Securing Global Routing System and Operators Approach

Fakrul Alam
Senior Training Officer
fakrul@apnic.net

SANOG 28
01 August - 09 August, 2016, Mumbai, India

Issue Date: [5th July 2016]
Revision: [1.0]

APNIC

Incidents

YouTube Hijacking: A RIPE NCC RIS case study

You're viewing an archived page. It is no longer being updated.

Introduction

On Sunday, 24 February 2008, Paki
announcement of the prefix 208.65
providers, PCCW Global (AS3491) f
Internet, which resulted in the hija

In this report we show how the eve
Service (RIS) and how, in general, o
network events.

Large scale BGP hijack out of India

Posted by Andree Toonk - November 6, 2015 - Hijack - 1 Comment

BGP hijacks happen every day, some of them affect more ne
and then there's a major incident that affects thousands of
keep an eye out for our users and if you would like to have
the world of BGP incidents, keep an eye on BGPstream.com.
those major incidents that affected thousands of networks.

Starting at 05:52 UTC, AS9498 (BHARTI Airtel Ltd.) started to claim ownership for thousan
prefixes by originating them in BGP. This affected prefixes for over two thousand unique
organizations (Autonomous systems).

Our systems detected origin AS changes (hijacks) for 16,123 prefixes. The scope and imp
different per prefix but to give you an idea, about 7,600 of these announcements were se
five or more of our peers (unique peers ASns) and 6,000 of these were seen by more than
our peers.

One of the reasons this was so widespread is because large networks such as AS174 (Cog
Communications) and AS52320 (GlobeNet Cabos Submarinos VZLA) accepted and propaga
these prefixes to their peers and customers.

Dyn Research
@DynResearch

Yesterday, Guam's Choice Phone (AS55863) leaked 45 /8's for a grand total of 754MM IPv4 routed addresses

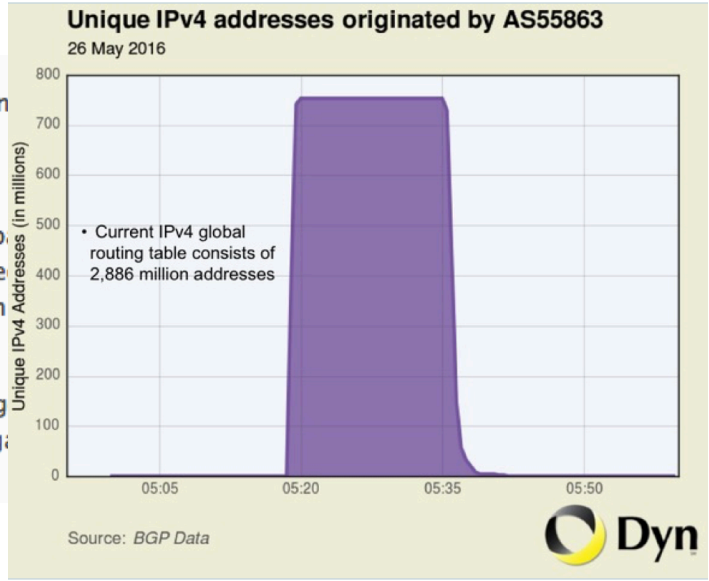

Motivations!

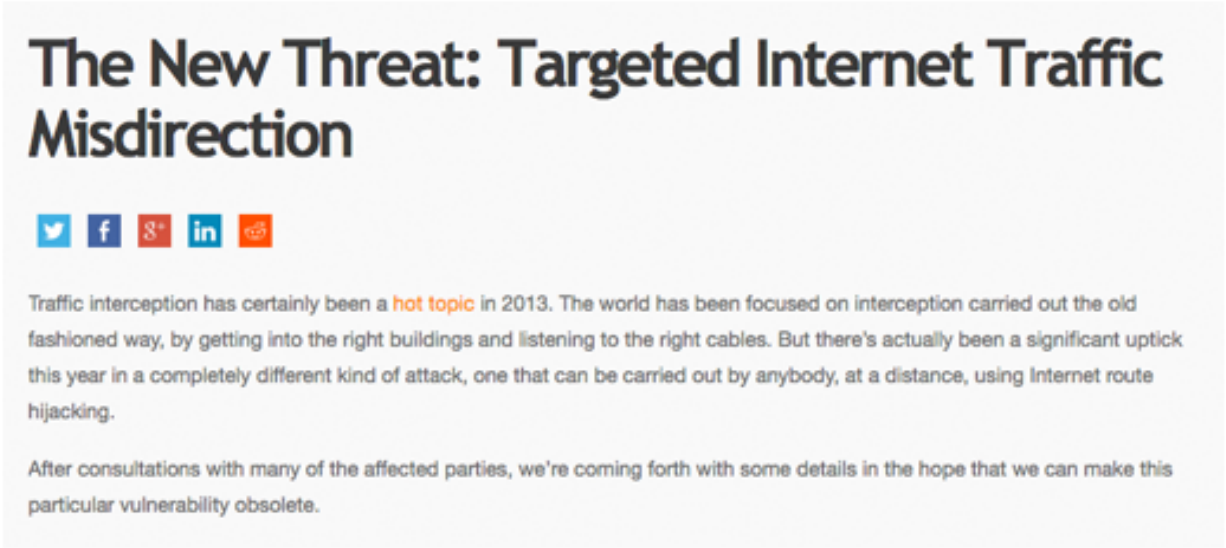

The New Threat: Targeted Internet Traffic Misdirection

Traffic interception has certainly been a hot topic in 2013. The world has been focused on interception carried out the old fashioned way, by getting into the right buildings and listening to the right cables. But there's actually been a significant uptick this year in a completely different kind of attack, one that can be carried out by anybody, at a distance, using Internet route hijacking.

After consultations with many of the affected parties, we're coming forth with some details in the hope that we can make this particular vulnerability obsolete.

some **spammers** are currently using **short-lived bogus BGP announcements** to send spam from hijacked parts of the IPv4 address space. Such a spammer would use BGP to **announce some address space**, then **send spam** from those addresses, and then **withdraw the announcement**.

