Let’s Fix the Internet Routing Security Problem

Aftab Siddiqui
Internet Technology Manager – APAC
Internet Society

The Problem

A Routing Security Overview

The Honor System: Routing Issues

Border Gateway Protocol (BGP) is based entirely on trust between networks

- No built-in validation that updates are legitimate
- The chain of trust spans continents
- Lack of reliable resource data

Recent Events

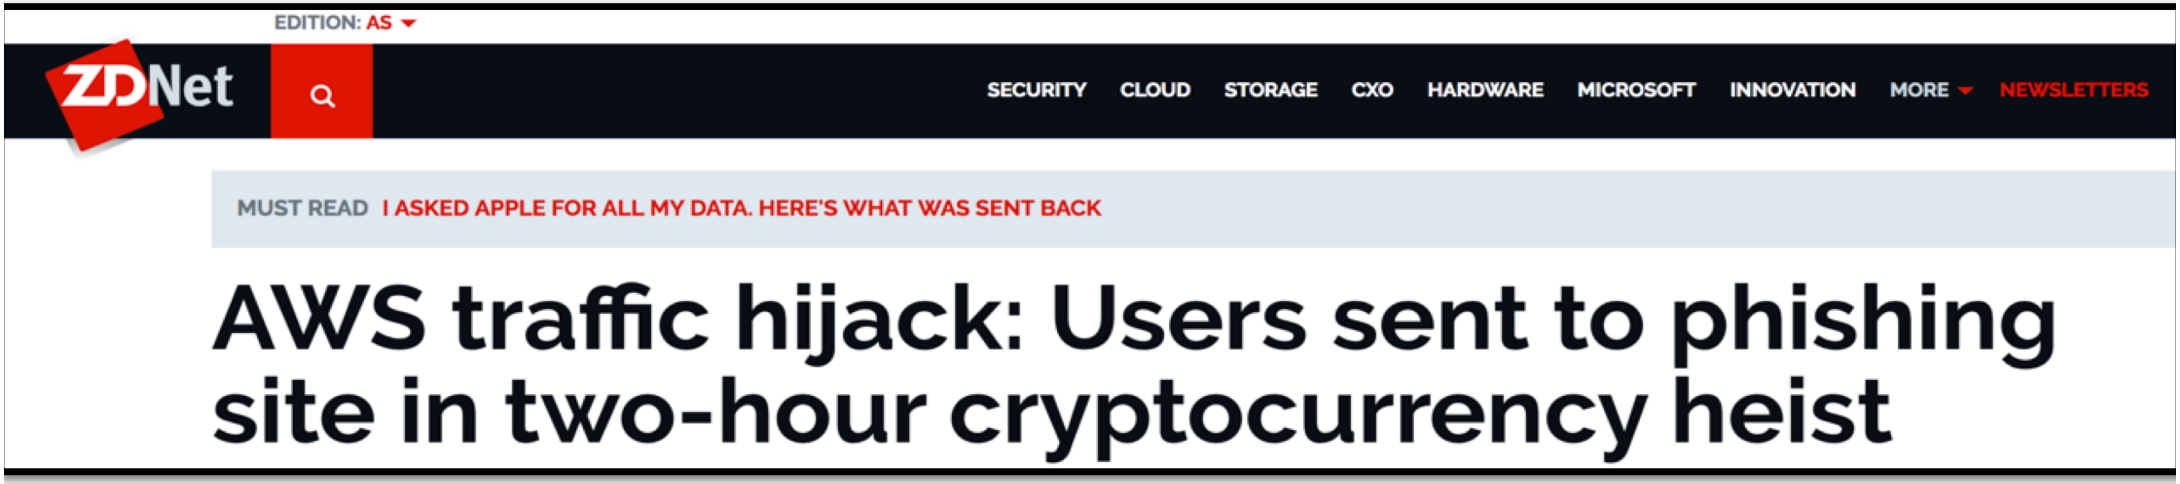

Telegram traffic from around the world took a detour through Iran

The Threats: What's Happening?

