

## Information Quality Analysis

### What is the Problem?

- Current registry data is considered:
  - Stale, insecure, inaccurate and incomplete
  - Registration is voluntary
  - Routing policy system is not adequately secure
- Despite weaknesses, data is used for:
  - Local route filtering
  - Debugging purposes
- Integrity of registry data is crucial to help maintain global secure routing infrastructure.
- No comprehensive investigations to date
- Improving quality and completeness of routing data could enable new BGP robustness mechanisms

1

### What are NIST's Goals?

- Perform comprehensive investigations and analysis of the integrity of the registry data
  - Characterize correctness and completeness of existing RIR/IRR databases
    - Quality measurement of IRR data
    - Analysis of syntactic correctness of IRR data
    - Analysis of IRR content changes
  - Characterize consistency of global IRR data with regard to BGP trace data
    - Quality analysis of registry data (route and NetHandle objects) vs. BGP route announcements (Updates)
  - Contribute to the improvement of the quality of global routing information infrastructure
- Understand and gain some insights by analyzing both registry and trace data, to help improve BGP routing robustness

2

### Registry Data Object Counts by Source

RIR	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%
RADB+	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  
+ includes all mirrored data EXCLUDING ARIN, APNIC, JPIRR  
Note that route objects can be registered any RIR regardless of where the address spaces are allocated.

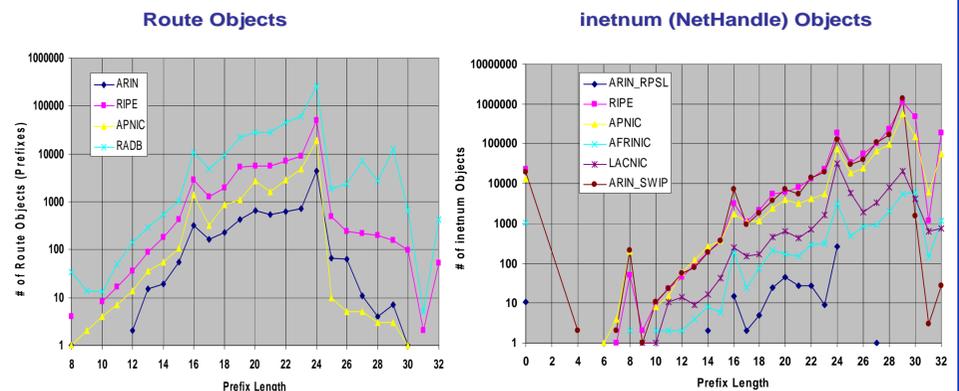
#### Data used for Consistency Checks

- Registry data: 2008-10-18
- BGP Updates & RIBs:
  - Collector: Oregon from Routeviews
  - All prefixes from Updates from 2008-6-1 to 2008-9-15: 1,159,558,753
  - Unique (prefix,origin) pairs: 362,649
- BGP RIBs on 2008-11-3: 283,035
  - Unique (prefix,origin) pairs other than those in Updates prefixes above: 14,197
- All unique (prefix,origin) pairs from both Updates and RIBs: 376,846

