of
Standards and Technology
U.S. Department of Commerce

Draft NIST Special Publication 800-189

Secure Interdomain Traffic Exchange

BGP Robustness and DDoS Mitigation

Kotikalapudi Sriram
Doug Montgomery
Advanced Network Technology Division
Information Technology Laboratory

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

December 2018

U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Institute of Standards and Technology
Walter Copan, NIST Director and Under Secretary of Commerce for Standards and Technology

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.

National Institute of Standards and Technology Special Publication 800-189
Natl. Inst. Stand. Technol. Spec. Publ. 800-189, 70 pages (December 2018)
CODEN: NSPUE2

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

Public comment period: *December 17, 2018* through *February 15, 2019*

National Institute of Standards and Technology
Attn: Advanced Network Technologies Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8920) Gaithersburg, MD 20899-8920
Email (for submission of reviewers' comments): sp800-189@nist.gov

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Abstract

In recent years, numerous routing control plane anomalies such as Border Gateway Protocol (BGP) prefix hijacking and route leaks have resulted in Denial of Service (DoS), unwanted data traffic detours, and performance degradation. Large-scale Distributed Denial of Service (DDoS) attacks on servers using spoofed Internet Protocol (IP) addresses and reflection-amplification in the data plane have also been frequent, resulting in significant disruption of services and damages. This special publication on Secure Interdomain Traffic Exchange (SITE) includes initial guidance on securing the interdomain routing control traffic, preventing IP address spoofing, and certain aspects of DoS/DDoS detection and mitigation.

Many of the recommendations in this publication focus on the Border Gateway Protocol (BGP). BGP is the control 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 the interdomain routing control traffic include Resource Public Key Infrastructure (RPKI), BGP origin validation (BGP-OV), and prefix filtering. Additionally, technologies recommended for mitigating DoS/DDoS attacks focus on prevention of IP address spoofing using Source Address Validation (SAV) with Access Control Lists (ACLs) and unicast Reverse Path Forwarding (uRPF). Other technologies (including some application plane methods) such as Remotely Triggered Black Hole (RTBH) filtering, Flow Specification (Flowspec), and Response Rate Limiting (RRL) are also recommended as part of the overall security mechanisms.

Keywords

Routing security and robustness; Internet infrastructure security; Border Gateway Protocol (BGP) security; prefix hijacks; IP address spoofing; Distributed Denial of Service (DDoS); Resource Public Key Infrastructure (RPKI); BGP origin validation (BGP-OV); prefix filtering; BGP path validation (BGP-PV); BGPsec; route leaks; Source Address Validation (SAV); unicast Reverse Path Forwarding (uRPF); Remotely Triggered Black Hole (RTBH) filtering; Flow Specification (Flowspec).

Acknowledgements

The authors are grateful to William T. Polk, Scott Rose, Okhee Kim, Oliver Borchert, Susan Symington, William C. Barker, William Haag, Allen Tan, and Jim Foti for their review and comments.

Audience

This document gives technical guidance and recommendations for secure interdomain traffic exchange. The primary audience include 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 will also be useful for enterprise and transit network operators and equipment vendors in general.

It is expected that the guidance and applicable recommendations from this publication will be incorporated in the security plans and operational processes of federal enterprise networks. Likewise, it is expected that applicable recommendations will be incorporated into the service agreements for federal contracts for hosted application services and Internet transit services.

Trademark Information

Executive Summary

There have been numerous incidents in recent years involving routing control plane anomalies such as Border Gateway Protocol (BGP) prefix hijacking, route leaks, and other forms of misrouting resulting in Denial of Service (DoS), unwanted data traffic detours and performance degradation. Large scale Distributed DoS (DDoS) attacks on servers using spoofed Internet Protocol (IP) addresses and reflection-amplification in the data plane have also been frequent, resulting in significant disruption of services and damages.

This document provides technical guidance and recommendations for technologies that improve the security and robustness of interdomain traffic exchange. The primary focus of these recommendations are the points of interconnection between enterprise networks, or hosted-service providers, and the public Internet. In other words, between what 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). These points of interconnection between stub and transit networks are often referred to as the Internet's edge. There is usually a contractual relationship between the transit networks and the stub networks that they service, and the technical procedures and policies defined in that relationship is commonly called its "peering policy".

Many of the recommendations in this document also apply to the points of interconnection between two transit networks. There are instances in which the recommendations for interdomain traffic exchange between transit networks will vary from those for exchanges between stub and transit networks.

The provided recommendations reduce the risk of accidental attacks (caused by misconfiguration) and malicious attacks in the routing control plane, and they help detect and prevent IP address spoofing and resulting DoS/DDoS attacks. These recommendations primarily cover technologies (for security and robustness) to be used in border routers that operate the Border Gateway Protocol (commonly called BGP routers). However, they also extend to other systems that support reachability in the Internet, e.g., Domain Name Servers (DNS) and other open Internet services, and Resource Public Key Infrastructure (RPKI) repositories.

It is expected that the guidance and applicable recommendations from this publication will be incorporated in the security plans and operational processes of federal enterprise networks. Likewise, it is expected that applicable recommendations will be incorporated into the service agreements for federal contracts for hosted application services and Internet transit services. This document may also be helpful in the ongoing efforts by NIST and NTIA [NIST2018] [Botnet-Roadmap] in response to the Presidential Executive Order 13800 [PEO-13800].

Technologies recommended in this document for securing the interdomain routing control traffic include Resource Public Key Infrastructure (RPKI), BGP origin validation (BGP-OV), and prefix filtering. Additionally, technologies recommended for mitigating DoS/DDoS attacks include prevention of IP address spoofing using Source Address Validation (SAV) with Access Control Lists (ACLs) and unicast Reverse Path Forwarding (uRPF). Other technologies (including some application plane methods) such as Remotely Triggered Black Hole (RTBH) filtering, Flow Specification (Flowspec), and Response Rate Limiting (RRL) are also

186 recommended as part of the overall security mechanisms.

187 **Table of Contents**

List of Appendices

List of Figures

List of Tables

1 Introduction

1.1 What This Guide Covers

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

Technologies recommended in this document for securing the interdomain routing control traffic include RPKI, BGP origin validation (BGP-OV), and prefix filtering. Additionally, technologies recommended for mitigating DoS/DDoS attacks include prevention of IP address spoofing using Source Address Validation (SAV) with Access Control Lists (ACLs) and unicast Reverse Path Forwarding (uRPF). Other technologies (including some application plane methods) such as Remotely Triggered Black Hole (RTBH) filtering, Flow Specification (Flowspec), and Response Rate Limiting (RRL) are also recommended as part of the overall security mechanisms.

1.2 What This Guide Does Not Cover

BGP origin validation relies on a global RPKI system (e.g., certificate authorities, publication repositories, etc.) as the source of trusted information about Internet address holders and their route origin authorization statements. Each RIR operates trusted root CA in the RPKI system and publishes a Certificate Practice Statement [RFC7382] describing the security and robustness properties of each implementation. 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 the scope of this document.

Transport layer security is key to integrity of messages communicated in BGP sessions. Making security recommendations for the underlying transport layer is also outside the scope of this document.

DDoS attacks using spoofed IP addresses make use of connectionless query-response services, e.g., DNS, NTP (Network Time Protocol), SSDP (Simple Service Discovery Protocol) servers, to "reflect" and amplify the impact of the attacks on the intended targets. This document addresses some but not all aspects of security hardening of the servers that are exploited for reflection and amplification. Security measures such as limiting packet rate of outlier source addresses, limiting IP connections, syn proxy, etc. may be effectively employed at servers that are used for reflection and amplification of DoS/DDoS attacks, but this document does not cover them.

1.3 Document Structure

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

- **Section 2:** Routing control plane attacks such as BGP prefix hijacking, AS path modification, and route-leaks are described.
- **Section 3:** Data plane attacks involving source IP address spoofing and reflection-amplification are described.
- **Section 4:** Solutions are described, and security recommendations are made for routing control plane/BGP security. The solution technologies that are discussed include RPKI, BGP origin validation (BGP-OV), prefix filtering, BGP path validation (BGP-PV), and route-leak detection and mitigation.
- **Section 5:** Solutions are described, and security recommendations are made for detection and mitigation of source IP address spoofing and reflection-amplification attacks. The solution technologies that are discussed include ACLs, various uRPF methods, response rate limiting (RRL), RTBH, and Flowspec.

1.4 Conventions Used in this Guide

Throughout this guide, the following format conventions are used to denote special use text:

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

URLs are included in the text and references to guide readers to a given website or online tool designed to aid administrators. This is not meant to be an endorsement of the website or any product/service offered by the website publisher. All URLs were considered valid at the time of writing.

2 Control Plane / BGP Vulnerabilities

2.1 Prefix Hijacking and Announcement of Unallocated Address Space

A BGP prefix hijack occurs when an Autonomous System (AS) accidentally or maliciously originates a prefix that it is not authorized (by the prefix owner) to originate. This is also known as false origination (or announcement). In contrast, if an AS is authorized to originate/announce a prefix by the prefix owner, then such a route origination/announcement is called legitimate. In the example illustrated in Figure 1, 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 ASes in the Internet, and this subset of ASes will install the false route in their routing table or Forwarding Information Base (FIB). That is, ASes for which AS64510 is closer (i.e., shorter AS path length) would choose the false announcement and thus data traffic from clients in those ASes destined for the network 192.0.2/24 will be misrouted to AS64510.

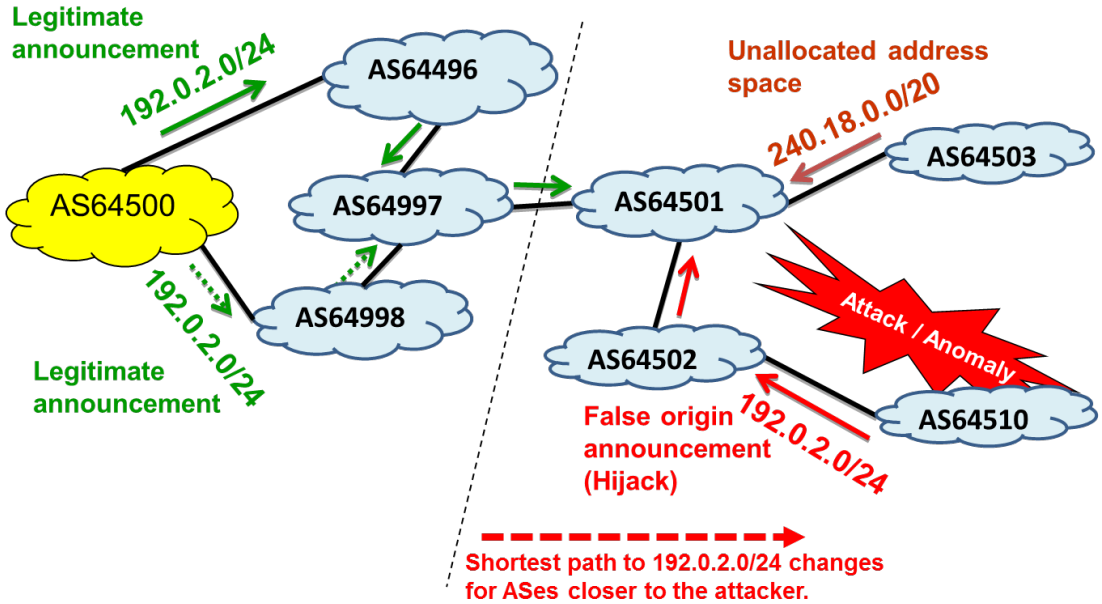

Figure 1: Illustration of Prefix Hijacking and Announcement of Unallocated Address Space.

The rules for IP route selection in 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 a prefix 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 (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 also falsely originate allocated but currently unused address space. This is referred to as *prefix squatting*, where someone else's unused prefix is temporarily announced and used for sending spam or other malicious purpose.

The various types of unauthorized prefix originations described above are called *prefix hijacks* or *false-origin announcements*. The unauthorized announcement of a prefix longer than the

legitimate announcement is called a "sub-prefix hijack". The consequences of such adverse actions can be serious, resulting in denial of service, eavesdropping, misdirection to imposter servers (e.g., to steal login credentials or inject malware), defeat of IP reputation systems to launch spam email, etc. There have been numerous incidents involving prefix hijacks in recent years. There are several commercial services and research projects that track and log anomalies in the global BGP routing system [BGPmon] [ThousandEyes] [BGPStream] [ARTEMIS]. Many of these sites provide detailed forensic analysis of observed attack scenarios.