Event	Explanation	Repercussions	Solution
Prefix/Route Hijacking	A network operator or attacker impersonates another network operator, pretending that a server or network is their client.	Packets are forwarded to the wrong place, and can cause Denial of Service (DoS) attacks or traffic interception.	Stronger filtering policies
Route Leak	A network operator with multiple upstream providers (often due to accidental misconfiguration) announces to one upstream provider that is has a route to a destination through the other upstream provider.	Can be used for traffic inspection and reconnaissance.	Stronger filtering policies
IP Address Spoofing	Someone creates IP packets with a false source IP address to hide the identity of the sender or to impersonate another computing system.	The root cause of reflection DDoS attacks	Source address validation

Prefix/Route Hijacking

Possible BGP hijack

Beginning at 2019-07-31 19:05:13 UTC, we detected a possible BGP hijack.
Prefix 185.82.215.0/24, is normally announced by AS174 COGENT-174 - Cogent Communications, US.

But beginning at 2019-07-31 19:05:13, the same prefix (185.82.215.0/24) was also announced by ASN 201187.

This was detected by 7 BGPMon peers.

Expected
Start time: 2019-07-31 19:05:13 UTC
Expected prefix: 185.82.215.0/24
Expected ASN: 174 (COGENT-174 - Cogent Communications, US)

Event Details
Detected advertisement: 185.82.215.0/24
Detected Origin ASN 201187 ()
Detected AS Path 29430 13030 42831 201187
Detected by number of BGPMon peers: 7

Source: bgpstream.com

Route Leak

BGP Leak

Beginning at 2019-07-31 20:11:25 UTC, we detected a possible BGP Leak
Prefix 212.154.218.0/23, Normally announced by AS50482 KAZAKHTELECOM-AS, KZ
Leaked by AS31500 GLOBALNET-AS, RU

This was detected by 12 BGPMon peers.

Leak Details

Start time: 2019-07-31 20:11:25 UTC

Leaked prefix: 212.154.218.0/23 (AS50482 KAZAKHTELECOM-AS, KZ)

Leaked By: AS31500 (GLOBALNET-AS, RU)

Leaked To:

- 6461 (ZAYO-6461 - Zayo Bandwidth, US)
- 3267 (RUNNET, RU)

Example AS path: 63956 703 7473 7473 7473 7473 7473 6461 3267 31500 205540 205540 205540 205540 205540 20764 20485 1299 3356 9002 9198 50482

Number of BGPMon peers that saw it: 12

Source: bgpstream.com

IP Address Spoofing

IP address spoofing is used to hide the true identity of the server or to impersonate another server. This technique can be used to amplify an attack.

Example: DNS amplification attack. By sending multiple spoofed requests to different DNS resolvers, an attacker can prompt many responses from the DNS resolver to be sent to a target, while only using one system to attack.

Fix: Source address validation: systems for source address validation can help tell if the end users and customer networks have correct source IP addresses (combined with filtering).

