SCTE·ISBE CABLE-TEC EXPO 2019 NEW ORLEANS, LA SEPT. 30-OCT. 3

2019 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS

Securing Interdomain Network Routing with Resource Public Key Infrastructure (RPKI)

A Technical Paper prepared for SCTE•ISBE by

Mark Goodwin
IP Design Engineer
Cox Communications, Inc
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328
404-269-8267
Mark.Goodwin@cox.com

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Table of Contents

List of Figures

© 2019 SCTE•ISBE and NCTA. All rights reserved.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

1. Introduction

In 2018, 1,300 IP addresses were hijacked from Amazon Web Services (Route 53). This malicious attack resulted in service disruption for about two hours and theft of approximately $150,000 in cryptocurrency. Further (at no fault of Amazon), this attack exposed both ISP peers and customers to fraudulent routes leaving them susceptible to attacks. The root cause of this incident—and hundreds of others alike—was the lack of security in Border Gateway Protocol (BGP), the protocol used for Interdomain Network Routing.

As BGP announces IP reachability information between domains, there is no way to validate the ownership of the IP information. This vulnerability, which arises from RFC 4272, creates opportunities for inadvertent advertisements and malicious theft of IP resources. Thus it potentially impacts network services and stability. As the Internet of Things (IoT) continues to penetrate customer devices and increases reliability expectations on the ISP networks, ISPs have a responsibility to deploy industry best security practices to protect the IP network for customers.

This document introduces Resource Public Key Infrastructure (RPKI) with BGP Route Origin Validation (ROV) to mitigate the security issues of BGP associated with origin attacks. RPKI is an out-of-band security infrastructure that uses public key cryptography to validate ownership of IP resources for a given Autonomous System (AS). This paper first details the existing vulnerabilities with BGP. Second, it shows how RPKI—as well as deployment of BGP ROV—mitigate BGP hijacks, and route leaks. Third, it outlines a deployment strategy for BGP RPKI. Finally, this paper demonstrates how the attacker power to hijack prefixes decreases after top ISPs adopt RPKI.

2. Motivations

2.1. BGP Security Analysis

BGP details the connection and exchange of IP prefixes between thousands of Autonomous Systems (ASes) that collectively make up the Internet as we know it. BGP operates on a transitive trust model, meaning, as IP prefixes are exchanged between ASes, there is a high degree of 'trust' that the IP routing information is correct i.e. IP prefixes originate from the correct AS or origin. Unfortunately this is not always the case. This vulnerability enables an AS to announce IP prefixes that it is does not own. This can be accidental or with malicious intent, in either case it can lead to service disruption, instability across sectors of the Internet, and loss of revenues. This fundamental vulnerability or security risk, is expressed in RFC 4272 as follows:

I. No mechanism specified within BGP to validate the authority of an AS to announce Network Layer Reachability Information (NLRI). In other words, a lack of origin validation.
II. No mechanism specified within BGP to ensure the authenticity of the path attributes announced by an AS. In other words, a lack of path validation.

These vulnerabilities enable a bad actor to originate/claim an IP prefix or sub-prefix for traffic interception, eavesdropping, and manipulation. Prefix hijacking occurs when an offending AS announces the exact same prefixes as its intended victim. These prefixes are preferred if the offender has an AS path shorter than the owner. When successful, traffic is directed to a malicious domain or network. This is represented in figure 1 and illustrates the vulnerability of BGP relating to a lack of origin validation. Note the attacker, AS6, originates the 21.0.0.0/10 prefix belonging to AS3. In this example, the attacker's announcement is preferred by AS2 due to a shorter AS path, which is essential to the BGP selection process.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

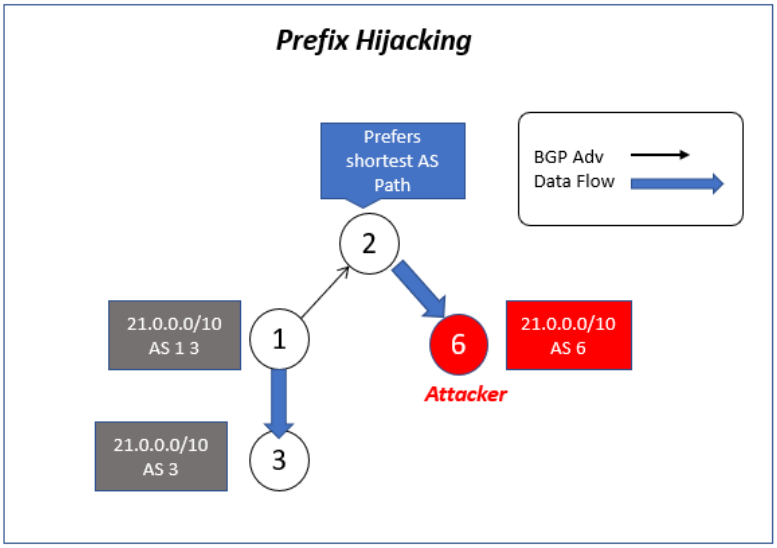

Figure 1 - Prefix Hijacking

Sub-prefix hijacking is another form of prefix hijacking in which the unauthorized announcement has a greater chance of success due to the BGP selection process. This hijacking occurs when an offending AS announces a more specific prefix than the potential victim. In this scenario, the attacker always wins due to the rules of the BGP route selection process. In figure 2 an attacker announces a subset (/16) of the 21.0.0.0/10 prefix and effectively hijacks traffic destined for AS 1.