2.2 AS Path Modification

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 that have been negotiated between networks. BGP is also vulnerable to modification of the AS_PATH information that it conveys. As an example, a malicious AS which receives a BGP update may illegitimately remove some of the preceding ASes in the AS_PATH attribute of the update to make the path length seem shorter. When the update modified in this manner is propagated, the ASes upstream can be deceived to believe that the path to the advertised prefix via the adversary AS is shorter. By doing this, the adversary AS may seek to increase (illegitimately) 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 that an adversary AS replaces a prefix in a received update by a more specific prefix (subsumed by the prefix), and then forwards the update to neighbors. This attack is known as 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 wins over a covering less specific prefix advertisement. This means that ASes in the Internet would widely accept and use the adversary AS's advertisement for the more specific prefix. The exceptions are the ASes that are in the AS path from the adversary to the prefix. These exception ASes 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 very serious. The adversary would be able to force almost all traffic for the more specific prefix to be routed via their AS. Thus, they can eavesdrop on the data (destined for the more specific prefix) while channeling it back to the legitimate destination to avoid detection.

2.3 Route Leaks

Previously we noted that the interconnections of networks in the Internet are dictated by contracted business relationships that express the policies and procedures for the exchange of control 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 a "customer", "transit provider", and/or "lateral peer" business relationship between networks.

Definitions of Peering Relations: 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 multihomed to more than one transit providers. A "stub customer" AS has no customer ASes or lateral peer ASes of its own. A "leaf customer" is a stub customer that is single-homed to one transit provider and not connected to any other AS. The term "customer cone prefixes" of an AS refers to the union of the prefixes received from all directly connected customers and the prefixes originated by the AS itself. Naturally, this set recursively includes customers' customers' prefix advertisements (down the hierarchy). "Lateral peer" ASes typically announce their customer-cone prefixes to each other, and subsequently they announce the lateral-peer's customer-cone prefixes to their respective customers but not to other lateral peers or transit providers.

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 is implemented by insuring that stub 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, and further 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) MainOne (a Nigerian ISP) leak of Google prefixes and outage caused for Google services for over an hour in November 2018 [Naik], (2) the Dodo-Telstra incident in March 2012 that caused outage of Internet services nationwide in Australia [Huston2012], (3) the massive Telekom Malaysia route leaks, which in turn Level3 accepted and propagated [Toonk-B], etc..

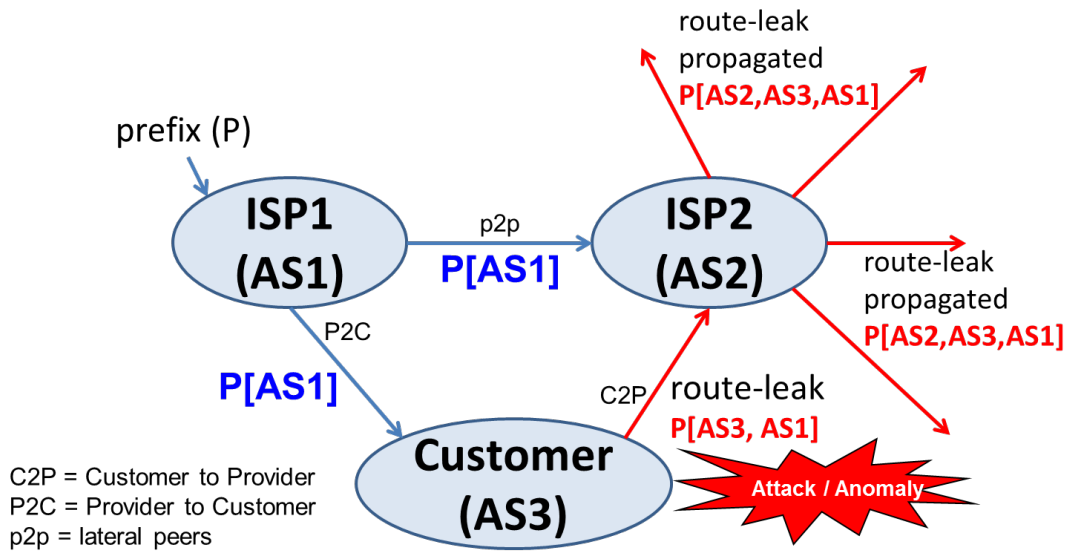

Figure 2: Illustration of the basic notion of a route leak.

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

3 IP Address Spoofing & Reflection-Amplification Attacks

3.1 Spoofed Source Addresses

Distributed Denial of Service (DoS) is a form attack where the attack traffic is generated from many distributed sources (to achieve a high-volume attack) and directed towards an intended victim (system/server) [ISOC] [Huston2016] [Mirai1] [Kaeo]. To conduct a direct DDoS attack, the attacker normally makes use of a few powerful computers or alternately a vast number of unsuspecting compromised third-party computers/devices (laptops, tablets, cell phones, Internet of Things (IoT) devices, etc.). The latter scenario is usually implemented through botnets [Arbor] [Huston2016] [NIST2018]. 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 address. For example, in the Mirai attacks [Mirai1] [Mirai2] [Winward] [TA16-288A], a very large number of compromised bots (IoT devices) that sent 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] (also see Section 5.1.6). If a hijacked prefix is being used, then the source addresses appearing in the DDoS attack packets is 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 multiply the attack traffic volume by a factor of 50 or more [ISOC]. The way this works can be explained with help of the illustration shown in Figure 3. The attacker normally makes use of a botnet consisting of many compromised devices to send query requests to high-performance Internet servers. The attacking systems insert the IP address of the target (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 contained in a single packet, and the exchange does not require the establishment of a connection (unlike Transmission Control Protocol (TCP)) between the source and the server. 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, amplifying the effect of the DoS attack (see Table 1 in Section 5.4). Such reflection and amplification attacks can result in massive DDoS with attack volumes in the range of hundreds of Gbps [Symantec] [ISTR-2015] [ISTR-2016] [ISTR-2017] [ISOC] [Verisign1] [Verisign2] [Bjarnason]. In Q1 2018, there was

an increase of 100% quarter-over-quarter and 700% year-over-year in DNS amplification attacks [HelpNet]. The attack volumes may still rise significantly if the Mirai-scale attacks are combined with reflection-amplification attacks.

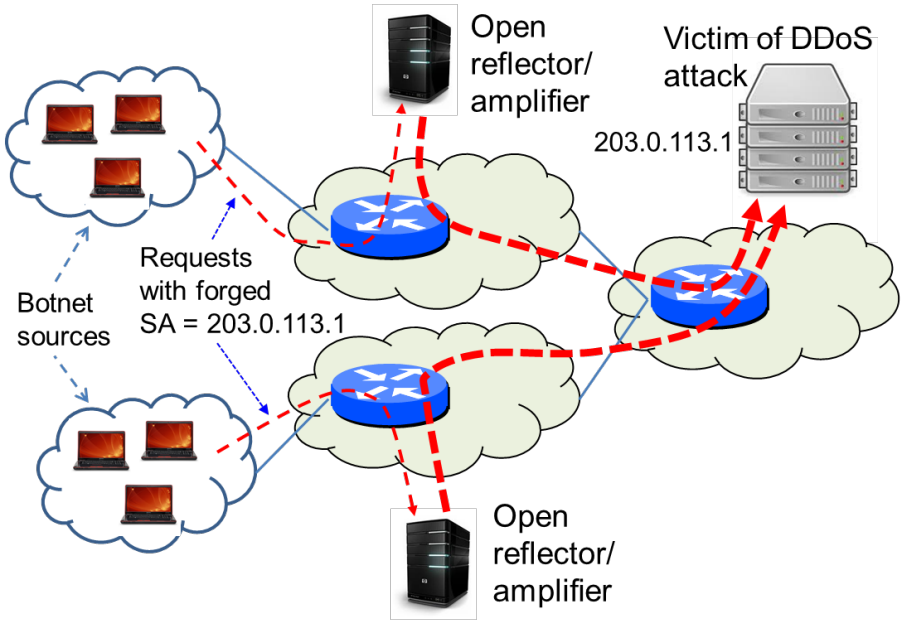

Figure 3: DDoS by IP source address spoofing, and reflection and amplification.

4 Control Plane / BGP Security – Solutions and Recommendations

BGP security vulnerabilities and mitigation techniques have been of interest for several years within the networking community (e.g., [IETF-SIDR] [RFC7454] [NIST800-54] [NANOG] [Murphy] [MANRS] [Quilt] [Levy] [CSRIC-WG6] [RFC6811] [RFC8205] [NSA-BGP]). This section highlights key BGP security technologies that have emerged from such efforts and makes related security recommendations. Many of the solution technologies discussed here have been developed and standardized in the IETF [IETF-SIDR] [IETF-SIDROPS] [IETF-IDR] [IETF-OPSEC] [IETF-GROW]. It is worth mentioning here that the [MANRS] document can be thought as complementary to this document since it provides implementation guidance for some of the solution technologies described in this section as well as Section 5.

4.1 Registration of Route Objects in Internet Routing Registries

Declarative data about Internet resource allocations and routing policies has 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 (ARIN, RIPE, etc.) 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 which may originate them. Routing Policy Specification Language (RPSL) [RFC4012] [RFC7909] and Shared Whois Project (SWIP) [SWIP] are two formats in which the data in RIRs/IRRs are presented. ARIN predominantly uses SWIP but some RPSL as well. The rest of the RIRs and ISPs' IRRs use only RPSL.

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

Security Recommendation 1: All Internet Number Resources (e.g., address blocks and ASNs) should be properly registered in the appropriate RIR registration database and all appropriate 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, branch-offices, etc.) that operate their own network services (e.g., Internet access, DNS, etc.).

Security Recommendation 2: Route objects corresponding to the BGP routes originated from an Autonomous System 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 directly and by their outsourced IT systems and services.

4.2 Certification of Resources in Resource Public Key Infrastructure

Resource Public Key Infrastructure (RPKI) is a standards-based approach for providing cryptographically-secured registries of Internet resources and routing authorizations [RFC6480] [RFC6482] [NANOG] [Murphy]. The IPv4/IPv6 address and AS number resource allocations follow a hierarchy. Internet Assigned Numbers Authority (IANA) allocates resources to the Reginal Internet Registries (RIRs) such as ARIN, RIPE, etc., and the RIRs suballocate resources to ISPs and enterprises. The ISPs may further suballocate to other ISPs and enterprises. In some regions, RIRs suballocate to Local Internet Registries (LIRs) which in turn suballocate to ISPs and enterprises. RPKI is a global certificate authority (CA) and registry service offered by all Reginal Internet Registries (RIRs). The RPKI certification chain follows the same allocation hierarchy (see Figure 4). Although RPKI certifications are illustrated only under ARIN in Figure 4, a similar pattern is found in all other RIRs. Ideally there should be a single root or Trust Anchor (TA) at the top of the hierarchy. But currently each of the five RIRs (AFRINIC, APNIC, ARIN, LACNIC, and RIPE) maintains an independent TA for RPKI certification services in its respective region. Thus, the global RPKI is currently operating with five TAs (see, for example, [ARIN1] [ARIN2] [RIPE1] [RIPE2]).

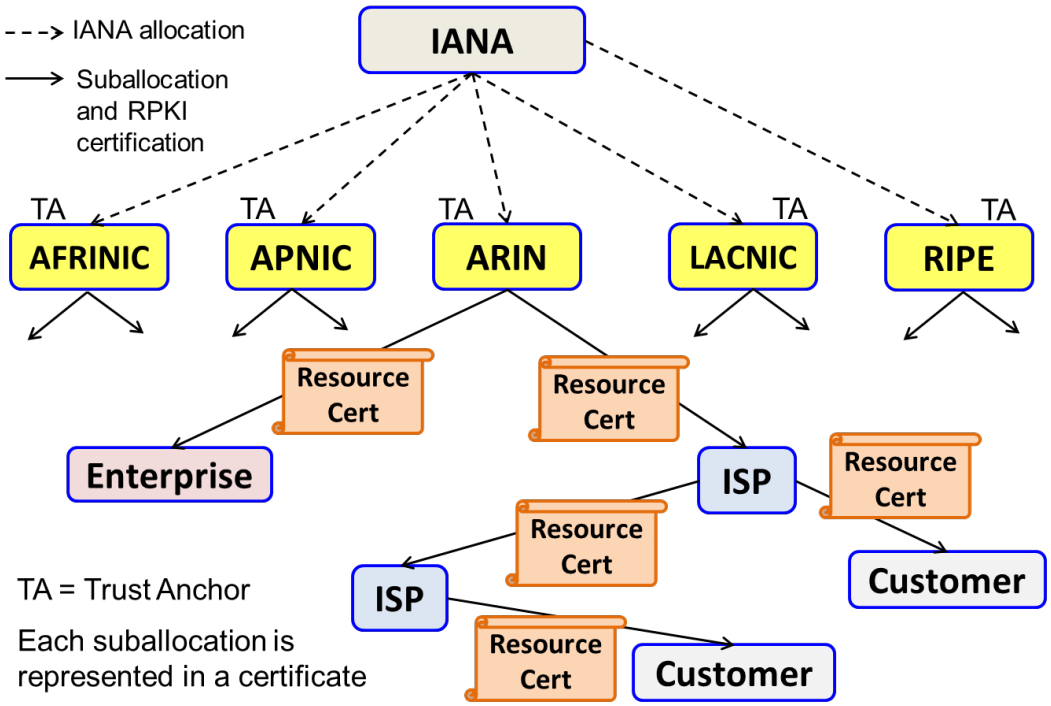

Figure 4: Illustration of resource allocation and certificate chain in RPKI.

