

# A BCP for the use of IRR DBs by IXP Route Servers



S. KONSTANTARAS – AMS-IX

M. D'ITRI - MINAP

W. van GULIK – ROMANDIX

A. DINU – INTERLAN IX

K. BLUMBERG – TORIX



ROIX (ROMANDIX)

RIPE 88 2024  
Krakow – PL

# BCP proposal

- IXP operator must trust and use only the following IRR databases for building and maintaining Route Server filters:

- AFRINIC
- ARIN
- RIPENCC
- LACNIC
  - NIC.MX
  - NIC.BR

- APNIC
  - APJJ
  - CNNIC
  - JPNIC
  - KRNIC
  - TWNIC
  - VNNIC
  - IRINN

# Grace Period 1/2

- The authors understand that the adoption of this policy will probably result in a massive transfer of RPSL objects from non-supported databases to the supported ones.
- Therefore, the policy introduces a grace period of 12 months in which the list of allowed databases is supplemented by these:
  - RADB
  - RIPE-NONAUTH
  - NTT
  - LEVEL3

# Grace Period 2/2

- At the end of the grace period, IXP operators must stop supporting these additional IRRs and operate filter generation tools that will only query the 5 RIR databases.
- During the grade period, IXPs will make multiple best effort attempts to warn their members relying on these IRRs about the need to replicate their policies in the appropriate RIR IRR.

# What's the impact if we apply this BCP now?

- We know the theoretical impact
  - Valid aut-num/route/route6 objects need to move back to their official DBs
  - However, Global-Ops might get affected with ARIN legacy space
- Can we measure the impact on ARIN Legacy space?
  - How many prefixes are affected?
  - What type of prefixes might get lost?

# Practical Analysis

- Marco wrote a script
  - You feed the script with your master table (or multiple BIRD tables)
  - You feed the script with ARIN's legacy space from ARIN's FTP server
  - Execute the script & get the results in a txt file

# Analysis – Prefix amounts

- Discovered around 5880 prefixes
  - Represents around 1.9% of AMS-IX full table
- InterLAN 4.3K prefixes (around 3.5% of their full table)
- MINAP 3.9K prefixes
- DE-CIX after a miss-calculation decided to do a deeper analysis and contribute results