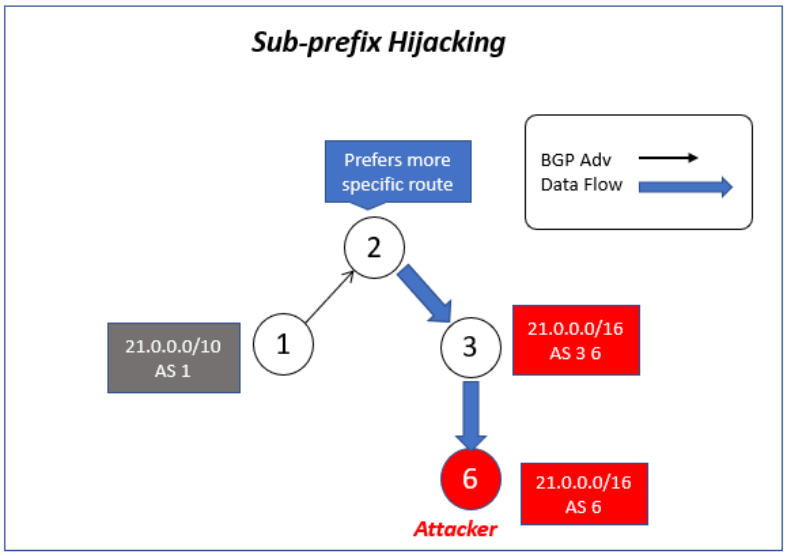

Figure 2 – Sub-Prefix Hijacking

These security risks stem from BGP's inability to validate a route's origin ASN. The Resource Public Key Infrastructure (RPKI) has become widely accepted by the Internet community as an industry standard in the mitigation of these threats.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

Man in the middle attacks (MITM) exploit the security weakness of BGP relating to lack of path validation. These are not as common as BGP hijacks, but they are seen out in the wild. With this method the attacker simply 'spoofs' the BGP packet by prepending the victim's ASN. This gives the impression that the attacker is connected to the victim and more importantly that the packet is originating from the victim's ASN. This strategy passes the RPKI check and effectively *overcomes* the protection RPKI was intended to provide. Effectively combating MITM attacks requires validating the authenticity of AS path attributes. Technologies such as BGPSec per IETF RFC 8205 are being developed to accomplish this. In the interim, to effectively detect and combat MITM attacks, operators should use an external BGP monitoring service which provides alerting capabilities in the event a more specific IP prefix is announced to the Internet. Industry tools with this capability include ThousandEyes and BGPStream. Figure 3 illustrates a MITM attack where AS 6 (attacker) announces the same prefix as its victim and prepends the victims ASN. AS4 and AS5 run the RPKI check and accepts the route because the orgin validation 'appears' to be correct.

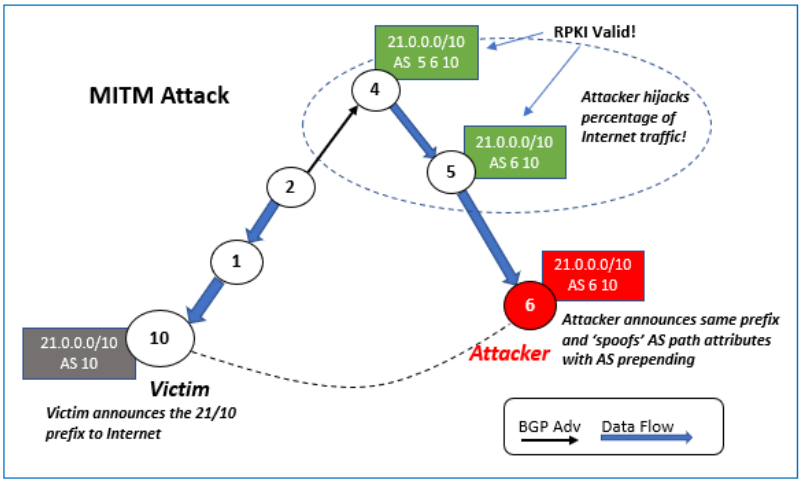

Figure 3 - MITM Attack

The next section presents RPKI components and their use in the mitigation of BGP hijacking relating to lack of origin validation.

3. RPKI Components

RPKI is an out of band framework designed to enhance security for the BGP interdomain routed network. It provides AS number to IP prefix mappings which enable cryptographic validation for the origin of IP prefixes. This practice provides a network operator with a method to effectively combat BGP origin attacks. This section details the RPKI components and their dependencies.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

3.1. Certificate Authority (CA)

The foundation of the RPKI system is the certificate authority (CA). It is the authoritative owner for all IP number resources and holds a root digital certificate to represent these. The CA issues other certificates to network operators and in doing so they convey the authority to use specified IP resources. The certificates contains digital signatures which can be cryptographically verified back to the issuer. In this manner a chain of trust is created from the issuer to the owner of IP resources. In the context of IP address allocation the regional internet registries (RIRs) naturally serve the role of a CA. As RIRs delegate IP resourses to network operators they can also issue a corresponding digital certificate representing ownership of these resources. The RIRs in the role of the CA will create, issue, and store digital certificates. They also act as a trusted 3rd party or trust anchor (TA) between the owner of certificates and parties who rely on the certificates for origin validation. Figure 4 illustrates the hierarchy for distribution of certificates. The RIRs sit at the top of the hierarchy and allocate certificates to network operators. The following sections cover digital certificates and the different forms they assume as they traverse the RPKI system.

