

An Analysis of ARIN NetHandles with OriginAS Data and Analysis of RIR/IRR Registry Data

O. Kim, K. Sriram, O. Borchert, P. Gleichmann, and D. Montgomery

Presentation at ARIN XXIII, San Antonio, TX
April 26-29, 2009

Contacts: okim@nist.gov, ksriram@nist.gov, doug@nist.gov
Project Website: www.antd.nist.gov/bgp_security

Outline

- Problem statement
- Analysis of ARIN NetHandles with OriginAS
- Analysis of Global Registries
(comparisons with what is announced in BGP)

What is the Problem?

- Current registry data is considered inaccurate, incomplete
- Despite weaknesses, data is used for:
 - Local route filtering
 - Debugging purposes
- No comprehensive investigations to date
- Improving quality and completeness of routing data could enable new BGP robustness mechanisms

Registry Data Object Counts by Source

RIR/IRR	route			inetnum (ARIN NetHandle)			aut-num (ARIN ASHandle)		
	06/18/2007	10/18/2008	Incr	06/18/2007	10/18/2008	Incr	06/18/2007	10/18/2008	Incr
ARIN	7,330	8,201	12%	338 (1,618,197)	434 (1,924,454)	28% 19%	758 (18,050)	890 (19,678)	17% 9%
RIPENCC	71,569	89,957	26%	2,044,536	2,458,119	20%	14,106	16,969	20%
APNIC*	23,616	35,515	50%	822,891	1,080,999	31%	4,559	5,347	17%
AFRINIC	0	0		13,948	22,706	63%	342	445	30%
LACNIC**	0	0		45,346	83,036	83%	1,219	1,339	10%
Standalone IRRs+	345,129	497,124	44%	1	1		3,785	4,643	23%
Total:	447,644	630,797	41%	2,927,060 (1,618,197)	3,645,295 (1,924,454)	25% 19%	24,769 (18,050)	29,633 (19,678)	20% 9%

* Includes TWNIC, JPIRR, JPNIC and APNIC


** RIR only

+ Independent IRR databases that are mirrored via **the RADB website** including RADB, but **EXCLUDING ARIN, APNIC, JPIRR and RIPE**

Note that route objects can be registered at any IRR regardless of where the address spaces are allocated.

Distribution of Prefix Length of inetnum (RPSL) and NetHandle (SWIP)