RPKI is based on the X.509 standard with RFC 3779 extensions that describe special certificate profiles for Internet number resources (prefixes and ASN numbers) [RFC5280] [RFC6487] [RFC3779]. As shown in Figure 4, the RIRs issue resource certificates, called 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 and hosts their own certificate authority and performs RPKI operations (e.g., signs ROAs, issues subordinate resource certificates to their customers).

Security Recommendation 3: Internet number resource holders with IPv4/IPv6 prefixes and/or AS numbers (ASNs) should obtain RPKI certificate(s) for their resources.

Security Recommendation 4: Transit providers should provide a service where they create, publish, and manage subordinate resource certificates for address space and/or ASNs suballocated to their customers.

Currently, RPKI services based on the hosted model and offered by RIRs are common. The security recommendation immediately above can be implemented in the hosted or the delegated model based on service agreements with customers.

4.3 BGP Origin Validation (BGP-OV)

Once an address prefix owner obtains a CA certificate, 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) [RFC6482] [RFC6811]. A ROA declares a specific AS as an authorized originator of BGP announcements for the prefix (see Figure 5). A ROA specifies one or more prefixes, optionally a maxlength per prefix, and a single AS number. The meaning of maxlength is as follows. If a maxlength is specified for a prefix in the ROA, then any more specific (i.e., longer) prefixes (subsumed under the prefix) with a length not exceeding the maxlength are 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.

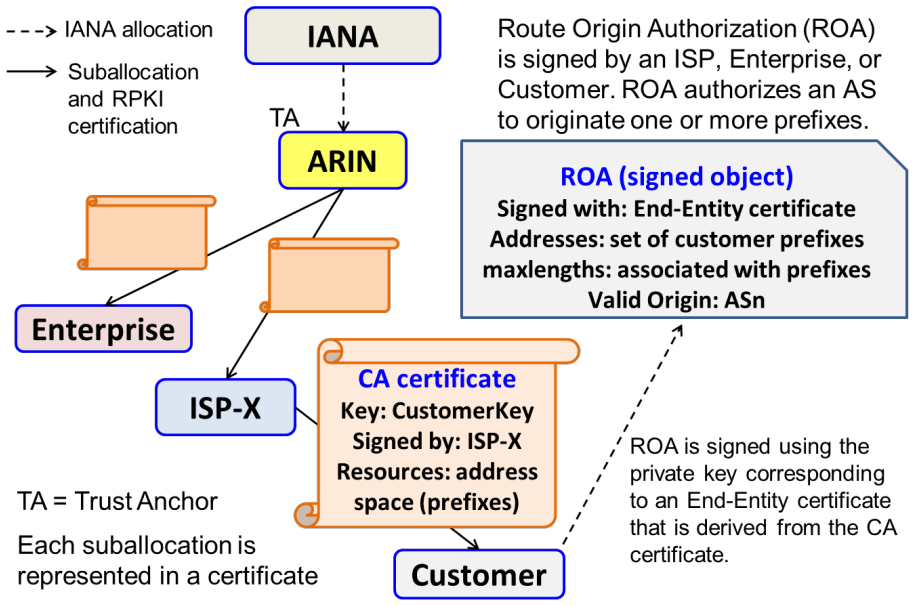

Figure 5: Creation of Route Origin Authorization (ROA) by prefix owner.

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. ISP can offer a service to its customers where the ISP creates and maintains 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* for short. 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 [RFC6810] [RFC8210]. More details regarding secure routing architecture based on RPKI can be found in [RFC6480].

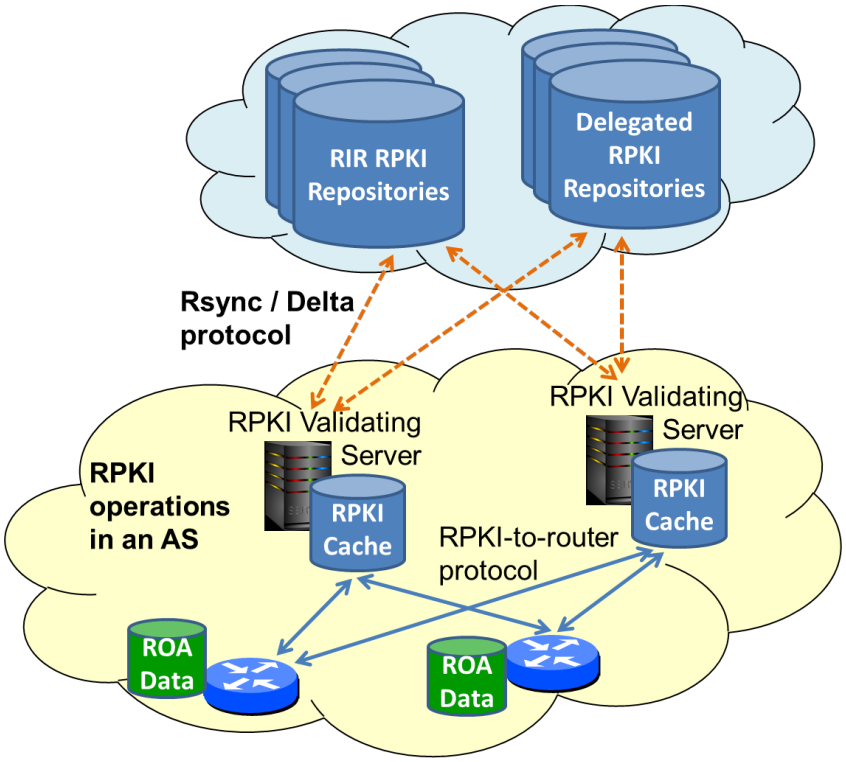

Figure 6: RPKI data retrieval, caching, and propagation to routers.

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 list of {prefix, maxlength, origin AS} tuples (derived from valid ROAs) from one or more RPKI cache servers. The router makes use of the list with the BGP origin validation (BGP-OV) 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) – Figure 7 provides additional details. Further, a route is deemed "NotFound" if the prefix announced is not covered by any prefix in the white 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 in BGP updates is discouraged in BCP 172 [RFC6472]. The RPKI-based origin validation may be supplemented by validation based on IRR data (see Section 4.1).

There are several implementations of RPKI-based BGP OV in both hardware and software-based router platforms [Juniper1] [Cisco1] [Patel] [Scudder] [NIST-SRx] [Parsons2] [goBGP] [RTRlib]. Deployment guidance and configuration guidance for many of these implementations are available from several sources [NCCoE-sidr] [RIPE1] [MANRS] etc. Although BGP-OV is already implemented in commercial BGP routers, the activation and ubiquitous use of RPKI and BGP-OV in BGP routers requires motivation and commitment on part of network operators.

588

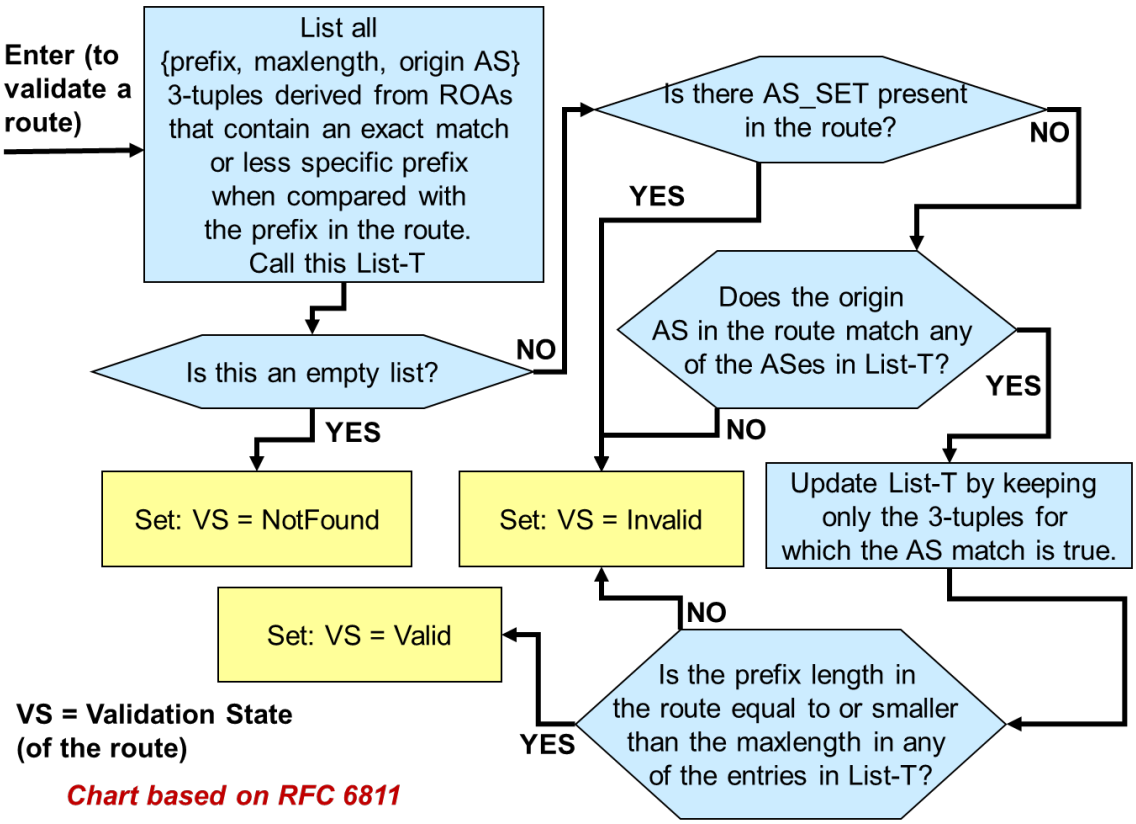

589

590 **Figure 7: Algorithm for origin validation (based on RFC 6811).**

591 **Security Recommendation 5:** Resource holders should register ROA(s) in the global
592 RPKI for all prefixes that are announced or intended to be announced in the public
593 Internet.

594 **Security Recommendation 6:** Transit providers should provide a service where they
595 create, publish, and maintain ROAs for their customers' prefixes.

596 Note: The security recommendation immediately above can be implemented in the hosted
597 or the delegated model based on service agreements with customers.

598 **Security Recommendation 7:** If a prefix that is announced (or intended to be
599 announced) is multihomed and originated from multiple ASes, then one ROA per
600 originating AS should be registered for the prefix (possibly in combination with other
601 prefixes which are also originated from the same AS).

602 **Security Recommendation 8:** When an ISP or enterprise owns multiple prefixes that
603 include less specific and more specific prefixes, they should ensure that the more specific
604 prefixes have ROAs before creating ROAs for the subsuming less specific prefixes.

605 **Security Recommendation 9:** An ISP should await until more specific prefixes that

606 are announced from within their customer cone have ROAs prior to the creation of its
607 own ROAs for subsuming less specific prefix(es).

608 AS 0 is a special AS number that is not allocated to any autonomous system. AS 0 is also not
609 permitted in routes announced in BGP. An AS0 ROA is one which has an AS 0 in it for the
610 originating AS [RFC6483] [APNIC1]. An address resource owner can create an AS 0 ROA for
611 their prefix to declare the intention that the prefix or any more specific prefix subsumed under it
612 must not be announced until and unless a normal ROA simultaneously exists for the prefix or the
613 more specific prefix.

614 **Security Recommendation 10:** An ISP or enterprise should create an AS0 ROA for
615 any prefix that is currently not announced to the public Internet.

616 **Security Recommendation 11:** A BGP router should not send updates with AS_SET
617 or AS_CONFED_SET in them (in compliance with BCP 172 [RFC6472]).

618 **Security Recommendation 12:** ISPs and enterprises who operate BGP routers
619 should also operate one or more RPKI validating caches.

620 **Security Recommendation 13:** A BGP router should maintain an up-to-date white
621 list consisting of {prefix, maxlength, origin ASN} that is derived from valid ROAs in the
622 global RPKI.

623 Note: The white list of {prefix, maxlength, origin ASN} 3-tuples can be typically
624 obtained (and periodically refreshed) by a router from a local RPKI cache server. As
625 mentioned before, the RPKI-to-router protocol [RFC6810] [RFC8210] is used for this
626 communication.

627 **Security Recommendation 14:** In partial/incremental deployment state of the RPKI,
628 the permissible {prefix, origin ASN} pairs should be generated by taking the union of
629 such data obtained from ROAs, IRR data, and customer contracts.

630 **Security Recommendation 15:** BGP-OV results should be incorporated into local
631 policy decisions to select BGP best paths.