# Analysis - Source

**Amount of Prefixes**

**Next Hop**

**Customer**

**Type**

**3331**

**80.249.209.150**

-

**Local**

188

80.249.211.6

-

Remote

170

80.249.208.124

-

Remote

103

80.249.212.8

-

Remote

97

80.249.213.7

-

Remote

85

80.249.211.191

-

Local

83

80.249.214.93

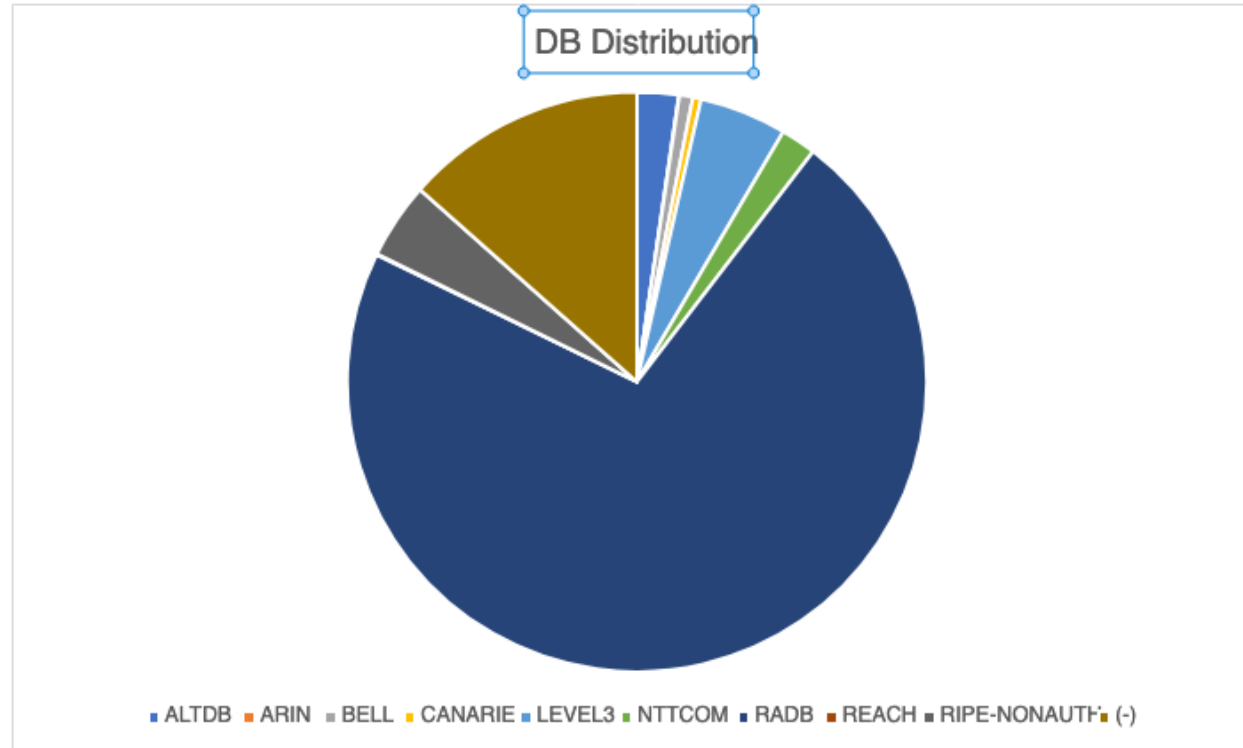
-

Remote



# Analysis – DB Distribution

DB	Number
ALTDB	136
ARIN	2
BELL	44
CANARIE	27
LEVEL3	287
NTTCOM	115
<b>RADB</b>	<b>4223</b>
REACH	3
RIPE-NONAUTH	250
(-)	792