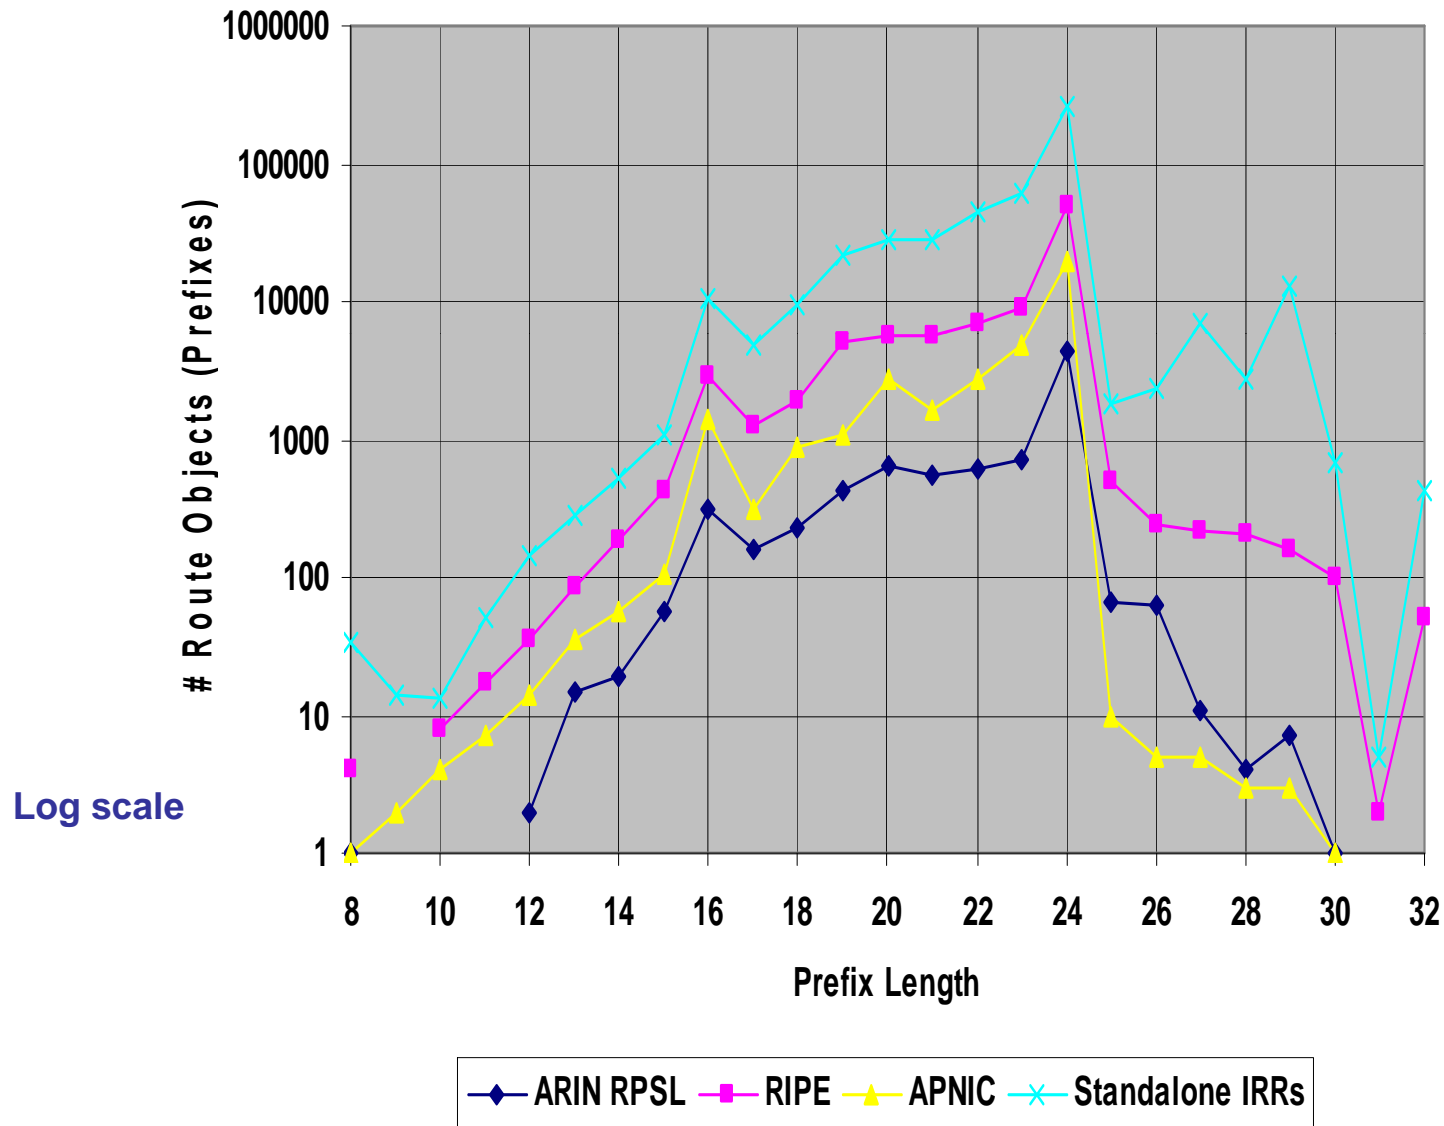
Registry Data Date: 2008-10-18




- Length 0 indicates that an address block cannot be represented by a single CIDR
- Length 4 specifies Multicast and Reserved Future Use blocks
- Some Legacy and ERX blocks may be included in one or more RIRs

Distribution of Prefix Length of Route Objects in IRR

Registry Data Date: 2008-10-18




Distribution of Sources of Prefix Allocations of Route Objects Registered to Standalone IRRs



All route objects registered in standalone IRRs on 2008-10-18: **497,124**

Growth of NetHandles with OriginAS



ARIN NetHandle Stats in Comparison to BGP Updates and RIBs

- Raw data
 - ARIN Registry data on 2008-10-18
 - * All NetHandle objects: 1,924,454
 - * Unique (NetHandle, OriginAS) pairs: 73,249 (4%)
 - * Unique (NetRange, OriginAS) pairs: 73,062
 - * Unique OriginASes: 2693
 - BGP Updates & RIB data:
 - * Collector: Oregon from Routeviews
 - * Updates (2008-06-01 to 2008-11-24)
 - Unique (prefix,origin) pairs: 531,820
 - * BGP RIBs on 2008-11-3: 283,035
 - unique (prefix,origin) pairs other than those in Updates prefixes above: 1
 - ALL Unique (prefix,origin) pairs from both Updates and RIBs: 531,821

Some Observations on ARIN NetHandles with OriginAS

- Multiple NetHandles that contain the exact same (NetRange, OriginAS) pairs with different allocation types:
 - Allocation types: allocation / reallocation / assignment / reassignment

# of instances with the following:	count
3 NetHandles containing the same (NetRange,OriginAS) pair	2
2 NetHandles containing the same (NetRange,OriginAS) pair	183
NetHandles with unique (NetRange,OriginAS) pair	72,877

Some Observations on ARIN NetHandles with OriginAS


- Two or more NetHandle objects contain the exact same (NetRange, OriginAS) pairs, but different *NetType*:
 - One Example: (66.97.96.0/20, 33125)

	NetHandle Object 1	NetHandle Object 2
NetHandle:	NET-66-97-96-0-1	NET-66-97-96-0-2
OrgID:	SNL-27	MCB-21
NetRange:	66.97.96.0 - 66.97.111.255	66.97.96.0 – 66.97.111.255
NetType:	<i>Allocation</i>	<i>Reassignment</i>
OriginAS:	AS33125	AS33125
Parent:	NET-66-0-0-0-0	NET-66-97-96-0-1
RegDate:	2006-10-10	2007-06-12
Updated:	2007-06-12	2007-06-12

ARIN NetHandles with OriginAS

Multiple OriginAS (MOAS) Distribution


Registry Data Date: 2008-10-18



- Some prefix owners register prefix with each of their ASes
- Some never remove old route registrations?

Distribution of NetHandles Associated with the Origin AS

Registry Data Date: 2008-10-18




- A large percentage of (NetHandle, Origin AS) pairs are associated with about 10 Origin ASes


Distribution of Prefix Length of NetHandles w/ OriginAS vs. BGP Trace Data

Registry Data Date: 2008-10-18

BGP Trace Data
from 2008-06-01 to 2008-11-24



Length 0 indicates that NetRange cannot be represented by a single CIDR




Length 0 indicates prefix 0.0.0.0/0

- Prefix 0.0.0.0/0 is announced by 15 Origin ASes (12956, 3561, 19151, 513, 9829, 3130, 293, 5602, 8546, 174, 47797, 28968, 31261, 47819, 18747).
- There exist 27 (prefix, origin) pairs with prefix length less than 8, excluding length 0 above.

Distribution of ARIN NetRange Address Block Allocations

Registry Data Date: 2008-10-18



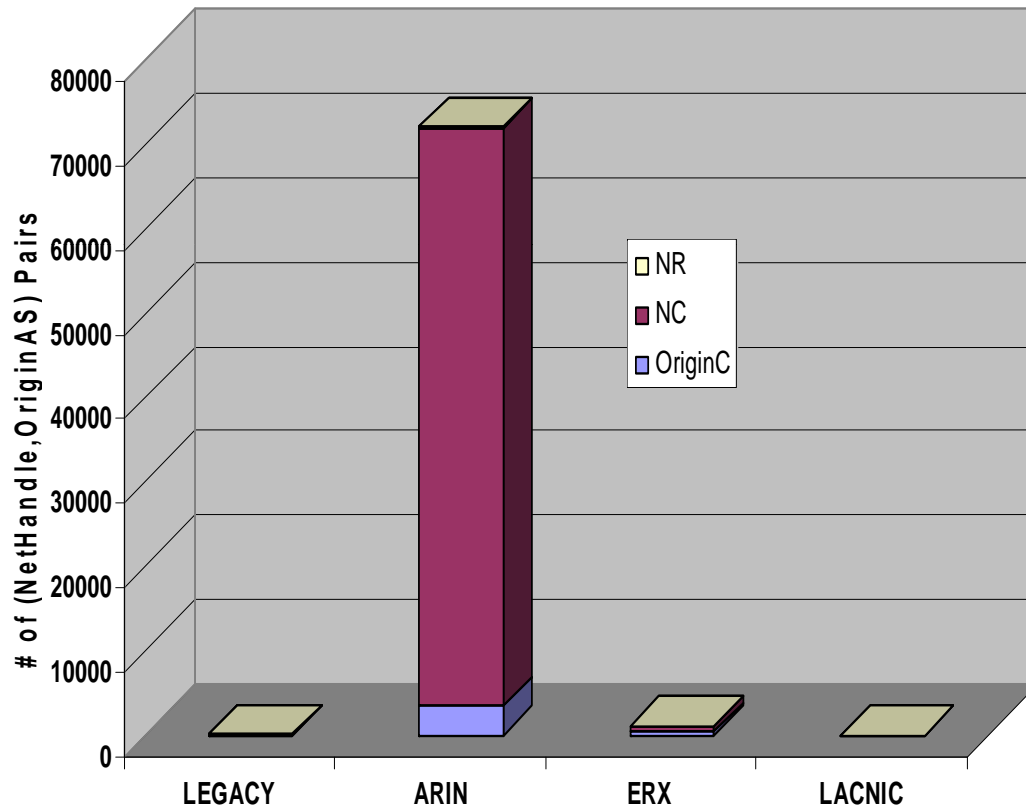
Note: Considering only NetHandles with Origin AS