632 Note (concerning the security recommendation immediately above): Exactly how BGP-
633 OV results are used in path selection is strictly a local policy decision for each network
634 operator. Typical policy choices include:

635 • Tag-Only – BGP-OV results are only used to tag/log data about BGP routes for
636 diagnostic purposes.
637 • Prefer-Valid – Use local preference settings to give priority to Valid routes. Note
638 this is only a tie breaking preference among routes with the exact same prefix.
639 • Drop-Invalid – Use local policy to ignore Invalid routes in the BGP decision
640 process.

641 Careful planning and thought should be given in the application of such policies. In
642 general, it is important that BGP-OV local policies be consistent throughout an individual

AS, both in terms of which peering sessions BGP-OV is enabled on, and in terms of how the results are used to influence the BGP decision process. It is recommended that network operators proceed through an incremental deployment process of adopting more stringent policies over time and after gaining experience and confidence in the system. The three example polices above, can be viewed as recommended stages of an incremental adoption plan.

It should be noted that enterprises should require their hosted-service providers (e.g., cloud, CDN, DNS, email, etc.) to follow the security recommendations stated here concerning certification of resources and creation of ROAs for the prefixes that are used in providing the hosted services and belong to the providers. An enterprise can do this themselves if the hosted-service provider is using the enterprises own address space for the hosted services.

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 a ROA for the target prefix onto his own unauthorized BGP announcement. In conjunction with forging the origin, for greater impact, 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. The security recommendations that follow are useful to minimize forged-origin attacks. (Note: BGP path validation (i.e., BGPsec [RFC8205]) described in Section 4.7 is required for full protection against prefix and/or path modifications.)

The following recommendation provides some degree of robustness against forged-origin attacks:

Security Recommendation 16: The maxlength in the ROA should preferably 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.

The following recommendation provides an even greater degree of robustness against forged-origin attacks.

Security Recommendation 17: 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 multiple ROAs (i.e., one ROA per prefix or more specific prefix) [maxlength].

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) [maxlength].

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] [NIST800-54]. Prefix filtering differs from BGP-OV 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 for access control lists. One option is to list ranges of IP prefixes that are to be denied, 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. Normally, 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] [IPv4-addr]. The RIRs have also nearly fully allocated their IPv4 address space [IPv4-addr]. (Some of the prefixes are designated for special use as discussed in Section 4.4.2.) 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 process, 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].

> **Security Recommendation 18:** 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.
>
> Note: If prefix resource owners regularly register AS 0 ROAs (see Section 4.3) for allocated (but possibly currently unused) prefixes, then those ROAs could be a complementary source for update of prefix filters mentioned above.

4.4.2 Special-Purpose Prefixes

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

> **Security Recommendation 19:** 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.

An AS may originate one or multiple prefixes. In the inbound direction, the AS should (in most cases) reject routes for the prefixes it originates if received from any of its eBGP peers (transit provider, customer, or lateral peer). In general, the data traffic destined for these prefixes should stay local and should not be leaked over external peering. However, if the AS operator is uncertain whether a prefix they originate is single-homed (or multihomed), then the AS should accept the prefix advertisement from an eBGP peer (and assign a lower local preference value) so that the desired redundancy is maintained.

Security Recommendation 20: For single-homed prefixes (subnets) that are owned and originated by an AS, any routes for those prefixes received at that AS from eBGP peers should be rejected.

4.4.3 Prefixes that Exceed a Specificity Limit

Normally, ISPs neither announce nor accept routes for prefixes that are more specific than a certain level of specificity. For example, maximum acceptable prefix lengths are mentioned in existing practices as /24 for IPv4 [RIPE-399] and /48 for IPv6 [RIPE-532]. The level of specificity that is acceptable is decided by each AS operator and communicated with peers. In instances when Flowspec (see Section 5.5) [RFC5575] [Hares] [Ryburn] is used between adjacent ASes for DDoS mitigation, the two ASes may mutually agree to accept longer prefix lengths (for example, a /32 for IPv4) but only for certain pre-agreed prefixes. That is, the announced more specific prefix must be contained within a pre-agreed prefix.

Security Recommendation 21: It is recommended that an eBGP router should set 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.

4.4.4 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 otherwise.

Security Recommendation 22: The default route (0.0.0.0/0 in IPv4 and ::/0 in IPv6) should be rejected except when a special peering agreement exists that permits accepting it.

4.4.5 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.

Security Recommendation 23: An Internet Exchange Provider (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 in turn should accept this prefix and reject any more specifics prefixes (of the IXP announced prefix) from any of its eBGP peers.

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

4.5 Prefix Filtering for Peers of Different Types

The inbound and outbound prefix filtering recommendations vary based on the type of peering relationship that exists between networks: lateral peer, transit provider, customer, and leaf customer (see definitions in Section 2.3). The different types of filters that apply are from the list described in Sections 4.4.1 through 4.4.5.

The security recommendations that follow apply to enterprises when they have eBGP peering with neighbor ASes. When an enterprise procures transit service from an ISP or hosted services (e.g., cloud, CDN, DNS, email, etc.) from hosted-service providers, the security recommendations should be included in the respective service contracts.

4.5.1 Prefix Filtering with Lateral Peer

Security Recommendation 24: Inbound prefix filtering (facing Lateral Peer): 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
- Default Route
- IXP LAN Prefixes

Security Recommendation 25: Outbound prefix filtering (facing Lateral Peer): The appropriate 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

Unallocated Prefixes may be omitted from the list of outbound prefix filters above if there is confidence that the inbound prefix filters are not letting them in.

4.5.2 Prefix Filtering with Transit Provider

Security Recommendation 26: Inbound prefix filtering (facing Transit Provider): 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

Not that the default route is not included in the above list. In some cases, a customer network prefers to receive the default route from a transit provider in addition to the full routing table.

Security Recommendation 27: Inbound prefix filtering (facing Transit Provider): 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.

Security Recommendation 28: 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 above Outbound prefix filtering security recommendation, some policy rules may also be applied if a transit provider is not contracted (or not chosen) to provide transit for some subset of outbound prefixes.

4.5.3 Prefix Filtering with Customer

Inbound prefix filtering: There are two scenarios that need consideration. **Scenario 1** is when there is full visibility of the customer and its cone of customers (if any), and there is knowledge of prefixes originated from such a customer and its cone. The knowledge of prefixes can be based on direct customer knowledge, IRR data and/or RPKI data (if that data is known to be in 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.

Security Recommendation 29: Inbound prefix filtering (facing Customer, Scenario 1): 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.

Scenario 2 is when there is not a reliable knowledge of all prefixes originated from the customer

and its cone of customers.

Security Recommendation 30: Inbound prefix filtering (facing Customer, Scenario 2): The same set of inbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1).

Security Recommendation 31: 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 full routing table. If it is the former, then the 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 filter may be added in the above list if the customer requires the full routing table but not the default route.

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.

Security Recommendation 32: Inbound prefix filtering (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 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

Security Recommendation 33: Outbound prefix filtering (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) to observe extra caution.

4.6 Role of RPKI in Prefix Filtering

An ISP can retrieve (from RPKI registries) all available Route Origin Authorizations (ROAs) corresponding to autonomous systems (ASes) that are known to belong in their customer cone. From the available ROAs, it is possible to determine the prefixes that can be originated from the corresponding ASes in the customer cone. Based on a knowledge of the tree structure of the customer cone, it is further possible to list all the prefixes that could be received on any given

customer interface (see Section 3.8 in [RouteLeak3]). As the RPKI registries become mature (with increasing adoption), the prefix lists derived from ROAs will become useful for prefix filtering. Even in the early stages of RPKI adoption, the prefix lists (from ROAs) can help cross-check and/or augment the prefix filter lists that an ISP constructs by other means.

Note: The list of ASes in an AS's customer cone can be determined by forming the list of unique origin ASes in all BGP announcements received (i.e., currently in the Adj-RIB-ins [RFC4271]) on all customer interfaces at the AS in consideration. This can be done in the network management system (off the router).

Security Recommendation 34: The ROA data (available from RPKI registries) should be used to construct and/or augment prefix filter lists for customer interfaces.

4.7 AS Path Validation (Emerging/Future)

Note: The IETF standard for BGP path validation (BGP-PV), namely BGPsec [RFC8205], is available but commercial vendor implementations are not currently available. Hence, this section briefly describes the technology and standards but does not make any security recommendations concerning BGP-PV.

As observed in Sections 4.3 and 4.3.1, BGP origin validation (BGP-OV) is necessary but by itself it is insufficient for fully securing the prefix and AS path in BGP announcements. BGP path validation (BGP-PV) is additionally required to protect against prefix modifications and forged-origin attacks (see Section 4.3.1) as well as other AS-path attacks such as path shortening and Kapela-Pilosov attacks (see Section 2.2). There is significant interest in the networking community to secure the AS path in BGP updates so that a more comprehensive protection can be provided to BGP updates [RFC8205] [RFC8208] [RFC7353] [Huston2011] [RFC8374]. RFC 8205 is the IETF standard that specifies the BGPsec protocol, i.e., the protocol for BGP path validation. Open source prototype implementations of BGP-PV are available [NIST-SRx] [Parsons2] [Adalier2].

The basic principles of BGP-PV are illustrated in Figure 8. (Please see [RFC8205] for a detailed protocol specification.) A ROA signed by the owner the prefix 10.1.0.0/16 attests that AS1 is authorized to originate the prefix. Further, each network operator that has deployed BGP-PV gets a resource certificate for their AS number, and the BGP-PV routers within the AS get router certificates and private keys for signing updates. The certificates for all BGP-PV routers are retrieved by all participating ASes, and the public keys of all BGP-PV routers are expected to be available at each BGP-PV 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 BGP-PV 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 will 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.

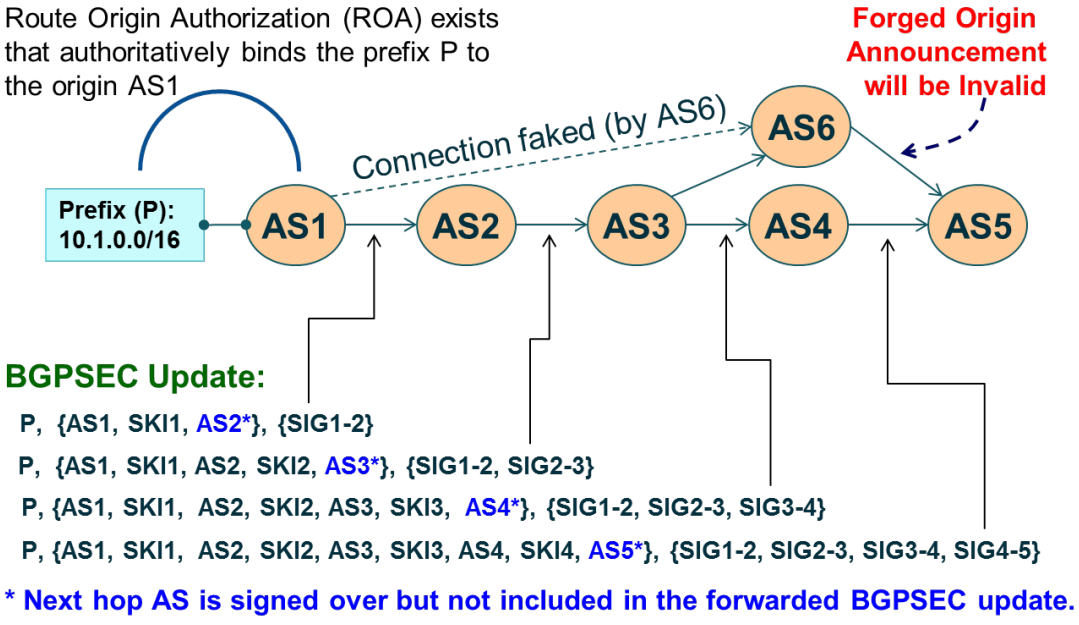

Figure 8: Basic principle of signing/validating AS paths in BGP updates.

ECDSA-P256 algorithm is currently recommended for signing BGPsec updates between ASes that peer with each other [RFC8208]. 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 BGP-PV routers will be required to perform additional processing due to signing and verification of path signatures. The performance characterization of BGP-PV 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].

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

The standards for BGP-PV are documented in IETF RFC's #8205 through #8210. When implementations based on these standards start to become available in commercial products, this document may be updated to recommend BGP-PV.

4.8 Route Leak Solution (Emerging/Future)

Section 2.3 described the route leaks problem space and noted that in RFC 7908 [RFC7908] the various types of route leaks are enumerated. Route leak solutions fall in two categories: (1) Intra-AS and (2) 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 is was received in eBGP from a customer, a lateral peer, a transit provider, etc. 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 – 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).

Security Recommendation 35: An AS operator should have 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.

Security Recommendation 36: An AS operator should have egress policy to utilize the tagged information (in the preceding Security Recommendation) to prevent route leaks when routes are forwarded on the egress.

