

Internet Routing Registry Tutorial

July 15, 2012, Karachi, Pakistan

In conjunction with

The logo for SANDOG, consisting of the word "SANDOG" in a bold, white, sans-serif font centered within a black rectangular box.

SANDOG

APNIC



Presenters

- Champika Wijayatunga
 - Training Unit Manager, APNIC
 - champika@apnic.net

- Vivek Nigam
 - Internet Resource Analyst, APNIC
 - vivek@apnic.net

Objectives

- To provide an introduction to the APNIC Routing Registry
 - Explain concepts of the global RR
 - Outline the benefits of the APNIC Routing Registry
 - Discuss Routing Policy Specification Language (RPSL)

Overview

- Whois DB Recap
- What is IRR?
- APNIC database and the IRR
- Using the Routing Registry
- Using RPSL in practice
- Benefit of using IRR

Whois Database Recap

APNIC Database

- Public network management database
 - APNIC whois database contains:
 - Internet resource information and contact details
 - APNIC Routing Registry (RR) contains:
 - routing information
- APNIC RR is part of IRR
 - Distributed databases that mirror each other

Database Object

- An object is a set of attributes and values
 - Each attribute of an object...
 - Has a value
 - Has a specific syntax
 - Is mandatory or optional
 - Is single- or multi-valued
 - Some attributes ...
 - Are primary (unique) keys
 - Are lookup keys for queries
 - Are inverse keys for queries
- Object “templates” illustrate this structure

Person Object Example

- Person objects contain contact information

Attributes

Values

person:	Test Person
address:	ExampleNet Service Provider
address:	2 Pandora St Boxville
address:	Wallis and Futuna Islands
country:	TC
phone:	+680-368-0844
fax-no:	+680-367-1797
e-mail:	tperson@example.com
nic-hdl:	TP17-AP
mnt-by:	MAINT-ENET-TC
changed:	tperson@example.com 20090731
source:	APNIC

Database Queries

- Flags used for inetnum queries

None find exact match

- l find one level less specific matches
- L find all less specific matches
- m find first level more specific matches
- M find all More specific matches
- x find exact match (if no match, nothing)
- d enables use of flags for reverse domains
- r turn off recursive lookups

Database Protection

- Authorisation
 - “mnt-by” references a mntner object
 - Can be found in all database objects
 - “mnt-by” should be used with every object!
- Authentication
 - Updates to an object must pass authentication rule specified by its maintainer object



Prerequisite for Updating Objects

- Create person objects for contacts
 - To provide contact info in other objects
- Create a mntner object
 - To provide protection of objects
- Protect your person object

What is IRR?

What is a Routing Registry?

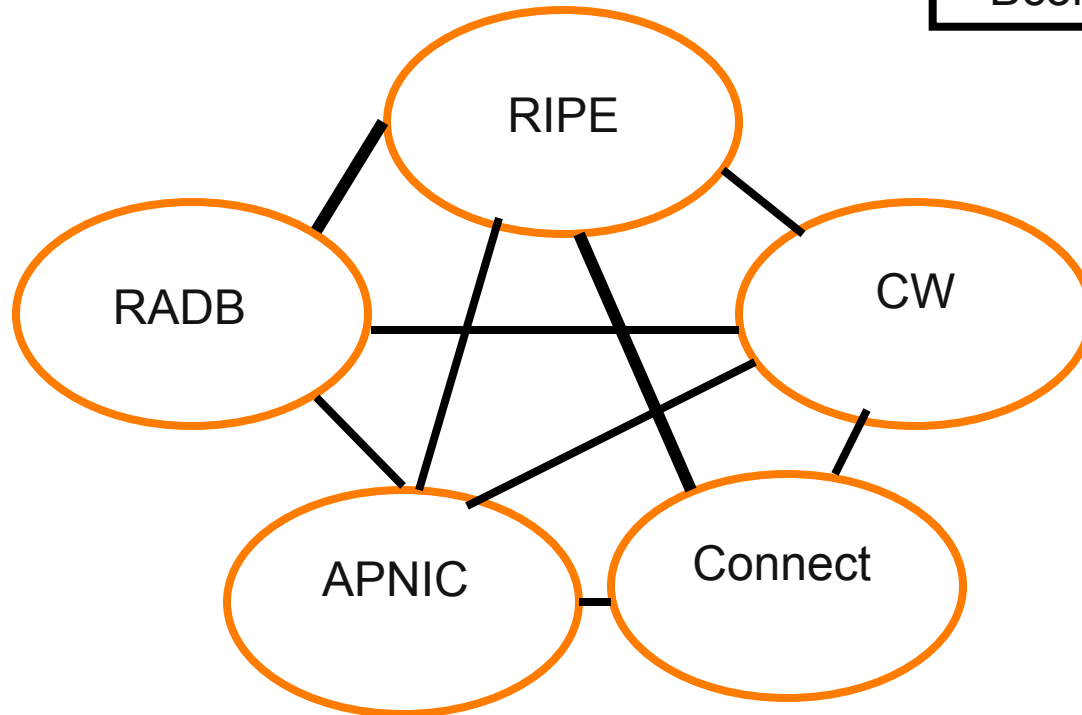
- A repository (database) of Internet routing policy information
 - Autonomous Systems exchanges routing information via BGP
 - Exterior routing decisions are based on policy based rules
 - However BGP does not provides a mechanism to publish/communicate the policies themselves
 - RR provides this functionality
- Routing policy information is expressed in a series of objects