Methodology for Consistency Checks

- Mntner, OrgID, Contact information (tech-c, admin-c, etc.) are compared across corresponding registered objects
- Origin Consistent: For {prefix, OriginAS} pair in NetHandle, ASHandle is consistent
- Not Registered: No ASHandle Exists
- NC: ASHandle is not consistent

Consistency Checks for ARIN NetHandles with OriginAS

Registry Data Date: 2008-10-18

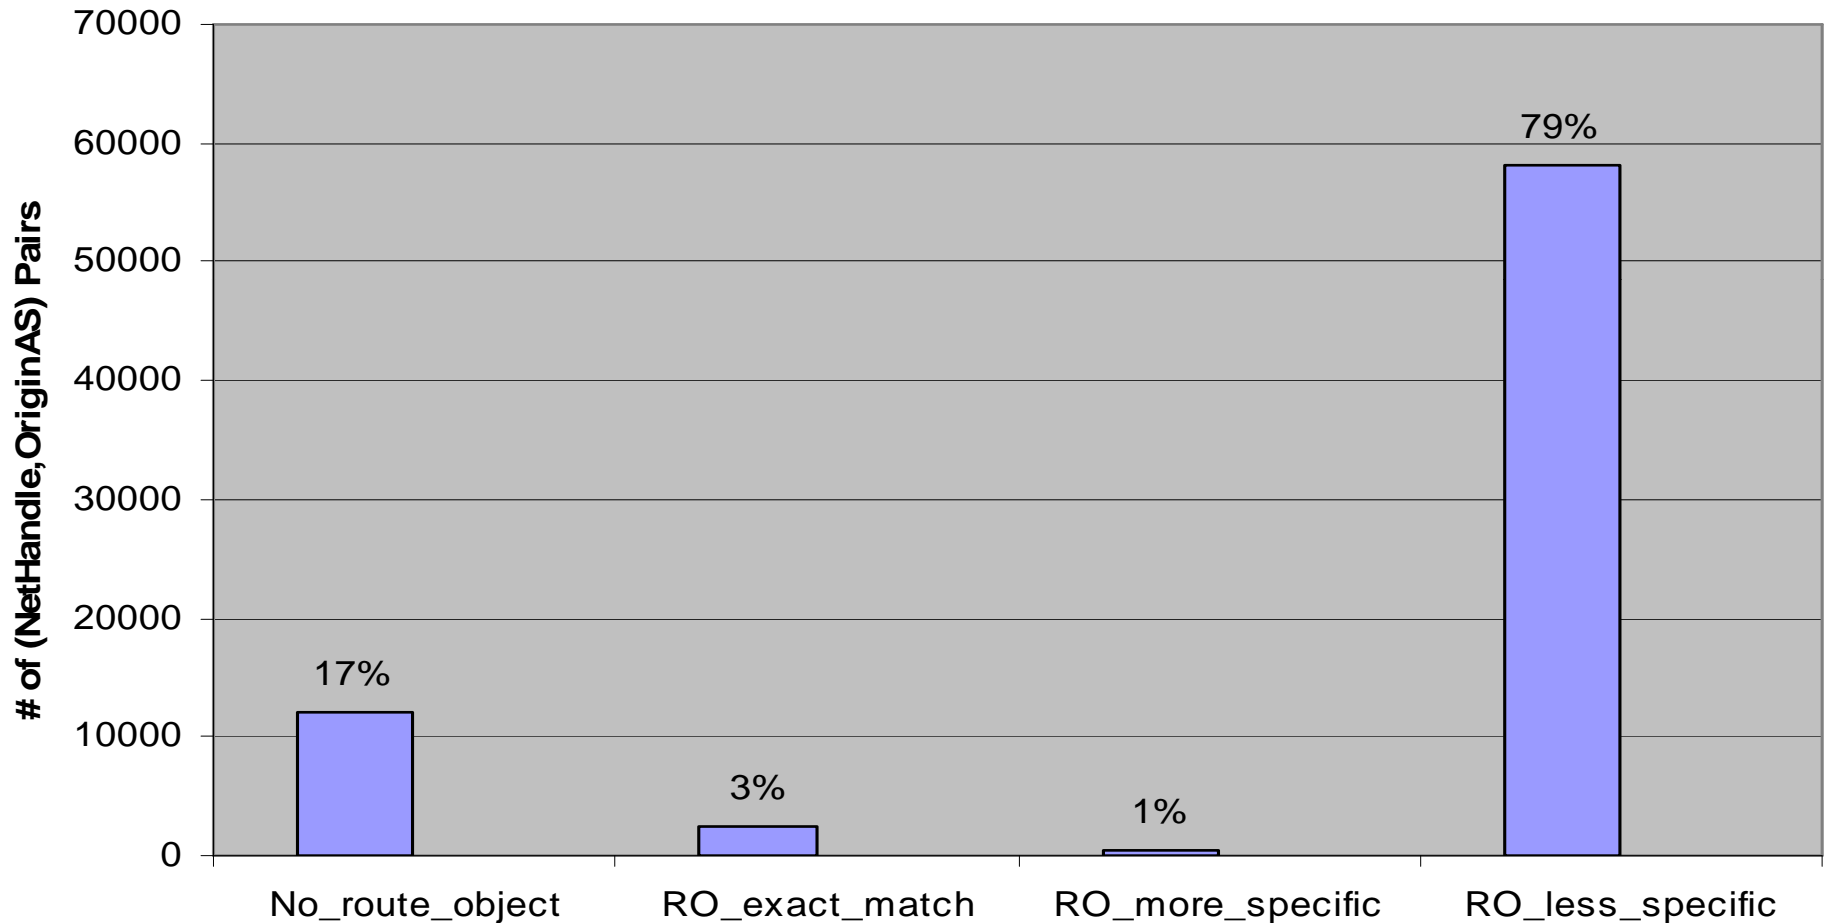


Region	OriginC	NC	NR	Total
Legacy	3	73	1	77
ARIN	3519	68437	113	72069
ERX	391	697	12	1100
Lacnic	1	2	0	3
Total	3914	69209	126	73249

