IRR Hygiene in the RPKI Era

Ben Du

GAIA Workshop

13 Oct 2021

BGP is Vulnerable to Prefix Origin Hijacking

- BGP provides no mechanism to prevent unauthorized re-routing of Internet traffic
- To secure BGP, researchers and operators have developed routing information databases so networks can verify BGP prefix origin information
- The Internet Routing Registry (IRR) was deployed in 1995
- Resource Public Key Infrastructure (RPKI) was deployed in 2012
- Both databases are still in use

IRR and RPKI may disagree

- Networks can choose to use either IRR or RPKI to implement route filtering.
- Networks need accurate routing information to properly implement routing filters
- IRR information can be inaccurate due to improper hygiene
- RPKI may have misconfigurations
- Such disagreement may introduce vulnerabilities – incorrect route filtering

Research Questions

- How much inconsistency is there between the information in IRR and RPKI?
- What are the causes of such inconsistency?
  - How many ASes are causing inconsistency?
  - Are ASes contributing more consistent or inconsistent information?
- Do ASes participating in routing security intiatives have more consistency maintaining IRR and RPKI records?
  - Mutually Agreed Norms for Routing Security (MANRS)

Background – IRR and RPKI

IRR

- The IRR is a distributed Internet routing information database. Networks voluntarily register their routing policies in one or more IRR databases.
- Inside a *route* object, the *route* and *origin* fields represent the IP prefixes of a network and the AS numbers that originates them in BGP

RPKI

- RPKI holds cryptographically attested routing information. The five RIRs are the trust anchors, who distribute RPKI certificates and Route Origin Authorization (ROA) objects to their members.
- Inside a ROA, the *IP prefix*, *ASN*, and *Max Length* fields are used to check
  1. If a BGP prefix's origin match the *ASN* in a matching ROA
  2. If the prefix length is less than the *Max Length* value

Dataset

- IRR dataset: Routing Assets Database (RADB) archive, August 2016 – October 2021 monthly snapshots. Only route objects are used.
- RPKI dataset: Validated ROA archive from the RIPE RPKI Validator, August 2016 – October 2021 monthly snapshots.

Research Questions

- How much inconsistency is there between the information in IRR and RPKI?
- What are the causes of such inconsistency?
    - How many ASes are causing inconsistency?
    - Are ASes contributing more consistent or inconsistent information?
- Do ASes participating in routing security intiatives have more consistency maintaining IRR and RPKI records?
    - Mutually Agreed Norms for Routing Security (MANRS)

