

A Survey On Utilization of IRR's Route Objects

Kengo NAGAHASHI
kenken@elab.ic.i.u-tokyo.ac.jp
The University of Tokyo
JPNIC IRR Planning Team

Agenda

1. Background
2. Goal
3. Related Work
4. Approach

Background

- What benefit IRR offers?
 - Getting Contact Information
 - Router Configuration
 - Others..
- However;
 - How much IRR objects registered?
 - Is all routing information registered in IRR ?

Our Goal

- To understand divergence between IRRs and BGP prefixes
 1. How many prefixes are registered in IRRs?
 2. What differences in IRRs (e.g. RADB, RIPE, APNIC)?
 - Region, History, Operation
 3. Is IRR practical for router configuration?

Related Work

□ RIPE RRCC/RIS

■ Similarity;

- Compared with BGP routing table and IRR database

■ Difference;


- We analyze RADB, RIPE, APNIC, JPIRR database and unified database

Approach


□ How to match IRR and BGP routing table?

- Two ways of matching method;

1. Exact Match



2. Best Match



- BGP Routing Table

- eBGP multi-hop from 2 ISPs

IRR Database

□ IRR Database

- # of route objects (as of 2003/08/11)
 - RADB 69,493 records
 - RIPE 31,087 records
 - APNIC 811 records
 - JPIRR 227 records

Definition of unified database

□ What?

- Combination of RADB, RIPE, APNIC, JPIRR
- Removed duplicated records

□ Why ?

- Routing Information -> world wide spread
- IRR route objects -> regional
- Ideal Database which covers all regions
- We made unified db

Unified Database

of duplicated objects (exact match)