Scores for Consistency Checks for ARIN NetHandle w/ OriginAS

ARIN NetHandles w/ OriginAS and the Existence of Corresponding Route Objects in RPSL


Registry Data Date: 2008-10-18



- For origin validation, ARIN RPSL route objects provide superior coverage than NetHandles with origin AS

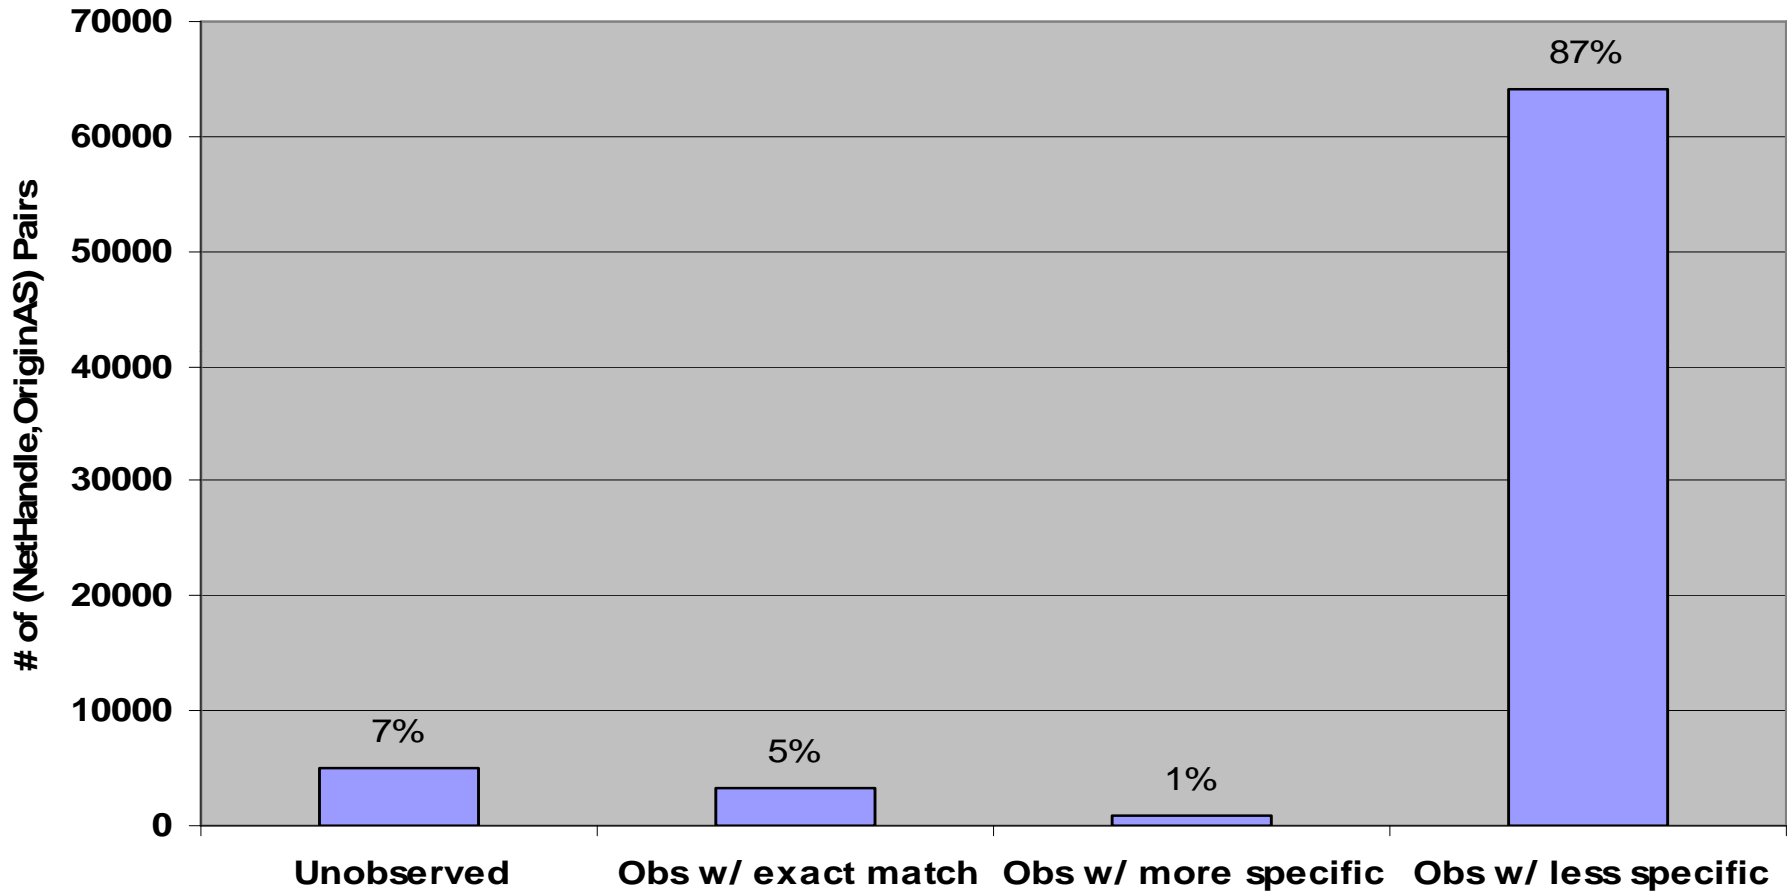
ARIN NetHandles w/ OriginAS and Existence and Quality of Corresponding Route Objects in RPSL

Registry Data Date: 2008-10-18



- **No_RO**: No Route objects exist
- **NR**: No Referenced objects exist (ie., ASHandle or aut-num)
- **NC**: (referenced objects exist, but) Not Consistent
- **FC**: Fully (Prefix & Origin) Consistent
- **PrefixC**: Only Prefix Consistent
- **OriginC**: Only Origin Consistent