What is a Routing Registry?

- Global Internet Routing Registry database
 - <http://www.irr.net/>
 - Uses RPSL
- Stability and consistency of routing
 - network operators share information
- Both public and private databases
 - These databases are independent
 - But some exchange data
 - Only register your data in one database

What is a Routing Registry?

ARIN, ArcStar, FGC, Verio,
Bconnex, Optus, Telstra, ...

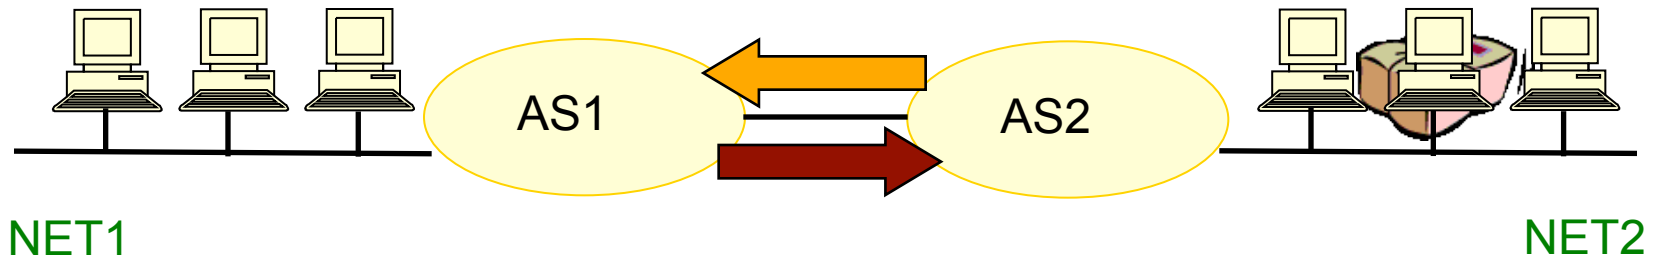


IRR = APNIC RR + RIPE DB + RADB + C&W + ARIN + ...

What is Routing Policy?

- Description of the routing relationship between autonomous systems
 - Who are my BGP peers?
 - Customer, peers, upstream
 - What routes are:
 - Originated by each neighbour?
 - Imported from each neighbour?
 - Exported to each neighbour?
 - Preferred when multiple routes exist?
 - What to do if no route exists?
 - What routes to aggregate?

Representation of Routing Policy



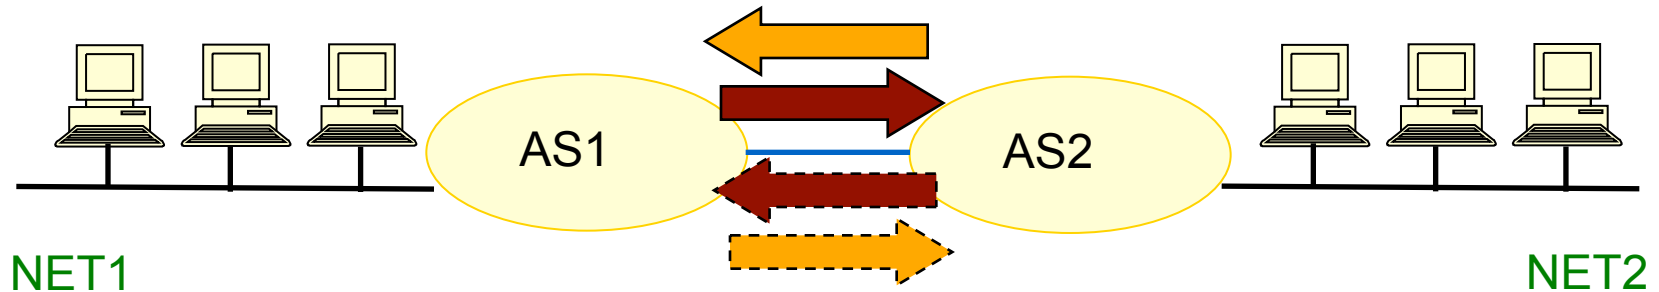
In order for traffic to flow from NET2 to NET1 between AS1 and AS2:

AS1 has to announce NET1 to AS2 via BGP

And AS2 has to accept this information and use it

Resulting in packet flow from NET2 to NET1

Representation of Routing Policy (cont.)



