APNIC 50

Prop132 deployment report

A policy implementation report for
Prop132 “AS0 for unallocated and unassigned resources”

George Michaelson

Product Manager: Registry

ggm@apnic.net

Virtual Policy Sig update!

Prop132 deployment report

A policy implementation report for
Prop132 “AS0 for unallocated and unassigned resources”

George Michaelson

Product Manager: Registry

ggm@apnic.net

ONLINE
8 - 10 September 2020

Prop132 “AS0 for unallocated and unassigned resources”

- We have implemented Prop132
  - APNIC now publishes and maintains an AS0 “ROA” for all un-delegated resources in our registry
  - These are the IPv4 and IPv6 resources listed as “available” or “reserved” in our daily published delegated statistics files
  - The AS0 ROA is defined in RFC6483 as “a disavowal of routing origination”

```
A ROA with a subject of AS 0 (AS 0 ROA) is an attestation by the holder of a
prefix that the prefix described in the ROA, and any more specific prefix,
should not be used in a routing context.
```

Deployment status

- This is now a fully deployed service
  - With systems monitoring 24/7 integrated into our operations platforms
  - Deployed in the cloud for the publication point (data repository)
  - At this stage, deployed in a stand-alone Trust Anchor Locator (TAL)

Grafana Status checks

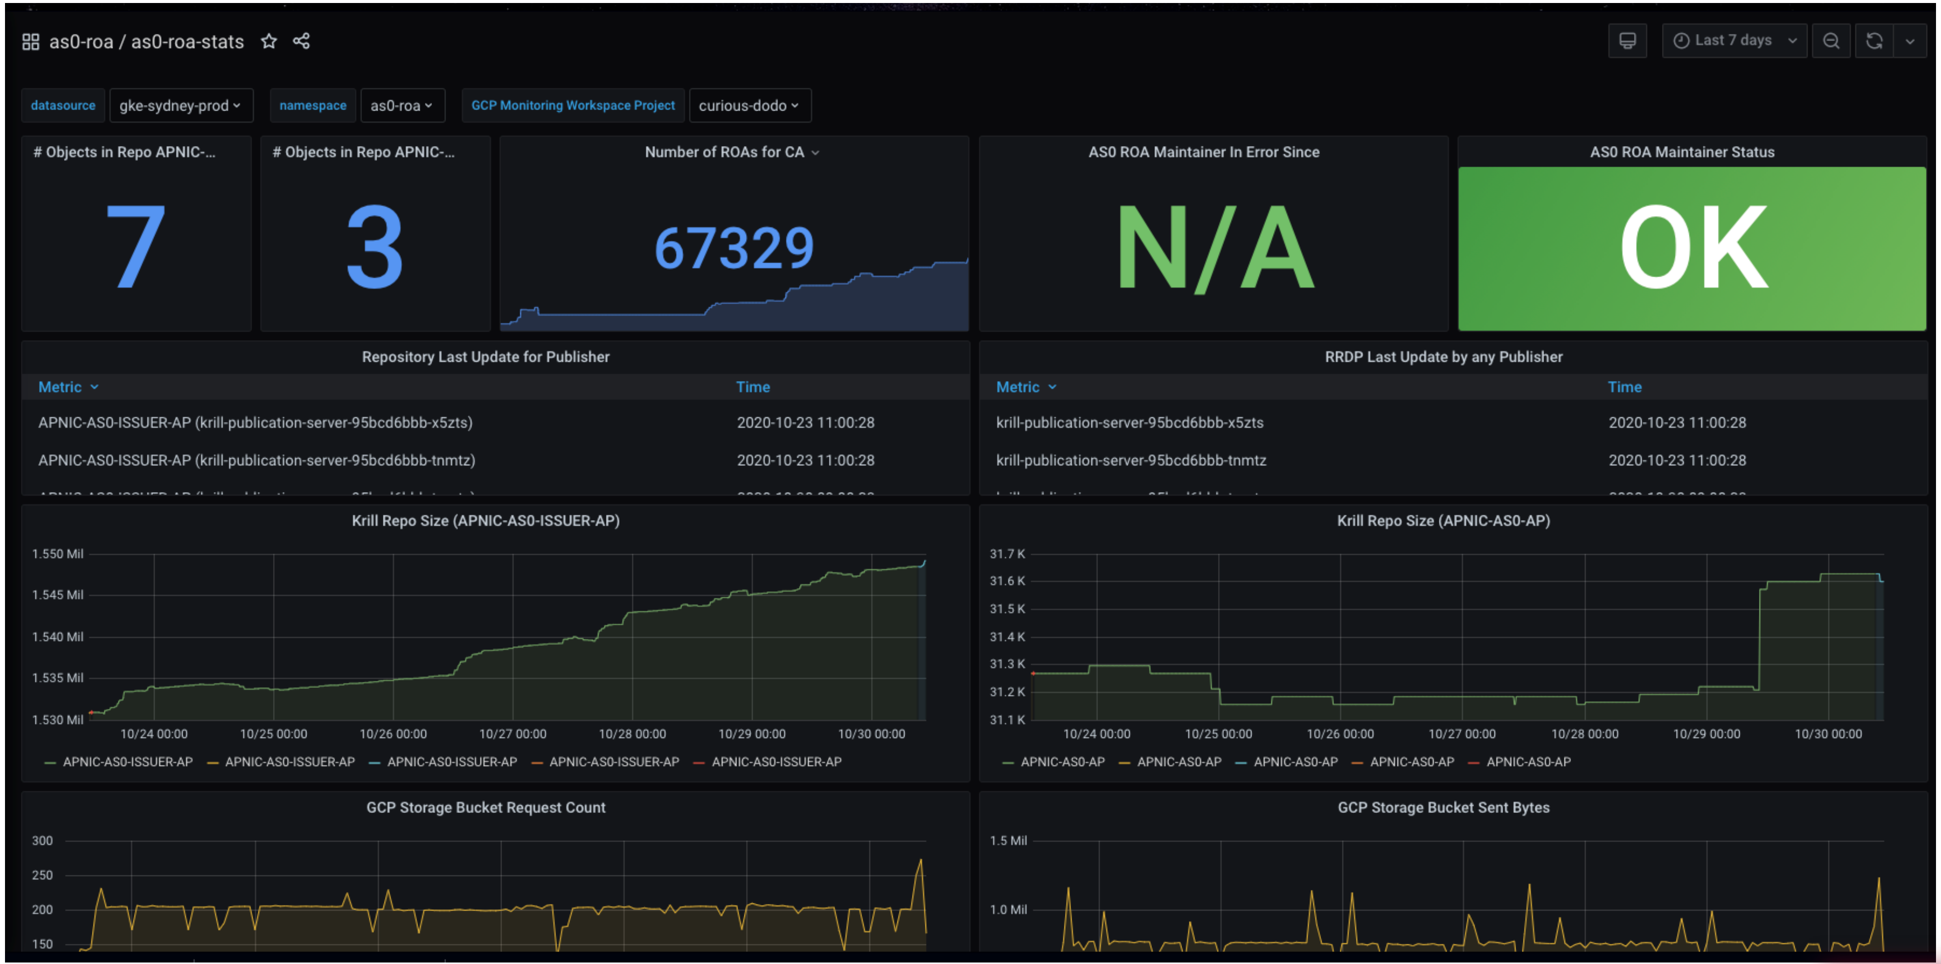