IRR and RPKI Inconsistency – Record Classification

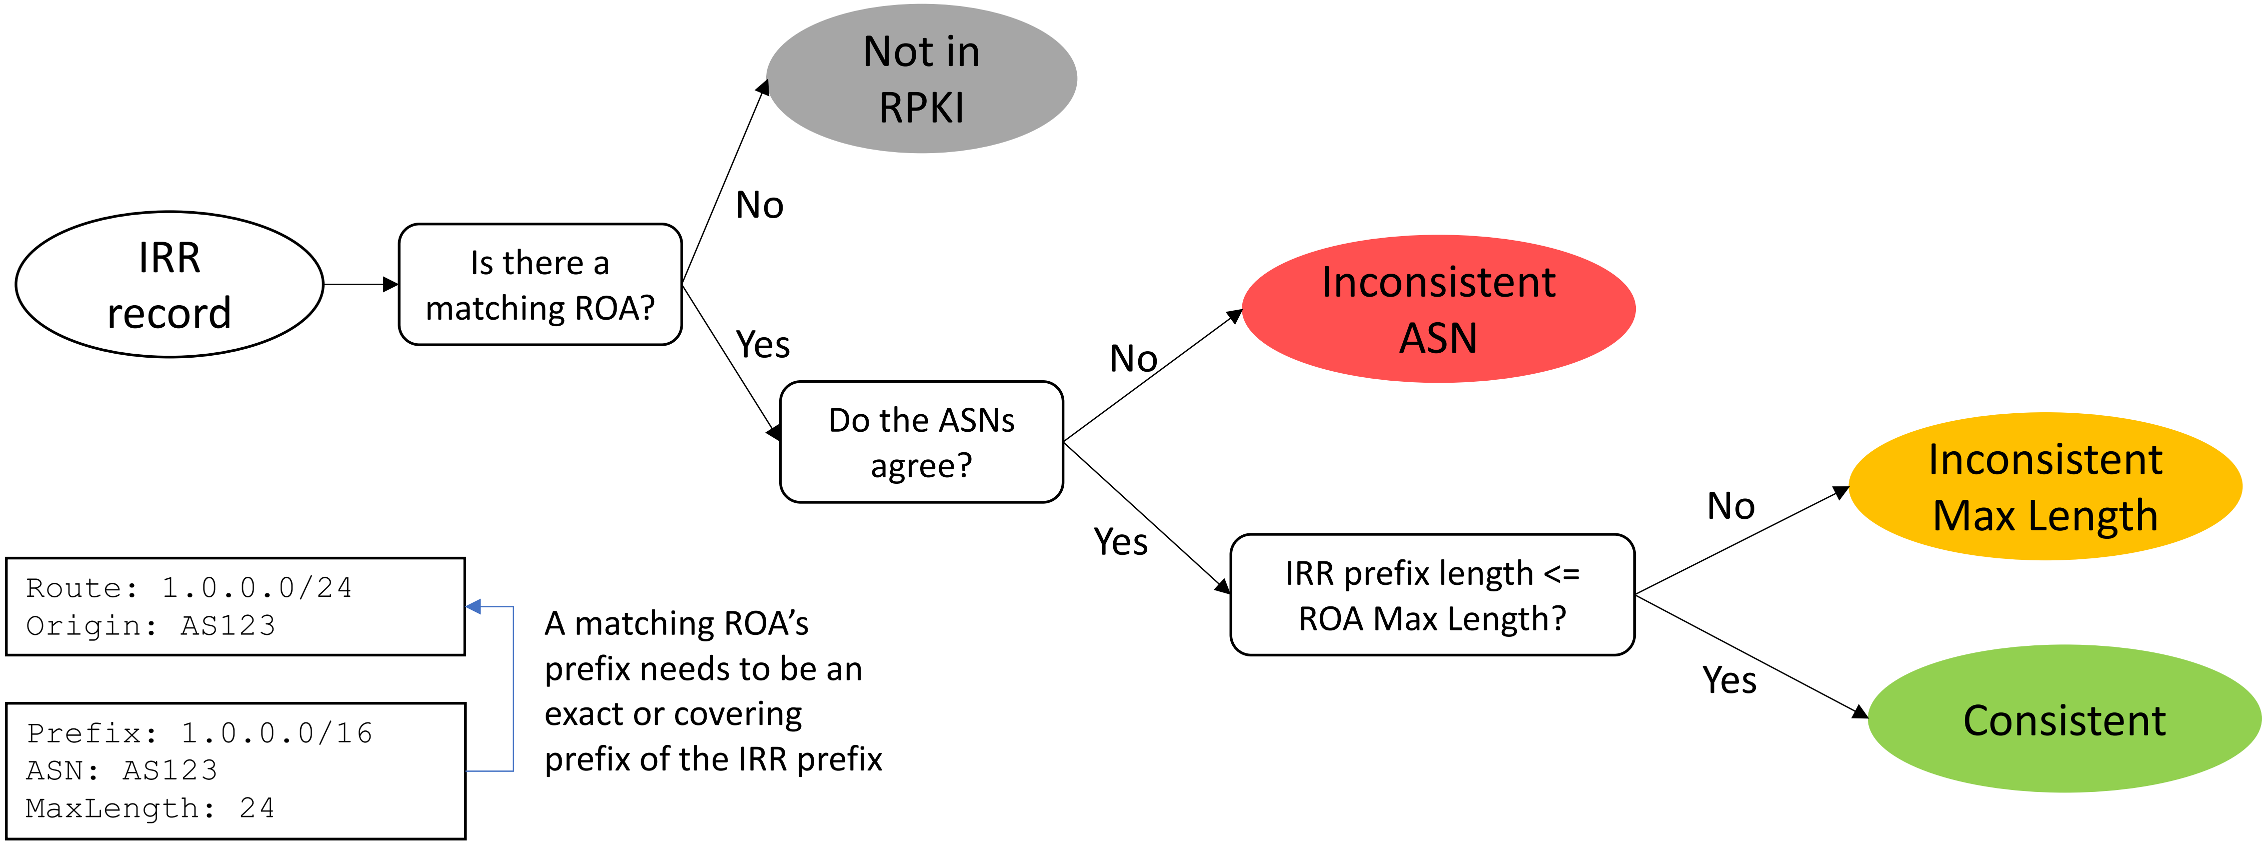

```
Route: 1.0.0.0/24
Origin: AS123
```

```
Prefix: 1.0.0.0/16
ASN: AS123
MaxLength: 24
```