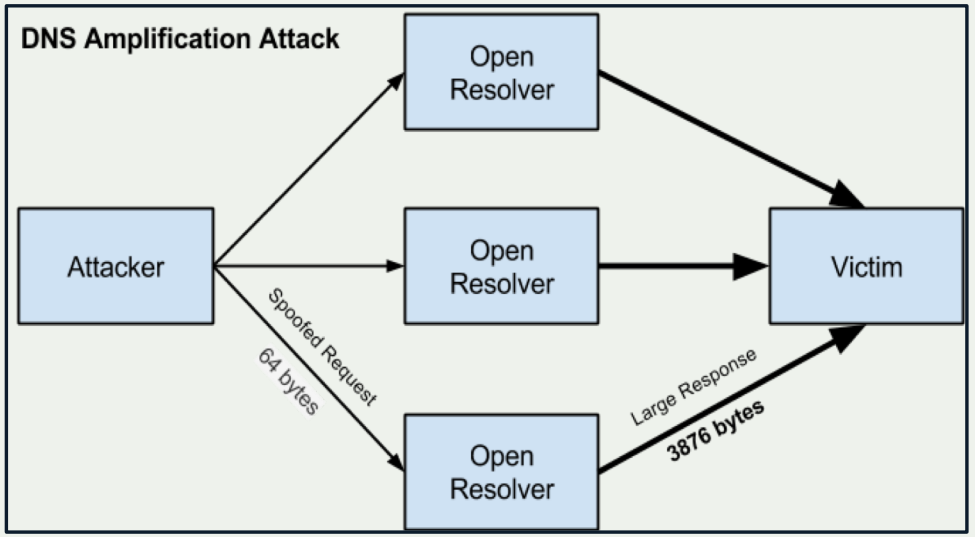

IP Address Spoofing

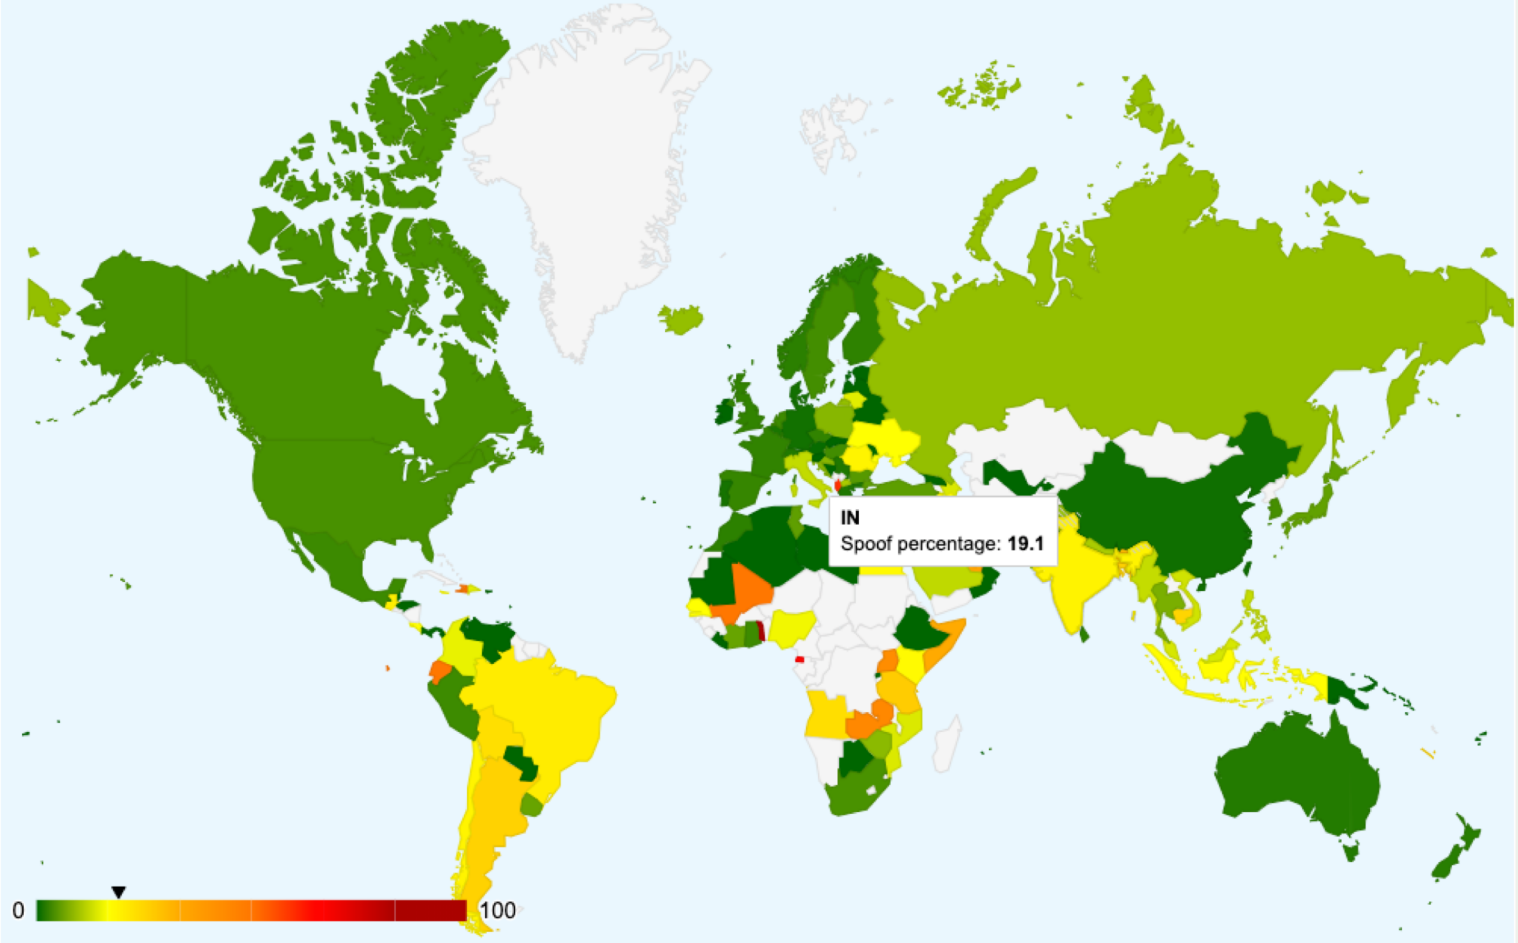

IP Address Spoofing - Thailand

Session	Timestamp (UTC)	Client IP Block	ASN	Country	NAT	Outbound Private Status	Outbound Routable Status	Adj Spoof Prefix Len
747021	2019-09-09 00:36:55	184.22.61.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
747021	2019-09-09 00:36:55	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64
746839	2019-09-08 14:37:36	58.137.176.x/24	4750 (CSLOXINFO-AS-AP)	tha (Thailand)	yes	unknown	unknown	none
746815	2019-09-08 13:24:40	184.22.143.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
746815	2019-09-08 13:24:40	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64
746778	2019-09-08 11:35:35	125.24.250.x/24	23969 (TOT-NET)	tha (Thailand)	yes	rewritten	received	none
746773	2019-09-08 11:12:36	125.24.250.x/24	23969 (TOT-NET)	tha (Thailand)	yes	rewritten	rewritten	none
746663	2019-09-08 06:05:30	184.22.114.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
746591	2019-09-08 01:21:35	184.22.141.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
746591	2019-09-08 01:21:35	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64
746466	2019-09-07 17:17:20	118.172.226.x/24	23969 (TOT-NET)	tha (Thailand)	yes	rewritten	received	none
746397	2019-09-07 14:15:37	58.137.104.x/24	4750 (CSLOXINFO-AS-AP)	tha (Thailand)	yes	unknown	unknown	none
746372	2019-09-07 13:04:37	184.22.141.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
746372	2019-09-07 13:04:37	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64
746178	2019-09-07 00:09:41	184.22.141.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
746178	2019-09-07 00:09:41	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64
745993	2019-09-06 15:51:06	58.137.104.x/24	4750 (CSLOXINFO-AS-AP)	tha (Thailand)	yes	unknown	unknown	none
745844	2019-09-06 10:55:57	184.22.121.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
745777	2019-09-06 07:51:18	184.22.115.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
745769	2019-09-06 07:35:44	184.22.140.x/24	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	yes	rewritten	rewritten	none
745769	2019-09-06 07:35:44	2405:9800:b5xx::/40	133481 (AIS-Fibre-AS-AP)	tha (Thailand)	no	blocked	blocked	/64

The Solution: Mutually Agreed Norms for Routing Security (MANRS)

Provides crucial fixes to eliminate the most common routing threats

Mutually Agreed Norms for Routing Security

MANRS defines four simple but concrete actions that network operators must implement to dramatically improve Internet security and reliability.

- The first two operational improvements eliminate the root causes of common routing issues and attacks, while the second two procedural steps improve mitigation and decrease the likelihood of future incidents.