# Analysis - Validity

Number of Prefixes

2090

IRR Valid

725

IRR Invalid

1365

RPKI Unknown

2080

RPKI Invalid

0

RPKI Valid

10

IRR VALIDATION



■ IRR Valid ■ IRR Invalid

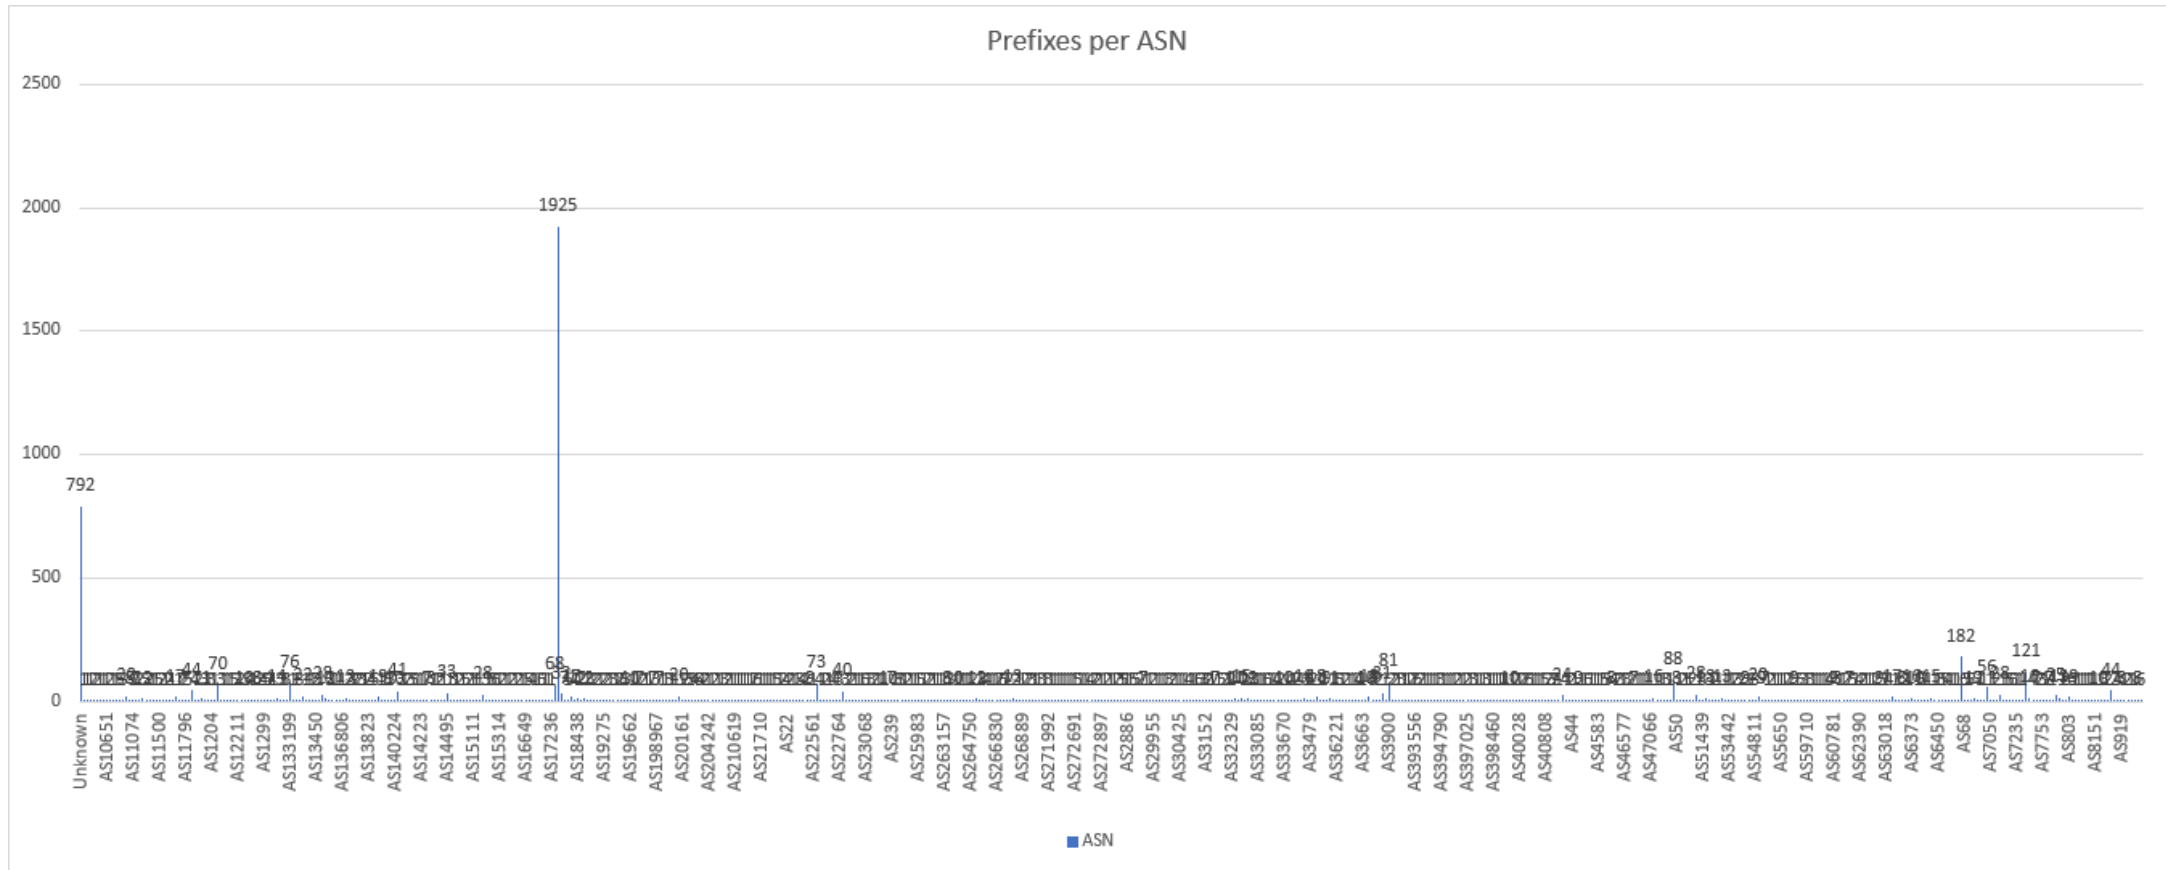
RPKI VALIDATION



■ RPKI Unknown ■ RPKI Invalid ■ RPKI Valid

# Analysis – Origin

Origin	Occurence
Unknown	792
<b>AS174</b>	<b>1925</b>
AS668	182
AS7377	121



# The problem with ARIN

- Legacy resources\* are the networks allocated before ARIN was established in 1997.
- Unlike the other RIRs, ARIN has decided to not provide IRR and RPKI services to legacy resources holders.
- There is no authoritative IRR server for these IP networks and they cannot be validated using RPKI either.