In order for traffic to flow towards from NET1 to NET2:

AS2 must announce NET2 to AS1

And AS1 has to accept this information and use it

Resulting in packet flow from NET 1 to NET2

RPSL

- Routing Policy Specification Language
 - Object oriented language
 - Based on RIPE-181
 - Structured whois objects



- Higher level of abstraction than access lists

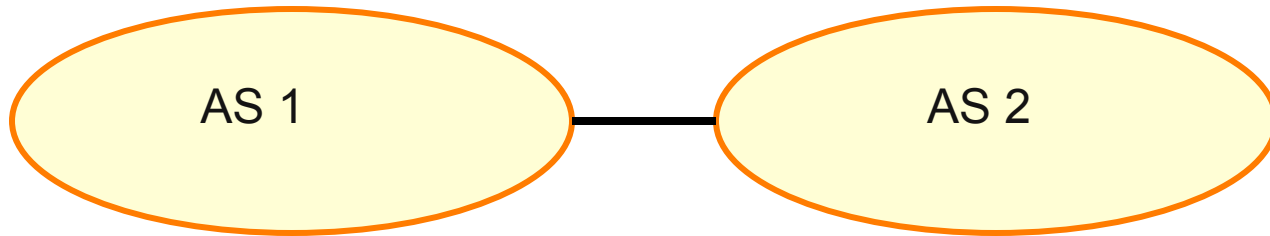


- Describes things interesting to routing policy:
 - Routes, AS Numbers ...
 - Relationships between BGP peers
 - Management responsibility



Routing Policy - Examples

Basic concept



“action pref” - the lower the value, the preferred the route

```
aut-num: AS1
...
import: from AS2
      action pref= 100
      accept AS2
export: to AS2 announce AS1
```

```
aut-num: AS2
...
import: from AS1
      action pref=100;
      accept AS1
export: to AS1 announce AS2
```