Current Practice

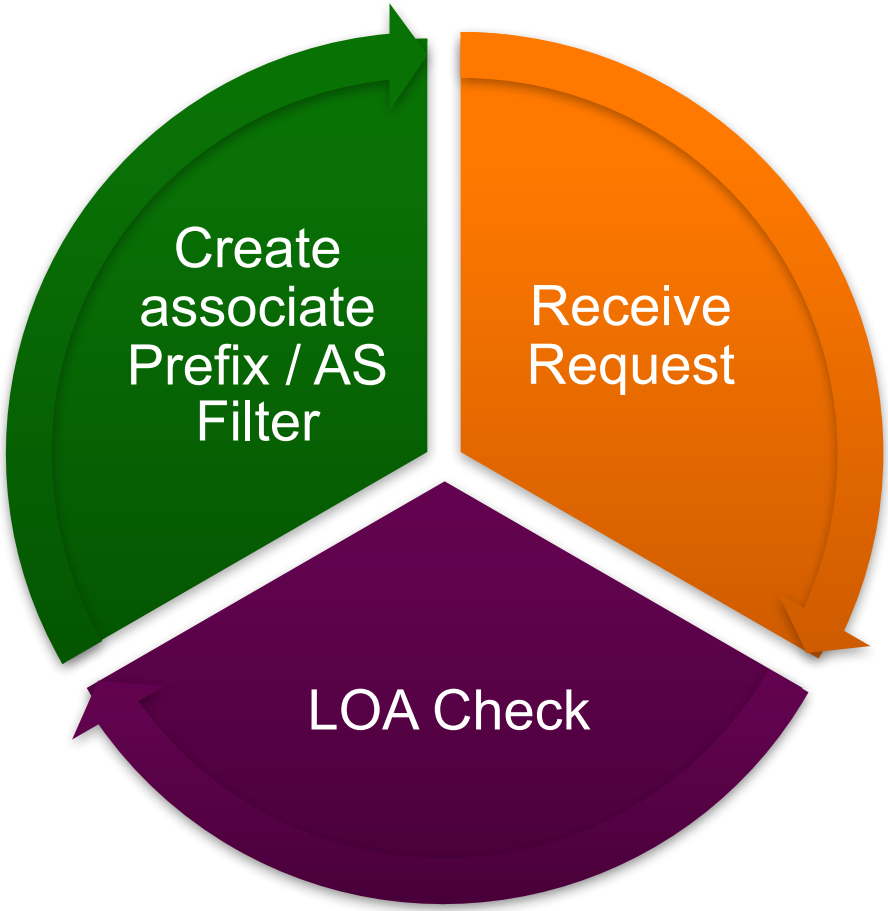

Tools & Techniques

- Manual LoA Check
  - Whois search on the customer's IP address from the IRR database
  - Find the admin-c / tech-c contact e-mail address from the database search and email them for verification
  - Check corresponding "route objects"
- Automated LoA Check
  - Fetch the routing policy from IRR Database
  - Generate associate prefix/as filter
  - Mostly done using RPSL
- RPKI
  - Check & validate prefix origin cryptographically

LoA Check

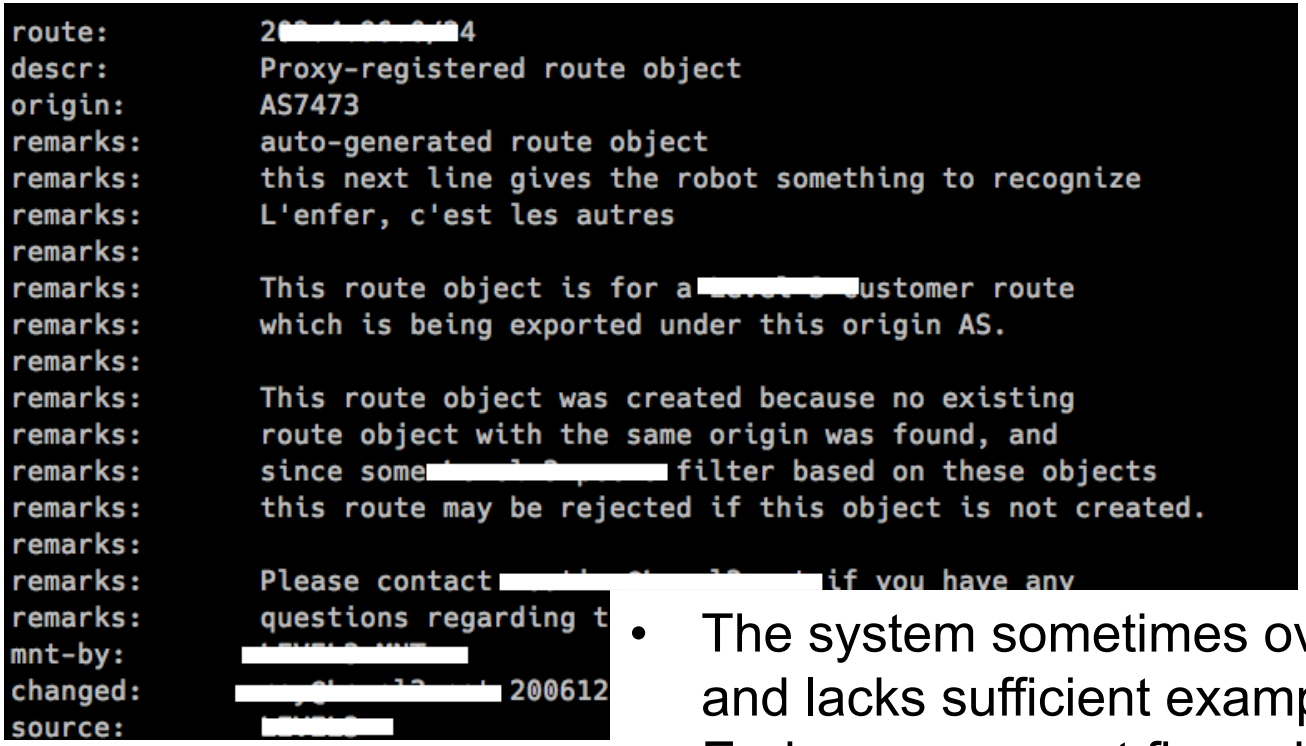

- The system sometimes overly complicated, and lacks sufficient examples.
- End users can not figure it out, which means another layer of support structure must be added, or proxy registration must be implemented.

LoA Check & RPSL

```
~ whois -h whois.radb.net AS1299 | more
aut-num:        AS1299
org:            ORG-TA45-RIPE
as-name:        TELIANET
import:         from AS57 action pref=50; accept AS-NLG-TO-TRANSIT
import:         from AS62 action pref=50; accept AS-c1
import:         from AS109 action pref=50; accept AS109
import:         from AS174 action pref=100; accept AS-PSINET
import:         from AS209 action pref=100; accept AS209
import:         from AS286 action pref=100; accept AS-KPN
import:         from AS293 action pref=100; accept AS-ESNET
import:         from AS577 action pref=50; accept AS577:AS-CUSTOMERS
import:         from AS612 action pref=50; accept AS612
import:         from AS701 action pref=100; accept AS701 AS701:AS-CUSTOMERS
import:         from AS702 action pref=100; accept AS702:RS-EURO AS702:RS-CUSTOMER
import:         from AS714 action pref=50; accept AS714
import:         from AS786 action pref=50; accept AS-JANETUS
import:         from AS812 action pref=50; accept AS-ROGERS:AS-CUSTOMERS
import:         from AS852 action pref=50; accept AS-TELUS
import:         from AS855 action pref=50; ac
import:         from AS1239 action pref=100;
import:         from AS1248 action pref=50; a
import:         from AS1257 action pref=100;
import:         from AS1267 action pref=50; a
import:         from AS1273 action pref=50; a
import:         from AS1280 action pref=50; a
```

A publicly accessible description of every import and export policy to every transit, peer, and customer, is difficult to maintain, and is not in the best business interests of many ISPs.

```
~ whois -h whois.radb.net AS1299 | wc -l
4924
~ _
```