\*<https://www.arin.net/resources/guide/legacy/>

# Legacy networks commonly seen in Europe

- Access operators leasing networks from Cogent.
- US universities (also: USG HEP labs).
- US local and central government.
- Enterprise networks of local branches of US companies.
- Amateur packet radio.
- USG-managed DNS root servers (C E G H).

But also:

- DoS scrubbing centers (mostly Prolexic (Akamai)).
- Companies actually based in the US and announced in Europe by global carriers.

# A better classification of prefixes (IXPs)

IXP	Total amount	Cogent	US EDU	US MIL	US GOV	AMPR	DoS	Other
MIX	44552	116				3	451	46
MINAP	6537	19						
NAMEX	6467	30						
<b>AMS-IX</b>	<b>214961</b>	<b>1432</b>	<b>93</b>		<b>1</b>	<b>61</b>		<b>456</b>
LONAP	32882	65				18		37
BCIX	13275	3				3		10
BNIX	2695	14				2		

# A better classification of prefixes (providers)

Carrier	Total amount	Cogent	US EDU	US MIL	US GOV	AMPR	DoS	Other
<b>HE</b>	<b>92963</b>	<b>1579</b>	<b>430</b>	<b>293</b>	<b>107</b>	<b>56</b>		<b>684</b>
NetIX	6259	40						
PacketFabric	2924	43	2					
Zayo	25497	171	5					

# Conclusions

- The BCP enhances routing security and helps towards this direction
- We have very strong signs that not supporting anymore filtering with RADB is feasible (but annoying).
- Losing the peering routes of US-based networks is usually (?) not a big deal. Anecdotal evidence suggests that most carry negligible traffic in Europe.
- Major casualties:
  - Local networks leasing Cogent IP space.
  - NASA's E root server