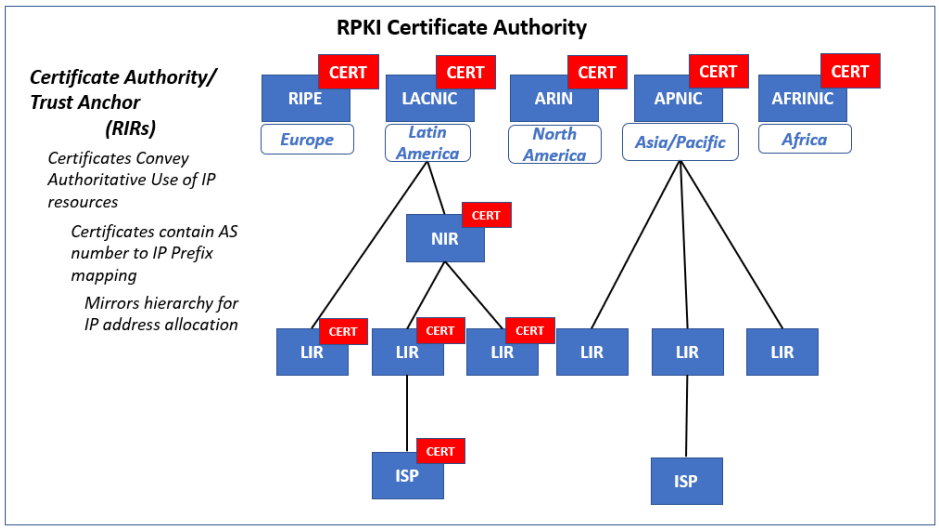

Figure 4 - RPKI Certificate Structure

3.2. Resource Certificate

The resource certificate (RC) within the RPKI system is created by the RIRs and issued to network operators. It is a X.509 certificate with important extensions for IP number resources (IPv4/v6 addresses, AS number). It includes AS number to IP prefix mappings and authoritatively certifies a network operator's ownership of IP prefixes. Distribution of RCs naturally mirror that of IP allocation from the RIRs. For a given IP assignment from an RIR there will be a corresponding resource certificate. The resource certificate is signed with a digital private key by the issuing RIR. Further, it binds the assigned IP prefixes to the a public key provided by the network operator. This essentially conveys ownership of the allotted IP resources to the network operator. Figure 5, noted below shows the high-level format for an RC.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

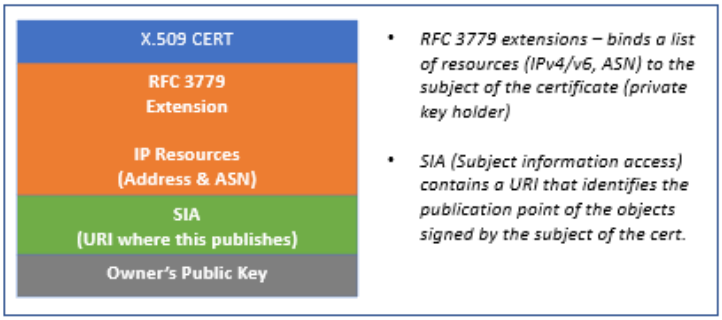

Figure 5 – Resource Certificate

3.3. Route Origin Authorizations (ROAs)

Network operators use resource certificates to create Route Origin Authorizations (ROAs). Here the network operator specifies which AS is authorized to originate its IPv4/v6 prefixes. The ASN can be that of the network operator or its customers. Specifically, within the ROA the network operator must specify the IP prefix, the longest prefix match, and the AS number to originate the stated IP prefixes. The network operator must then digitally sign the ROA with its digital private key. This action further conveys authority to the specified AS to originate prefixes. The ROAs are then distributed to the RPKI repository and now allow for cryptographic verification by referencing the public and private keys in the resource certificate and the ROA respectively. Per RFC 6482 the ROA is considered valid if the following conditions are met:

- Its corresponding **resource certificate** is valid
- The IP address prefixes of the ROA match the IP address prefixes in the resource certificate

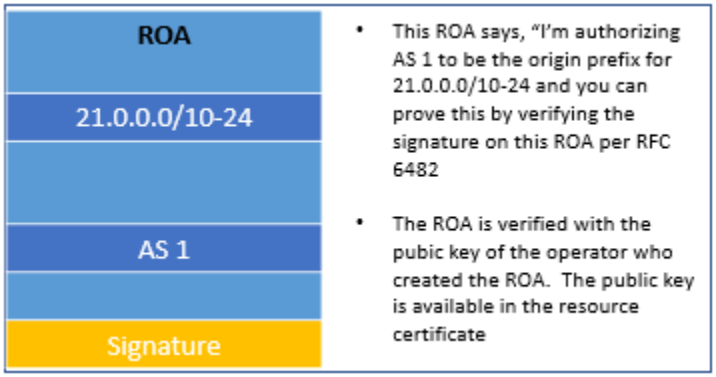

Figure 6 – Route Origin Authorization (ROA)

3.4. RPKI Validating Cache

The RPKI validating cache (VC) is a software entity used to retrieve resource certificates and ROAs from each of the five RIRs. The VC then uses mathematical computations in the decryption and verification of ROAs. This verification process spans the RPKI chain of trust. This feature of the VC relaxes the requirement of the router for any cryptographic operations. The RPKI VC produces a prefix-to-AS mapping database in the form of a validated ROA payload (VRP), which is consumable by the router. The VRP contains the prefix, maximum length, and route origin (ASN).