MANRS

MANRS Actions

Filtering

Prevent propagation of incorrect routing information

Ensure the correctness of your own announcements and announcements from your customers to adjacent networks with prefix and AS-path granularity

Anti-spoofing

Prevent traffic with spoofed source IP addresses

Enable source address validation for at least single-homed stub customer networks, their own end-users, and infrastructure

Coordination

Facilitate global operational communication and coordination between network operators

Maintain globally accessible up-to-date contact information in common routing databases

Global Validation

Facilitate validation of routing information on a global scale

Publish your data, so others can validate

Action 1: Filtering
BCP 194 – RFC7454

BGP Operations and Security

BCP 194 – Prefix Filtering

- IPv4 and IPv6 Special-Purpose Prefixes

The IANA IPv4 and IPv6 Special-Purpose Address Registry maintains the list of special-purpose prefixes and their routing scope, and it SHOULD be used for prefix-filter configuration.

- **Unallocated Prefixes**

IANA allocates prefixes to RIRs that in turn allocate prefixes to LIRs (Local Internet Registries). It is wise not to accept routing table prefixes that are not allocated by IANA and/or RIRs. This section details the options for building a list of allocated prefixes at every level. It is important to understand that filtering unallocated prefixes requires constant updates, as prefixes are continually allocated. Therefore, automation of such prefix filters is key for the success of this approach.

BCP 194 – Prefix Filtering

Inbound Filtering (Loose – Strict)

Loose - where no check will be done against RIR allocations

Strict - where it will be verified that announcements strictly conform to what is declared in routing registries.

BCP 194 – Prefix Filtering

Inbound Filtering Loose Option

In this case, the following prefixes received from a BGP peer will be filtered:

- prefixes that are not globally routable
- prefixes not allocated by IANA (IPv6 only)
- routes that are too specific
- prefixes belonging to the local AS
- IXP LAN prefixes
- the default route

BCP 194 – Prefix Filtering

Inbound Filtering Strict Option

In this case, filters are applied to make sure advertisements strictly conform to what is declared in routing registries. This varies across the registries and regions of the Internet.

In addition to this, apply the following filters beforehand in case the routing registry that's used as the source of information by the script is not fully trusted:

- prefixes that are not globally routable
- routes that are too specific
- prefixes belonging to the local AS
- IXP LAN prefixes and the default route

BCP 194 – Prefix Filtering

Outbound Filtering

The configuration should ensure that only appropriate prefixes are sent. These can be, for example, prefixes belonging to both the network in question and its downstream. This can be achieved by using BGP communities, AS paths, or both. Also, it may be desirable to add the following filters before any policy to avoid unwanted route announcements due to bad configuration:

- Prefixes that are not globally routable
- Routes that are too specific
- IXP LAN prefixes
- The default route

BCP 194 – Max Prefix Filtering

It is RECOMMENDED to configure a limit on the number of routes to be accepted from a peer. The following rules are generally RECOMMENDED:

- From peers, it is RECOMMENDED to have a limit lower than the number of routes in the Internet. This will shut down the BGP peering if the peer suddenly advertises the full table.
- From upstreams that provide full routing, it is RECOMMENDED to have a limit higher than the number of routes in the Internet. A limit is still useful in order to protect the network (and in particular, the routers' memory) if too many routes are sent by the upstream.

BCP 194 – AS Path Filtering

Following are the RECOMMENDED practices when processing BGP AS paths.

- Network administrators SHOULD accept from customers only 2-byte or 4-byte AS paths containing ASNs belonging to (or authorized to transit through) the customer.
- Network administrators SHOULD NOT accept prefixes with private AS numbers in the AS path unless the prefixes are from customers.
- Network administrators SHOULD NOT accept prefixes when the first AS number in the AS path is not the one of the peer's unless the peering is done toward a BGP route server
- Network administrators SHOULD NOT advertise prefixes with upstream AS numbers in the AS path to their peering AS unless they intend to provide transit for these prefixes.