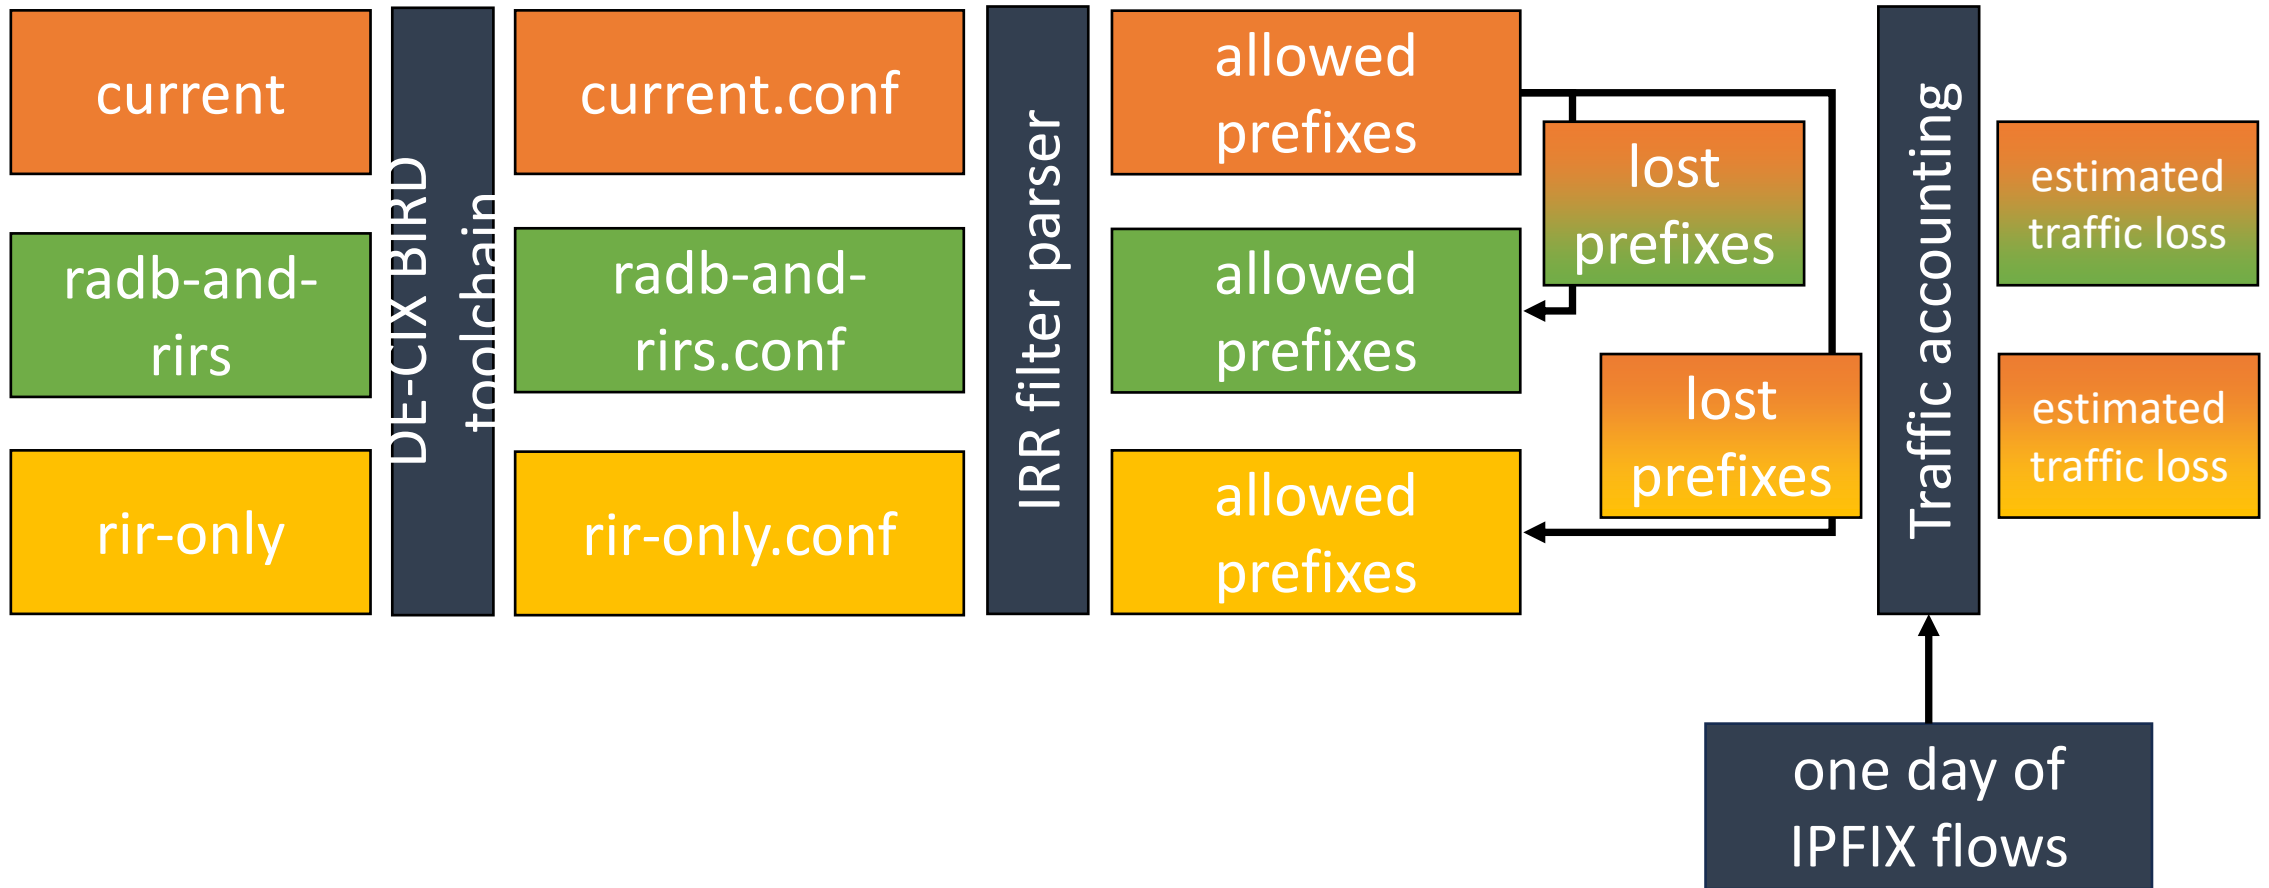
# Collaboration with DE-CIX

- An impact analysis of the BCP at the DE-CIX platform
  - Performed by Matthias Wichtlhuber and Daniel Wagner

# Simulated Scenarios

- **current:** default IRR db search order
  - peers prioritize their preferred db to the front
- **radb-and-rirs:** default IRR db search order without non-RIR dbs but with RADB
  - peers prioritize their preferred db to the front
  - if the preferred db is a non-RIR db, customer preference is ignored
- **rir-only:** same as radb-and-rirs, but RIR dbs only

# Simulation Method

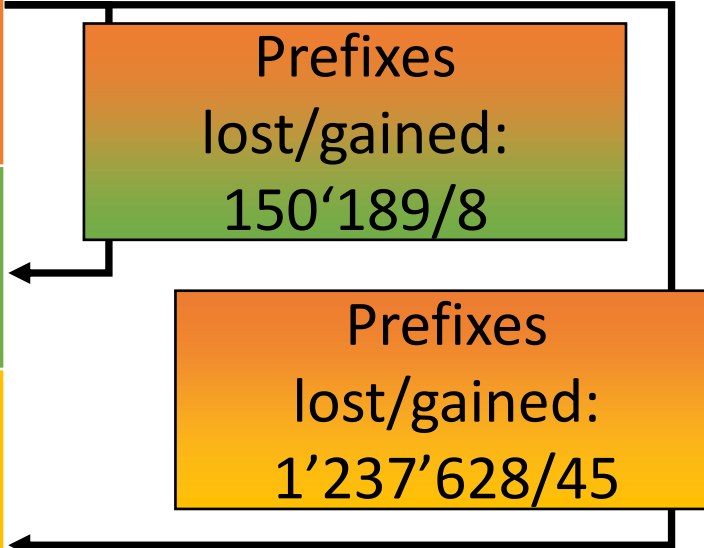


# DE-CIX toolchain AS filter/prefix loss

Scenario	Resolved IRR prefix lists	Deaggregated unique /24s
current	98'253 (100%)	10'997'437 (100%)
radb-and-rirs	96'424 (98.1%)	10'847'256 (98.6%)
rir-only	65'870 (67%)	9'759'854 (88.7%)


Prefixes  
lost/gained:  
150'189/8

Prefixes  
lost/gained:  
1'237'628/45