as0-roa / as0-roa-stats

datasource: gke-sydney-prod
namespace: as0-roa
GCP Monitoring Workspace Project: curious-dodo
Last 7 days

Objects in Repo APNIC-...: 7

Objects in Repo APNIC-...: 3

Number of ROAs for CA: 67329

AS0 ROA Maintainer In Error Since: N/A

AS0 ROA Maintainer Status: OK

Repository Last Update for Publisher

Metric	Time
APNIC-AS0-ISSUER-AP (krill-publication-server-95bcd6bbb-x5zts)	2020-10-23 11:00:28
APNIC-AS0-ISSUER-AP (krill-publication-server-95bcd6bbb-tnmtz)	2020-10-23 11:00:28

RRDP Last Update by any Publisher

Metric	Time
krill-publication-server-95bcd6bbb-x5zts	2020-10-23 11:00:28
krill-publication-server-95bcd6bbb-tnmtz	2020-10-23 11:00:28

Krill Repo Size (APNIC-AS0-ISSUER-AP)

Krill Repo Size (APNIC-AS0-AP)

GCP Storage Bucket Request Count

GCP Storage Bucket Sent Bytes

Implementation report:testbed

- An initial Testbed was deployed for APRICOT/APNIC49
  - Based on the “Krill” system from NLNet Labs
  - Operating on the delegated files as a daily view of registry
  - Using a temporary, soft-keyed Trust Anchor (TA) in a TAL file
  - Publishing the repository inside APNIC VM on the test network
  - This service was used by a small number of people (<10)
    - We were able to confirm issues with discrete ROA per prefix
    - We understood our operational needs to manage the ROA as resources are issued by APNIC