© 2019 SCTE•ISBE and NCTA. All rights reserved.

3.5. Relying Party (RP)

The relying party (RP) in the case of RPKI is represented by network operators who *rely* upon the outcome of the validating cache process (*the VRPs*) for origin validation. The VRPs serve as a representation of the certificate authority for the RPKI system and effectively provides proof for the origin of IP prefixes. The network operator consume the VRPs using edge routers or peering points in their networks. These edge routers utilize the **RPKI to router protocol** as detailed by RFC 6810 to connect to the validating cache. Figure 6 details the interaction of the router and validating cache utilizing the RPKI to router protocol. For initial startup, the router sends the cache a request or reset query. The cache responds by sending all its data records or VRPs followed by a end of data PDU. From here, the router must periodically query the validating cache for updated records.

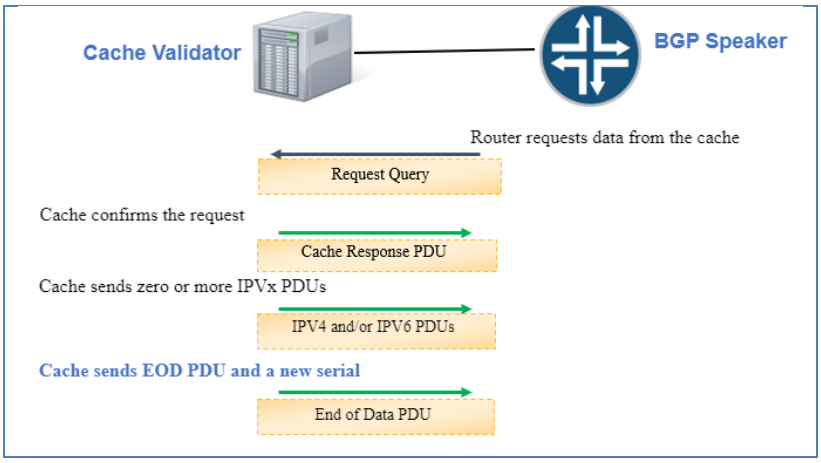

Figure 7 - RPKI to Router startup exchange. CiscoLive 2014 BRKST-2446

The edge routers then use the **BGP route origin validation** (ROV) per RFC 6811 to ingest and store the VRPs in a route validation database. This database of validated routes can now be compared to against incoming routes from each external BGP session. Per RFC 6811, incoming routes will have on of the following validation states:

- **Notfound**: No VRP entry covers the route prefix.
- **Valid**: At least one VRP matches the route prefix.
- **Invalid**: At least one VRP Covers the route prefix, but no VRP matches it.

Note: a 'match' refers to a route prefix that is covered by a VRP i.e its prefix length is less than or equal to the VRP maximum length, and the route origin ASN is equal to the VRP ASN.

As illustrated in figure 8, the first criteria for BGP selection will now be origin validation using a inbound filter that matches on the validation database. From there, the BGP selection process will run as usual and distribute the routing information to the iBGP network. Options are available to apply origin validation extended communities.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

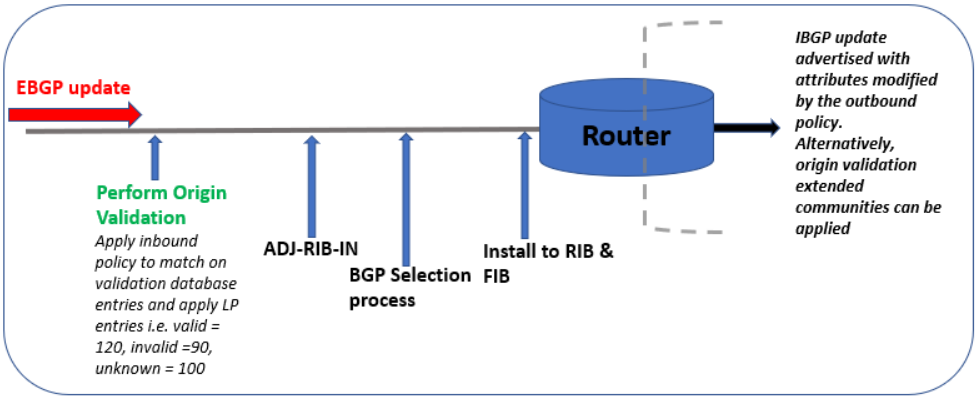

Figure 8 - BGP ROV processing

4. RPKI Architecture

4.1. Overview

This section illustrates a high-level architecture for RPKI with the requirements and functions for each component.