The above intra-AS solution for prevention of route leaks can also be implemented using a BGP Attribute (instead of BGP Community). The Attribute-based solution [RouteLeak2] has the advantage that it can be made available in commercial routers as a standard feature, which in turn minimizes manual network operator actions. However, such a solution involves an update to the BGP protocol [RFC4271] and requires standardization. Such an effort takes time and is currently in progress in the IETF [RouteLeak2].

The second type of solution that is inter-AS is intended to work in eBGP across AS hops. With the inter-AS solution, the focus shifts to detection and mitigation in case a route leak has already occurred and started to propagate. The idea is that if a leak indeed propagates out of an AS, then the peer AS or any AS along the subsequent AS path should be able to detect and stop it. Solution for inter-AS route leak detection and mitigation is also work in progress in the IETF [RouteLeak1] [RouteLeak3].

For robustness of the Internet routing infrastructure, inter-AS route-leak detection and mitigation capability will also need to be implemented in addition to the intra-AS prevention capability. When mechanisms for route-leak detection and mitigation capability are standardized and become available in products, this document may be updated to include appropriate security recommendations to reflect the same.

5 Securing Against DDoS & Reflection-Amplification – Solutions and Recommendations

There are various existing techniques and recommendations for deterrence against DDoS attacks with spoofed addresses [BCP38] [BCP84] [NABCOP] [CSRIC-WG5]. There are also some techniques used for prevention of reflection-amplification attacks [RRL] [TA14-017A], which are used in achieving greater impact in DDoS attacks. Employing a combination of these preventive techniques in enterprise and ISP border routers, hosted-service provider networks, DNS/NTP servers, broadband and wireless access networks, and data centers provides the necessary protections against DDoS attacks.

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), Digital Subscriber Line Access Multiplexers (DSLAM), and Packet Data Network (PDN) gateways in mobile networks. Ingress/egress Access Control List (ACL) and unicast Reverse Path Forwarding (uRPF) are techniques employed for implementing SAV [BCP38] [BCP84] [ISOC] [RFC6092; REC-5, REC-6]. Ingress SAV applies to incoming (received) packets and egress SAV applies to outgoing (transmitted) packets.

5.1.1 SAV using Access Control List

Ingress/egress Access Control Lists (ACLs) are maintained which list acceptable (or alternatively, unacceptable) prefixes for the source addresses in the incoming/outgoing Internet Protocol (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 multihomed, 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) 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 disallowed prefix blocks, RFC 1918 prefixes, or provider's/enterprise's own prefixes).

5.1.2 SAV using Strict Unicast Reverse Path Forwarding

In the strict unicast Reverse Path Forwarding (uRPF) method, an ingress packet on an interface at the border router is accepted only if (1) the Forwarding Information Base (FIB) contains a prefix that encompasses the source address, and (2) 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 in consideration. It is well known that this method has limitations when a network or autonomous system is multi-homed and there is asymmetric routing of packets. Asymmetric routing occurs (see Figure 9) 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.

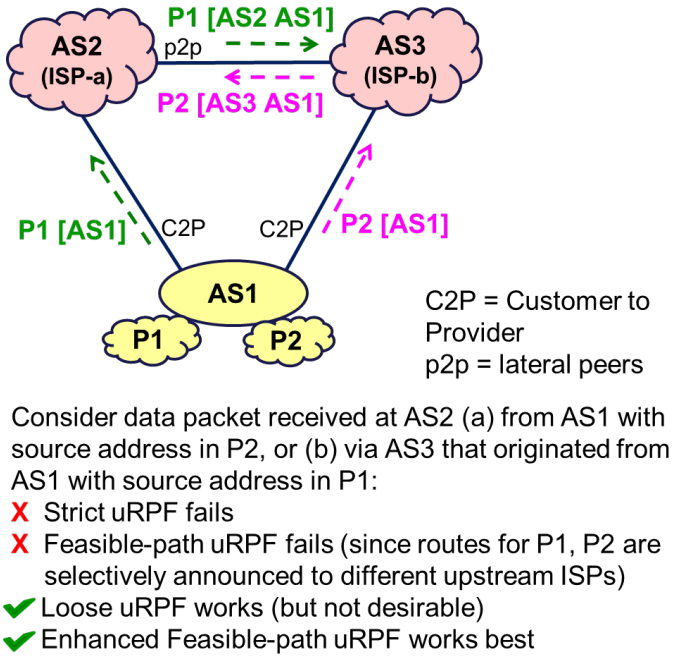

Figure 9: Scenario 1 for illustration of efficacy of uRPF schemes.

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 the difference is that instead of inserting one best route in the FIB (or an equivalent Reverse Path Forwarding (RPF) table), alternative routes are also added there. This method relies on announcements for the same prefixes (albeit some may be prepended to effect lower preference) propagating to all the eBGP-peer routers performing feasible-path uRPF check. So, in the multi-homing scenario, 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 (see Figure 10). Alternatively, it also works if the customer AS announces the aggregate of P1 and P2 (if possible) to each transit provider in addition to announcing P1 to one provider and P2 to the other provider. It should be mentioned that the feasible-path uRPF works in this scenario only if customer route is preferred at AS2 and AS3 over the shorter path.

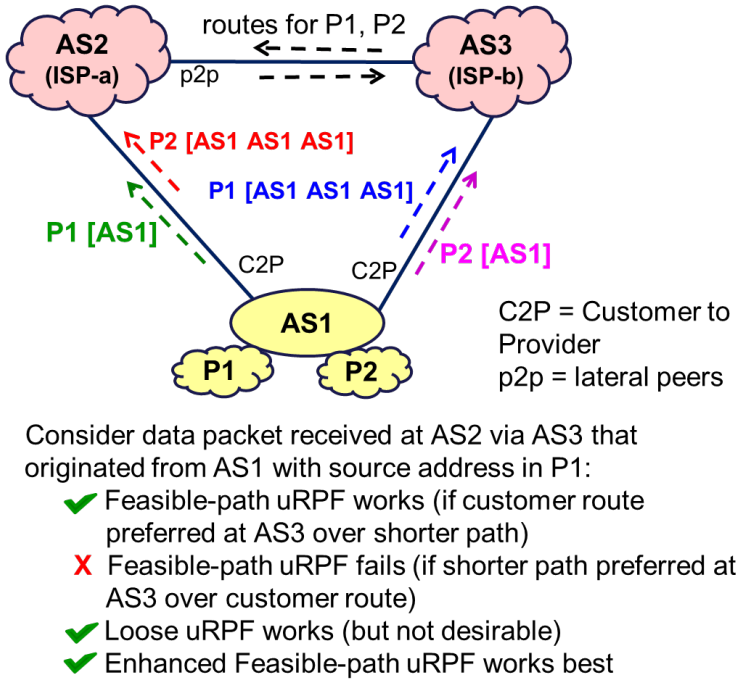

Figure 10: Scenario 2 for illustration of efficacy of uRPF schemes.

However, the feasible-path uRPF method has limitations as well. One form of limitation naturally occurs when the recommendation of propagating the same prefixes to all routers is not heeded. Another form of limitation can be described as follows. In Scenario 2 (described above, illustrated in Figure 10), it is possible that the second transit provider (ISP-b) does not propagate the prepended route (i.e., P1 [AS1 AS1 AS1]) to the first transit provider (ISP1). 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 source address in prefix P1 that traverses via AS3 (ISP-b) will get dropped at AS2 (ISP-a) despite the flexibility accorded by feasible path uRPF.

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 encompass the source address. That is, a packet is dropped if no route exists in the FIB for the source address. Loose uRPF sacrifices directionality. In most cases, this method is not effective for prevention of address spoofing. Nearly all IPv4 address space already appears in the global routing table. Hence, for IPv4, loose uRPF only drops packets if the spoofed address is non-routable (e.g., RFC 1918, unallocated, allocated but currently not routed). It may be noted that the method is more useful for IPv6 than IPv4.

5.1.5 SAV using Enhanced Feasible-Path uRPF

Note: The status of the Enhanced Feasible-Path uRPF (EFP-uRPF) is that it is currently work in progress in the IETF [EFP-uRPF]. It holds promise for providing a significant improvement in

effectiveness and deployability over the Feasible Path uRPF. Hence, this section briefly describes the technology and standards effort but does not make a security recommendation concerning use of EFP-uRPF.

Enhanced feasible-path uRPF (proposed in [EFP-uRPF]) adds greater flexibility and accuracy to uRPF operation than the three uRPF methods discussed above in Sections 5.1.2 through 5.1.4. The basic principle of EFP-uRPF method for enhancing the efficacy in multi-homing and asymmetric routing scenarios is as follows. 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 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 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.

Looking back at Scenarios 1 and 2 (Figure 9 and Figure 10), the EFP-uRPF provides comparable or better performance than the other uRPF methods for those scenarios. Scenario 3 (Figure 11) further illustrates that of EFP-uRPF method works best even in a much more complex asymmetric routing scenario. In Scenario 3 (Figure 11), the focus is on AS4 receiving data packets with source address in {P1, P2, P3}. If EFP-uRPF is used, the operator (at AS4) can be assured that DDoS mitigation would work effectively while none of those data packets would be subject to denial of service. The details concerning EFP-uRPF can be found in [EFP-uRPF]. It is still work in progress, so no security recommendations involving EFP-uRPF are offered here.

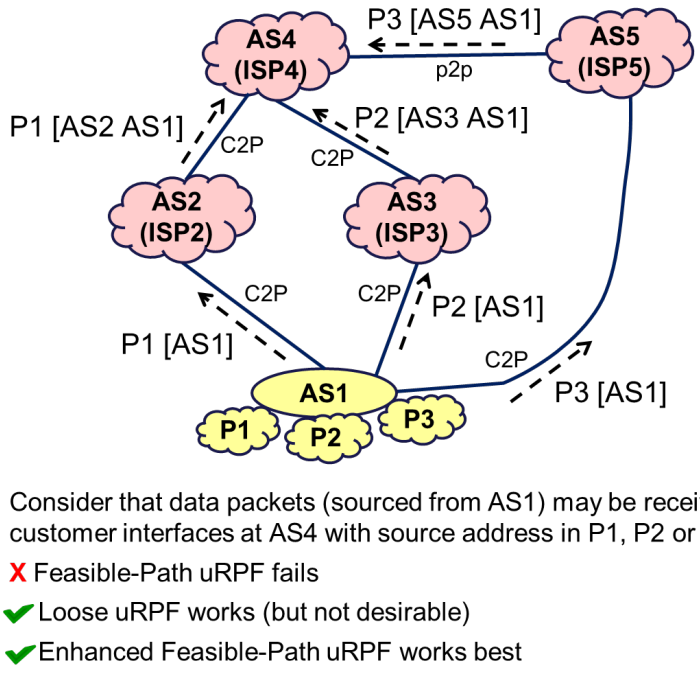

Figure 11: Scenario 3 for illustration of efficacy of uRPF schemes.

5.1.6 More Effective Mitigation with Combination of Origin Validation and SAV

It is worth noting that with the combination of BGP origin validation (BGP-OV) (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 12. In the scenario in Figure 12, the attacker first compromises routers (or perhaps owns some of them) at AS98 and AS99, 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). The feasible-path uRPF filters at AS5 and AS6 are effectively deceived, and the attacker stays under the radar because the hijacked prefix is a less specific prefix. Then the attacker would be able to successfully perform address spoofing and DDoS with reflection-amplification. To protect against this type of multi-pronged attack, the combination of BGP-OV (to prevent the hijacking) and feasible-path uRPF (to prevent the address spoofing) should be employed. For this to work, the target prefix (10.1.0.0/22) owner should create a ROA for the prefix and all ASes (especially, AS5 and AS6) in Figure 12 should be performing BGP-OV in addition to employing uRPF.

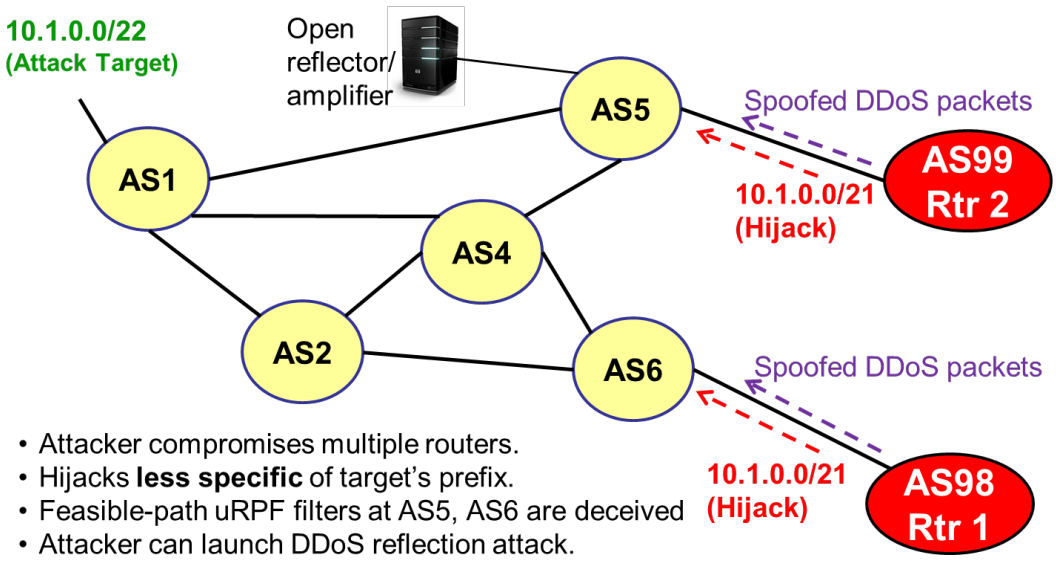

Figure 12: Illustration of how origin validation complements SAV.