Action 2: Anti-Spoofing

BCP 38 – RFC2827

Network Ingress Filtering

Source Address Validation

Check the source IP address of IP packets

- filter invalid source address
- filter close to the packets origin as possible
- filter precisely as possible

If no networks allow IP spoofing, we can eliminate these kinds of attacks

Action 3: Coordination

Facilitating global operational communication and coordination between network operators

Coordination

Maintaining Contact Information in Regional Internet Registries (RIRs): AFRINIC, APNIC, RIPE NCC, LACNIC, ARIN

For Example: whois –h whois.apnic.net AS9541

```
aut-num:        AS9541
as-name:        CYBERNET-AP
descr:          Cyber Internet Services (Pvt) Ltd.
country:        PK
org:            ORG-CISP3-AP
mnt-routes:     MAINT-PK-CYBERNET
admin-c:        AQ84-AP
tech-c:         AQ84-AP
mnt-by:         APNIC-HM
mnt-irt:        IRT-CYBERNET-PK
mnt-lower:      MAINT-PK-CYBERNET
last-modified:  2019-06-09T22:39:23Z
source:         APNIC
```

Coordination

```
irt:            IRT-CYBERNET-PK
address:        A904, 9th Floor,Lakson Bldg 3,Sarwar Shaheed Rd,Karachi-74200
e-mail:         noc-abuse@cyber.net.pk
abuse-mailbox:  noc-abuse@cyber.net.pk
admin-c:        AQ84-AP
tech-c:         AQ84-AP
auth:           # Filtered
mnt-by:         MAINT-PK-AQ
last-modified:  2016-01-05T10:59:53Z
source:         APNIC

organisation:   ORG-CISP3-AP
org-name:       Cyber Internet Services Pakistan
country:        PK
address:        A - 904 9th Floor Lakson Square Building No. 3
address:        No. 3, Sarwar Shaheed Road Karachi-74200 Pakistan
phone:          +92-21-38400654
fax-no:         +92-213-5680842
e-mail:         noc-abuse@cyber.net.pk
last-modified:  2019-04-25T12:55:55Z
source:         APNIC
```

Action 4: Global Validation

Facilitating validation of routing information on a global scale

Global Validation

Routing information should be made available on a global scale to facilitate validation, which includes routing policy, ASNs and prefixes that are intended to be advertised to third parties. Since the extent of the internet is global, information should be made public and published in a well known place using a common format.

Object	**Source**	Description
aut-num	IRR	Policy documentation
route/route6	IRR	NLRI/origin
as-set	IRR	Customer cone
ROA	RPKI	NLRI/origin

Global Validation

There are 2 ways to provide the validation information (IRR and/or RPKI)

Providing information through the IRR system

Internet Routing Registries (IRRs) contain information—submitted and maintained by ISPs or other entities—about Autonomous System Numbers (ASNs) and routing prefixes. IRRs can be used by ISPs to develop routing plans.

The global IRR is comprised of a network of distributed databases maintained by Regional Internet Registries (RIRs) such as APNIC, service providers (such as NTT), and third parties (such as RADB).

Global Validation

```
$ whois -h whois.apnic.net 1.1.1.0/24

route:          1.1.1.0/24
origin:         AS13335
descr:          APNIC Research and Development, 6 Cordelia
St
mnt-by:         MAINT-AU-APNIC-GM85-AP
last-modified:  2018-03-16T16:58:06Z
source:         APNIC
```

Global Validation

```
$ whois -h whois.radb.net 1.1.1.0/24

route:          1.1.1.0/24
origin:         AS13335
descr:          APNIC Research and Development, 6 Cordelia
St
mnt-by:         MAINT-AU-APNIC-GM85-AP
last-modified:  2018-03-16T16:58:06Z
source:         APNIC

route:          1.1.1.0/24
descr:          Cloudflare, Inc.
descr:          101 Townsend Street, San Francisco,
California 94107, US
origin:         AS13335
mnt-by:         MNT-CLOUD14
```

Global Validation

Providing information through the RPKI system

The RPKI repository can store information about prefixes originated by your network in the form of Route Origin Authorization (ROA) objects. Note, that these do not include your customer announcements, but only prefixes that belong to your ASN. Only the origin ASN is verified, not the full path.