ARIN NetHandles w/ OriginAS that are Observed in BGP Trace Data



- About 6% of the NetHandles with origin AS are usable for direct verification of origin in BGP update messages; that is less than 5K NetHandles (in Oct. 2008)

Comparison of ARIN NetHandles with OriginAS vs. Announced (p, OAS) Pairs for Prefix Length ≥ 25

		Prefix length ≥ 25	
		# of (p, OAS)	percentage
ARIN NetHandles with OriginAS	73k	60k	82.2%
Announced (p, OAS) that correspond to ARIN Address Space	186k	6.5	3.5%
Globally announced (p, OAS)	532k	29.3k	5.5%

ARIN NetHandle w/ OriginAS Consolidation of (NetHandle, OriginAS) Pairs

On 2008-10-18	Total	# sub-prefixes
All unique (NetRange,OriginAS) Pairs	73,062	
Distinct NH_OAS (NetHandle w/ OriginAS) with no super-prefixes	39,297	
Of these (39,297):		
# of NH_OAS with no sub-prefixes	38,693	0
# of NH_OAS with sub-prefixes (only one level below)	584	16828
# of NH_OAS with sub-prefixes (two levels below)	20	16937

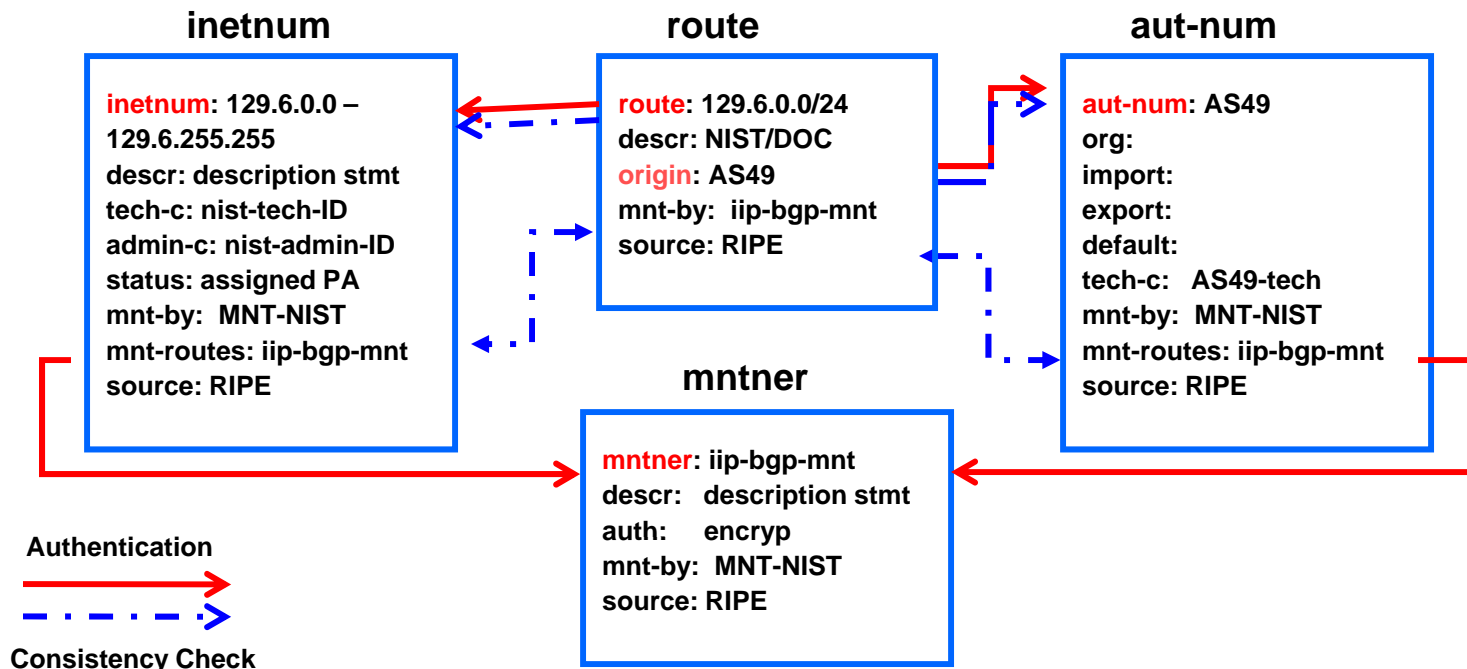
- Note: $38,693 + 584 + 20 + 16828 + 16937 = 73,062$
- Many of the consolidated 39,297 are also subprefixes of what are actually observed

Outline

- Problem statement
- Analysis of ARIN NetHandles with OriginAS
- Analysis of Global Registries
(comparisons with what is announced in BGP)

Registry Self-Consistency Check (Quality Analysis Algorithm)

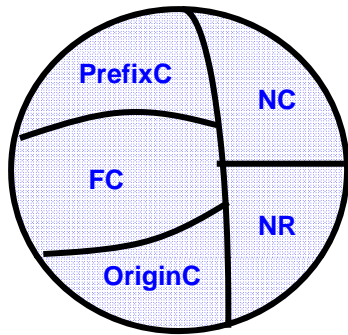
- Self-Consistency check criteria:
 - Check consistency between relevant objects by comparing the following attributes:
 - * 'mntner' related attributes: Used mainly for RPSL
 - * 'orgID' attribute: Used mainly for SWIP
 - * Contact information (i.e., tech-c / admin-c / TechHandle / AbuseHandle)
- A route object is considered as fully consistent if, based on the above criteria, it matches with both of these:
 - ⊕ the referenced **aut-num** for the **origin**; and
 - ⊕ the referenced **inetnum** for the **prefix**.