5.2 SAV Recommendations for Various Types of Networks

Three types of network scenarios are considered here and SAV security recommendations are provided for each scenario. The network types are: (1) Networks that have customers with directly-connected allocated address space such as broadband and wireless service providers, (2) Enterprise networks, and (3) Internet Service Providers (ISPs).

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

5.2.1 Customer with Directly-Connected Allocated Address Space: Broadband and Wireless Service Providers

SAV with ACLs (described above) is relatively easy when a network served by an ISP's edge device (e.g., border router, CMTS, DSLAM, PDN gateway) is directly connected (without multi-homing) and using an IP address space that is suballocated by the ISP. Hence, SAV using ACL method should be always used in such cases. For the egress packets (i.e., packets transiting via the edge device into the Internet), the source address must be within the allocated space. As an example, the DOCSIS 3.0 specification for CMTS already incorporates this security check [DOCSIS][Comcast].

Security Recommendation 37: BGP routers that have directly-connected customers with suballocated address space, CMTS (or equivalent) in broadband access networks, and PDN gateways (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).

5.2.2 Enterprise Border Routers

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

Security Recommendation 38: An enterprise border router that is multi-homed should always announce all its prefixes to each of its upstream transit providers (albeit with appropriate AS prepending for traffic engineering). It should avoid selectively announcing some prefixes to one transit ISP and other prefixes to another transit ISP.

Note: By following the above recommendation, the enterprise border router ensures that that the transit ISPs' border routers discard (due to uRPF) only those data packets from the enterprise that do not have source addresses belonging in any of the enterprise's announced prefixes. Thus, it also ensures that data packets from the enterprise that have source addresses belonging in any of the enterprise's announced prefixes are never denied.

Security Recommendation 39: This is the exception case when the enterprise border router does not adhere to the above recommendation and instead selectively announces some prefixes to one upstream transit ISP and other prefixes to another upstream transit ISP. In this case, it should ensure (by appropriate internal routing) that the source addresses in the data packets towards each upstream transit ISP belong in the prefix or prefixes announced to that ISP.

Security Recommendation 40: 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, RFC 1918 prefixes, or enterprise's own prefixes).

5.2.3 Internet Service Providers

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

Security Recommendation 41: On customer facing interfaces, ISPs should do 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, especially in the case of multi-homed customers.

Note: In the future, the enhanced feasible-path uRPF (see Section 5.1.5) may be considered (based on progress with its standardization and availability of commercial implementation).

Security Recommendation 42: For feasible-path uRPF to work appropriately, the ISPs (at least those near the Internet edge) should propagate all their customer routes to their upstream transit ISPs (albeit with appropriate AS prepending for traffic engineering).

Security Recommendation 43: 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 the above recommendation facilitates a basis for adhering to the preceding recommendation as well. (The above recommendation is also one of the stability conditions on BGP policy for ensuring stable convergence of routing information [Gao-Rexford].)

Security Recommendation 44: On interfaces with lateral (i.e., non-transit) peers, ISPs should do SAV on ingress packets by deploying the feasible-path uRPF (see Sections 5.1.3). They should avoid using strict or loose uRPF as they are not very effective for SAV on the lateral peer interfaces.

Security Recommendation 45: On interfaces with transit providers, ISPs should do SAV on ingress packets by deploying 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, RFC1918 prefixes, ISP's own prefixes).

Security Recommendation 46: On the egress side towards customers, lateral (i.e., non-transit) peers and transit providers, the ISP's border routers should deploy ACLs (Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks, RFC 1918 prefixes, ISP's internal-use only prefixes).

5.3 Role of RPKI in Source Address Validation

A method was described in Section 4.6 on how ISPs can use the ROAs in RPKI registries to assist with construction of prefix filters. The same technique can be applied to construct ACLs for SAV on each customer facing interface. These ACLs can be used to cross-check and/or augment entries in the RPF lists corresponding to each customer facing interface.

Security Recommendation 47: The ROA data (available from RPKI registries) should be used to construct and/or augment ACLs/RPF lists for SAV on customer interfaces

5.4 Monitoring UDP/TCP Ports with Vulnerable Applications and Employing Traffic Filtering

DDoS threats involving vulnerable applications using various UDP/TCP ports and IoT devices are continually evolving and varied, e.g., memcached DDoS reflection attacks and SSDP diffraction, etc. [Bjarnason]. Hence, traffic filtering methods mentioned in this section are not meant to be exhaustive.

Traffic monitoring and filtering based on specific User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) ports is done to deny traffic of certain application types that are not expected on a given interface in consideration [TA14-017A] [Acunetix] [ISC2] [Arbor]. In some cases, the applications may be legitimate but the observed traffic volumes may be suspiciously high, in which case response rate limiting is applied [Redbarn] [ISC1].

In the case of the DNS (Port 53), the DNS resolver can limit the scope of clients from which it will accept requests. The clients normally come from within the same network where the DNS resolver resides. Hence, the DNS resolver can maintain access lists in the configuration so that an otherwise open DNS resolver can be effectively ‘closed’ [ISOC]. Another effective measure is for the authoritative DNS resolvers to monitor the rate of queries per source address and apply Response Rate Limiting (RRL). The RRL dampens the rate at which authoritative servers respond to high volumes of malicious queries [Redbarn] [ISC1].

1190 Table 1 below lists application-layer protocols and their port numbers. The UDP-based
1191 applications have been identified as vulnerable to reflection/amplification attacks.

1192 **Table 1: Common Applications and their TCP/UDP Port Numbers.**

Application Protocol	Bandwidth Amplification Factor	Port #	Port Assignment Status
Domain Name System (DNS)	28 to 54	53, 853, 953	Official
Network Time Protocol (NTP)	557	123	Official
Simple Network Management Protocol (SNMP), SNMPv2	6	161	Official
NetBIOS Name/Datagram/Session	4	137/138/139	Official
Simple Service Discovery Protocol (SSDP); discovery of UPnP devices	31	1900	Official
Character Generation Protocol (CharGEN)	359	19	Official
Quote of the Day (QOTD)	140	17	Official
BitTorrent	4	6881-6887; 6889-90; 6891-6900; etc. various ranges	Unofficial
Kad network (Kademlia P2P overlay protocol)	16	6419, 6429	Unofficial
Quake Network Protocol	64	15, 28, 27500-27900, 27901-27910, 27950, 27952, 27960-27969, etc.	Unofficial
Streaming Protocols (e.g., QuickTime)		6970-9999, etc.	Unofficial
Real-Time Streaming Protocol (RTSP); ms-streaming		554, 1755	Official
Routing Information Protocol (RIP, RIPng)	131	520, 521	Official
Multicast DNS (mDNS)	2 to 10	5353	Official
Portmap/RPC	7 to 28	369	Official

1193 In Table 1, the amplification factor listed for each protocol is the traffic volume multiplier that
1194 can be achieved by exploiting the reflection/amplification effect of that protocol run on UDP
1195 [TA14-017A]. Port assignment status is called ‘Official’ if officially assigned by IANA;
1196 otherwise it is ‘Unofficial’ [TCP-UDP-port]. The following set of security recommendations

pertain to vulnerable applications such as those listed in Table 1.

Security Recommendation 48: Port 0 is a reserved port. Hence, deny TCP/UDP traffic on Port 0 on all interfaces.

Security Recommendation 49: In BGP routers, allow peers to connect to only port 179. The standard port for receiving BGP session OPEN messages is port 179, so attempts by BGP peers to reach other ports are likely to indicate faulty configuration or potential malicious activity.

Security Recommendation 50: Disable applications or services that are unwanted in the network or system in consideration.

Security Recommendation 51: Deny traffic for any TCP/UDP ports for which the network or system in consideration does not support the corresponding applications. In some cases, an application or service is supported on some interfaces (e.g., customer or internal facing interfaces) but not others (e.g., Internet facing interfaces). In such cases, the traffic with port ID specific to the application in consideration should be denied on interfaces on which the application is not supported.

Security Recommendation 52: This recommendation is aimed at detection of traffic overload and mitigating actions. The relevant mitigation techniques are (a) Response Rate Limiting (RRL) [ISC1] [Redbarn], and (b) Source-based Remote Triggered Black Hole (S/RTBH) filtering enabled with Flowspec [RFC5575] (see Section 5.5 for details). These techniques are applicable to open services/protocols such as those listed in Table 1 which are themselves vulnerable to DoS/DDoS attacks or may be exploited for reflection/amplification. The recommendation consists of multiple steps as follow [TA14-017A]:

- Monitor the rate of queries/requests per source address and detect if abnormally high volume of responses is headed to the same destination (i.e., same IP address).
- Apply the Response Rate Limiting (RRL) technique to mitigate the attack.
- Using BGP messaging (Flowspec), create a Remotely Triggered Black Hole (RTBH) filter. This can be coordinated with the upstream ISP.
- Maintain emergency contact information for the upstream provider to coordinate response to the attack.
- An upstream ISP should actively coordinate response with downstream customers.

Note: The RRL technique is commonly used in DNS and dampens the rate at which authoritative servers respond to high volumes of malicious queries. It can also be applied in other applications (shown in Table 1) for dampening the response rate.

The security recommendations that follow below are specific to NTP and DNS.

Security Recommendation 53: Deny NTP monlist request traffic (by disabling the monlist command) altogether, or at least enforce that the requests come from valid (permitted) source addresses.

Security Recommendation 54: To limit exploitation, a DNS recursive resolver should limit the scope of clients from which it accepts requests. The clients normally come from within the same network where the DNS resolver resides. Hence, the DNS resolver can maintain access lists in the configuration so that the recursive resolver is not open to the entire network (or Internet) [ISOC] [TA14-017A].

Security Recommendation 55: Deny all traffic with a source or destination address that matches a DNS anycast address. An exception should be made for internal recursive resolvers that are used to do outbound recursion.

Security Recommendation 56: Block all inbound/outbound Port 53 UDP messages at DNS recursive resolvers except those from designated recursive resolvers.

5.5 BGP Flow Specification (Flowspec)

Destination-based Remote Triggered Black-Holing (D/RTBH) [RFC3882] [RFC7999] and Source-based Remote 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 [RFC5575] [RFC7674] [Hares] [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. 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 a 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 a 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, a flow specification NLRI is defined and it is used to convey information about traffic filtering rules for traffic that should be discarded [RFC5575]. This mechanism allows an upstream AS to perform inbound filtering in their edge routers of traffic that a given downstream AS wishes to drop. Table 2 shows the information that can be included in BGP Flowspec [RFC5575].

Table 2: BGP Flowspec 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

Table 3 shows the extended community values that are defined to specify various types of actions [RFC5575] requested at the upstream AS.

Table 3: Extended community values defined in Flowspec to specify various types of actions.

type	extended community	encoding
0x8006	traffic-rate (set to 0 to drop all traffic)	2-byte as#, 4-byte float
0x8007	traffic-action (sampling)	bitmask
0x8008	redirect to VRF (route target)	6-byte Route Target
0x8009	traffic-marking	DSCP value

In the table above VRF stands for Virtual Routing and Forwarding, and DSCP stands for Differentiated Services Code Point (DSCP). As evident from the discussion above and Table 2 and Table 3, Flowspec facilitates flexible specification and communication (by downstream AS) of rules and actions for DDoS mitigation to be executed at edge routers in the upstream AS.

Security Recommendation 57: Edge routers should be equipped to perform Destination-based Remote Triggered Black Hole (D/RTBH) filtering and Source-based Remote Triggered Black Hole (S/RTBH) filtering.

Security Recommendation 58: Edge routers should be equipped to make use of BGP flow specification (Flowspec) to facilitate DoS/DDoS mitigation (in coordination between upstream and downstream autonomous systems).

Security Recommendation 59: Edge routers – in an AS providing RTBH filtering – should have ingress policy towards RTBH customers to accept routes more specific than /24 in IPv4 and more specific than /64 in IPv6. Also, the edge routers should accept such 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 (e.g., has a ROA for the less specific route). 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.

Security Recommendation 60: A customer AS should make sure that the routes announced for RTBH filtering have NO_EXPORT, NO_ADVERTISE, or similar communities.

Security Recommendation 61: An ISP providing RTBH filtering service to customers must have egress policy that denies routes that have community tagging meant for triggering RTBH filtering. This is an additional safeguard in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails to work for some reason.