A matching ROA's prefix needs to be an exact or covering prefix of the IRR prefix

IRR and RPKI Inconsistency

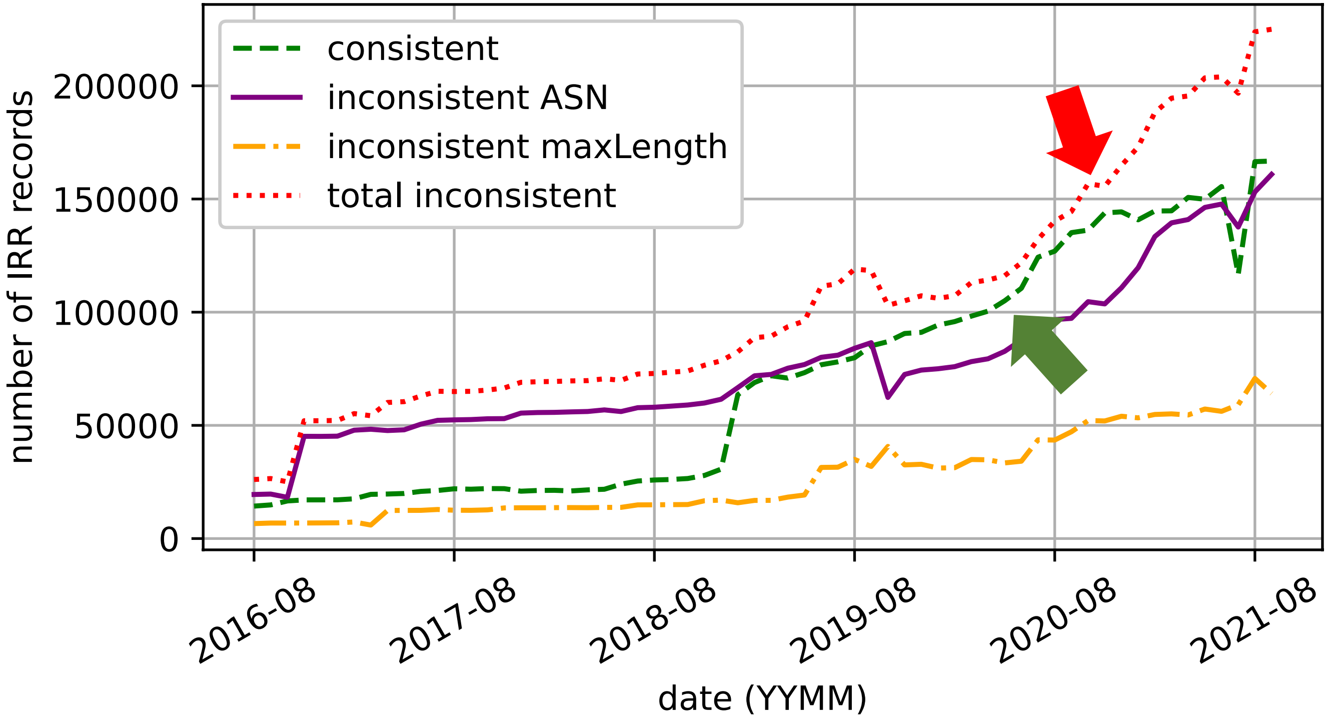

- Red line = purple line + yellow line
- There has always been more inconsistent records than consistent ones