	RADB	RIPE	APNIC	JPIRR
RADB	5925	1906	139	209
RIPE	1906	934	5	0
APNIC	139	5	0	0
JPIRR	209	0	0	0

RADB	69,493 records
RIPE	31,087 records
APNIC	811 records
JPIRR	227 records

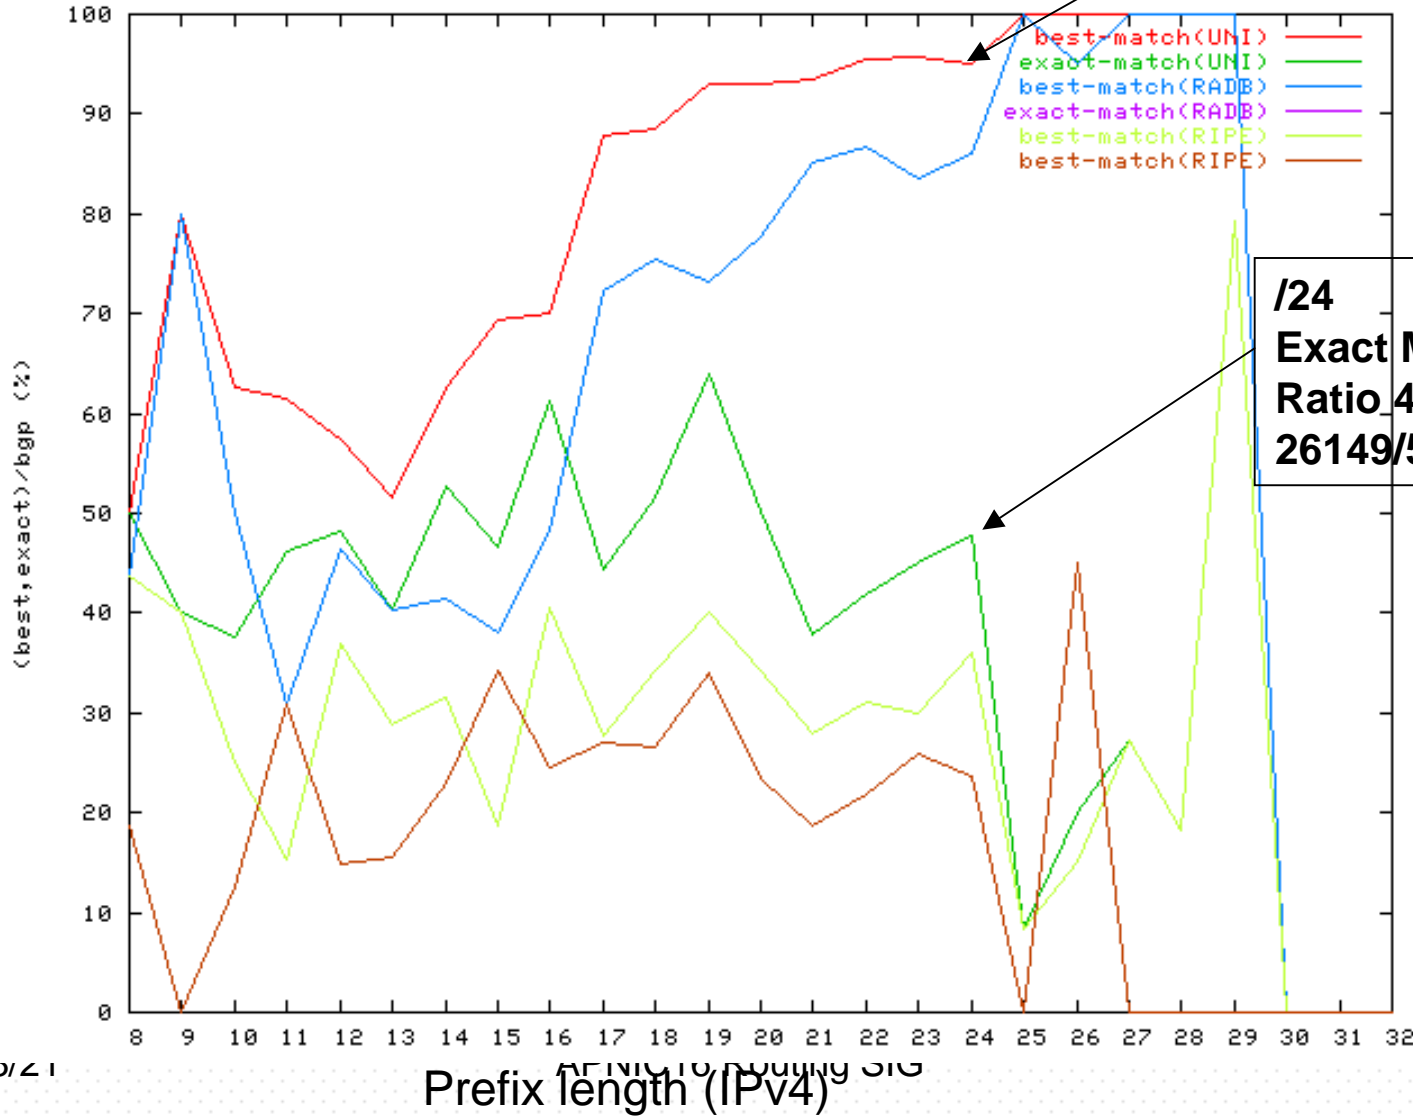
of Route Objects in unified DB is:

Total # of All Route Objects – dup. of (RADB+RIPE+APNIC+JPIRR)
= 101,618 – 9,118 = **92,500** records

Snapshot 1 (2003/08/11)