3

### Distribution of Prefix Length of Route Objects vs. inetnum (NetHandle) Objects

Registry Data Date: 2008-10-18

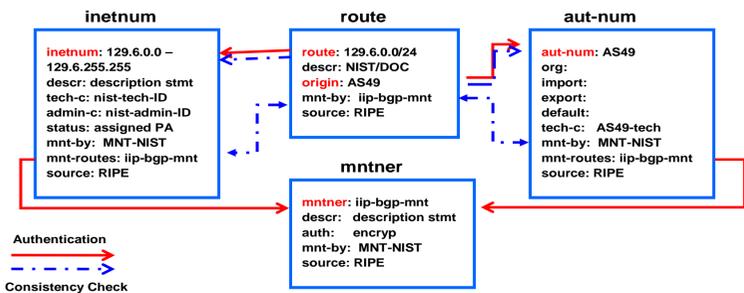


Length 0 indicates that an address block cannot be represented by a single CIDR

4

### Registry Self-Consistency Check (Quality Analysis Algorithm)

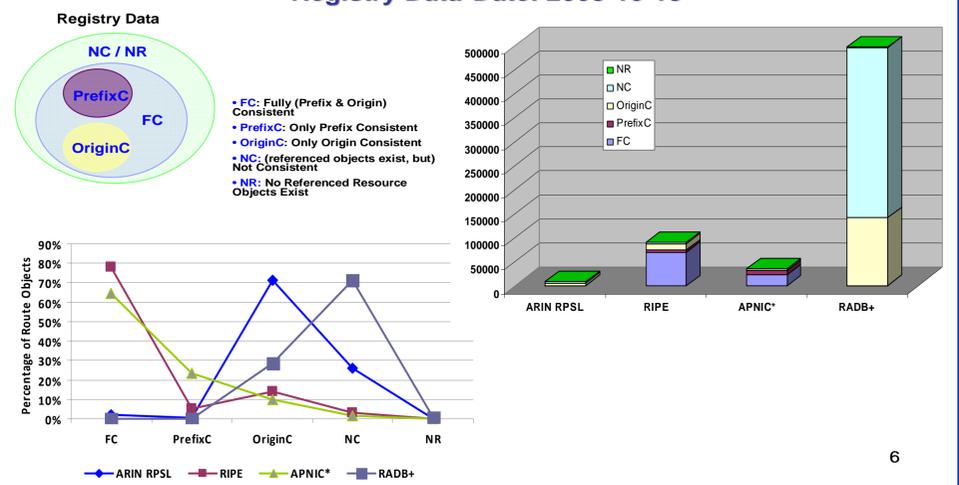
- Self-Consistency check criteria:
  - Check consistency between relevant objects by comparing the following attributes:
    - 'mntner' related attributes: Used mainly for RPSL
    - 'orgID' attribute: Used mainly for SWIP
    - Contact information (i.e., tech-c / admin-c)
- A route object is considered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



5

### Characterization of IRR Consistency Based on Route Object Registrations

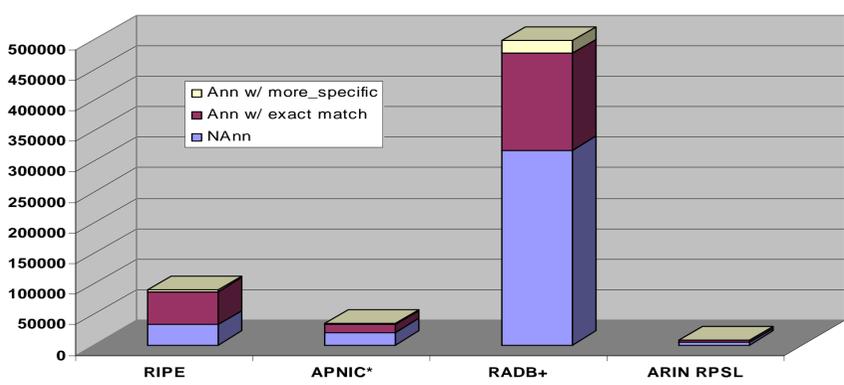
Registry Data Date: 2008-10-18



6

### IRR Route Objects WRT BGP Trace Data

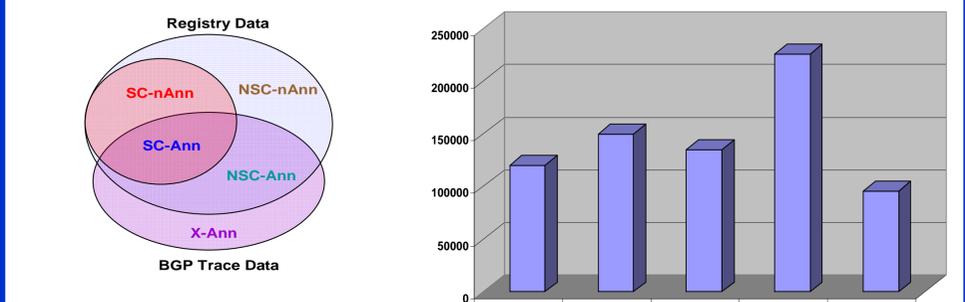
Registry Data Date: 2008-10-18  
Trace Data from 2008-06-01 to 2008-09-15



Route objects are observed with exact match or more specific prefixes with the same originAS in BGP trace data

7

### Measuring Correctness / Completeness of Registry Data WRT BGP Trace Data



Metrics	Registered route objects in the Registry	(prefix, originAS) Pairs in Trace Data
SC-Ann	Self-consistent	Observed matching exact match or more specific prefixes with same origin
SC-nAnn	Self-consistent	Not observed
NSC-Ann	Not self-consistent	Observed matching exact match or more specific prefixes with same origin
NSC-nAnn	Not self-consistent	Not observed
X-Ann	No registered route objects	Observed (prefix,originAS) pairs

All global registered route objects: 630,797  
All unique observed (prefix, origin) pairs: 376,836

8