RPKI Implementation

- Origin Validation
- Hosted CA
  - Easy to deploy, but have to trust a third party with your private key
- Delegated
  - Complexity in installing CA, generate ROAs, publish URI & point TA
- Upgrade at least ASBRs to RPKI capable code

Technology & Learning Curve

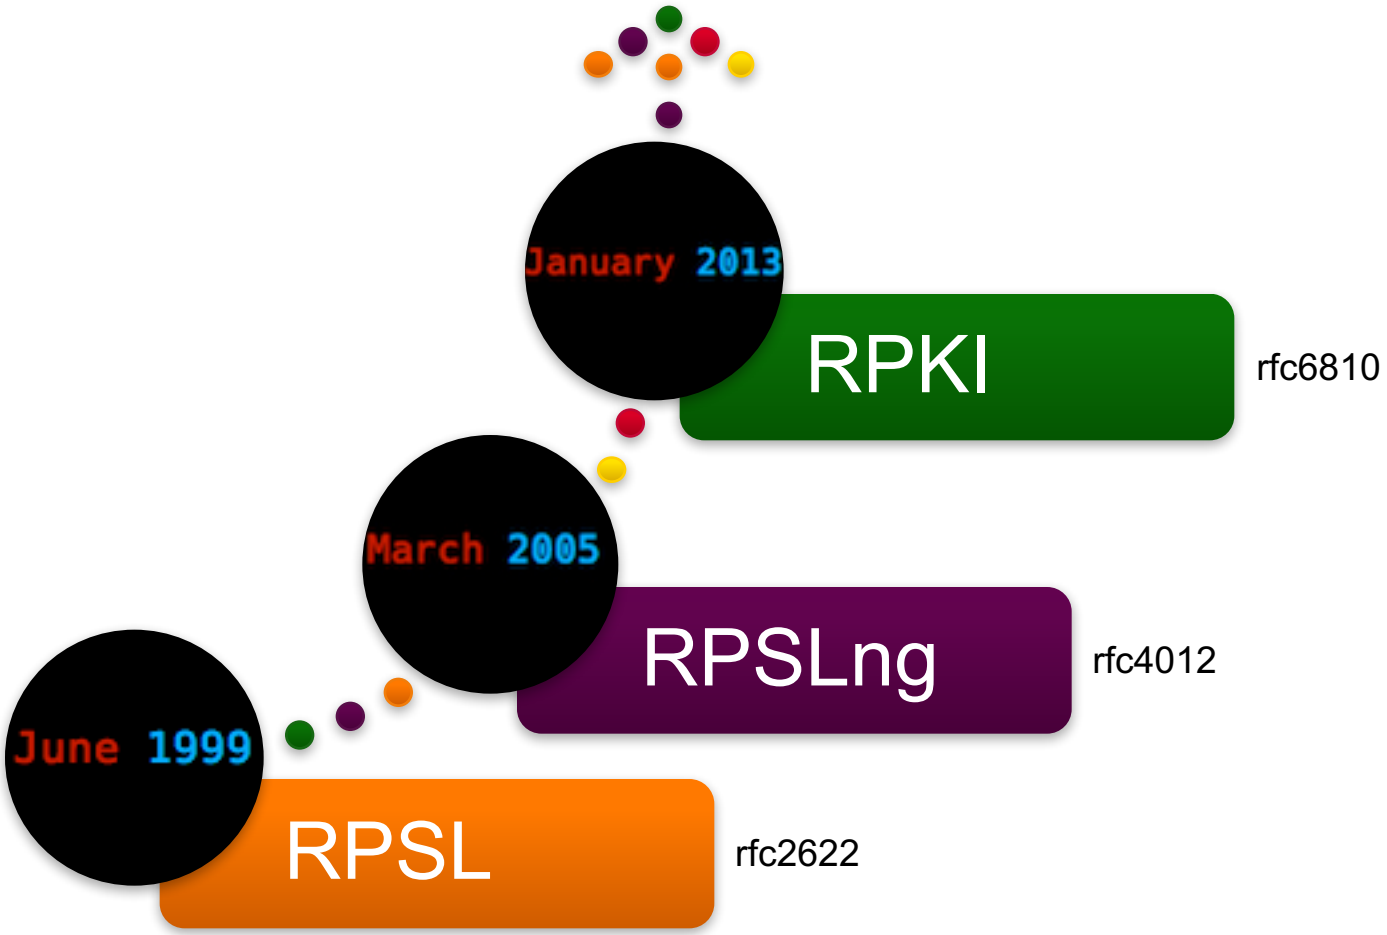

But how Operators

are

Adopting & Implementing?