/24
Best Match (Uni)
Ratio 94%
51927/54664

Percentage of IRR/BGP ratio



/24
Exact Match (Uni)
Ratio 47%
26149/54664


2003/8/11

10

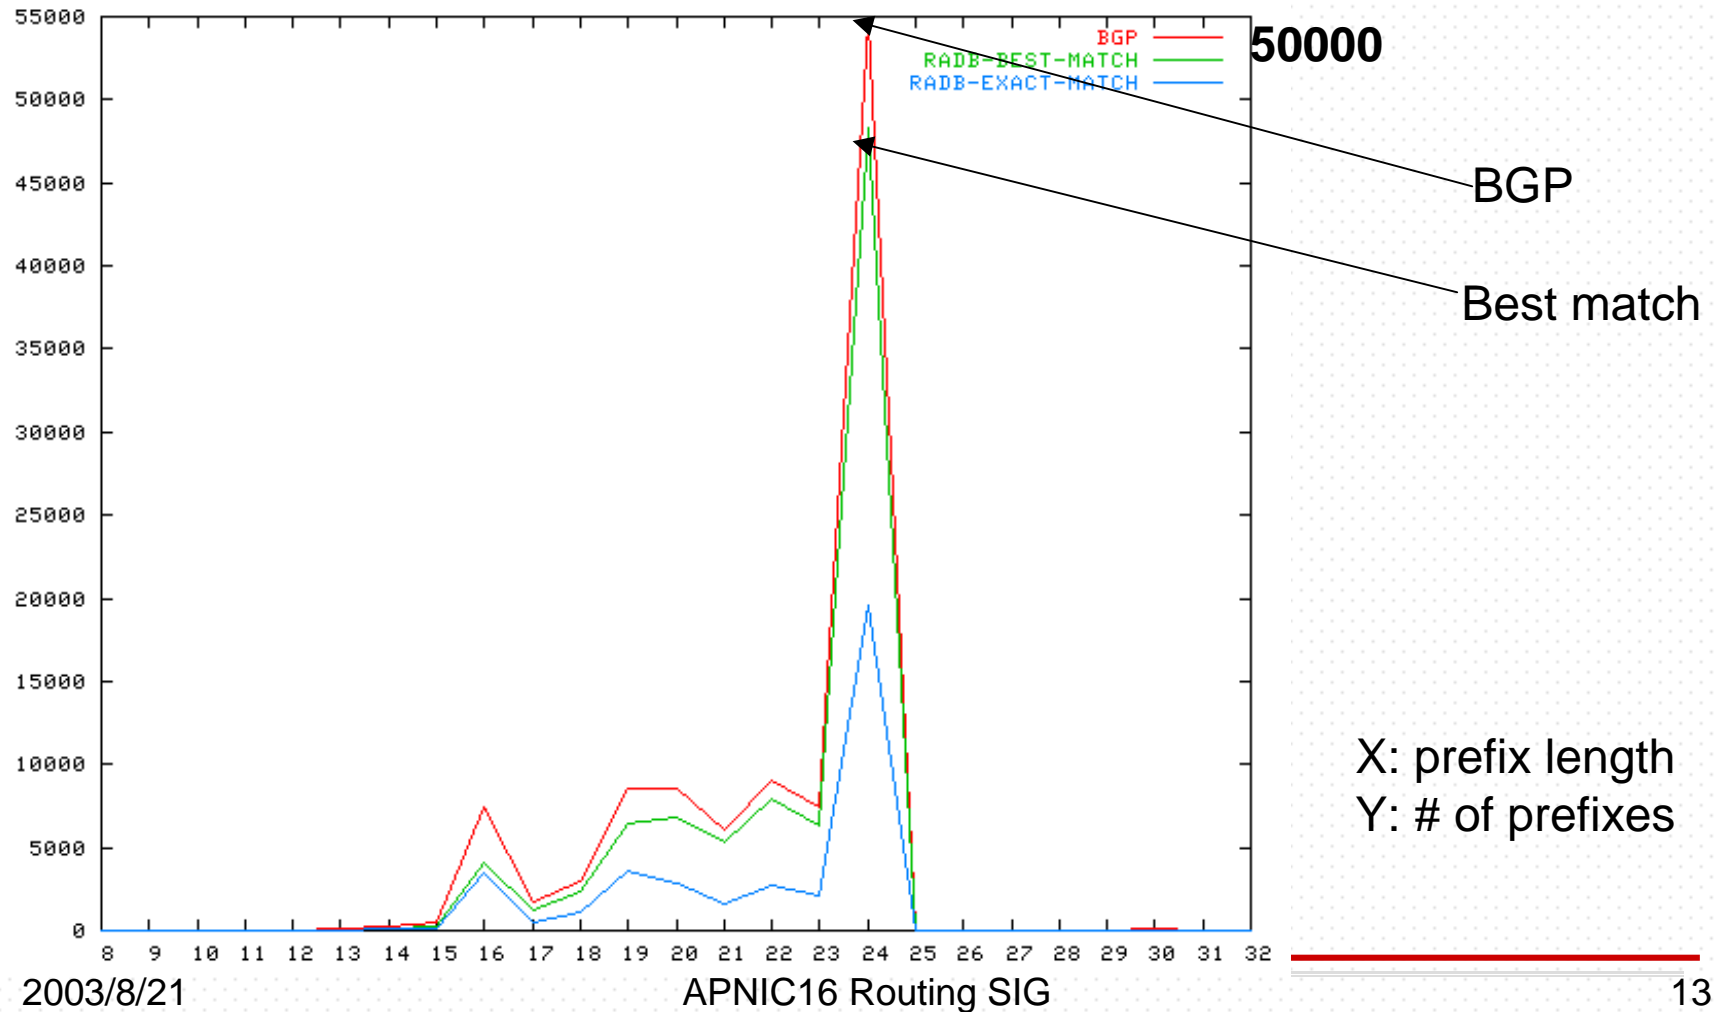
Snapshot 2 (2003/08/11)

