hierarchical vs non-hierarchical as-sets

Aftab Siddiqui
siddiqui@isoc.org

The Problem

When setting up a BGP session, it’s important to decide on:

- The prefixes you’re willing to accept from your peer —whether it’s the full global routing table or a sub-set of it based on certain internal policies.
- Which prefixes are you authorized to advertise to them.

Because you can’t trust whether everyone is advertising their prefixes correctly, the IRR developed Routing Policy Specification Language (RPSL), which automatically verifies routing policies and announcements that are uploaded to the IRR

RPSL – RFC2622

RPSL was designed so that a view of the global routing policy can be contained in a single cooperatively maintained distributed database to improve the integrity of Internet's routing. RPSL is not designed to be a router configuration language.

RPSL is designed so that router configurations can be generated from the description of the policy for one autonomous system combined with the description of a router (**inet-rtr** class), mainly providing router ID, autonomous system number of the router, interfaces and peers of the router, and combined with a global database mappings from AS sets to ASes (**as-set** class), and from origin ASes and route sets to route prefixes (**route** and **route-set** classes).

The Problem…..cont..

The IRR (Internet Routing Registry) consists of several globally distributed routing information databases and can be put in to two distinct categories i.e. Authenticated and non-authenticated. Some notable non-authenticated IRR instances include RADb, ALTDB,

Whereas, those run by the Regional Internet Registries (RIRs) are all authenticated IRR instances. RIRs allow you to create objects defined by RPSL in their databases, which can be used to verify the information and can generate BGP filters and prefix lists using whois queries.

The Problem…..cont..

Each record in an IRR database contains — a set of attributes with corresponding values, which describe things such as people, organization, IP addresses, AS numbers, routing policy, and network/abuse contact information.

Object name	Description
as-set	It provides a mechanism for publicly documenting the relationship between Autonomous Systems (ASes).
aut-num	Contains details of the registered holder of an AS number and their routing policy for that AS.
inetnum/inet6num	Contains details of an allocation or assignment of IPv4/IPv6 address space.
irt (APNIC) abuse-c (other RIRs)	Incident Response Team. It’s used to provide information about the contact details of the abuse handling team.

AS-SET

An as-set provides a way to document the relationship between ASes which can then be publicly verified.

This object defines a group of ASNs that are peers in the routing network and through which traffic can be routed. The as-set members can include ASNs as well as the names of other as-sets .

AS-SET

The as-set attribute defines the name of the set. It is an RPSL name that starts with “as-“. The member’s attribute lists the members of the set. The members attribute is a list of AS numbers, or other as-set names.

Attribute	Value	Type
as-set	<object-name>	Mandatory, single-valued, class key
members	List of <as-numbers> or <as-set-names>	Optional, multi-valued
mbrs-by-ref	List of <mntner-names>	Optional, multi-valued

RFC 2622 Section 5.1

AS-SET

An as-set serves the purpose of describing which networks compose the so-called ‘customer cone’ of an AS you peer with. By adding a ‘member’ attribute pointing to either an ASN or to another as-set, an entity is indicating which prefixes should be accepted by their BGP Neighbors. Here are two examples from Australian networks:

```
% Information related to 'AS-IPTRANSIT'

as-set:          AS-IPTRANSIT
descr:           IP Transit
tech-c:          ITPL4-AP
admin-c:         ITPL4-AP
mnt-by:          MAINT-IPTRANSIT-AU
members:         AS64098
members:         AS134720
last-modified:   2017-05-10T17:50:45Z
source:          APNIC
```

Back to ”The Problem”

```
% Information related to 'AS-IPTRANSIT'

as-set:         AS-IPTRANSIT
descr:          IP Transit
tech-c:         ITPL4-AP
admin-c:        ITPL4-AP
mnt-by:         MAINT-IPTRANSIT-AU
members:        AS64098
members:        AS134720
last-modified:  2017-05-10T17:50:45Z
source:         APNIC
```