Security Recommendation 62: An ISP providing RTBH filtering service to customers must have egress policy that denies prefixes that are longer than expected. This provides added safety in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails to work for some reason.

Appendix A— Consolidated List of the Security Recommendations

Table 4 provides a consolidated list of the Security Recommendations (copied from the various sections throughout the document). If "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). Similar statement applies for ISPs when the ISP column is checked. The "Open Servers" column pertains to providers of open Internet services such as DNS, DNSSEC, NTP, etc. When an enterprise outsources services, then the feature/service corresponding to a security recommendation that applies to them would in turn apply to their hosting 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 4 corresponding to Internet Exchange Point (IXP), but the BGP (control plane) security recommendations for ISPs also apply to opaque IXPs (i.e., IXPs that insert their ASN in the AS path and operate BGP).

Table 4: Consolidated List of the Security Recommendations

	Applicable to		
Security Recommendation	**Enter-prise**	**ISP**	**Open Servers**
BGP Origin Validation:			
Security Recommendation 1: All Internet Number Resources (e.g., address blocks and ASNs) should be properly registered in the appropriate RIR registration database and all appropriate 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, branch-offices, etc.) that operate their own network services (e.g., Internet access, DNS, etc.).	X	X	
Security Recommendation 2: Route objects corresponding to the BGP routes originated from an Autonomous System 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 directly and by their outsourced IT systems and services.	X	X	
Security Recommendation 3: Internet number resource holders with IPv4/IPv6 prefixes and/or AS numbers (ASNs)	X	X	

should obtain RPKI certificate(s) for their resources.			
Security Recommendation 4: Transit providers should provide a service where they create, publish, and manage subordinate resource certificates for address space and/or ASNs suballocated to their customers.		X	
Security Recommendation 5: Resource holders should register ROA(s) in the global RPKI for all prefixes that are announced or intended to be announced in the public Internet.	X	X	
Security Recommendation 6: Transit providers should provide a service where they create, publish, and maintain ROAs for their customers' prefixes. Note: The security recommendation immediately above can be implemented in the hosted or the delegated model based on service agreements with customers.		X	
Security Recommendation 7: If a prefix that is announced (or intended to be announced) is multihomed and originated from multiple ASes, then one ROA per 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 8: When an ISP or enterprise owns multiple prefixes that include less specific and more specific prefixes, they should ensure that the more specific prefixes have ROAs before creating ROAs for the subsuming less specific prefixes.	X	X	
Security Recommendation 9: An ISP should await until more specific prefixes that are announced from within their customer cone have ROAs prior to the creation of its own ROAs for subsuming less specific prefix(es).		X	
Security Recommendation 10: An ISP or enterprise should create an AS0 ROA for any prefix that is currently not announced to the public Internet.	X	X	
Security Recommendation 11: A BGP router should not send updates with AS_SET or AS_CONFED_SET in them (in compliance with BCP 172 [RFC6472]).	X	X	

Security Recommendation 12: ISPs and enterprises who operate BGP routers should also operate one or more RPKI validating caches.		X	
Security Recommendation 13: A BGP router should maintain an up-to-date white list consisting of {prefix, maxlength, origin ASN} that is derived from valid ROAs in the global RPKI.	X	X	
Security Recommendation 14: In partial/incremental deployment state of the RPKI, the permissible {prefix, origin ASN} pairs should be generated by taking the union of such data obtained from ROAs, IRR data, and customer contracts.	X	X	
Security Recommendation 15: BGP-OV results should be incorporated into local policy decisions to select BGP best paths. Note (concerning the security recommendation immediately above): Exactly how BGP-OV results are used in path selection is strictly a local policy decision for each network operator. Typical policy choices include: • Tag-Only – BGP-OV 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 this is only a tie breaking preference among routes with the exact same prefix. • Drop-Invalid – Use local policy to ignore Invalid routes in the BGP decision process.	X	X	
Security Recommendation 16: The maxlength in the ROA should preferably 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 17: 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 multiple ROAs (i.e., one ROA per prefix or more specific prefix) [maxlength]. 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) [maxlength].	X	X	

Prefix (Route) Filtering:			
Security Recommendation 18: 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. Note: If prefix resource owners regularly register AS 0 ROAs (see Section 4.3) for allocated (but possibly currently unused) prefixes, then those ROAs could be a complementary source for update of prefix filters mentioned above.	X	X	
Security Recommendation 19: 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 20: For single-homed prefixes (subnets) that are owned and originated by an AS, any routes for those prefixes received at that AS from eBGP peers should be rejected.	X	X	
Security Recommendation 21: It is recommended that an eBGP router should set 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 22: The default route (0.0.0.0/0 in IPv4 and ::/0 in IPv6) should be rejected except when a special peering agreement exists that permits accepting it.	X	X	
Security Recommendation 23: An Internet Exchange Provider (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 in turn should accept this prefix and reject any more specifics prefixes (of the IXP announced prefix) from any of its eBGP peers.	X	X	
Security Recommendation 24: Inbound prefix filtering (facing Lateral Peer): The following prefix filters should be applied in the inbound direction: • Unallocated Prefixes • Special-Purpose Prefixes	X	X	

• Prefixes that the AS Originates • Prefixes that Exceed a Specificity Limit • Default Route • IXP LAN Prefixes			
Security Recommendation 25: Outbound prefix filtering (facing Lateral Peer): The appropriate 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	X	X	
Security Recommendation 26: Inbound prefix filtering (facing Transit Provider): 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 27: Inbound prefix filtering (facing Transit Provider): 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 28: 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 above Outbound prefix filtering security recommendation, some policy rules may also be applied if a transit provider is not contracted (or not chosen) to provide transit for some subset of outbound prefixes.	X	X	
Security Recommendation 29: Inbound prefix filtering (facing Customer, Scenario 1): Only the prefixes that are known to be originated from the customer and its customer cone should		X	

be accepted and all other route announcements should be rejected.			
Security Recommendation 30: Inbound prefix filtering (facing Customer, Scenario 2): The same set of inbound prefix filters should be applied as those for a lateral peer (see Section 4.5.1).		X	
Security Recommendation 31: 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 full routing table. If it is the former, then the 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 filter may be added in the above list if the customer requires the full routing table but not the default route.		X	
Security Recommendation 32: Inbound prefix filtering (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 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 33: Outbound prefix filtering (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) to observe extra caution.	X		
Security Recommendation 34: The ROA data (available from RPKI registries) should be used to construct and/or augment prefix filter lists for customer interfaces.		X	

Route Leak Mitigation:			
Security Recommendation 35: An AS operator should have 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 36: An AS operator should have egress policy to utilize the tagged information (in the preceding Security Recommendation) to prevent route leaks when routes are forwarded on the egress.	X	X	
DDoS Mitigation (Anti-spoofing):			
Security Recommendation 37: BGP routers that have directly-connected customers with allocated address space, CMTS (or equivalent) in broadband access networks, and PDN gateways (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	
Security Recommendation 38: An enterprise border router that is multi-homed should always announce all its prefixes to each of its upstream transit providers (albeit with appropriate AS prepending for traffic engineering). It should avoid selectively announcing some prefixes to one transit ISP and other prefixes to another transit ISP.	X		
Security Recommendation 39: This is the exception case when the enterprise border router does not adhere to the above recommendation and instead selectively announces some prefixes to one upstream transit ISP and other prefixes to another upstream transit ISP. In this case, it should ensure (by appropriate internal routing) 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 40: 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, RFC 1918 prefixes, or enterprise's own prefixes).	X		

Security Recommendation 41: On customer facing interfaces, ISPs should do 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, especially in the case of multi-homed customers.		X	
Security Recommendation 42: For feasible-path uRPF to work appropriately, the ISPs (at least those near the Internet edge) should propagate all their customer routes to their upstream transit ISPs (albeit with appropriate AS prepending for traffic engineering).		X	
Security Recommendation 43: 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 the above recommendation facilitates a basis for adhering to the preceding recommendation as well. (The above recommendation is also one of the stability conditions on BGP policy for ensuring stable convergence of routing information [Gao-Rexford].)		X	
Security Recommendation 44: On interfaces with lateral (i.e., non-transit) peers, ISPs should do SAV on ingress packets by deploying the feasible-path uRPF (see Sections 5.1.3). They should avoid using strict or loose uRPF as they are not very effective for SAV on the lateral peer interfaces.		X	
Security Recommendation 45: On interfaces with transit providers, ISPs should do SAV on ingress packets by deploying 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, RFC1918 prefixes, ISP's own prefixes).		X	
Security Recommendation 46: On the egress side towards customers, lateral (i.e., non-transit) peers and transit providers, the ISP's border routers should deploy ACLs (Section 5.1.1) to drop packets when the source address is spoofed (i.e., belongs to obviously disallowed prefix blocks, RFC 1918 prefixes, ISP's internal-use only prefixes).		X	
Security Recommendation 47: The ROA data (available from RPKI registries) should be used to construct and/or augment ACLs/RPF lists for customer interfaces.		X	

Traffic Filtering (Monitoring UDP/TCP Ports with Vulnerable Applications):			
Security Recommendation 48: Port 0 is a reserved port. Hence, deny TCP/UDP traffic on Port 0 on all interfaces.	X	X	X
Security Recommendation 49: In BGP routers, allow peers to connect to only port 179. The standard port for receiving BGP session OPEN messages is port 179, so attempts by BGP peers to reach other ports are likely to indicate faulty configuration or potential malicious activity.	X	X	
Security Recommendation 50: Disable applications or services that are unwanted in the network or system in consideration.			X
Security Recommendation 51: Deny traffic for any TCP/UDP ports for which the network or system in consideration does not support the corresponding applications. In some cases, an application or service is supported on some interfaces (e.g., customer or internal facing interfaces) but not others (e.g., Internet facing interfaces). In such cases, the traffic with port ID specific to the application in consideration should be denied on interfaces on which the application is not supported.			X
Security Recommendation 52: This recommendation is aimed at detection of traffic overload and mitigating actions. The relevant mitigation techniques are (a) Response Rate Limiting (RRL) [ISC1] [Redbarn], and (b) Source-based Remote Triggered Black Hole (S/RTBH) filtering enabled with Flowspec [RFC5575] (see Section 5.5 for details). These techniques are applicable to open services/protocols such as those listed in Table 1 which are themselves vulnerable to DoS/DDoS attacks or may be exploited for reflection/amplification. The recommendation consists of multiple steps as follow [TA14-017A]: • Monitor the rate of queries/requests per source address and detect if abnormally high volume of responses is headed to the same destination (i.e., same IP address). • Apply the Response Rate Limiting (RRL) technique to mitigate the attack. • Using BGP messaging (Flowspec), create a Remotely Triggered Black Hole (RTBH) filter. This can be coordinated with the upstream ISP.			X

• Maintain emergency contact information for the upstream provider to coordinate response to the attack. An upstream ISP should actively coordinate response with downstream customers.			
Security Recommendation 53: Deny NTP monlist request traffic (by disabling the monlist command) altogether, or at least enforce that the requests come from valid (permitted) source addresses.			X
Security Recommendation 54: To limit exploitation, a DNS recursive resolver should limit the scope of clients from which it accepts requests. The clients normally come from within the same network where the DNS resolver resides. Hence, the DNS resolver can maintain access lists in the configuration so that the recursive resolver is not open to the entire network (or Internet) [ISOC] [TA14-017A].			X
Security Recommendation 55: Deny all traffic with a source or destination address that matches a DNS anycast address. An exception should be made for internal recursive resolvers that are used to do outbound recursion.			X
Security Recommendation 56: Block all inbound/outbound Port 53 UDP messages at DNS recursive resolvers except those from designated recursive resolvers.			X
DDoS Mitigation (Remote Triggered Black Hole filtering, Flow specification):			
Security Recommendation 57: Edge routers should be equipped to perform Destination-based Remote Triggered Black Hole (D/RTBH) filtering and Source-based Remote Triggered Black Hole (S/RTBH) filtering.	X	X	
Security Recommendation 58: Edge routers should be equipped to make use of BGP flow specification (Flowspec) to facilitate DoS/DDoS mitigation (in coordination between upstream and downstream autonomous systems).	X	X	
Security Recommendation 59: Edge routers – in an AS providing RTBH filtering – should have ingress policy towards RTBH customers to accept routes more specific than /24 in IPv4 and more specific than /64 in IPv6. Also, the edge routers should accept such more specific route (in case of D/RTBH) only if it is subsumed by a less specific route that the customer is authorized		X	

to announce as standard policy (e.g., has a ROA for the less specific route). 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.			
Security Recommendation 60: 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 61: An ISP providing RTBH filtering service to customers must have egress policy that denies routes that have community tagging meant for triggering RTBH filtering. This is an additional safeguard in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails to work for some reason.		X	
Security Recommendation 62: An ISP providing RTBH filtering service to customers must have egress policy that denies prefixes that are longer than expected. This provides added safety in case NO_EXPORT, NO_ADVERTISE, or similar tagging fails to work for some reason.		X	

1312

Appendix B— Acronyms

Selected acronyms and abbreviations used in this paper are defined below.