IRR and RPKI Inconsistency - Verisign

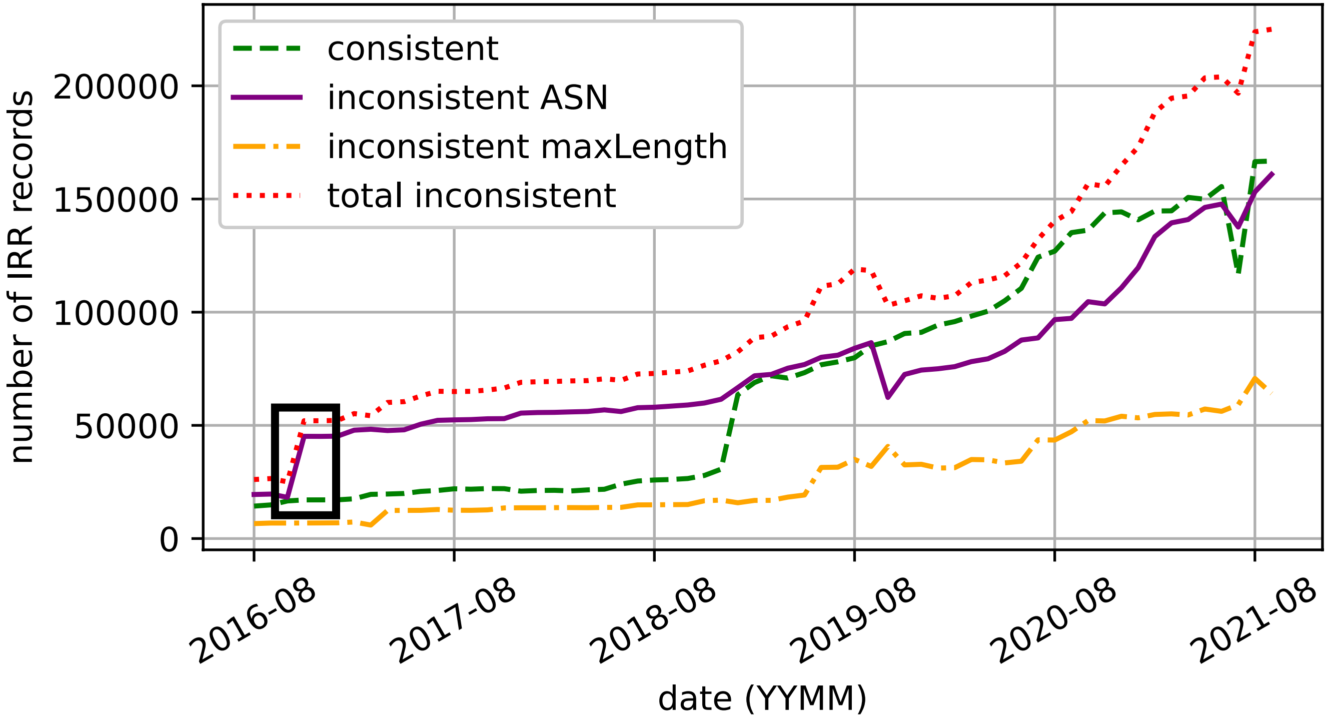

- Uptick in Oct 2016
- Verisign customers registered their prefixes and ASes in RPKI
- In IRR, the prefixes are under Verisign AS, but labeled as customer route, causing inconsistency
- +26,647 inconsistent IRR records