This as-set is created by the maintainer IP Transit as shown here.

```
% Information related to 'AS-AFTABSIDDIQUI'

as-set:         AS-AFTABSIDDIQUI
descr:          Aftab Siddiqui Testing AS-SET
tech-c:         ISAL1-AP
admin-c:        ISAL1-AP
mnt-by:         MAINT-ISAL-SG
last-modified:  2022-12-16T04:40:25Z
source:         APNIC
```

This as-set is created by the maintainer ISAL, which belongs to ISOC Asia Pacific and has no relationship with AS-AFTABSIDDIQUI

Back to ”The Problem”

How to name the AS-SET is also defined in the RFC2622 and has 2 categories Hierarchical and Non Hierarchical

The issue we saw in previous slide is that ISAL-SG technically can create this object because it’s permissible as per the RIR policy: it doesn’t contradict the standard outlined in the RFC for non-hierarchical objects as well.

Back to ”The Problem”

What is the Significance of Creating an as-set With a Bogus Name?

Network operators use filters based on the IRR and RPKI. While the valid RPKI ROA status of routes in the global routing table is growing but it is just above 40%, implementing IRR-based filtering is a great starting point.

Network operator who wants to build a prefix filter for AS58280 then I can simply use bgpq4 (A BGP Filter Generator). For example:

```
$ bgpq4 as58280
no ip prefix-list NN
ip prefix-list NN permit 45.129.224.0/22
```

Back to ”The Problem”

if I need to generate similar data for an operator with hundreds of downstream customers, then this method doesn’t scale and that’s why we use an as-set.

Instead of passing an ASN as the argument to bgpq4, I can pass as-set as an argument

$ bgpq4 -S APNIC AS-AFTABSIDDIQUI

no ip prefix-list NN

ip prefix-list NN permit 103.138.210.0/24

But remember, I didn’t create this AS-SET.

It worked because member “AS139038”

```
% Information related to 'AS-AFTABSIDDIQUI'

as-set:          AS-AFTABSIDDIQUI
descr:           Aftab Siddiqui Testing AS-SET
tech-c:          ISAL1-AP
admin-c:         ISAL1-AP
mnt-by:          MAINT-ISAL-SG
last-modified:   2023-02-28T23:47:26Z
members:         as139038
source:          APNIC
```

Back to ”The Problem”

But if ISAL-SG the maintainer of AS-AFTABSIDDIQUI just removes the AS139038 from the member attribute then this is what I will get.

```
$ bgpq4 -S APNIC AS-AFTABSIDDIQUI

ERROR:Key not found expanding !a4AS-AFTABSIDDIQUI

no ip prefix-list NN

! generated prefix-list NN is empty

ip prefix-list NN deny 0.0.0.0/0
```

Remember we are currently using “Authenticated IRR” i.e. APNIC and it is totally legal to this . There are several non-authenticated IRRs as well.

Back to ”The Problem”

By Maintainer (Amazon)	NOT by Maintainer
as-set: AS-AMAZON descr: Amazon ASNs members: AS-AMAZON-NA, AS-AMAZON-AP, AS-AMAZON-EU admin-c: AC6-ORG-ARIN tech-c: AC6-ORG-ARIN notify: noc@amazon.com mnt-by: MAINT-AS16509 changed: noc@amazon.com 20151027 #17:32:13Z source: RADB	as-set: AS-AMAZON tech-c: DUMY-RIPE admin-c: DUMY-RIPE mnt-by: KATERINA-MNT created: 2022-10-23T19:05:59Z last-modified: 2022-10-23T19:05:59Z source: RIPE

Amazon is maintaining their primary AS-SET in RADB (non-authenticated IRR) and someone in RIPE has created the same AS-SET. If I use RIPE as the source to generate prefix filters for Amazon assuming it’s the Authenticated source then I will get an empty prefix list.