How we took testbed to production

- We have now deployed this service into production
  - Still based on delegated files, but with a delay to prevent accidental exclusions if delegated files are out of synchronization with registry
  - Live updates to Registry (delegations) are applied within 5 minutes to both main RPKI and AS0 RPKI state
    - Delegations are removed from the AS0 ROA within 5 minutes of resources being assigned or allocated from the free pool.
  - We are collecting statistics on use, and the scale of BGP effects which will be presented to the Routing Security SIG

Implementation report: Production

- In-house deployment on VM under operations monitoring
  - HSM backed trust anchor keypair
    - Same level of assurance as main line TA
- Cloud deployment of repository (GCP/GKE)
  - Both rsync and RRDP supported
  - Will distribute in GCP
    - When 2nd and further nodes commissioned
    - When the main RPKI RRDP/rsync service is distributed

Implementation architecture

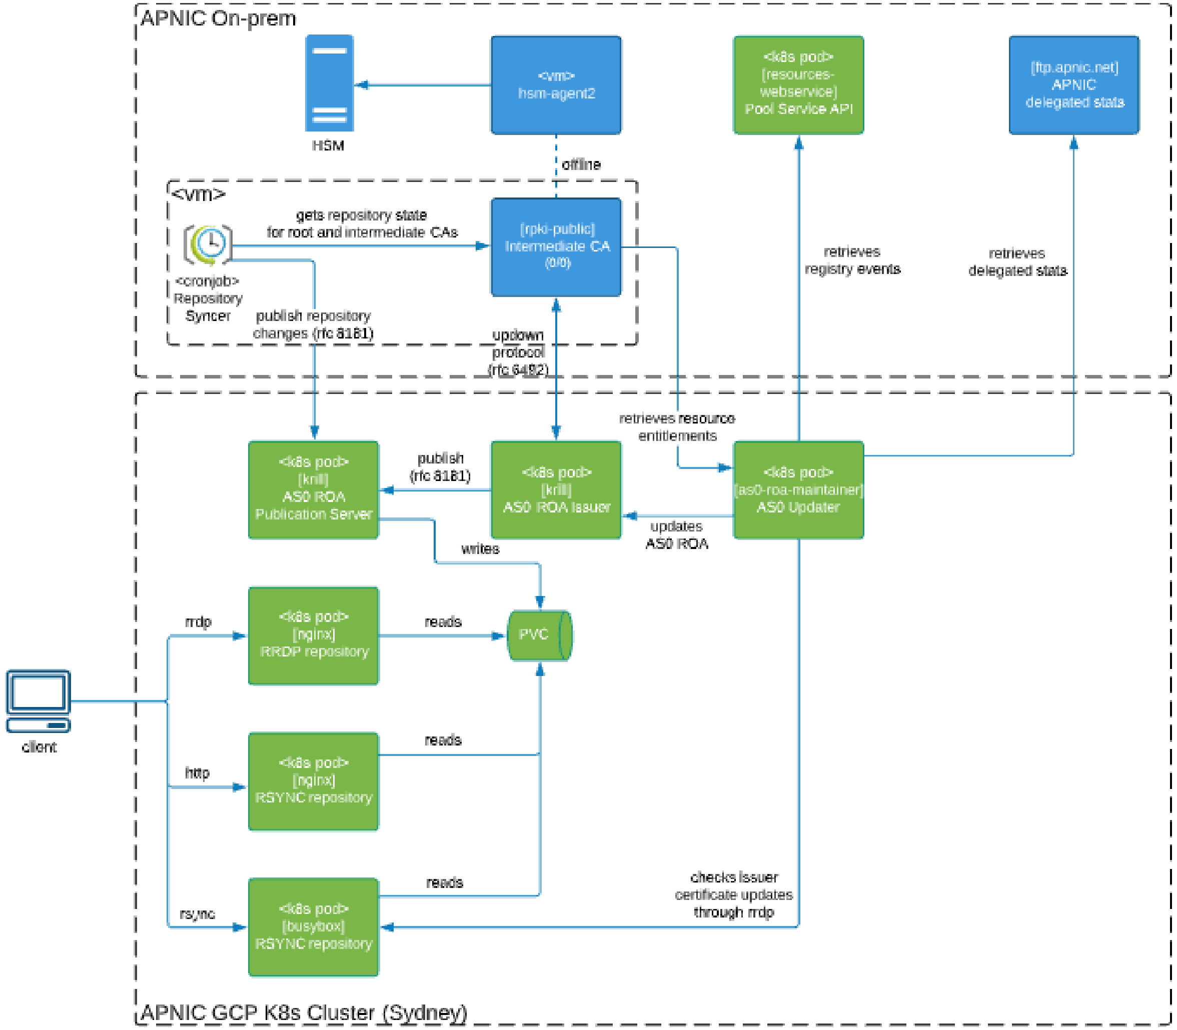

- On-premises and GKE Sydney deployments
- HSM backed TAL, follows main line RPKI
- Re-use of existing RPKI systems code
  - Actual signing carried out by Krill (NLNet)

- Repository structure served from GKE
  - Capable of being distributed in future
  - Using CloudFlare front-end

Where to from here?