IRR and RPKI Inconsistency - Verisign

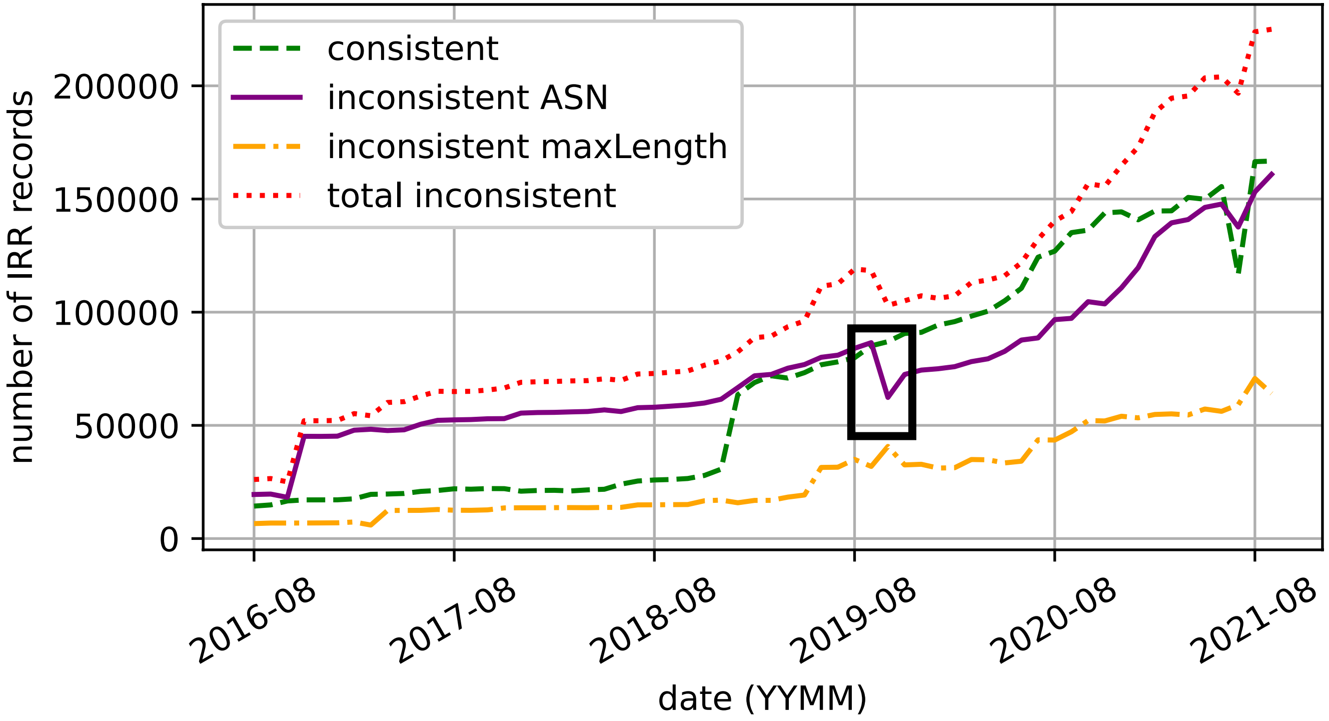

- Downtick in Sept. 2019
- Verisign deleted inconsistent records from RADB
- -26,682 inconsistent IRR records