# Estimated traffic loss current vs. radb-and-rirs/rirs-only

DE-CIX traffic loss under BCOP impact (current vs. radb-and-rirs)



DE-CIX traffic loss under BCOP impact (current vs. rir-only)



# Key Takeaways

- Phasing out all alternative IRR dbs except RADB is doable mid-term
- RADB has substantial impact
  - We loose one third of resolved filtering lists and 11% of /24s
  - Traffic loss could be up to 250 Gbps during peak
- Proposal
  1. Gather like-minded IXPs and find a date X (>6 months) for removing non-official IRR dbs except RADB
  2. Coordinate and prepare customer communications, monitor results
  3. Implement changes at date X
  4. Work jointly on a plan for RADB

# Thank you



Questions?

Let the discussion begin

URL for the Document:

[https://amsix-my.sharepoint.com/:b:/g/personal/stavros\\_konstantaras\\_amsix\\_net/EWvrrurCwi3VNnUc-ocF5ODwBd\\_OLGCbdhC9T8jNb1IQNZg?e=MregHo](https://amsix-my.sharepoint.com/:b:/g/personal/stavros_konstantaras_amsix_net/EWvrrurCwi3VNnUc-ocF5ODwBd_OLGCbdhC9T8jNb1IQNZg?e=MregHo)