4.2. Component Communication

Think of the RPKI architecture as two separate architectures. One architecture allows for the creation, signage, and distribution of resource certificates and ROAs; it is completely hosted in software. This is represented in figure 9 by the left side of the drawing with the RPKI repository, resource certificates, and ROAs. The functions for each of these components are offered as a service from the RIRs or they can be developed internally by network operators.

The next architecture is partially hosted in software and is represented in figure 9 on the right side of the drawing. The function of this architecture include decryption and filtering. The decryption functionality is implemented in the validating cache, and filtering is implemented in routing hardware with BGP ROV processing and route policy.

As illustrated in figure 9, the RPKI validator utilizes the rsync protocol to fetch resource certificates and ROAs from the RPKI repositories. The validator then cryptographically validates each ROA from the root of the RPKI hierarchy to the leaves. For the output, it generates a whitelist of all valid, invalid, and unknown ROAs in the form of VRPs. This RPKI Cache is queried by a BGP router using the RPKI-to-router protocol. Here the router consumes the VRP data into its validation database and constructs policy to compare all incoming routes against these entries. Based on the outcome of this comparison routes can be labeled valid, invalid, or unknown.

Note in figure 9 the RPKI validating cache receives a ROA from the repository. After validating this ROA it is consumed by the network operator's router. The router receives 2 incoming routes with different origin values and is able to reject the 'milicious' route based on comparison against a validated entry.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

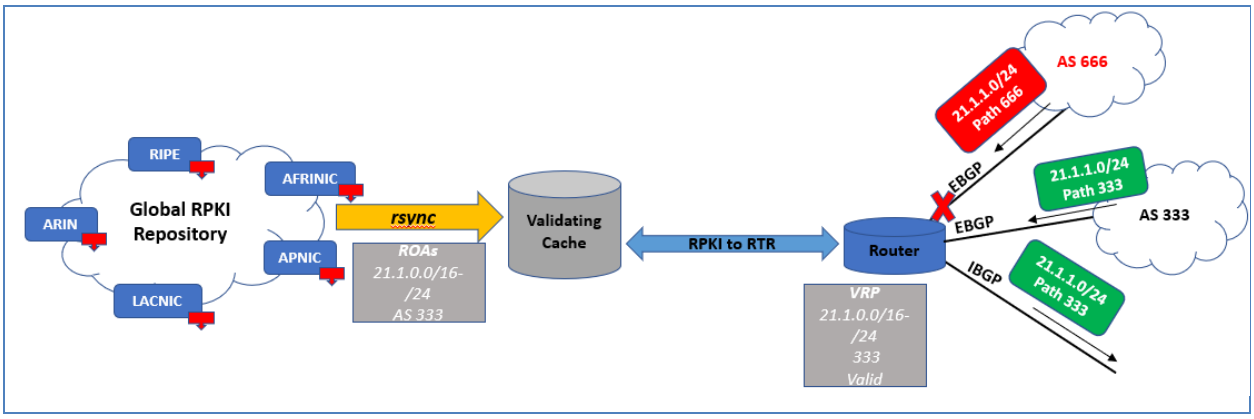

Figure 9 - RPKI Component Communication

4.3. General Requirements

Network operators can access the RPKI repository using a web base graphical user interface (GUI). The appropriate credentials associated with the company's authorized point of contact are required to retrieve resource certificates and create ROAs. The network operator can only create ROAs for the Registry to which they have been assigned IP prefixes.

4.4. Software Requirements

Implementation for the validating cache will require the following:

- a Unix like OS
- Java 8
- 4 GB of memory
- rsync protocol

Edge routers within the network must run software supportive of the BGP route origin validation.

4.5. Hosted RPKI

Each of the RIRs offer hosted RPKI as a service. The hosted service is straight forward and offers ease of use. With a hosted model, the Certification Authority (CA) and the publication components are built and maintained by each RIR. This removes the cost and operational functions such as coding, key storage, maintenance, and publication from operators looking to utilize RPKI. With the hosted model operators simply 'click around' the RIR GUI to receive RCs and issue ROAs. Hosted RPKI requires the operator IP resources be issued from the RIR providing the service.

4.6. Delegated RPKI

The deletated RPKI model offers operators the flexiblility to host their own CA to issue RCs. This requires operators to build their own CA and RPKI repository to address the functions of maintenance, encryption and decryption of ROAs. This model carries more operational complexity and overhead as operators are now required to manage software and operational functions, such as uptime. This model does, however, carry the benefit of better integration with the operator's system. The decision to implement hosted versus the deletated RPKI system is largely operator dependent and the Network Architect is responsible for understanding the pros and cons of each model.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

4.7. High Availability

The operator network is highly dependent upon the RIRs for the distribution and publication of RCs and ROAs. If the RPKI system is unavailable, the operator IP resources simply reverts to a status of unknown. Operators are encouraged to deploy redundant validating caches with appropriate priorities between the two. This prevents a single point of failure, thus ensuring stability in the event one VC is lost due to power or other operational issues.

4.8. Security

The RPKI validating cache should be housed in secure environments to prevent and defend against any malicious attempts at hacking the RPKI system.