IRR and RPKI Inconsistency - TWNIC

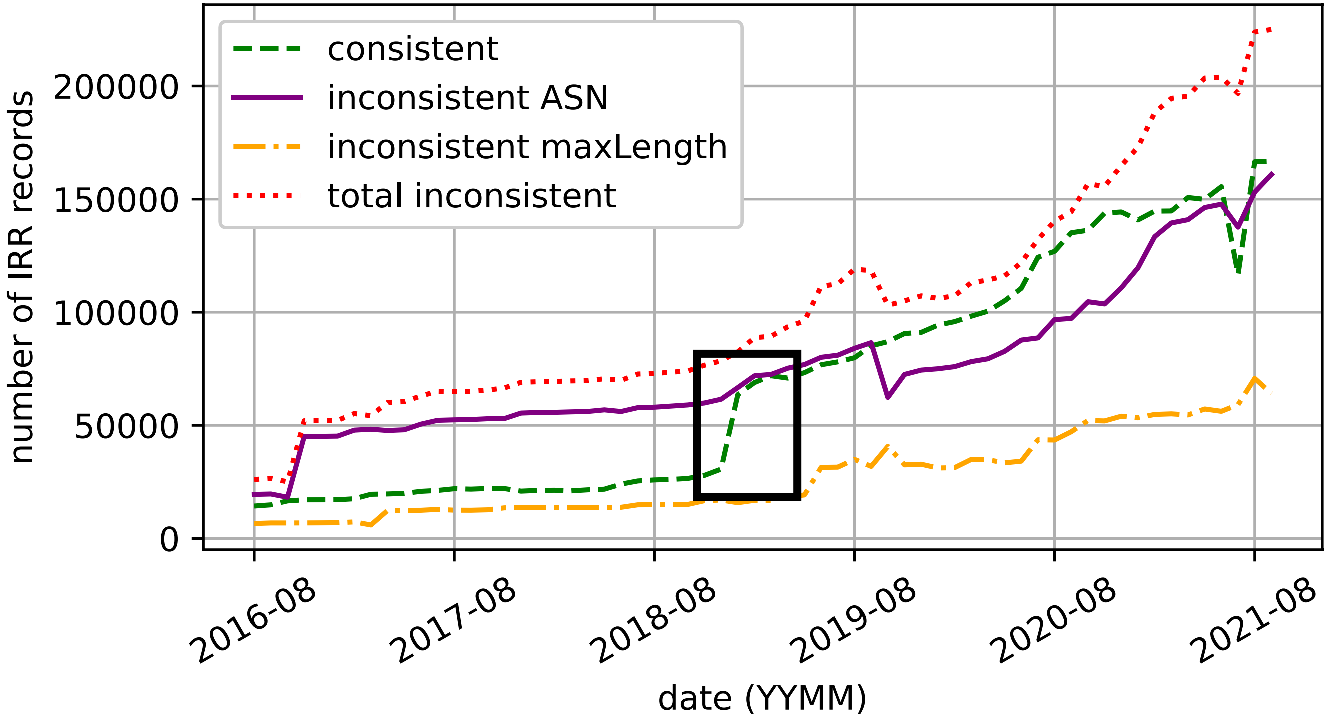

- Green line uptick in Nov. 2018
- TWNIC ASes bulk registered their prefixes in RPKI
- +34,430 consistent IRR records