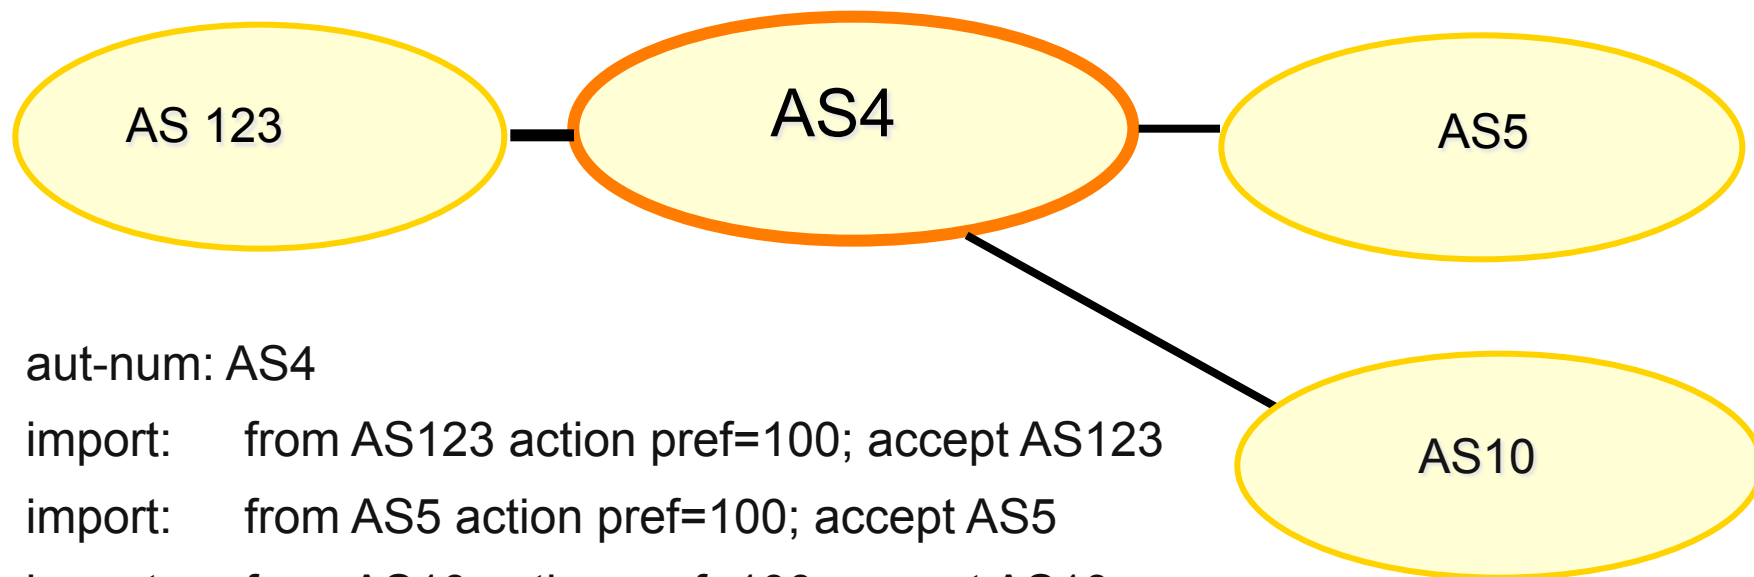
Routing Policy - Examples



More complex example

- AS4 gives transit to AS5, AS10
- AS4 gives local routes to AS123

Routing Policy - Examples



aut-num: AS4

import: from AS123 action pref=100; accept AS123

import: from AS5 action pref=100; accept AS5

import: from AS10 action pref=100; accept AS10