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
BGP-OV	BGP Origin Validation
BGP-PV	BGP Path Validation
BGPsec	Broder Gateway Protocol with Security Extensions
DA	Destination Address
DSCP	Differentiated Services Code Point
DHS	Department of Homeland Security
DoS	Denial of Service
DDoS	Distributed Denial of Service
DNS	Domain Name System
DNSSEC	Domain Name System Security Extensions
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
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
ROA	Route Origin Authorization
RPKI	Resource Public Key Infrastructure
RPKI-to-router protocol	RPKI cache to router protocol
RLP	Route Leak Protection
RRDP	RPKI Repository Delta Protocol
RTBH	Remotely Triggered Black-Holing

D/RTBH	Destination-based Remotely Triggered Black-Holing
S/RTBH	Source-based 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)
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

1315

1316 ## Appendix C— References

[Acunetix] "Prevention of NTP Reflection DDoS attacks based on CVE-2013-5211," Acunetix blog, September 2014. http://www.acunetix.com/blog/articles/ntp-reflection-ddos-attacks/

[Adalier1] M. Adalier, K. Sriram, O. Borchert, K. Lee, and D. Montgomery, "High Performance BGP Security: Algorithms and Architectures", North American Network Operators Group (NANOG69), Washington D.C, February 2017. https://nanog.org/meetings/abstract?id=3043

[Adalier2] M. Adalier, "Efficient and Secure Elliptic Curve Cryptography Implementation of Curve P-256," NIST Workshop on ECC Standards, June 2015. 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. https://blog.apnic.net/2018/11/09/myapnic-rpki-service-now-supports-as0-roa-creation/

[Arbor] "Worldwide Infrastructure Security Report," Vol. XI, Arbor Networks report (2016). https://www.arbornetworks.com/images/documents/WISR2016_EN_Web.pdf

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

[ARIN2] M. Kosters, "Securing Core Internet Functions – Resource Certification, RPKI," Presentation by ARIN at NANOG on The Road (NOTR), September 2018. http://www.cvent.com/events/notr-washington-dc/custom-17-2bd749ab3c1f46de9be85e47c942fd5d.aspx

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

[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. https://tools.ietf.org/html/bcp38

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

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

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

[Bjarnason] S. Bjarnason, "Withstanding the Infinite: DDoS Defense in the Terabit Era," Presentation at NANOG-74, October 2018. 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. https://www.commerce.gov/sites/default/files/2018-11/Botnet%20Road%20Map%20112918%20for%20posting_0.pdf

[Cisco1] "BGP—Origin AS Validation," 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

[Cisco2] "Understanding Unicast Reverse Path Forwarding," Cisco blog, http://www.cisco.com/c/en/us/about/security-center/unicast-reverse-path-forwarding.html

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

[Cisco4] "Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide, Release 5.2.x – Chapter: Implementing BGP Flowspec," http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-2/routing/configuration/guide/b_routing_cg52xasr9k/b_routing_cg52xasr9k_chapter_011.html

[Comcast] "Comcast network management: Preventing Network Spoofing," March 2014, http://networkmanagement.xfinity.com/index.php/faqs-on-preventing-network-spoofing

[CVE-2013-5211] "Vulnerability summary for CVE-2013-5211," (for vulnerability related to monlist feature in NTP), National Vulnerability Database, September 27, 2016. https://nvd.nist.gov/vuln/detail/CVE-2013-5211

[CSRIC-WG5] CSRIC Working Group 5 Final Report: Remediation of Server-Based DDoS Attacks. https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG5_Remediation_of_Server-Based_DDoS_Attacks_Report_Final_(pdf)_V11.pdf

[CSRIC-WG6] "Long-Term Core Internet Protocol Improvements," Working Group 6 presentation, September 2015. https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG6_Presentation_09242014.pdf

[DOCSIS] "DOCSIS 3.0: MAC and upper layer protocols interface specification," Cable Labs publication. h http://www.cablelabs.com/wp-content/uploads/specdocs/CM-SP-MULPIv3.0-I29-151210.pdf

[EFP-uRPF] K. Sriram, D. Montgomery, and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Filtering," IETF Internet Draft, April 2018. https://datatracker.ietf.org/doc/draft-ietf-opsec-urpf-improvements/

[FISMA2002] Federal Information Security Management Act of 2002, Pub. L. 107-347 (Title III), 116 Stat. 2946. 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. http://www.gpo.gov/fdsys/pkg/PLAW-113publ283/pdf/PLAW-113publ283.pdf.

[Gao-Rexford] Freedman, M., "Interdomain Routing Policy", Princeton University COS 461 Lecture Notes; Slides 25-27, Spring 2011, 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. https://github.com/osrg/gobgp/blob/master/docs/sources/rpki.md

[Hares] S. Hares, C. Loibl, R. Raszuk, D. McPherson, and M. Bacher, "Dissemination of Flow Specification Rules," IETF I.D. draft-ietf-idr-rfc5575bis (work in progress), June 2018. https://datatracker.ietf.org/doc/draft-ietf-idr-rfc5575bis/

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

[Huston2011] G. Huston and R. Bush, "Securing BGP," The Internet Protocol Journal, Volume 14, No. 2, June 2011. http://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-52/142-bgp.html

[Huston2012] G. Huston, "Leaking Routes," Asia Pacific Network Information Centre (APNIC) Blog, March 2012, 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. http://www.circleid.com/posts/20161026_closer_look_at_recent_ddos_attacks_and_what_it_means_for_dns/

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

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

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

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

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

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

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

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

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

[ISC1] "A Quick Introduction to Response Rate Limiting," ISC Knowledge Base blog. https://kb.isc.org/article/AA-01000/0/A-Quick-Introduction-to-Response-Rate-Limiting.html

[ISC2] "A Chargen-base DDoS? Chargen still a thing?" ISC blog, https://isc.sans.edu/forums/diary/A+Chargenbased+DDoS+Chargen+is+still+a+thing/15647

[ISOC] P. Vixie (Ed.), "Addressing the challenge of IP spoofing," ISOC report, September 2015. 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. 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. 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. https://www.symantec.com/content/dam/symantec/docs/reports/istr-22-2017-en.pdf

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

[Juniper2] "Configuring Unicast RPF," Juniper blog, https://www.juniper.net/documentation/en_US/junos14.2/topics/usage-guidelines/interfaces-configuring-unicast-rpf.html

[Juniper3] "Example: Configuring Unicast Reverse-Path-Forwarding Check," Juniper blog, http://www.juniper.net/documentation/en_US/junos15.1/topics/topic-map/unicast-rpf.html

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

[Kaeo] M. Kaeo, "Routing Security, DDoS and Route Hijacks," NANOG on The Road (NOTR), September 2018. http://www.cvent.com/events/notr-washington-dc/custom-17-2bd749ab3c1f46de9be85e47c942fd5d.aspx

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

[Levy] M. Levy, "RPKI - The required cryptographic upgrade to BGP routing," Cloudflare blog, September 2018. https://blog.cloudflare.com/rpki/

[MANRS] "Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide," Published by the Internet Society (ISOC), retrieved October 2018. https://www.manrs.org/isps/guide/

[maxlength] Y. Gilad, S. Goldberg, K. Sriram, J. Snijders, and B. Maddison, "The use of maxlength in the RPKI," IETF Internet Draft, April 2018. https://tools.ietf.org/html/draft-ietf-sidrops-rpkimaxlen-00

[Merit-RADb] "Merit RADb" (Merit Network Inc.) 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. 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. 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 (NORT), St. Louis, MO, November, 2015 https://www.nanog.org/sites/default/files/04-Murphy-StLouis.pdf

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

[Naik] A. Naik, "Internet Vulnerability Takes Down Google," ThousandEyes report, November 2018. 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, NANOG 67, Chicago, IL, June 2016. https://www.nanog.org/meetings/abstract?id=2846

[NANOG-list] "Intra-AS messaging for route leak prevention," NANOG Email List - Discussion Thread, June 2016. 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. https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14b-draft.pdf

[NIST2018] 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. https://csrc.nist.gov/publications/detail/white-paper/2018/05/30/enhancing-resilience-against-botnets--report-to-the-president/final

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

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

[NIST-SRx] BGP Secure Routing Extension (BGP-SRx): Open source Origin Validation and BGPsec Path Validation implementations in Quagga. https://www-x.antd.nist.gov/bgpsrx/

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

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

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

[Parsons1] "Secure Your Routing Infrastructure," Parsons blog. http://www.securerouting.net/

[Parsons2] Open source Origin Validation and BGPsec Path Validation implementations in BIRD, Parsons blog. http://www.securerouting.net/tools/bird/

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

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

[Redbarn] "Response Rate Limiting in the Domain Name System (DNS RRL)," Redbarn blog. http://www.redbarn.org/dns/ratelimits

[RFC3882] D. Turk, "Configuring BGP to Block Denial-of-Service Attacks," IETF RFC 3882, September 2004. 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. https://tools.ietf.org/html/rfc4012

[RFC4271] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF RFC 4271, January 2006.

[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. http://www.ietf.org/rfc/rfc5280.txt.

[RFC5575] P. Marques et al., "Dissemination of Flow Specification Rules," IETF RFC 5575, August 2009. https://tools.ietf.org/html/rfc5575

[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. https://tools.ietf.org/html/rfc5635

[RFC6092] J. Woodyatt, "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service," IETF RFC 6092, January 2011. https://tools.ietf.org/html/rfc6092

[RFC6472] W. Kumari and K. Sriram, "Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP," BCP 172 (RFC 6472), December 2011. https://tools.ietf.org/html/rfc6472

[RFC6480] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," RFC6480, February 2012. 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. https://tools.ietf.org/html/rfc6481

[RFC6482] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012. https://tools.ietf.org/html/rfc6482

[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. 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. 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. https://tools.ietf.org/html/rfc6492

[RFC6810] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol," RFC 6810, January 2013. 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. 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. 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. 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. https://tools.ietf.org/html/rfc7382

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

[RFC7674] J. Haas, "Clarification of the Flowspec Redirect Extended Community," IETF RFC 7674, October 2015. https://tools.ietf.org/html/rfc7674

[RFC7908] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, June 2016. 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. 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. https://tools.ietf.org/html/rfc7935

[RFC7999] T. King, et al., "BLACKHOLE Community," IETF RFC 7999, October 2016. 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. https://tools.ietf.org/html/rfc8182

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

[RFC8208] S. Turner and O. Borchert, "BGPsec Algorithms, Key Formats, & Signature Formats," IETF RFC 8208, September 2017. https://tools.ietf.org/html/rfc8208

[RFC8210] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," IETF RFC 8210, September 2017. https://tools.ietf.org/html/rfc8210

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

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

[RIPE2] RIPE NCC RPKI Validator, 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, 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. 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. 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, July 2018. https://datatracker.ietf.org/doc/draft-ietf-idr-route-leak-detection-mitigation/

[RouteLeak2] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and K. Sriram, "Route Leak Prevention using Roles in Update and Open Messages", IETF Internet Draft, June 2018. https://datatracker.ietf.org/doc/draft-ietf-idr-bgp-open-policy/

[RouteLeak3] K. Sriram (Ed.), "Design Discussion of Route Leaks Solution Methods", IETF Internet Draft, July 2018. https://datatracker.ietf.org/doc/draft-sriram-idr-route-leak-solution-discussion/

[Rsync] Wiki page on the Rsync protocol. 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. https://www.ietf.org/proceedings/86/slides/slides-86-sidr-1.pdf

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

[Ryburn] J. Ryburn, "DDoS Mitigation," NANOG-63, February 2015. https://www.nanog.org/meetings/abstract?id=2487

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

[SP800-53] Joint Task Force Transformation Initiative, "Security and Privacy Controls for Federal Information Systems and Organizations," (National Institute of Standards and Technology, Gaithersburg, MD) NIST Special Publication (SP) 800-53 Revision 4, April 2013 (includes updates as of 01-22-2015). https://doi.org/10.6028/NIST.SP.800-53r4

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

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

[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. 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

[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. http://www.symantec.com/connect/blogs/denial-service-attacks-short-strong

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

[Winward] R. Winward, "Mirai – Inside of an IoT Botnet," NANOG-69, February 2017. 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. https://pc.nanog.org/static/published/meetings//NANOG74/daily/day_2.html#talk_1767

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

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

[TCP-UDP-port] "List of TCP and UDP ports," https://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers .

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

[Cymru-UTRS] Unwanted traffic removal service (UTRS), Team Cymru blog, http://www.team-cymru.com/utrs.html

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

[Toonk-B] Toonk, A., "Massive route leak causes Internet slowdown", BGPMON Blog, June 2015, 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. 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. http://www.verisign.com/en_US/security-services/ddos-protection/ddos-report/index.xhtml

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

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

1317