IRR and RPKI Inconsistency - TWNIC

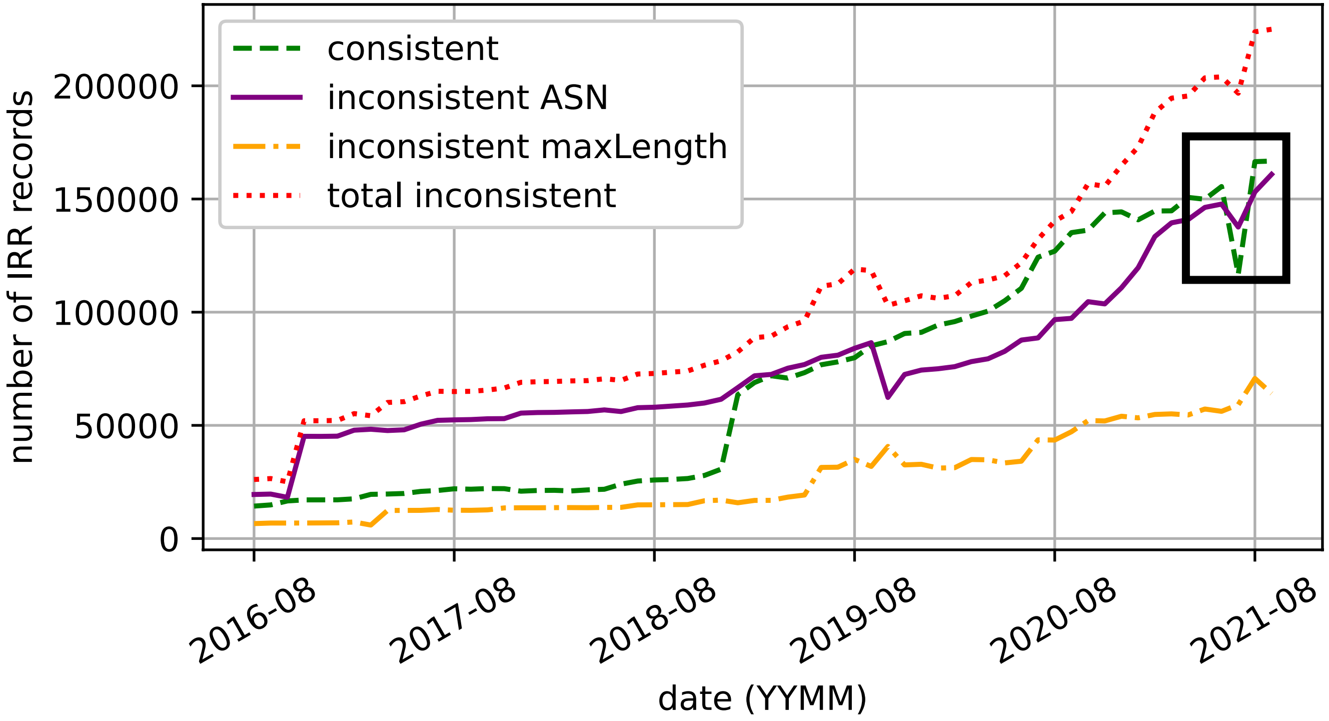

- Green line valley in July 2021
- TWNIC ASes temporarily removed their ROAs from RPKI
- 33,216 IRR records impacted

Research Questions

- How much inconsistency is there between the information in IRR and RPKI?
- What are the causes of such inconsistency?
  - How many ASes are causing inconsistency?
  - Are ASes contributing more consistent or inconsistent information?
- Do ASes participating in routing security intiatives have more consistency maintaining IRR and RPKI records?
  - Mutually Agreed Norms for Routing Security (MANRS)

ASes Behind the Inconsistency

- We classify the ASes based on their record classification
- E.g. An AS in category 5 have IRR records classified as consistent and inconsistent ASN

	Consistent	Inconsistent ASN	Inconsistent MaxLength
1. Entirely consistent	✓	✗	✗
2. Entirely inconsistent ASN	✗	✓	✗
3. Entirely inconsistent ML	✗	✗	✓
4. Entirely inconsistent, both	✗	✓	✓
5. Mixed, consistent + inconsistent ASN	✓	✓	✗
6. Mixed, consistent + inconsistent ML	✓	✗	✓
7. Mixed, all 3	✓	✓	✓

ASes Behind the Inconsistency

- Take the Oct 2021 snapshots in the IRR and RPKI datasets
- There are more ASes that keep their entirety of IRR records consistent with RPKI than those leave all their records inconsistent with RPKI

	Number of ASes	IRR Records of ASes
1. Entirely consistent	4749	32799
2. Entirely inconsistent ASN	3728	51900
3. Entirely inconsistent ML	26	47
4. Entirely inconsistent, both	9	106
5. Mixed, consistent + inconsistent ASN	1459	122849
6. Mixed, consistent + inconsistent ML	341	12112
7. Mixed, all 3	387	171924