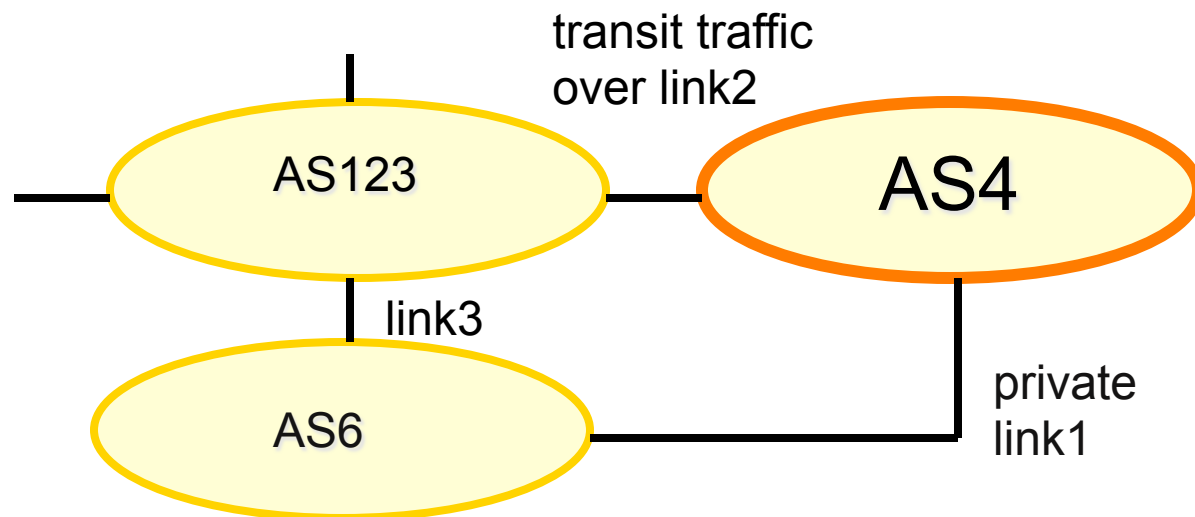
export: to AS123 announce AS4

export: to AS5 announce AS4 AS10

export: to AS10 announce AS4 AS5

← *Not a path*

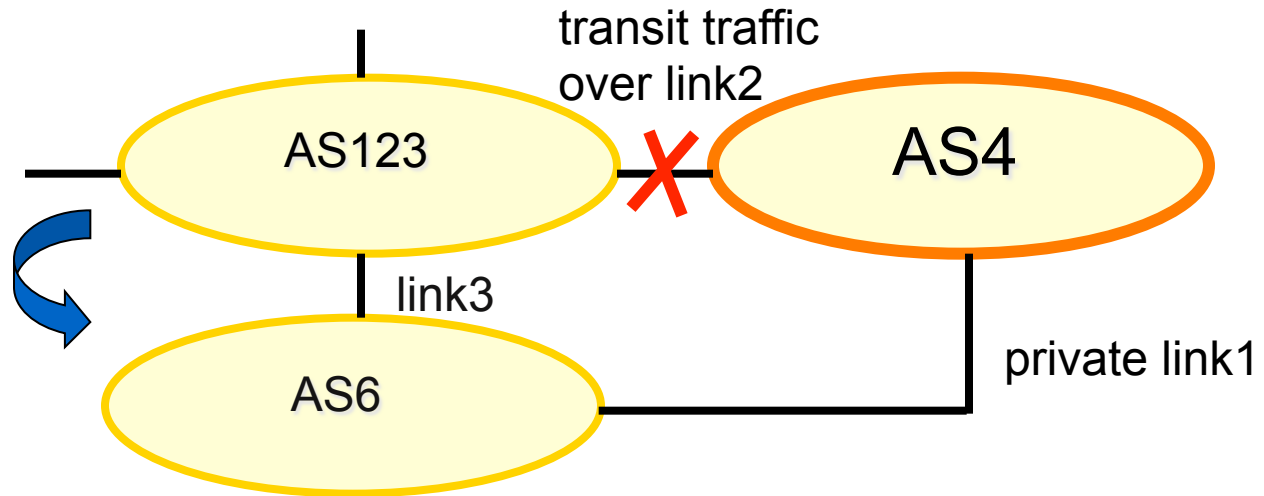
Routing Policy - Examples



More complex example

- AS4 and AS6 private link1
- AS4 and AS123 main transit link2
- backup all traffic over link1 and link3 in event of link2 failure