5. RPKI Adoption Considerations

Network operators have a responsibility to both customers and peers to provide reliable routing information and to effectively filter incorrect and malicious routes. Any case where incorrect routing information is propagated can result in service disruption and in some cases loss of revenues. This is of key concern to large network operators as they could expose millions of customers to threats. The following sections references a leading industry study on RPKI deployment and security.[1] It considers the benefits and consequences of RPKI adoption with BGP ROV.

5.1. Collateral Benefit

In adopting BGP RPKI, not only are threats mitigated for a given operator's network, but also for all customer networks sitting behind it. Per the simple network depicted in figure 10, note AS 2 has adopted BGP RPKI with ROV. In doing so it also protects the AS 3 network sitting behind it. Per the illustration, AS 2 detects that the prefix 1.1.1/24 originating from AS 666 as invalid because neither its origin nor prefix length matches its validation database entry for 1.1/16. AS 2 BGP ROV implementation is to drop the prefix. This represents a large ISP's capability to protect thousands of customer networks when implementing RPKI with BGP ROV.

[1] Are We There Yet? On RPKI's Deployment and Security, https://eprint.iacr.org/2016/1010.pdf. Yossi Gilad et al.

© 2019 SCTE•ISBE and NCTA. All rights reserved.

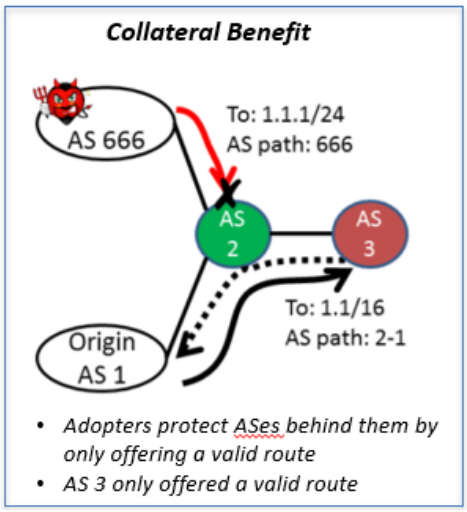

Figure 10 - Collateral Benefit. Source: Yossi Gilad et al

5.2. Collateral Damage

In contrast, collateral damage occurs by not adopting RPKI with BGP ROV. Essentially networks that don't adopt RPKI are subject to propagate faulty routing information, thus exposing their customer base to BGP hijacks. This is the case even when networks behind non-adopting RPKI networks implement BGP RPKI. Consider the example below. Even though AS 3 enforces ROV, because its provider network does not, it can still fall prey to an attacker as illustrated in figure 11. AS 3 implements RPKI with BGP ROV and drops the attacker's advertisement of 1.1.1/24. AS 3, however, accepts the legitimate advertisement of 1.1/16 from AS 2 1 and proceeds to forward traffic. The problem occurs as the traffic traverses it's upstream provider, AS 2. AS 2 doesn't implement BGP ROV and accepts the faulty route from the attacker. This causes AS 2 to forward traffic from AS 3 to the faulty destination of AS 666 for the traffic that falls within the subset of the /24. This represents a small network operator's inability to combat BGP hijacking when their upstream providers do not implement RPKI with BGP ROV.

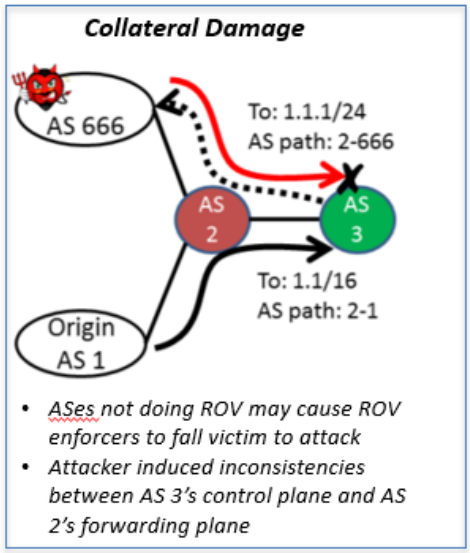

Figure 11 - Collateral Damage. Source: Gilad et al

© 2019 SCTE•ISBE and NCTA. All rights reserved.

5.3. Impact to Global Internet Routing Security

The study completed by Gilad et al. concludes that RPKI with BGP ROV is most effective once adopted by the 'core of the Internet' or the top 100 ISPs. This is essentially the ISPs with the largest customer bases (ASNs). In adoption of RPKI with BGP ROV, the top 100 ISPs provide the extreme collateral benefit in the reduction of BGP hijacks for the Internet at large. Figure 12 below shows that as a larger percentage of the top 100 ISPs adopt RPKI with BGP ROV the attacker power is diminished for prefix and subprefix hijacking respectively. For example, at 50% adoption rate of BGP ROV, the attacker success rate is redcuced by 20% (from 50% to 30%) for prefix hijacking. It is reduced by 10% (from 100% to 90%) for subprefix hijacking. Once the adoption rate is at 100%, the attacker success rate is reduced to about 5% for prefix hijacking and about 25% for subprefix hijacking.