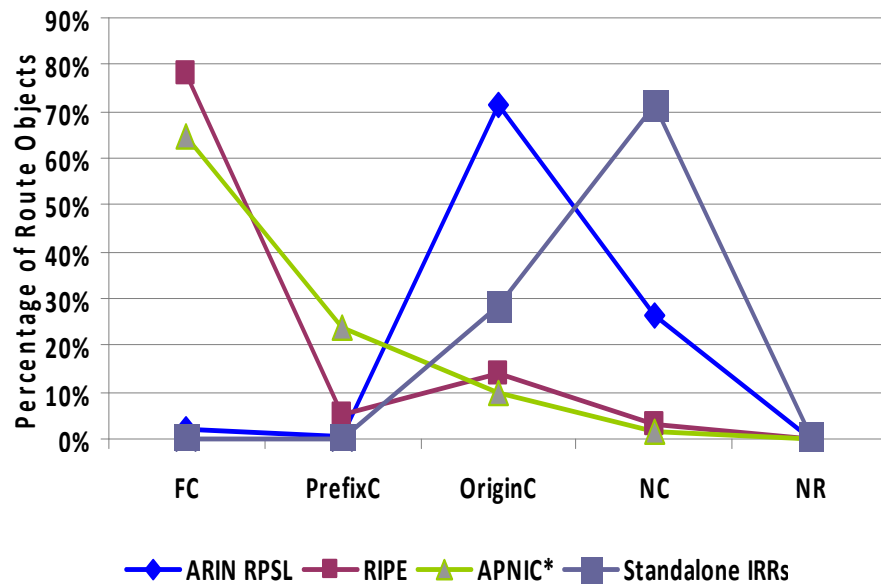
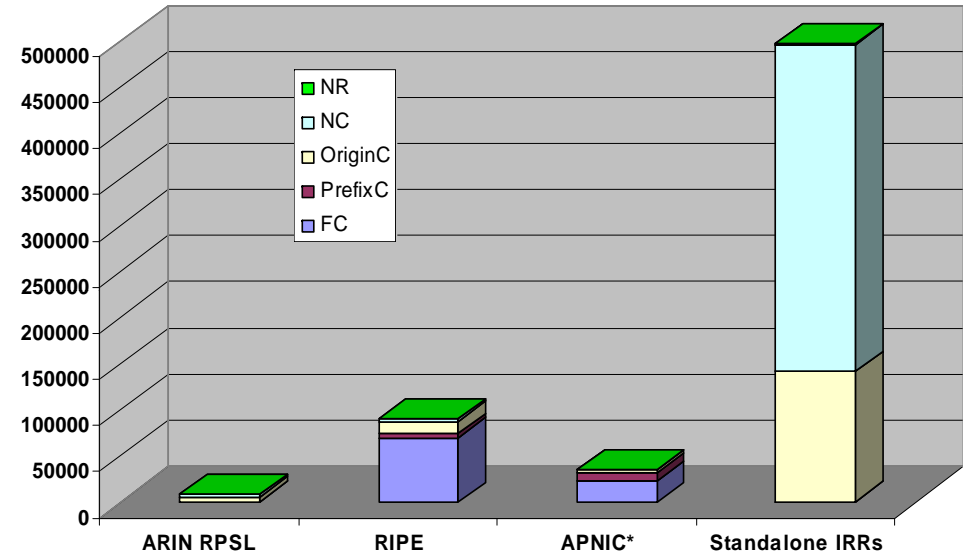
Characterization of IRR Consistency Based on Route Object Registrations

Registry Data Date: 2008-10-18

Registry Data




- **FC:** Fully (Prefix & Origin) Consistent
- **PrefixC:** Only Prefix Consistent
- **OriginC:** Only Origin Consistent
- **NC:** (referenced objects exist, but) Not Consistent
- **NR:** No Referenced Resource Objects Exist



	ARIN RPSL		RIPE		APNIC*		Standalone IRRs	
FC	169	2%	70057	78%	22981	65%	534	0%
PrefixC	27	1%	4458	5%	8364	23%	107	0%
OriginC	5845	71%	12627	14%	3562	10%	141323	29%
NC	2147	26%	2815	3%	608	2%	353598	71%
NR	13	0%	0	0	0	0	1534	0%
Total	8201		89957		35515		497096	