Routing Policy - Examples



aut-num: AS4

import: from AS123 action pref=100; accept ANY ← *full routing received*

import: from AS6 action pref=50; accept AS6

import: from AS6 action pref=200; accept ANY

export: to AS6 announce AS4

export: to AS123 announce AS4

higher cost for backup route

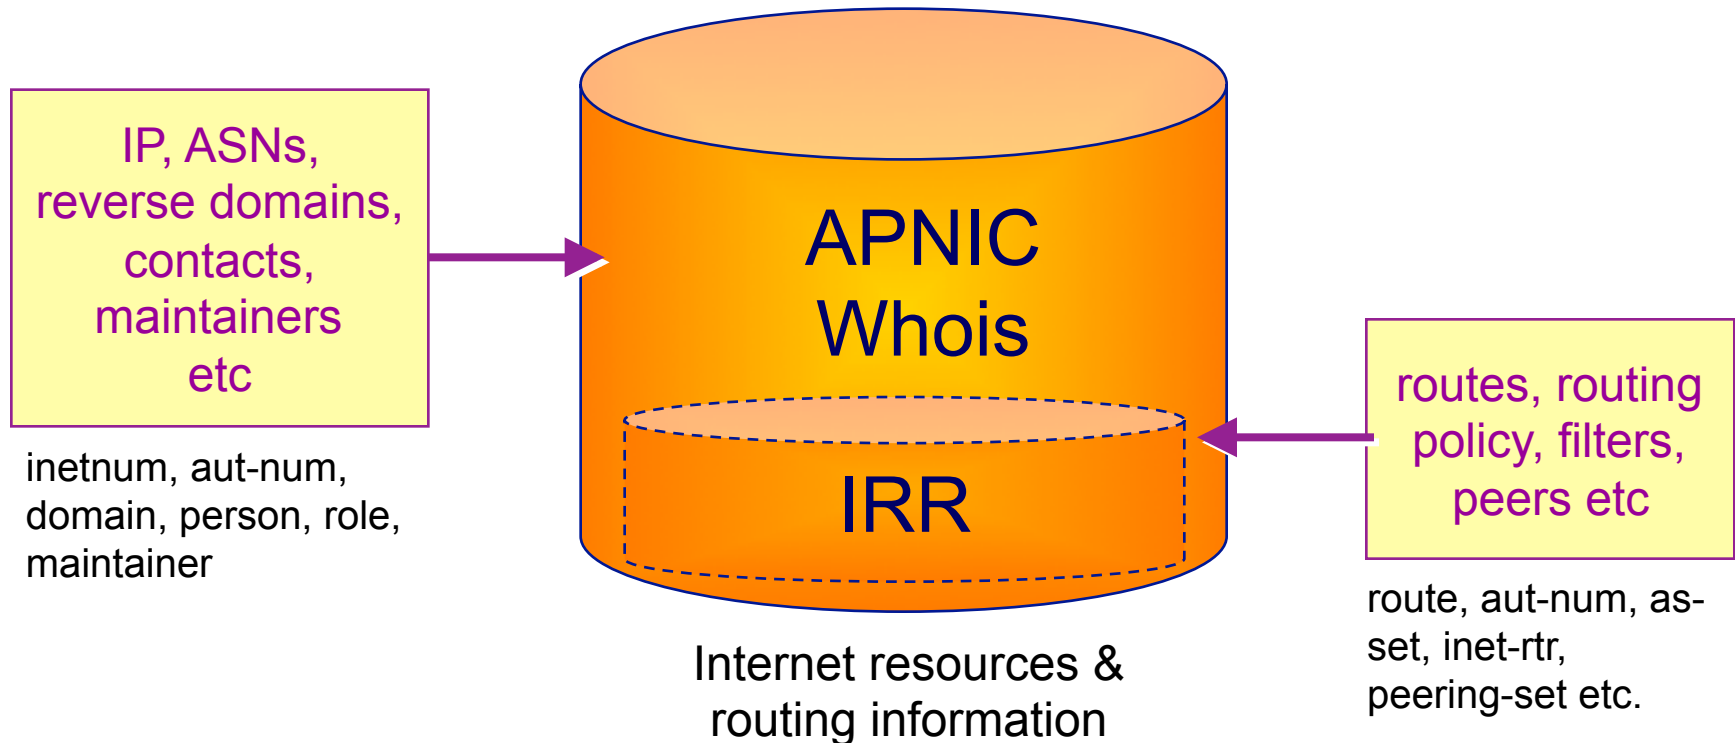
APNIC Database and the IRR

APNIC Database & the IRR

- APNIC whois Database
 - Two databases in one
- Public Network Management Database
 - “whois” info about networks & contact persons
 - IP addresses, AS numbers etc
- Routing Registry
 - contains routing information
 - routing policy, routes, filters, peers etc.
 - APNIC RR is part of the global IRR