prefixlen	BGP	Unified IRR		RADB		RIPE		APNIC		JPIRR	
		best	exact	best	exact	best	exact	best	exact	best	exact
8	16	8	8	7	7	3	3	0	0	0	0
9	5	4	2	4	2	0	0	0	0	0	0
10	8	5	3	4	2	1	1	0	0	0	0
11	13	8	6	4	2	4	4	0	0	0	0
12	54	31	26	25	20	8	6	0	0	0	0
13	97	50	39	39	28	15	12	0	0	0	0
14	260	163	137	108	82	60	56	2	2	1	1
15	476	330	222	181	89	163	135	0	0	9	8
16	7,390	5,180	4,529	3,567	2,990	1,820	1,658	10	10	96	64
17	1,684	1,480	748	1,218	465	455	309	46	4	127	13
18	3,005	2,663	1,553	2,270	1,026	798	581	22	4	175	22
19	8,544	7,949	5,467	6,261	3,418	2,914	2,474	78	41	447	25
20	8,564	7,962	4,293	6,647	2,936	2,009	1,487	223	28	476	64
21	6,032	5,639	2,285	5,134	1,690	1,133	660	236	7	172	6
22	9,060	8,644	3,792	7,865	2,819	1,973	1,148	314	16	196	12
23	7,448	7,132	3,351	6,217	2,236	1,933	1,303	271	7	140	8
24	54,664	51,927	26,149	47,034	19,717	12,986	7,597	2,437	784	637	32
25	24	24	2	24	2	0	0	0	0	0	0
26	20	20	4	19	3	9	1	0	0	0	0
27	11	11	3	11	3	0	0	0	0	0	0
28	11	11	2	11	2	0	0	0	0	0	0
29	24	24	19	24	19	0	0	0	0	0	0
30	4	0	0	0	0	0	0	0	0	0	0
31	0	0	0	0	0	0	0	0	0	0	0
32	19	0	0	0	0	0	0	0	0	0	0
total	107,433	99,265	52,640	86,674	37,558	26,284	17,435	3,639	903	2,476	255

Snapshot 3 (2003/08/11) Unified DB



Snapshot 4 (2003/08/11) RADB



Best Match

- ❑ Utilization of Best Match is high (94%, /24)
- ❑ Utilization of Exact Match is low (47%, /24)
- ❑ Best Match is indeed valid method?
- ⇒ Validation of Origin-AS

Validation of Origin-AS

IRR	total # of best match prefix	correct origin	incorrect origin	incorrect origin ratio(%)
RADB	86,674	29,845	56,829	65.57%
RIPE	26,284	17324	8960	34.09%

- Why high ratio of incorrect origin data?
 - Obsolete Data
 - E.g. invalid origin at RADB
 - AS11908 2749 best matched prefix
 - AS1 2003 best matched prefix
 - AS6389 1735 best matched prefix

Conclusion (1)

1. How many prefixes are registered in IRR?

- ❑ Unified DB (Average):

- ❑ Exact Match: 49%

- ❑ Best Match: 92.4%

- ❑ RADB (Average):

- ❑ Exact Match: 34.96%

- ❑ Best Match: 80.68%

- ❑ RIPE (Average):

- ❑ Exact Match: 24.47%

- ❑ Best Match: 16.23%

- ❑ APNIC

- ❑ Exact Match: 0.84%

- ❑ Best Match : 3.39%

- ❑ JPIRR

- ❑ Exact Match: 0.24%

- ❑ Best Match : 2.30%

Conclusion (2)

2. What differences in IRRs (e.g. RADB, RIPE, APNIC)?
 - ❑ RADB stores many un-maintained objects
 - ❑ RIPE stores more maintained objects than RADB
 - ⇒ difference in origin-AS ratio
 - ❑ RADB: incorrect origin-AS ratio: 65.57%
 - ❑ RIPE: incorrect origin-AS ratio: 34.09%

Conclusion (3)

3. Is IRR practical for router configuration?

Current:

- Accuracy of RIPE IRR is high
- Therefore, it is relatively practical to make router configuration using RIPE authoritative IRR objects

Future:

- Further investigation is needed to clarify these differences between RIPE and other databases

URL

<http://jpirr.nic.ad.jp/stat/>