Prefixes Distribution

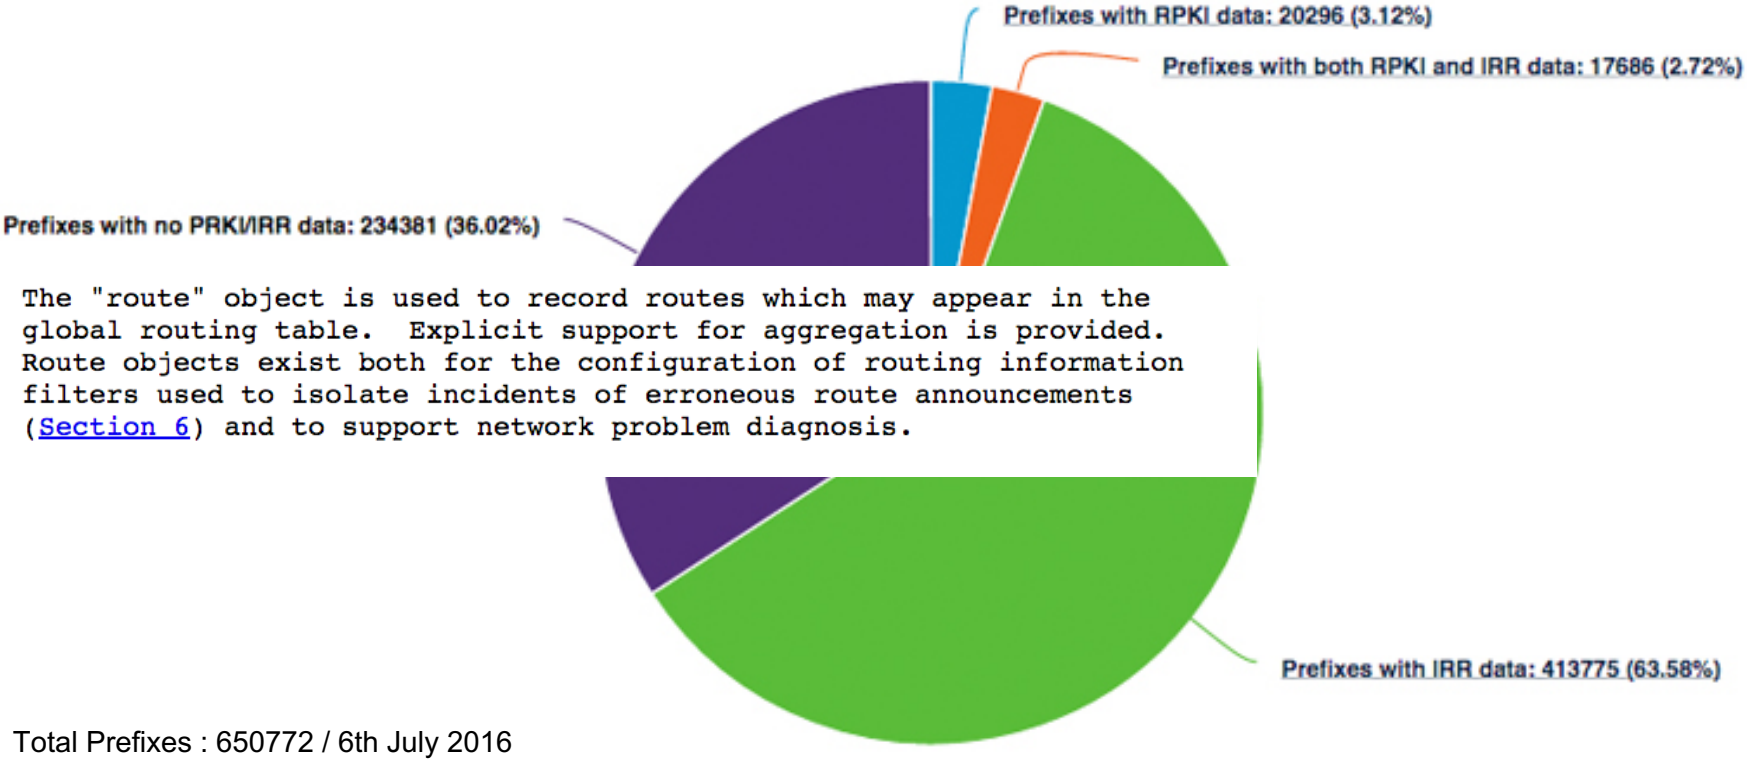

Total Prefixes : 650772 / 6th July 2016

Prefixes with IRR Data

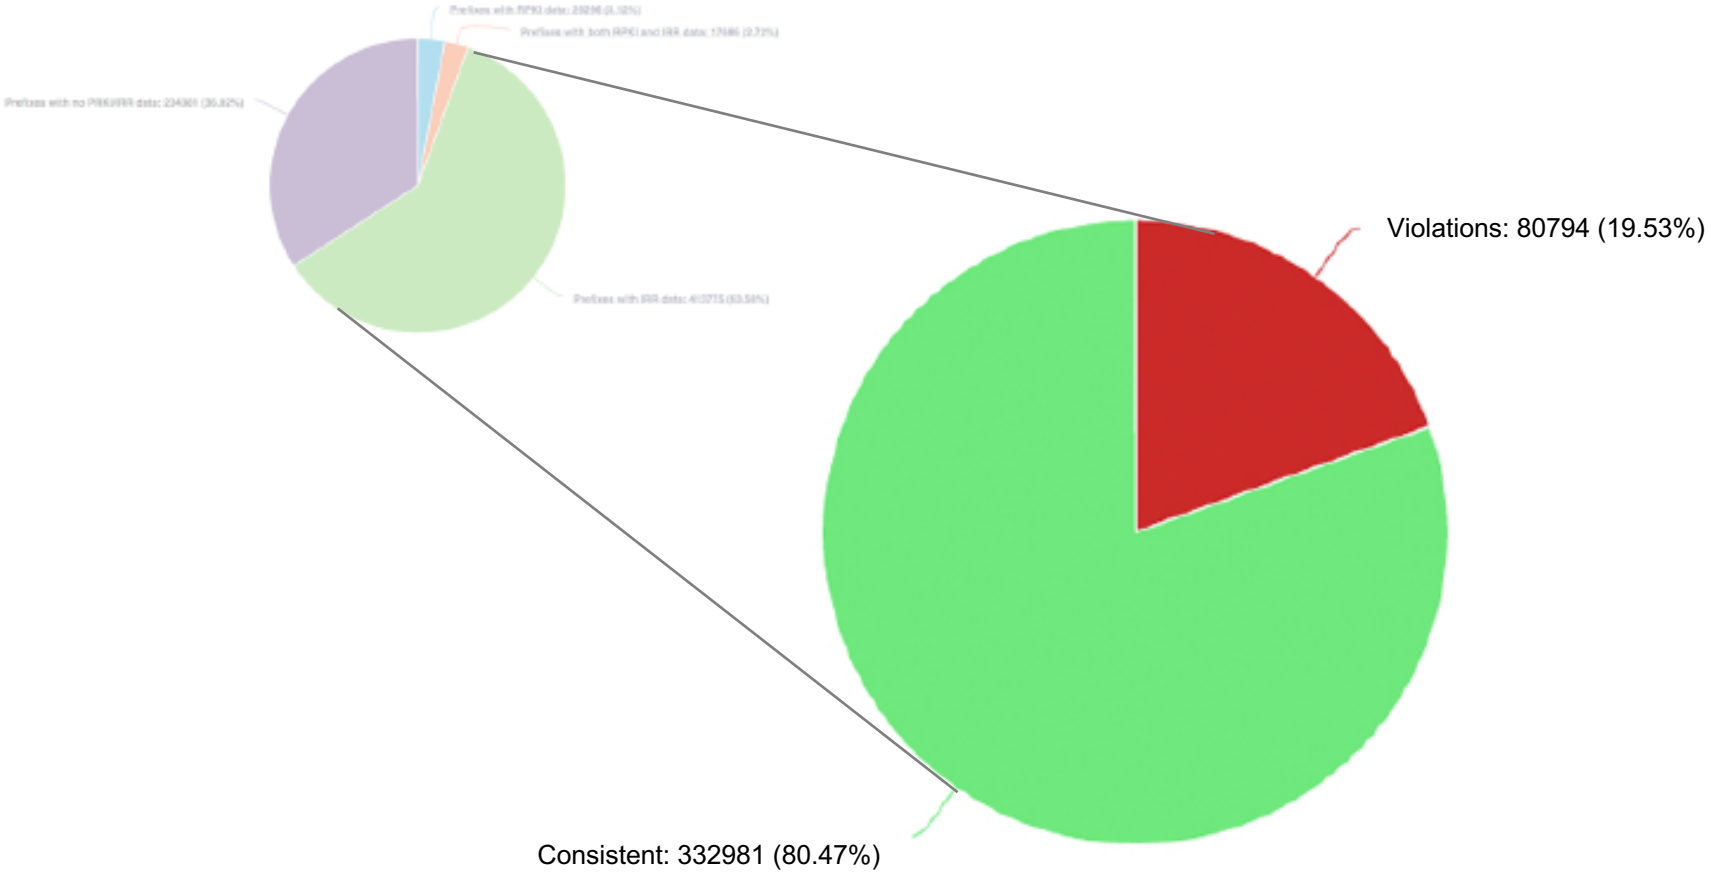

IRR Data Violations Example

Prefix/Len	Recv Origin AS	IRR Origin AS
203.27.30.0/24	4294836336	2147483647
103.62.28.0/24	4294836383	2147483647
103.62.29.0/24	4294836383	2147483647

```
~  Desktop  whois -h whois.radb.net 203.27.30.0/24
route:      203.27.30.0/24
descr:      Proxy route object registered by AS2764
origin:     AS4294836336
remarks:    This route object was created by AAPT on behalf of a customer.
remarks:    As some of AAPTs upstream networks filter based on IRR objects,
remarks:    this rou
remarks:    of this
notify:     routing.
mnt-by:     MAINT-AS
changed:    nobody@a
source:     RADB
```

```
a          .net>show ip bgp 103.62.29.0/24
BGP routing table entry for 103.62.29.0/24, version 198378
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     2
  Refresh Epoch 2
  4826 [redacted] 4294836383
    49.255.232.169 from 49.255.232.169 (114.31.194.12)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 316282439 316282836 316333767
      rx pathid: 0, tx pathid: 0x0
```