ASes Behind the Inconsistency

- However entirely consistent ASes have fewer IRR records
- Some ASes in Category 2 have 5000+ IRR records
- The largest AS in Category 1 have only ~1200 IRR records

	Number of ASes	IRR Records of ASes
1. Entirely consistent	4749	32799
2. Entirely inconsistent ASN	3728	51900
3. Entirely inconsistent ML	26	47
4. Entirely inconsistent, both	9	106
5. Mixed, consistent + inconsistent ASN	1459	122849
6. Mixed, consistent + inconsistent ML	341	12112
7. Mixed, all 3	387	171924

Research Questions

- How much inconsistency is there between the information in IRR and RPKI?
- What are the causes of such inconsistency?
    - How many ASes are causing inconsistency?
    - Are ASes contributing more consistent or inconsistent information?
- Do ASes participating in routing security intiatives have more consistency maintaining IRR and RPKI records?
    - Mutually Agreed Norms for Routing Security (MANRS)

ASes with Records of Mixed Consistency

- 744 ASes in Category 5 have more consistent records than inconsistent ones
- 195 in the Category 6
- Overall the mixed AS keep more consistent records than inconsistent ones
  - Green = consistent
  - Red = inconsistent ASN
  - Yellow = inconsistent Max Length

	Number of ASes	IRR Records of ASes
5. Mixed, consistent + inconsistent ASN	1459	70911 + 51938
6. Mixed, consistent + inconsistent ML	341	8401 + 3711
7. Mixed, all 3	387	55519 + 56069 + 60036

Research Questions

- How much inconsistency is there between the information in IRR and RPKI?
- What are the causes of such inconsistency?
  - How many ASes are causing inconsistency?
  - Are ASes contributing more consistent or inconsistent information?
- Do ASes participating in routing security intiatives have more consistency maintaining IRR and RPKI records?
  - Mutually Agreed Norms for Routing Security (MANRS)

MANRS ASes

- MANRS ASes are required to register in either in IRR or RPKI.
- In general, the majority of MANRS ASes that registered both in the IRR and RPKI keep all their records consistent

	MANRS ASes (741)	Records
1. Entirely consistent	171	3497
2. Entirely inconsistent ASN	39	1055
3. Entirely inconsistent ML	0	0
4. Entirely inconsistent, both	0	0
5. Mixed, consistent + inconsistent ASN	53	480 = 404 + 76
6. Mixed, consistent + inconsistent ML	13	2336 = 1121 + 1215
7. Mixed, all 3	36	21023 = 5466 + 8176 + 7381
Has IRR but No RPKI	429	-

MANRS ASes vs All ASes

- Compared to all ASes, a greater fraction of MANRS ASes practice good IRR Hygiene

	All ASes in IRR and RPKI (11709)	**MANRS ASes in IRR and RPKI (312)**
1. Entirely consistent	4749 (49.18%)	171 (54.81%)
2. Entirely inconsistent ASN	3728 (31.84%)	39 (12.5%)
3. Entirely inconsistent ML	26 (0.22%)	0
4. Entirely inconsistent, both	9 (0.08%)	0
5. Mixed, consistent + inconsistent ASN	1459 (12.46%)	53 (16.99%)
6. Mixed, consistent + inconsistent ML	341 (2.91%)	13 (4.17%)
7. Mixed, all 3	387 (3.31%)	36 (11.54)

Summary

1. There is significant disagreement between IRR and RPKI. Overall there are more IRR records that are inconsistent RPKI than ones that are consistent
2. More ASes keep all their IRR records consistent with RPKI than ones that leave all their records inconsistent. However, such ASes with good hygiene own fewer prefixes
3. ASes that participate in routing security initiatives are more likely to have good IRR hygiene

Future Work

- How does inconsistency between IRR and RPKI impact networks? Is there harm?
- Are there malicious actors who registered false records in the IRR?

Questions?

bendu@ucsd.edu