Characterization of IRR Consistency Based on Route Object Registrations


Registry Data Date: 2008-10-18



Stability of (p, OAS) in the Trace Data

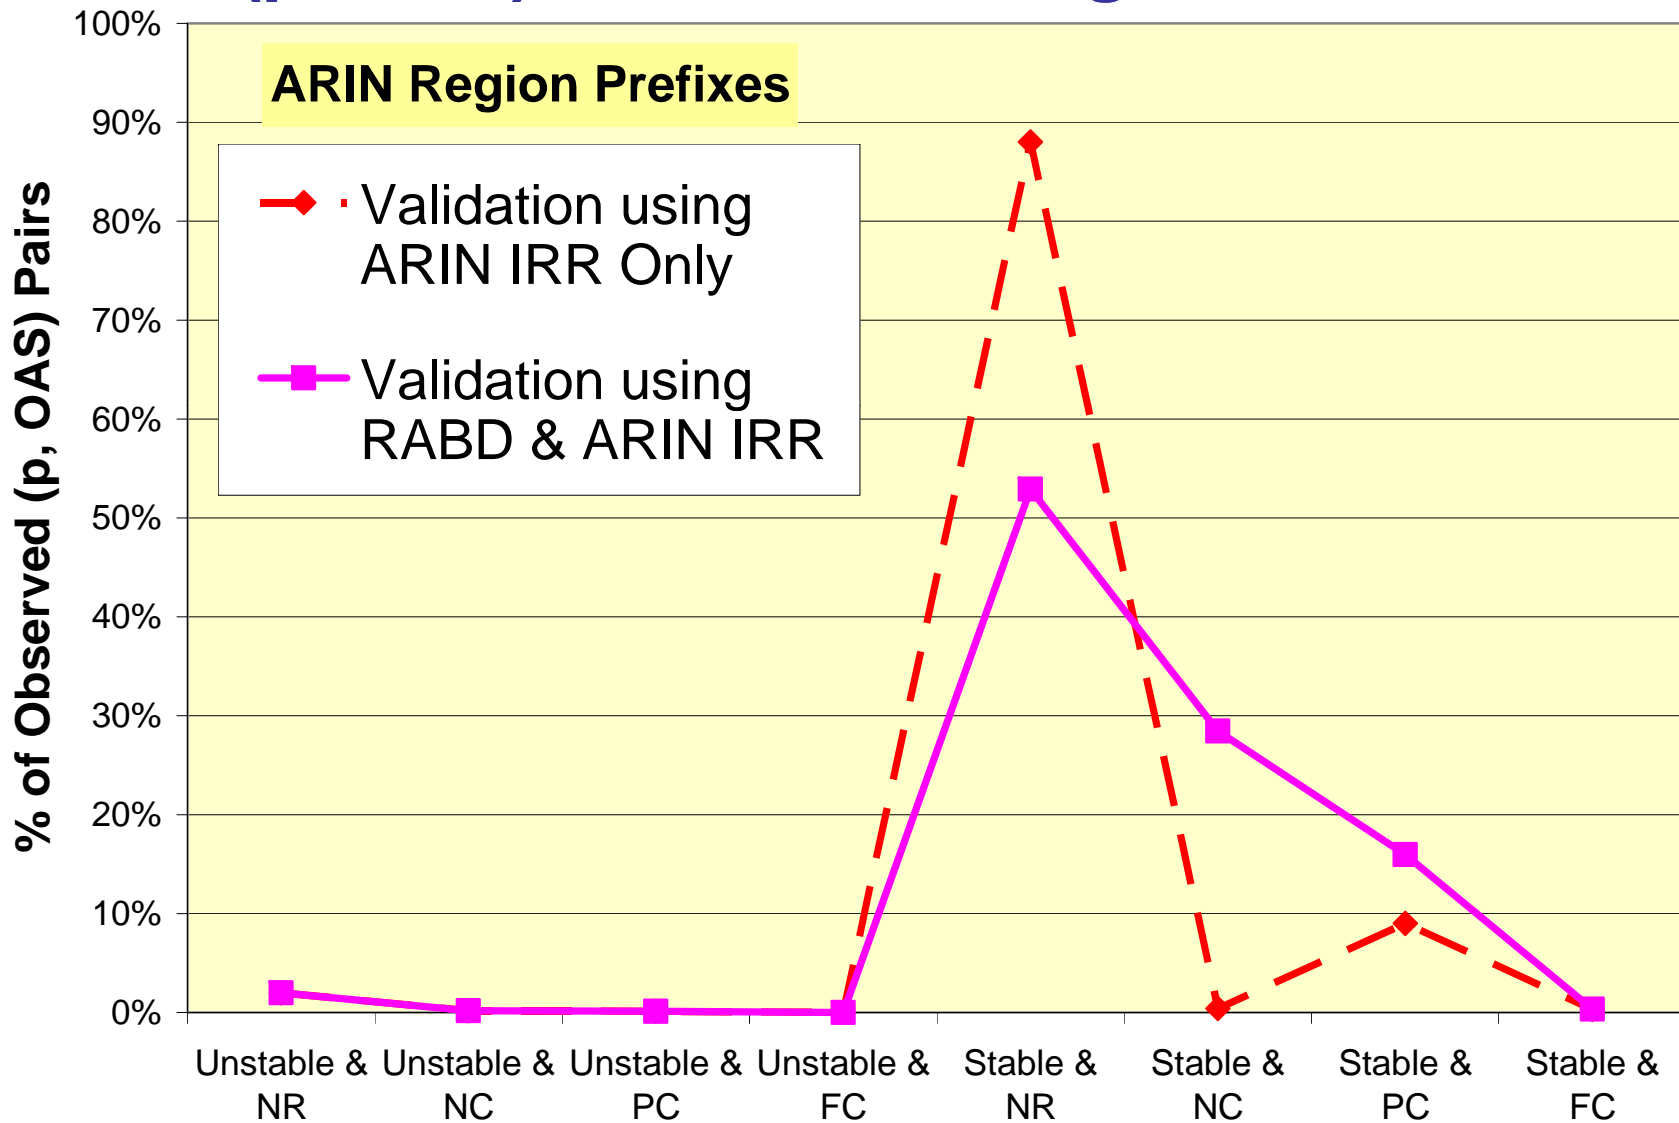
- If (p, OAS) pair remained in RIBs stably for 48 hours or more at least once during the observation period (6 months), then the (p, OAS) pair is considered stable
- Otherwise, the (p, OAS) pair is considered unstable (transient)

Classification of Observed (p, OAS) Pairs According to Stability / Consistency Scores



FC = Fully Consistent; PC = Partially Consistent; NC = Not Consistent; NR = Not Registered