Prefixes with RPKI

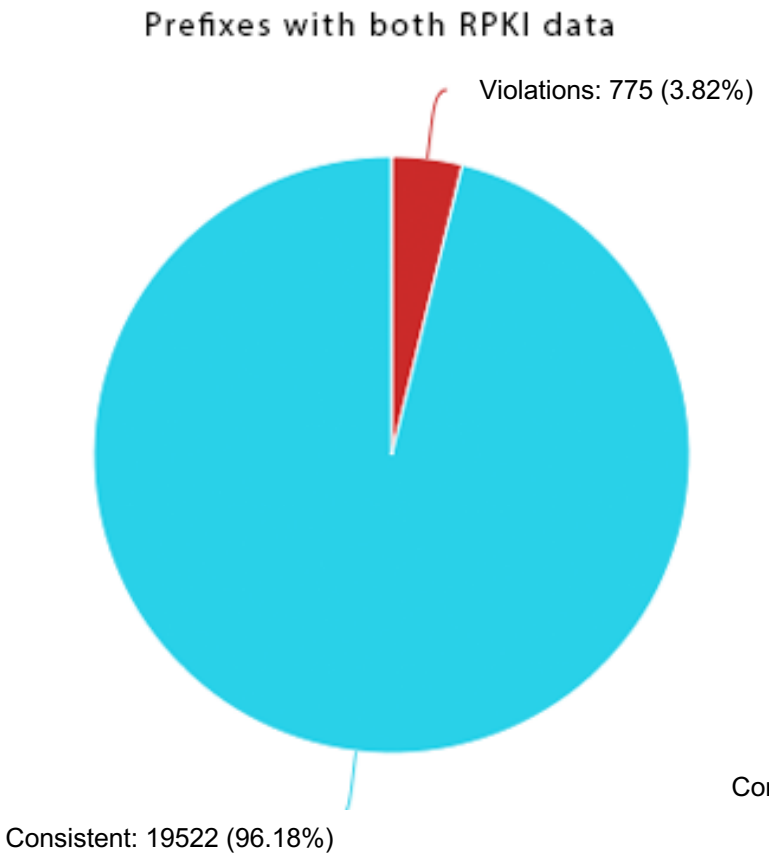

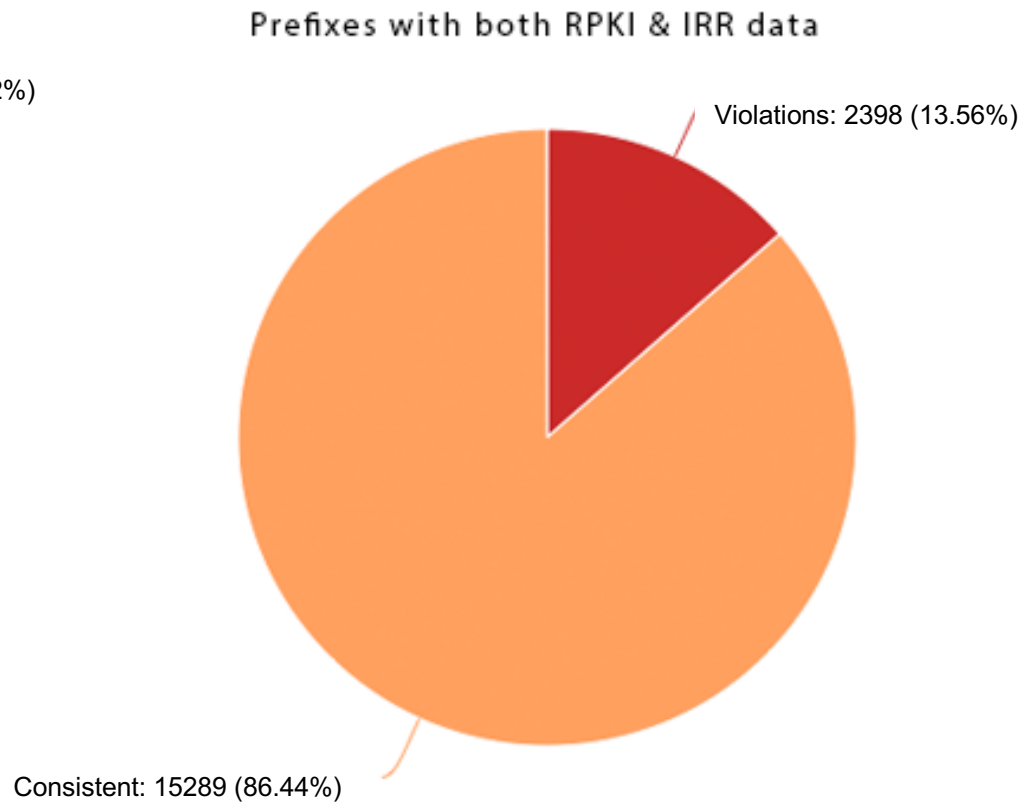

RPKI Data Violation Example

- Most of the cases Invalid Prefix (Fixed length mismatch)
  - Create ROA for /22 but announce 24
- Invalid origin AS is also visible

```
~ whois -h whois.bgpmon.net " --roa 14080 213.192.242.0/23"
2 - Not Valid: Invalid Origin ASN, expected 8903
```

RPKI Data Violation Example


```
{
  "validated_route": {
    "route": {
      "origin_asn": "AS58456",
      "prefix": "202.70.91.0/24"
    },
    "validity": {
      "state": "Invalid",
      "reason": "as",
      "description": "At least one VRP Covers the Route Prefix
      "VRPs": {
        "matched": [],
        "unmatched_as": [
          {
            "asn": "AS23752",
            "prefix": "202.70.64.0/19",
            "max_length": 19
          }
        ],
        "unmatched_length": []
      }
    }
  }
}
```