- Further discussion of this service is now conducted in the APNIC Routing Security SIG
  - Statistics on use,
  - Size of ROA,
  - Operational experiences,
  - Future directions.
- Initial outcome: 69 routes marked bad in DFZ from ~65k prefixes
  - (reported by Job Snijders during deployment testing)

Some initial statistics

- Initial outcome: 69 routes marked bad in DFZ from ~65k prefixes
  - (reported by Job Snijders during deployment testing)
- Usage: Released week of 1st September
  - 35 ASN now fetching from the service (as of end October)

ASN fetching the AS0 TAL (september)

ASN	Name	Economy
9443	VOCUS-RETAIL-AU Vocus Retail	AU
38345	ZDNS Internet Domain Name System Beijing	CN
4837	CHINA169-BACKBONE CHINA UNICOM	CN
4812	CHINANET-SH-AP China Telecom (Group)	CN
4847	CNIX-AP China Networks Inter-Exchange	CN
17621	CNCGROUP-SH China Unicom Shanghai network	CN
23910	CNGI-CERNET2-AS-AP China Next Generation Internet	CN
4134	CHINANET-BACKBONE No.31	CN
1136	KPN KPN National	EU

ASN	Name	Economy
3265	XS4ALL-NL Amsterdam	NL
8587	INFRACOM-AS	NL
15169	GOOGLE	US
20473	AS-CHOOPA	US
395747	CLOUDFLARENET-SFO05	US
8075	MICROSOFT-CORP-MSN-AS-BLOCK	US
14618	AMAZON-AES	US
14061	DIGITALOCEAN-ASN	US
132892	CLOUDFLARE Cloudflare	US

ASN fetching the AS0 TAL (October)

ASN	AS Name	CC
38195	SUPERLOOP-AS-AP	AU
4764	Aussie Broadband	AU
7545	TPG Telecom Limited	AU
51852	PLI-AS	CH
17621	China Unicom Shanghai network	CN
38345	ZDNS	CN
4134	CHINANET-BACKBONE	CN
4837	CHINA UNICOM Backbone	CN
4847	China Networks Inter-Exchange	CN
1136	KPN National	EU
12876	SAS	FR
16276	OVH	FR

ASN	AS Name	CC
9009	M247	GB
206238	FREEDOMINTERNET	NL
3265	XS4ALL-NL	NL
8587	INFRACOM-AS	NL
34665	PINDC-AS	RU
132892	Cloudflare	US
63949	Linode	US
13649	ASN-VINS	US
14061	DIGITALOCEAN-ASN	US
14618	AMAZON-AES	US
15169	GOOGLE	US
16509	AMAZON-02	US

ASN	AS Name	CC
16628	DEDICATED-FIBER-COMMUNICATIONS	US
174	COGENT-174	US
20473	AS-CHOOPA	US
209	CENTURYLINK-US-LEGACY-QWEST	US
36351	SOFTLAYER	US
394474	WHITELABELCOLO393	US
53667	PONYNET	US
54538	PAN0001	US
62874	WEB2OBJECTS	US
8075	MICROSOFT-CORP-MSN-AS-BLOCK	US

Some initial statistics: Size of ROA

- AS0 ROA is ~1mb (at present)
- 64k++ IPv4 and IPv6 prefixes encoded in one Object
  - Approx 1k IPv4
  - 64k IPv6
    - The IPv6 count is a function of “sparse” allocation
  - Size varies with allocation/assignment and returns

Recommendations on use

- We do not recommend integration in the BGP ROV process directly
  - We ask Relying Party (Validator) code developers not to integrate it into the code, it should be hand configured.
- We do recommend use of the AS0 TAL as a check against BGP state, but not directly integrated into BGP
  - Look at AS0 TAL rejected routes, understand why.

What should I do (ISP or IXP operator)?

- You don’t **need** to do anything.
  - this is an advisory/information service only.
- If you want to understand leakage of AS0 tagged ‘unused’ resources
  - run a validator which includes the TAL for AS0
  - look at what kinds of BGP announcements would be filtered,
    - especially ones which originate in your downstreams, or within the IX community

Questions?