All Regional Internet Registries offer a so-called hosted Resource Certification service where keys are kept and managed by the RIR and all operations are performed on the RIR's servers.

Global Validation

Providing information through the RPKI system

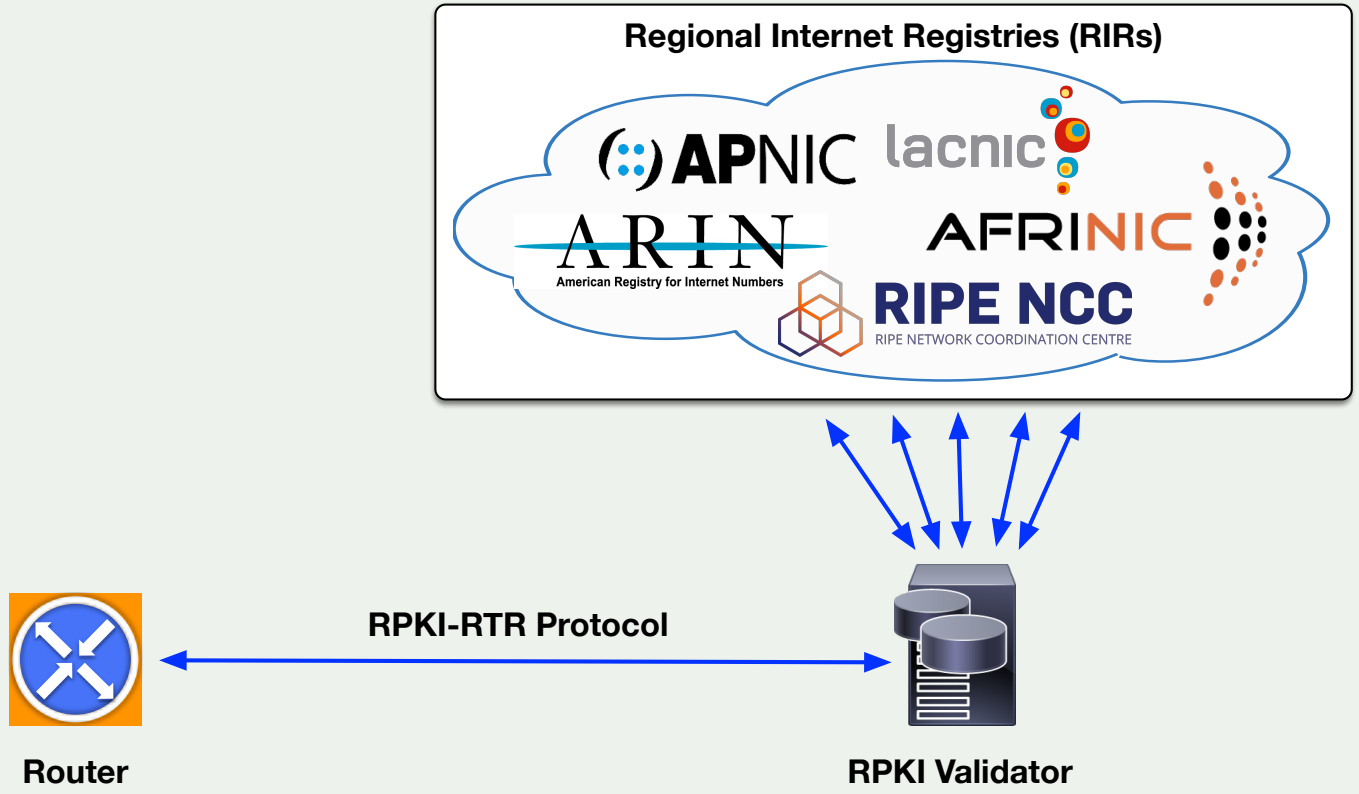

Global Validation

http://localcert.ripe.net:8088/api/v1/validity/AS13335/1.1.1.0/24


```
{
  "validated_route":{
    "route":{
      "origin_asn":"AS13335",
      "prefix":"1.1.1.0/24"
    },
    "validity":{
      "state":"Valid",
      "description":"At least one VRP Matches the Route Prefix",
      "VRPs":{
        "matched":[{
          "asn":"AS13335",
          "prefix":"1.1.1.0/24",
          "max_length":24
        }],
        "unmatched_as":[],
        "unmatched_length":[]
      }
    }
  }
}
```

Why join MANRS?