Stability/Consistency Scores of Observed (p, OAS) Pairs: ARIN Region Prefixes

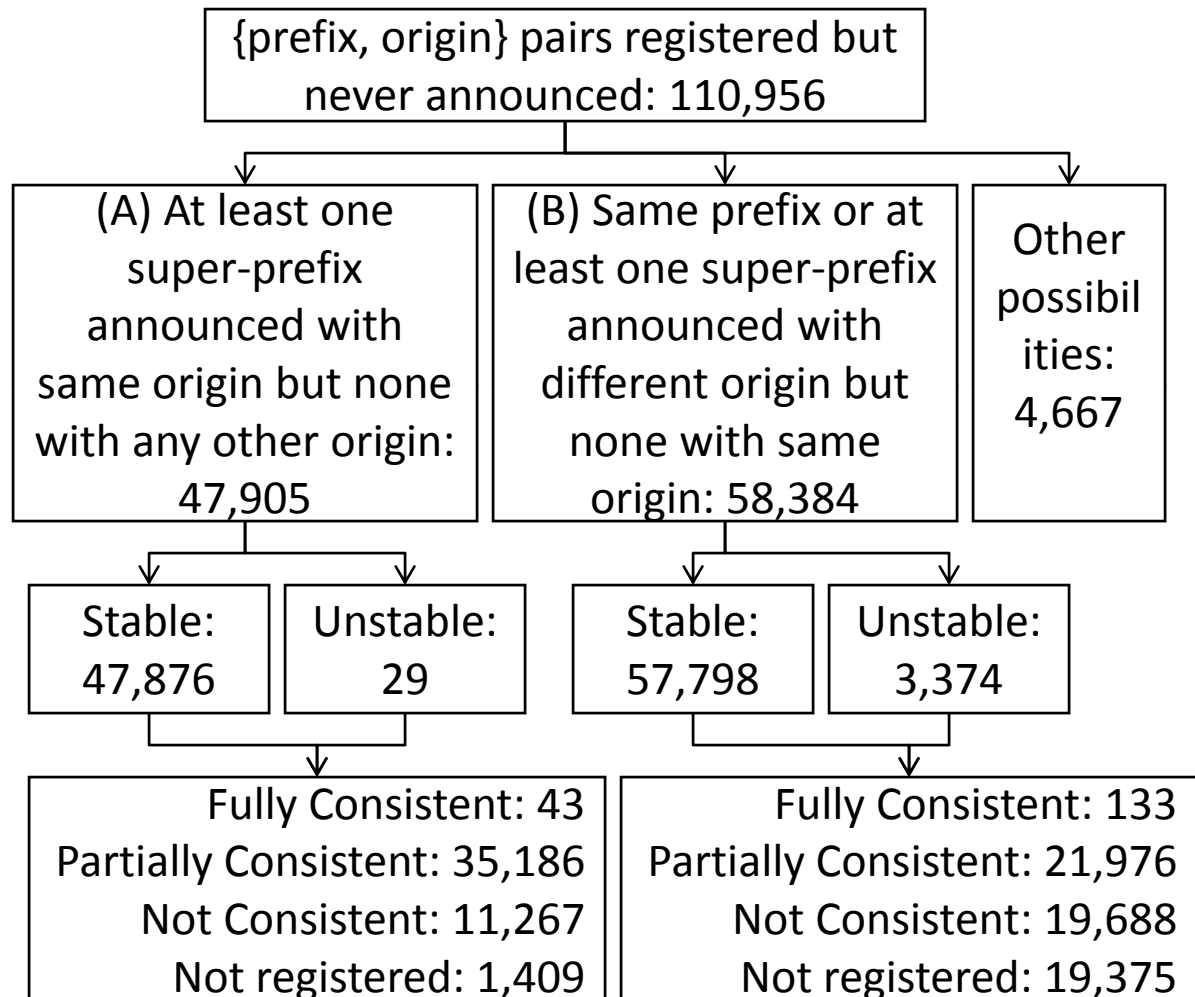


FC = Fully Consistent; PC = Partially Consistent; NC = Not Consistent; NR = Not Registered

Analysis of Registered But Unobserved Routes

ARIN Prefixes

- Large number of {prefix, origin} pairs registered but never announced
- In most cases, super-prefixes are announced with the same origin AS (as in registered route) or a different origin AS
- Re-origination type of aggregation by a higher tier ISPs and/or stale Route registrations?

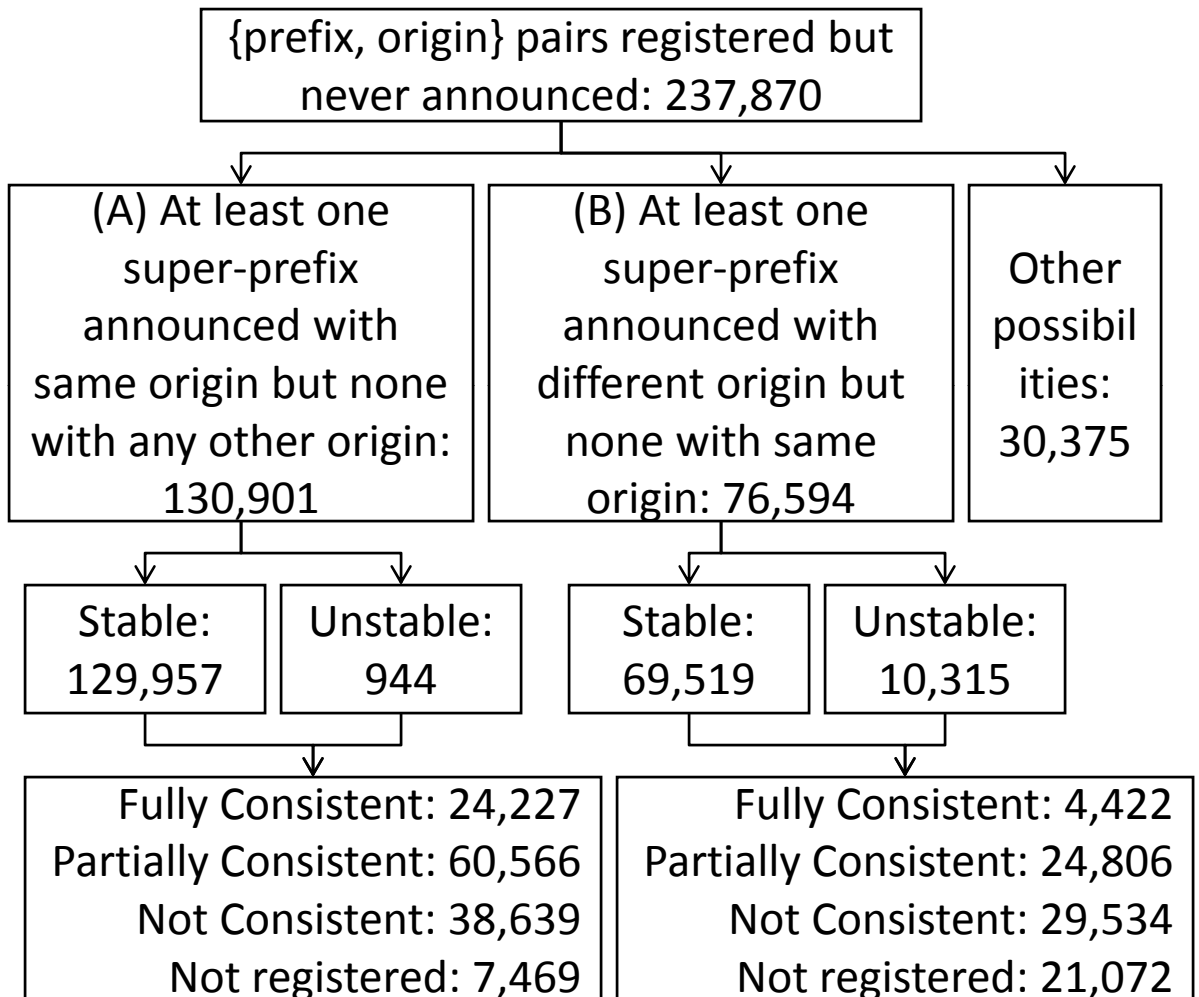


← For the super-prefixes with their observed origin ASes →

Analysis of Registered But Unobserved Routes

Global Prefixes

- Large number of {prefix, origin} pairs registered but never announced
- In most cases, super-prefixes are announced with the same origin AS (as in registered route) or a different origin AS
- Re-origination type of aggregation by a higher tier ISPs and/or stale Route registrations?



↖ For the super-prefixes with their observed origin ASes ↗

Conclusions and Future Work

- ARIN NetHandles with Origin AS -- dominantly for prefix lengths ≥ 25
- Announced prefixes are dominantly of length ≤ 24
- As it stands, ARIN RPSL routes (~10K) more useful than NetHandles with origin AS (~100K)
- Routes exist in standalone RABD but not enough and lacking consistency (Verizon alone has about 60 different OrgIDs *)
- It would be immensely helpful whatever RIRs / ISPs can do to encourage/support:
 - Route registrations
 - Using consistent OrgIDs
 - SIDR RPKI trials and testing

* Based on informal communication between NIST and Verizon