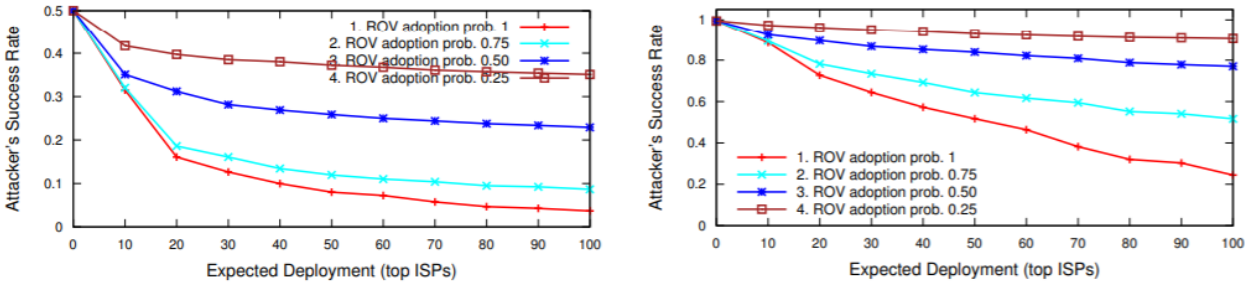

Figure 12 - Impact for Core Internet's adoption of ROV - Gilad et al

The Gilad el al study also shows, however, that RPKI with BGP ROV is least effective without adoption of the top ISPs. This holds true even when a significantly large number of small ISPs adopt RPKI with BGP ROV. Here the outcome of extensive simulations show that small network operators who adopt RPKI with BGP ROV gain minimal benefit over that gained by only the top 100 ISPs adoption. Put another way the security of small network operators is largely dependent on the core of the Internet's adoption of BGP ROV. Figure 13 below illustrates how the attacker's success rate is dependent the collateral benefit or damage provided by the top 100 ISPs adoption of RPKI with BGP ROV.

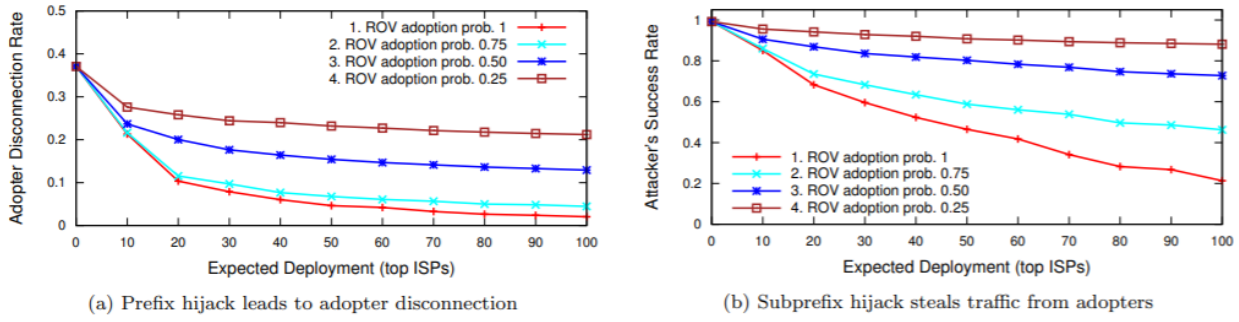

Figure 13 - Gilad et al - Collateral damage. Without adoption of top ISPs, smaller ISPs with adopt ROV are doomed

© 2019 SCTE•ISBE and NCTA. All rights reserved.

5.4. Current RPKI and BGP ROV adoption status

The National Institute of Standards and Technology (NIST) organization currently reports very low global adoption status for RPKI with BGP ROV. Per the illustration below, for IPv4 address space, of the 836,476 unique prefix/origin pairs, only around 15.6% show registry within the RPKI system. The unknown entries largely indicate around 85% of the IPv4 IP space is unregistered with RPKI.

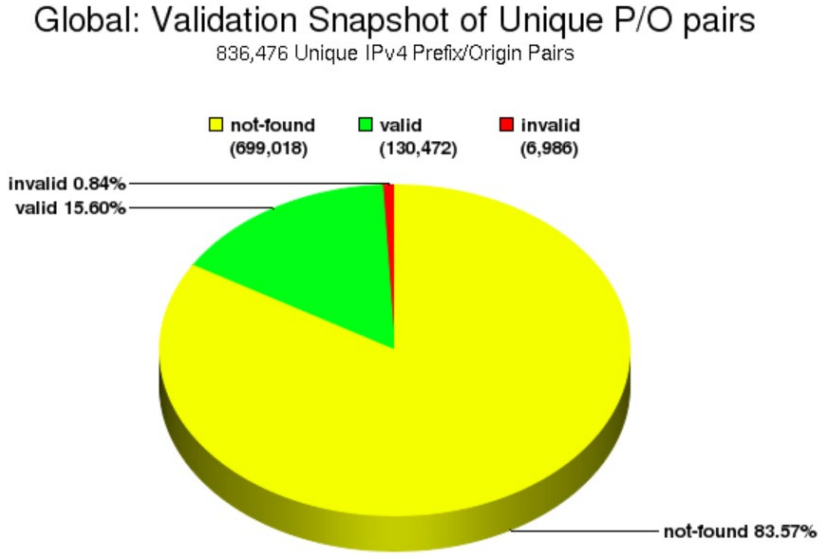

Figure 14 - NIST RPKI Monitor 8-20-19 rpki-monitor.antd.nist.gov

The North American adoption of RPKI as represented by the American Registry of Internet Numbers (ARIN) is slightly lower than the global adoption status. Of the 302,056 unique prefix to origin pairs for IPv4, NIST shows a 7.04% adoption rate with a roughly 93% nonadoption rate.