Join Us

Visit https://www.manrs.org

- Fill out the sign up form with as much detail as possible.
- We may ask questions and run tests

Get Involved in the Community

- Members support the initiative and implement the actions in their own networks
- Members maintain and improve the document and promote MANRS objectives

MANRS Implementation Guide

If you're not ready to join yet, implementation guidance is available to help you.

- Based on Best Current Operational Practices deployed by network operators around the world
- https://www.manrs.org/bcop/

Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide

MANRS

Version 1.0, BCOP series
Publication Date: 25 January 2017

MANRS Training Modules

6 training modules based on information in the Implementation Guide.

Walks through the tutorial with a test at the end of each module.

Working with and looking for partners that are interested in integrating it in their curricula.

https://academy.apnic.net/en/course/manrs/

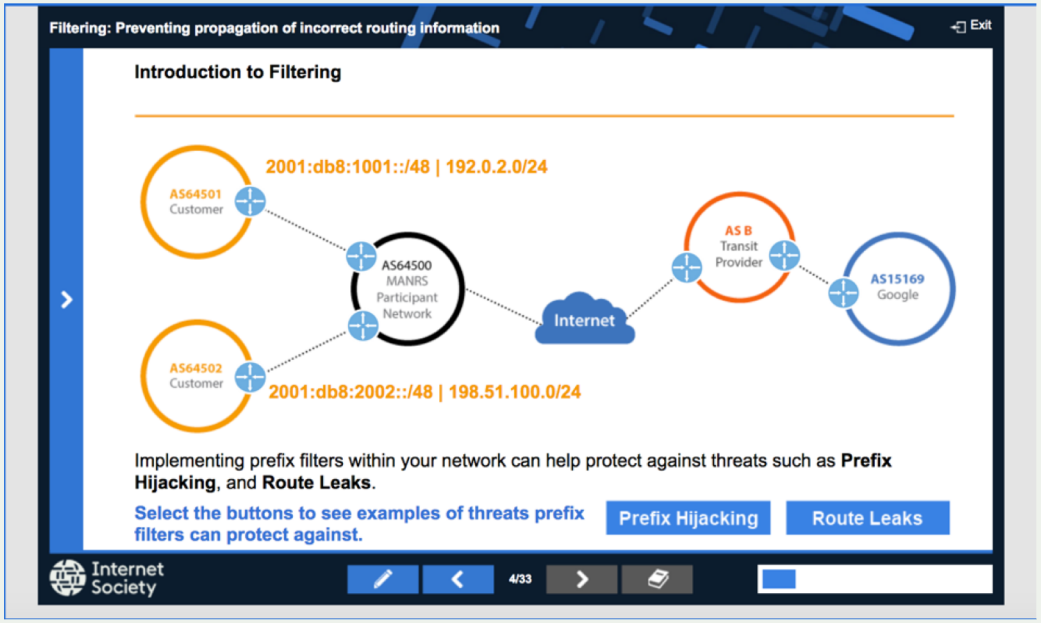

Thanks to APNIC for hosting MANRS Tutorial

MANRS IXP Partnership Programme

How can IXPs Join MANRS?

The IXP Programme Action Set (Actions 1 and 2 are mandatory, and the IXP must implement at least one additional Action.)

Action 1. Prevent propagation of incorrect routing information. (Mandatory)

Action 2. Promote MANRS to the IXP membership. (Mandatory)

Action 3. Protect the peering platform.

Action 4. Facilitate global operational communication and coordination between network operators.

Action 5. Provide monitoring and debugging tools to the members.

“The good we secure for ourselves is precarious and uncertain until it is secured for all of us and incorporated into our common life.”

— **Jane Addams** (Nobel Peace Prize Winner)

Thank you.

Aftab Siddiqui
siddiqui@isoc.org

manrs.org