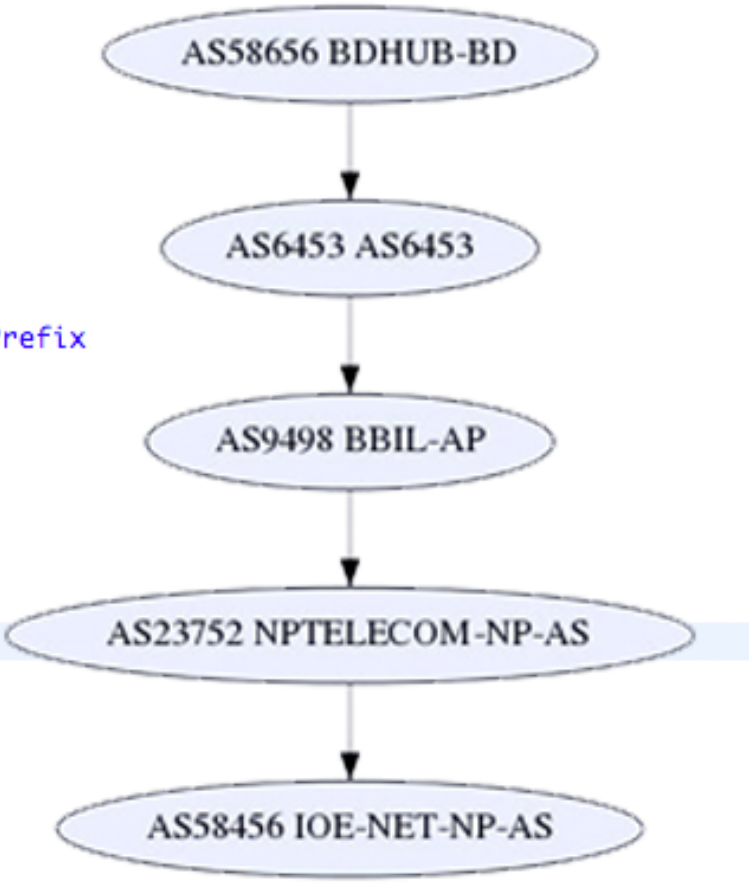

How About South Asia!

ROA in South Asia

Country	IPv4 Prefixes Covered	IPv4 Prefixes Valid
Afghanistan	0%	0%
Bangladesh	25.11%	24.05%
Bhutan	86.67%	86.67%
India	0.04%	0.03%
Nepal	55.3%	18.28%
Maldives	0%	0%
Pakistan	12.17%	12.14%
Sri Lanka	50.18%	40.57%

source : https://lirportal.ripe.net/certification/content/static/statistics/world-roas.html
date : 18th July 2016