Back to ”The Problem”

But if ISAL-SG the maintainer of AS-AFTABSIDDIQUI just removes the AS139038 from the member attribute then this is what I will get.

```
$ bgpq4 -S APNIC AS-AFTABSIDDIQUI

ERROR:Key not found expanding !a4AS-AFTABSIDDIQUI

no ip prefix-list NN

! generated prefix-list NN is empty

ip prefix-list NN deny 0.0.0.0/0
```

Remember we are currently using “Authenticated IRR” i.e. APNIC and it is totally legal to this . There are several non-authenticated IRRs as well.

AS-SET in APNIC

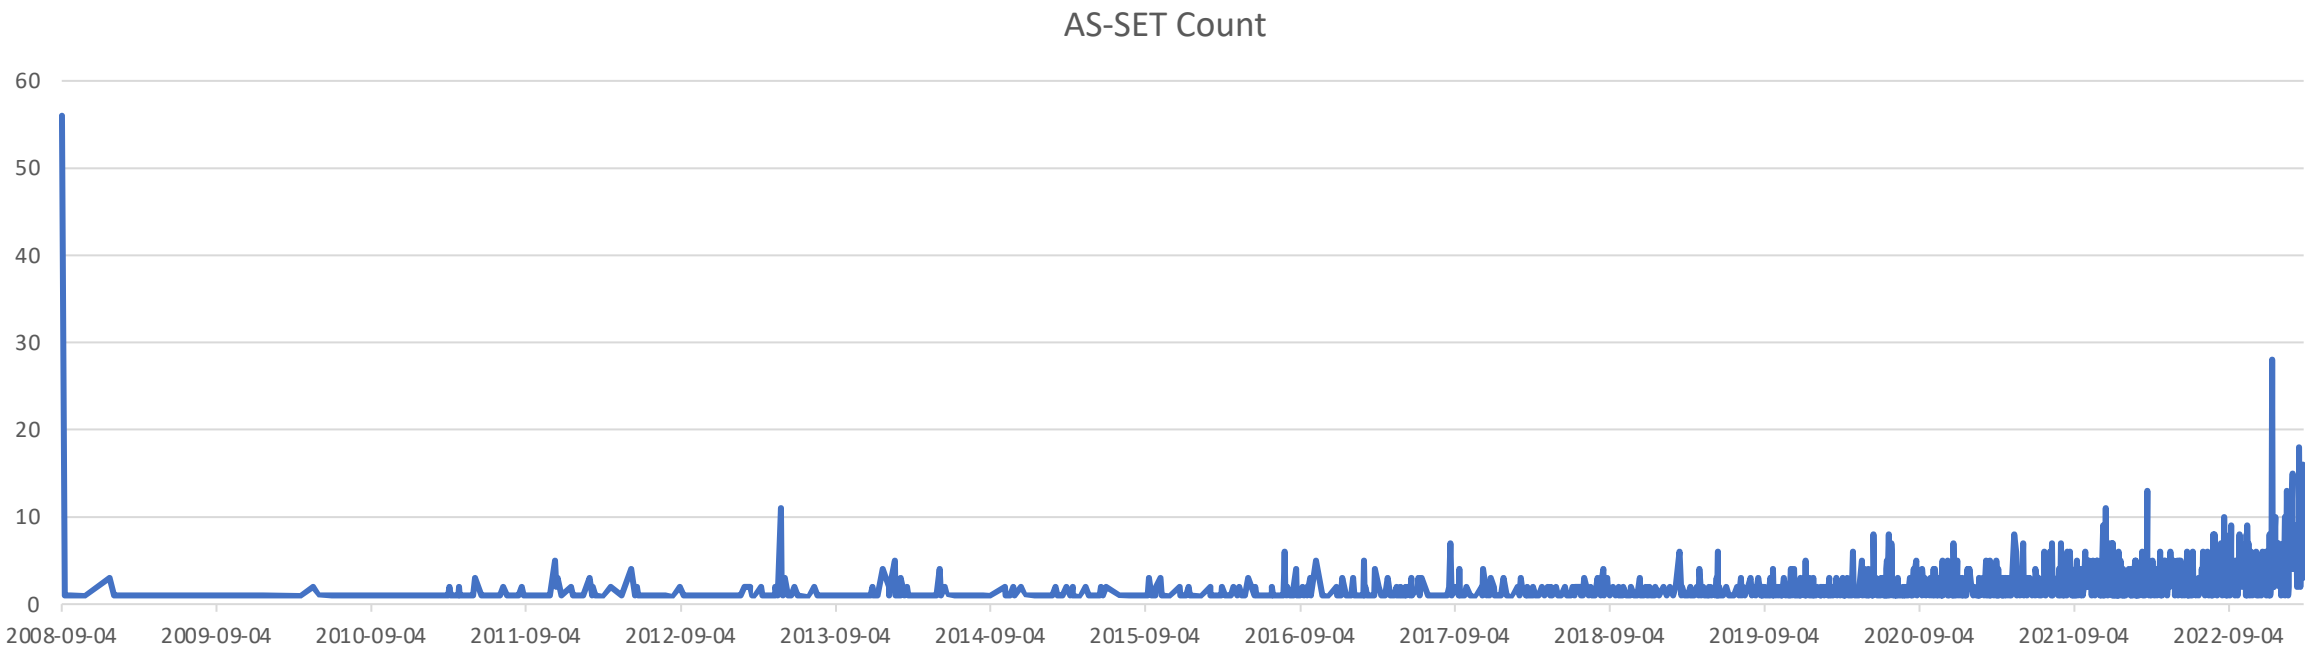

In last couple of years, network operators in APNIC region are creating AS-SET very consistently.