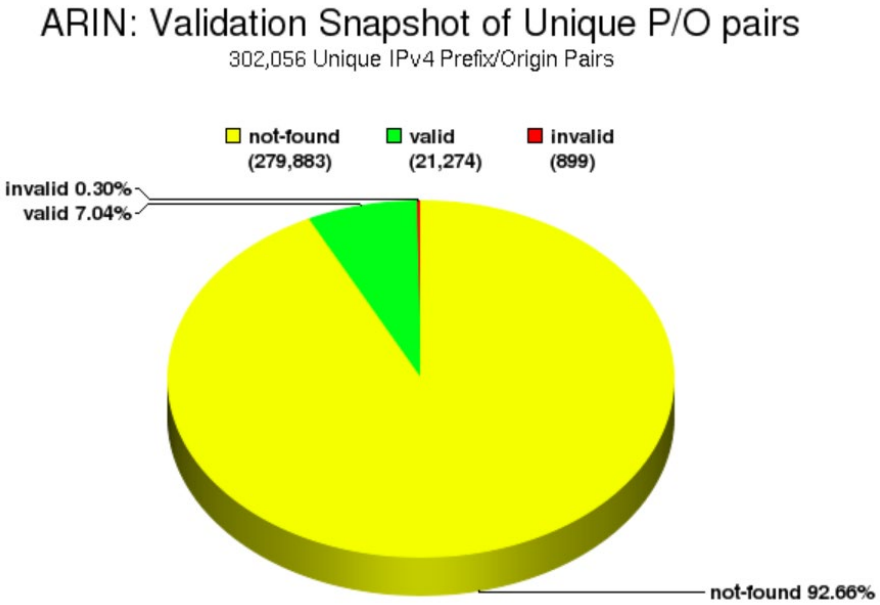

Figure 15 - rpki-monitor.antd.nist.gov

© 2019 SCTE•ISBE and NCTA. All rights reserved.

6. Conclusion

RPKI with BGP ROV is an additional security layer for the Internet infrastructure. It's main purpose is to defend against BGP hijacking which has been the root cause of service disruption for multiple operators around the global Internet. BGP hijacking occurs when IP resources (prefixes) are advertised by an operator who is unauthorized to do so. This enables the unauthorized operator to divert traffic from its intended destination thereby causing service disruption. RPKI with BGP ROV combats this by providing cryptographic validation of IP prefixes. RPKI presents a 'proof' model for the origin of IP prefixes to interdomain routing and can significantly reduce the occurance of incidential and intentional BGP hijacking. For RPKI to be effective, it is must be implemented by the top 100 ISPs who encompass the core of the Internet. Otherwise no substantial security benefits will be achieved.

Abbreviations

RPKI	resource public key infrastructure
CA	Certificate Authority
RC	Resource certificate
ROV	Route Origin Validation
AS	Autonomous System
RIR	Regional Internet Registry
ARIN	American Registry for Internet Numbers
ROA	Route origion authorization
RP	Relying party
VRP	Validated ROA payload
ISP	Internet Service Provider
IP	Internet Protocol
BGP	Border Gateway Protocol
RFC	Request For Comment
IoT	Internet of Things
NLRI	Network Layer Reachability Information
MITM	Man in the Middle
EE	End Entity
VC	Validating Cache
GUI	Graphical User Interface

Bibliography & References

Are We There Yet? On RPKI's Deployment and Security, https://eprint.iacr.org/2016/1010.pdf. Yossi Gilad et al.
RPKI Deployment, https://academy.apnic.net/wp-content/uploads/2019/04/slides_rpki_deployment.pdf. APNIC
RPKI, https://www.arin.net/resources/manage/rpki/ arin.net
BGP Security Vulnerabilities Analysis, RFC 4272, https://www.ietf.org/rfc/rfc4272.txt
An Infrastructure to Support Secure Internet Routing, RFC 6480, https://tools.ietf.org/html/rfc6480
Validation of Route Origination Using RPKI and ROA, RFC 6483, https://tools.ietf.org/html/rfc6483
A Profile for the Resource Certificate Repository Structure, RFC 6481, https://tools.ietf.org/html/rfc6481

© 2019 SCTE•ISBE and NCTA. All rights reserved.

A Profile for Route Origin Authorizations (ROAs), RFC 6483, https://tools.ietf.org/html/rfc6482
RIPE network coordination centre, Using published RPKI data. https://www.ripe.net/manage-ips-and-asns/resource-management/certification/using-published-rpki-data.
The Resource Public Key Infrastructure (RPKI) to router protocol, RFC 6810, https://tools.ietf.org/html/rfc6810
https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14-draft.pdf
https://www.blackhat.com/presentations/bh-dc-09/Zmijewski/BlackHat-DC-09-Zmijewski-Defend-BGP-MITM.pdf

© 2019 SCTE•ISBE and NCTA. All rights reserved.