Bangladesh

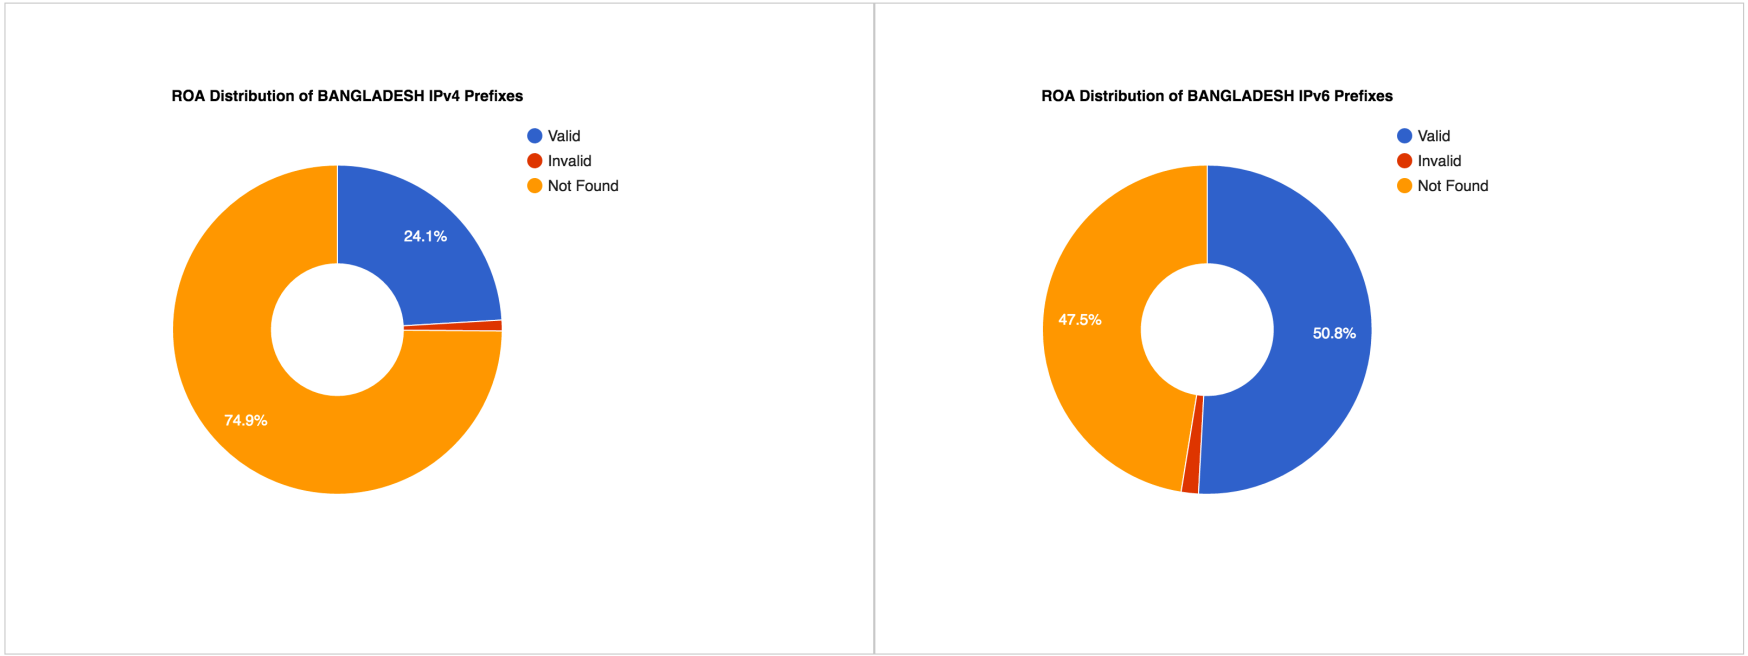

ref link : http://rpki.apnictraining.net/output/bd.html

Bhutan

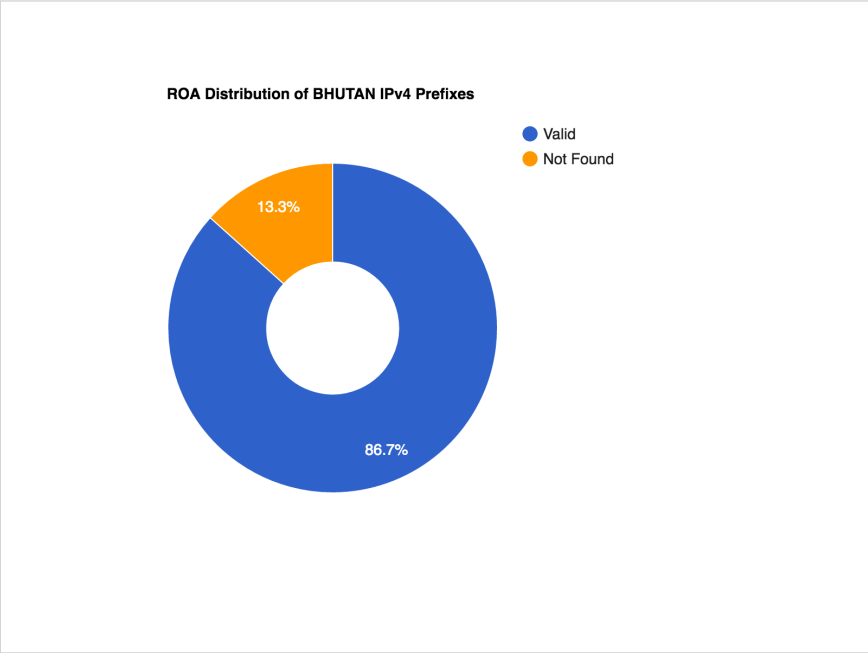

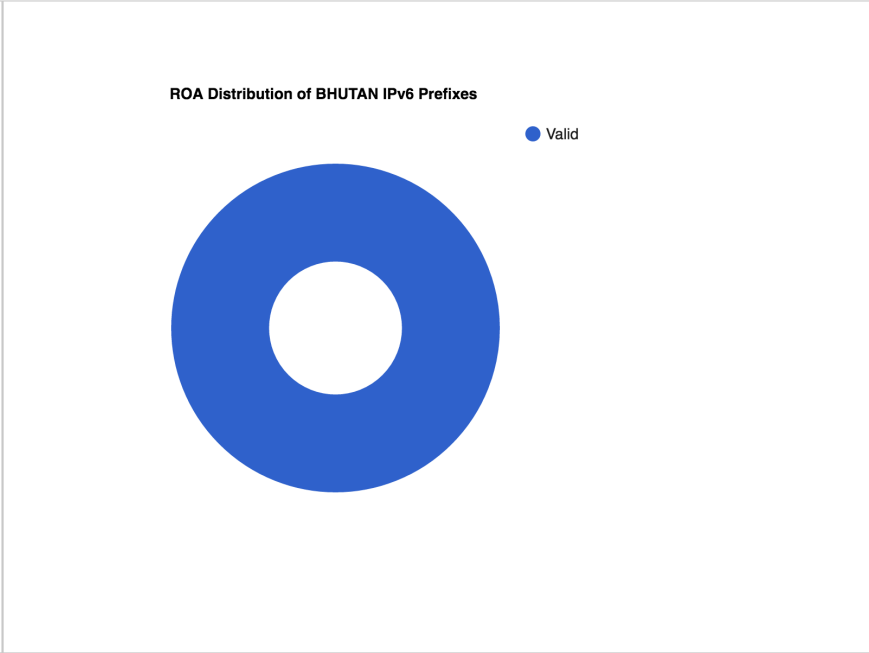

ref link : http://rpki.apnictraining.net/output/bt.html

India

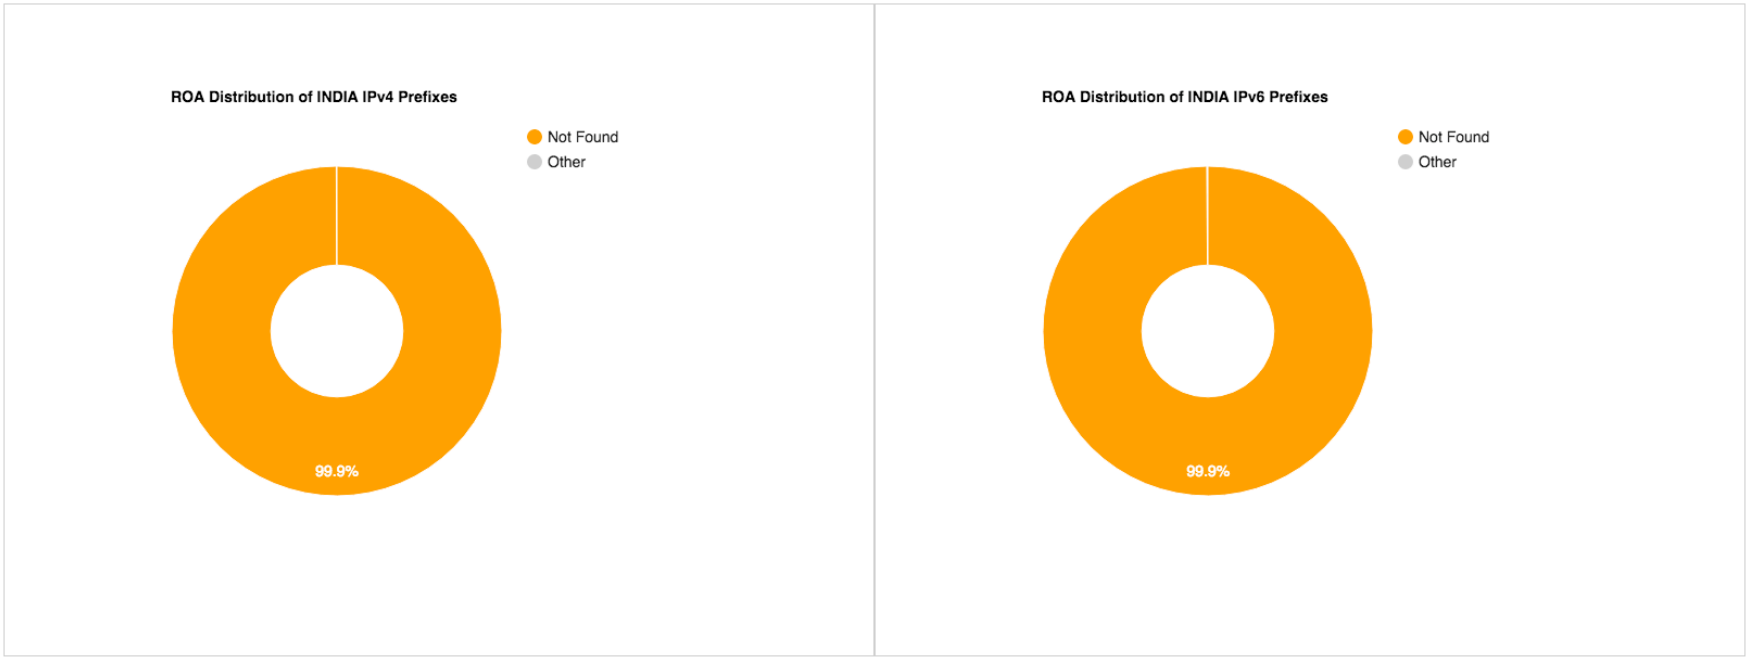

ref link : http://rpki.apnictraining.net/output/in.html

Nepal

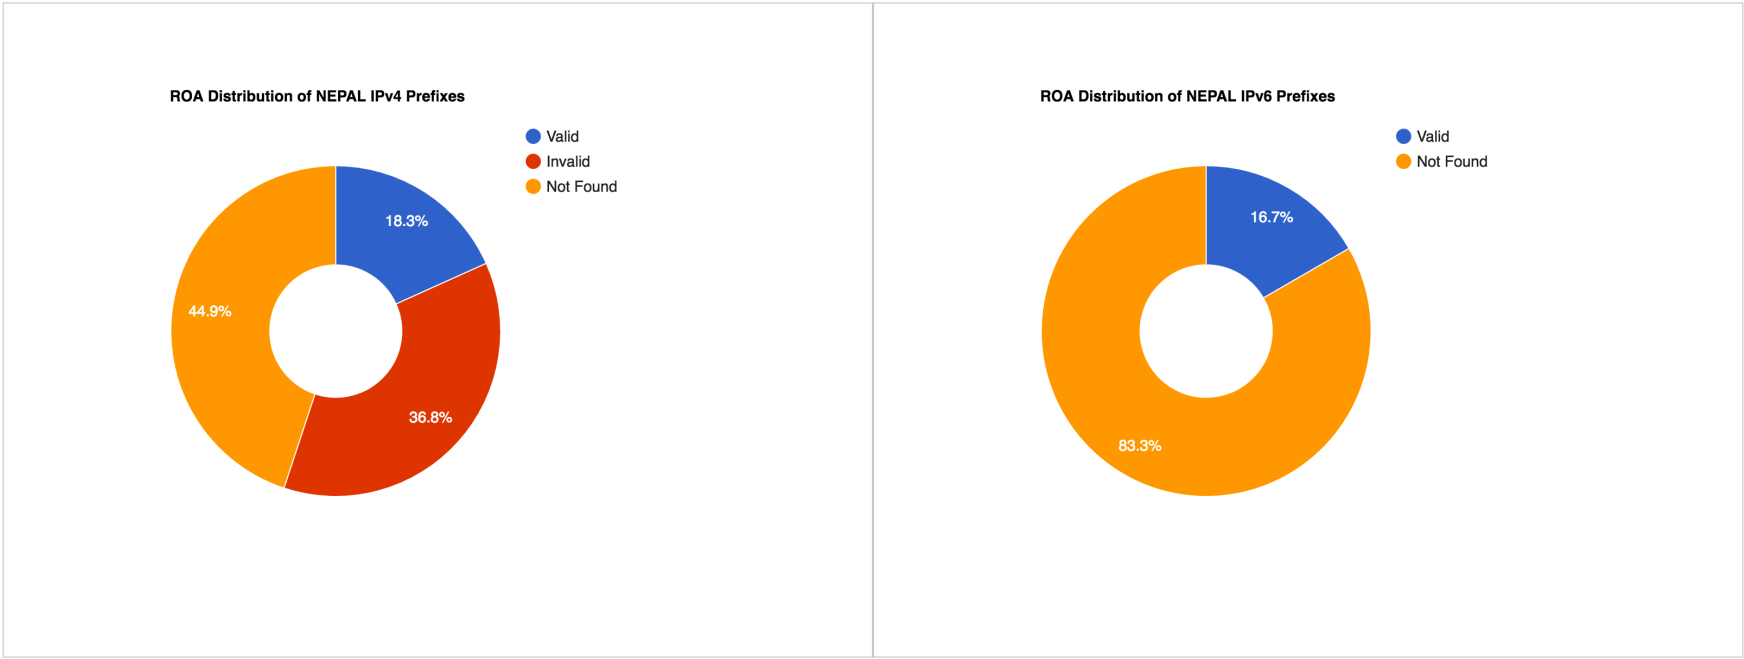

ref link : http://rpki.apnictraining.net/output/np.html

Pakistan

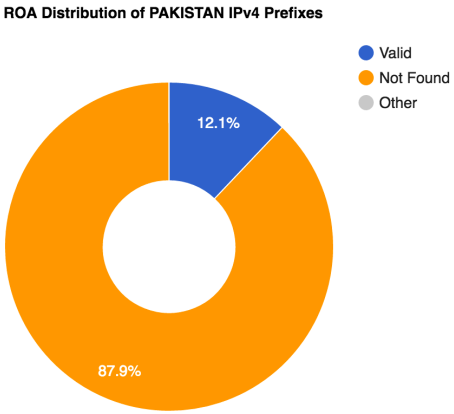

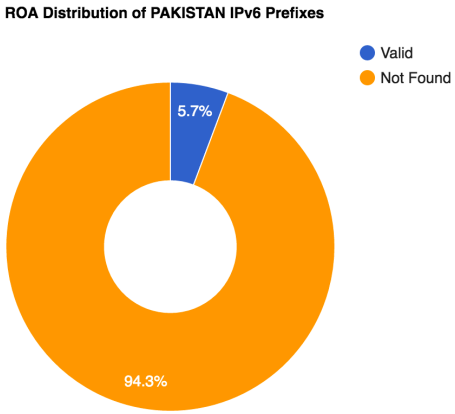

This graph generated on Wed 20 Jul 2016 23:30:36 AEST

ref link : http://rpki.apnictraining.net/output/pk.html

Sri Lanka

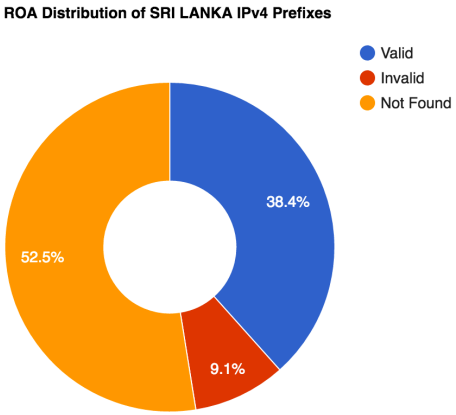