AS-SET in APNIC – non-hierarchical

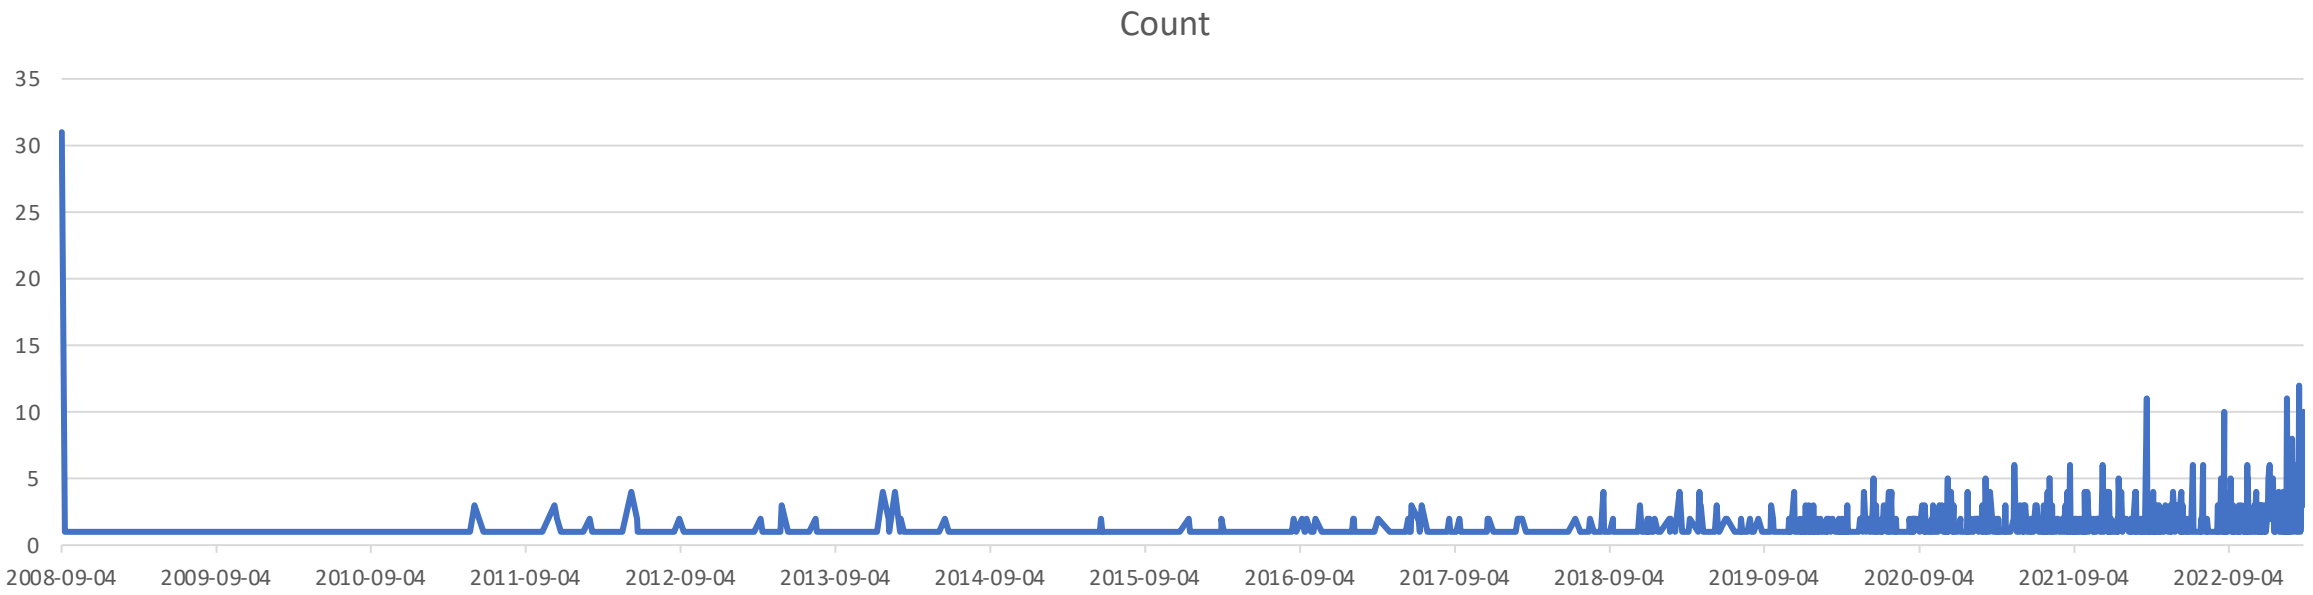

These are the numbers of non-hierarchical as-set in APNIC whois database

AS-SET in APNIC – non-hierarchical

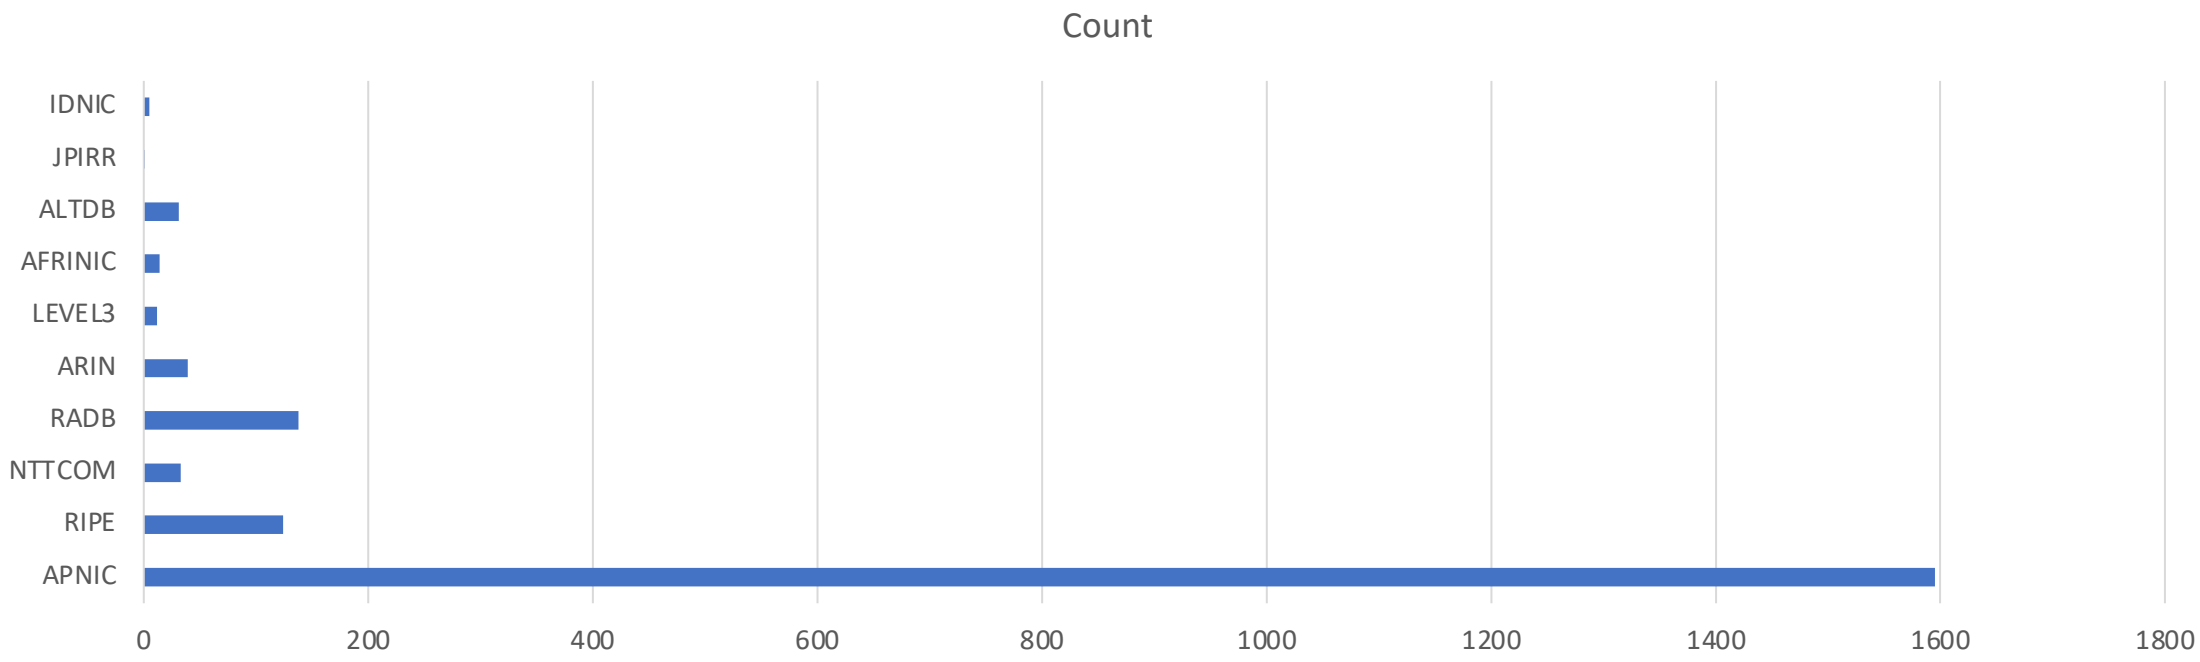

APNIC non-hierarchical as-set which also exist in other IRR databases

Stop non-hierarchical as-set in APNIC

Proposal 151: This proposal would restrict APNIC account holders from creating a non-hierarchical as-set, and notify all Members who already have non-hierarchical as-set that it is recommended they move to a hierarchical as-set.

https://www.apnic.net/community/policy/proposals/prop-151/