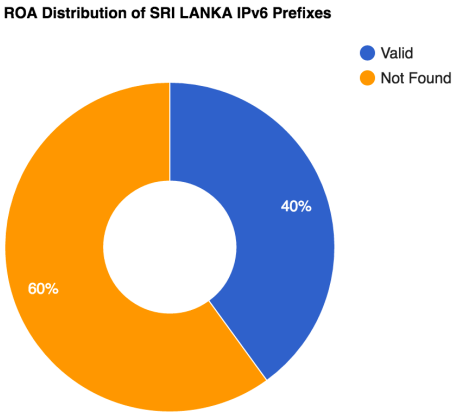

This graph generated on Wed 20 Jul 2016 23:42:25 AEST

ref link : http://rpki.apnictraining.net/output/lk.html

Summary

- RPKI adoption is growing
  - Most of the cases operators create ROA for min length and advertise longest prefix.
  - Some invalid ROA due to further allocation to customers.
- BGP operations and security
  - draft-ietf-opsec-bgp-security-07

Data Collection

- OpenBMP
  - https://github.com/OpenBMP/openbmp
- RPKI Dashboard
  - https://github.com/remydb/RPKI-Dashboard
- RIPE RPKI Statistics
  - https://lirportal.ripe.net/certification/content/static/statistics/world-roas.html
- RIPE Cache Validator API
  - http://rpki-validator.apnictraining.net:8080/export

Thank You

APNIC