Integration of Whois and IRR

- Integrated APNIC Whois Database & Internet Routing Registry

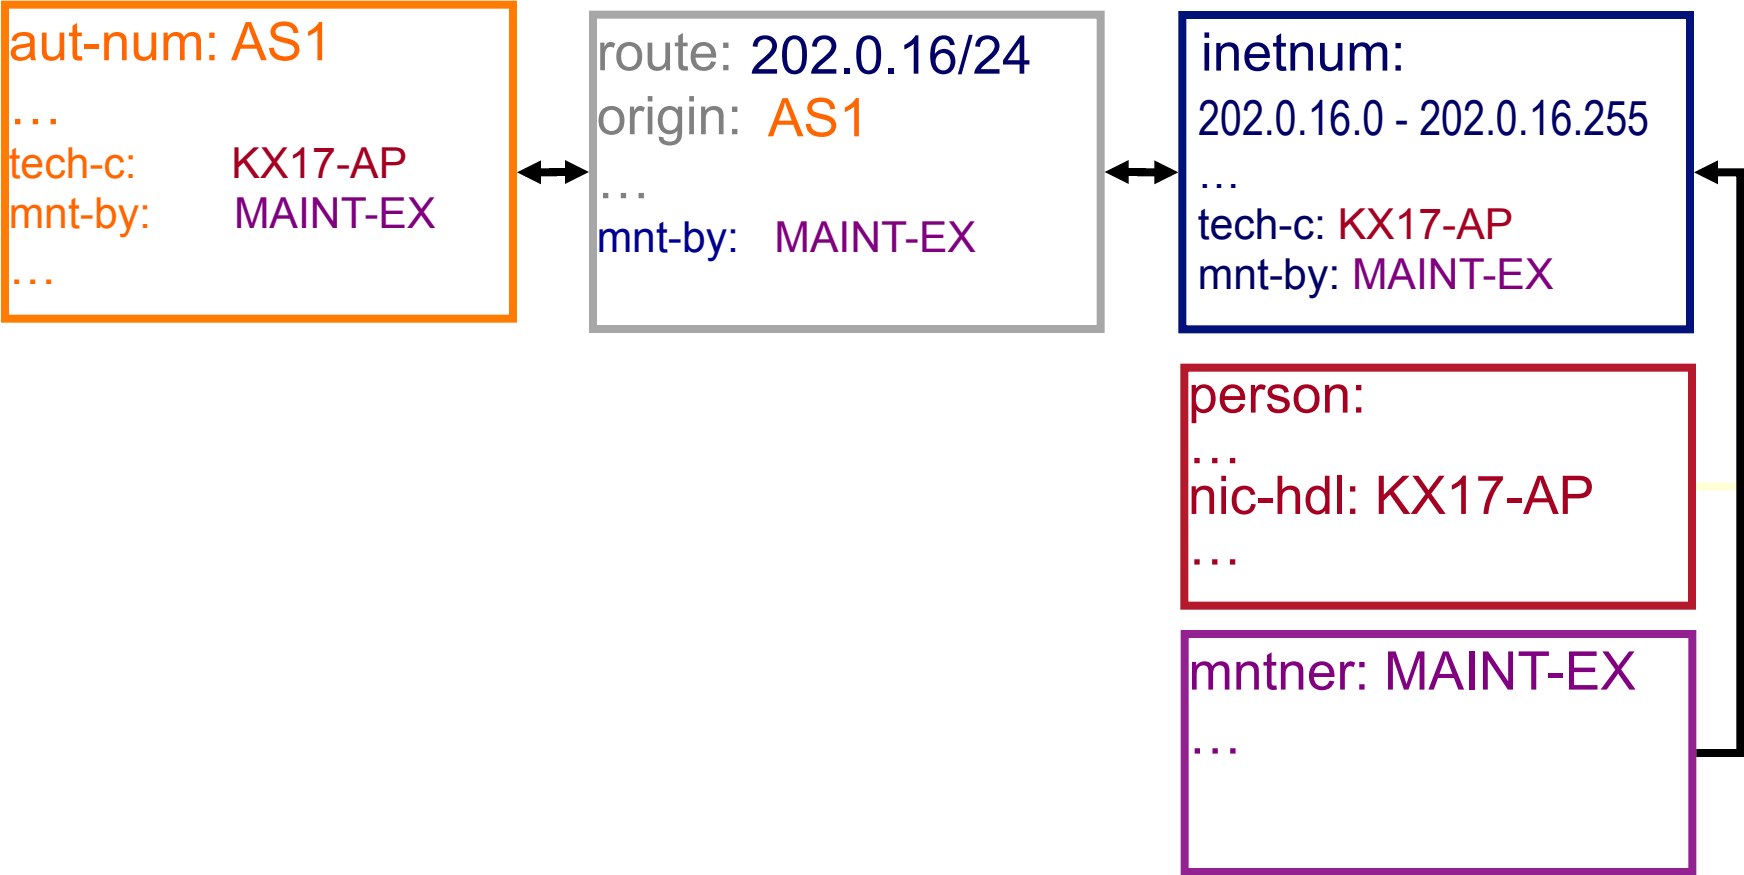


IRR Objects

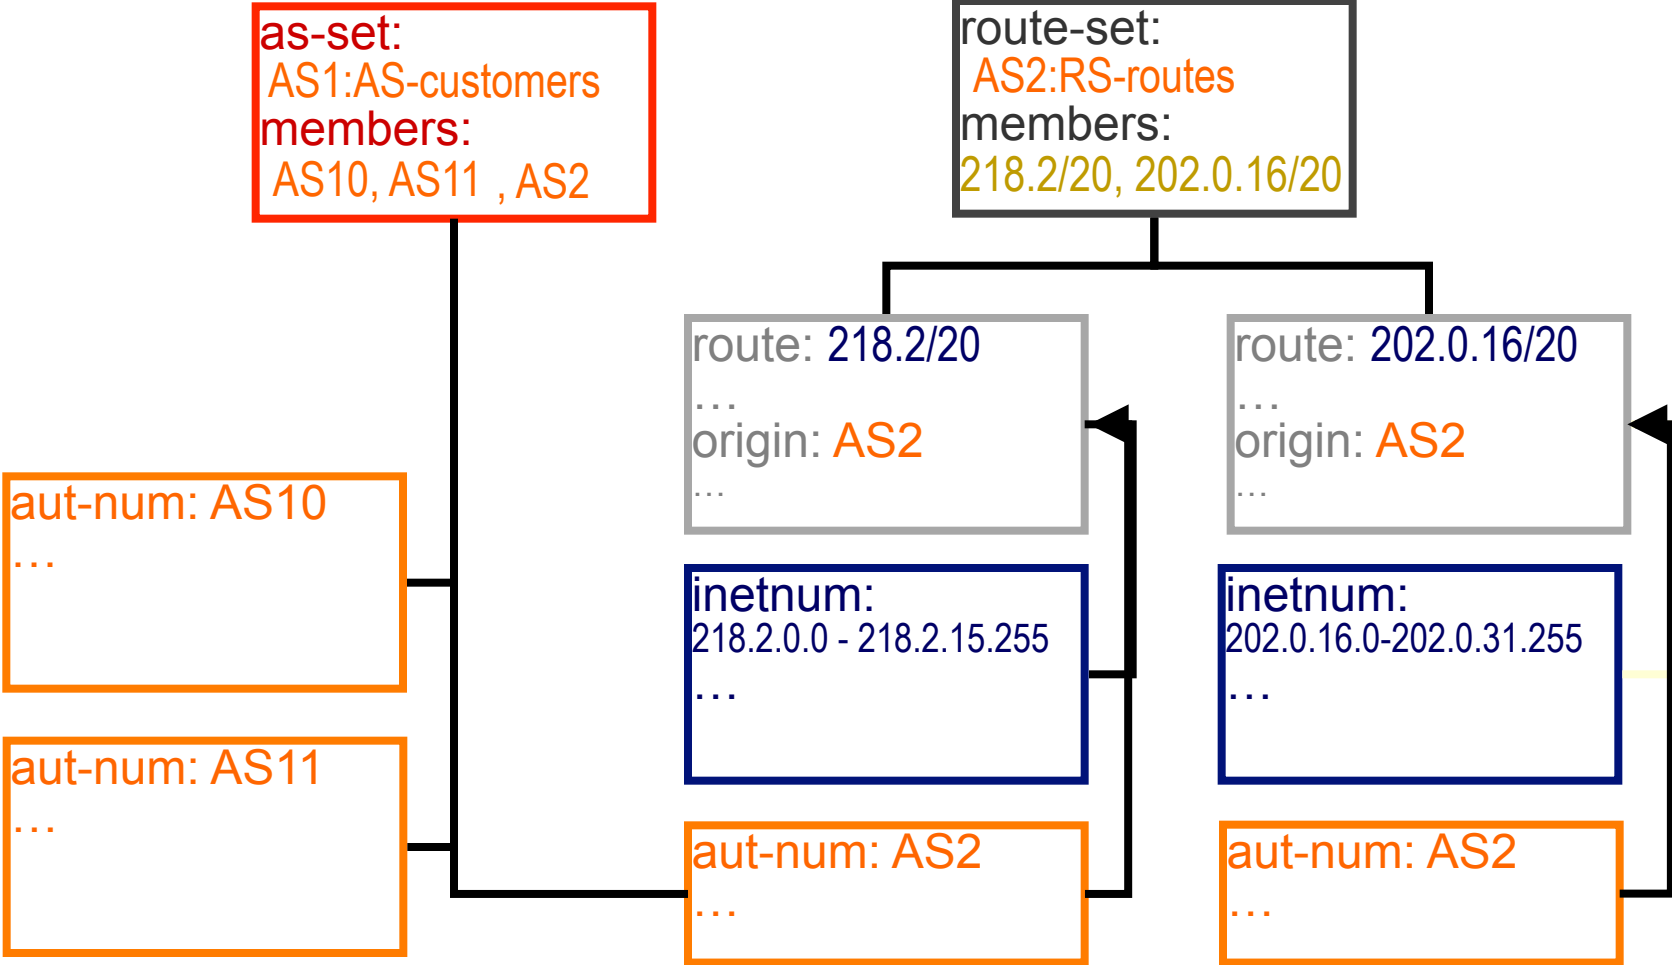
- route
 - Specifies interAS routes
- aut-num
 - Represents an AS. Used to describe external routing policy
- inet-rtr
 - Represents a router
- peering-set
 - Defines a set of peerings
- route-set
 - Defines a set of routes
- as-set
 - Defines a set of **aut-num** objects
- rtr-set
 - Defines a set of routers
- filter-set
 - Defines a set of routes that are matched by its filter

www.apnic.net/db/ref/db-objects.html

Inter-related IRR Objects



Inter-related IRR Objects



Hierarchical Authorisation

- **mnt-routes**

- authenticates *creation* of route objects

- creation of route objects must pass authentication of mntner referenced in the mnt-routes attribute

- Format:

- `mnt-routes: <mntner>`

In:

`inetnum`

`aut-num`

`route`

Authorisation Mechanism

```
inetnum:      202.137.181.0 - 202.137.196.255
netname:      SPARKYNET-TC
descr:        SparkyNet Service Provider
...
mnt-by:       APNIC-HM
mnt-lower:    MAINT-SPARKYNET1-TC
mnt-routes:   MAINT-SPARKYNET2-TC
```

This object can only be modified by APNIC

Creation of more specific objects (assignments) within this range has to pass the authentication of MAINT-SPARKYNET

Creation of route objects matching/within this range has to pass the authentication of MAINT-SPARKYNET-WF

Creating Route Objects

- Multiple authentication checks:
 - Originating ASN
 - mntner in the mnt-routes is checked
 - If no mnt-routes, mnt-lower is checked
 - If no mnt-lower, mnt-by is checked
 - AND the address space
 - Exact match & less specific route
 - mnt-routes etc
 - Exact match & less specific inetnum
 - mnt-routes etc
 - AND the route object mntner itself
 - The mntner in the mnt-by attribute

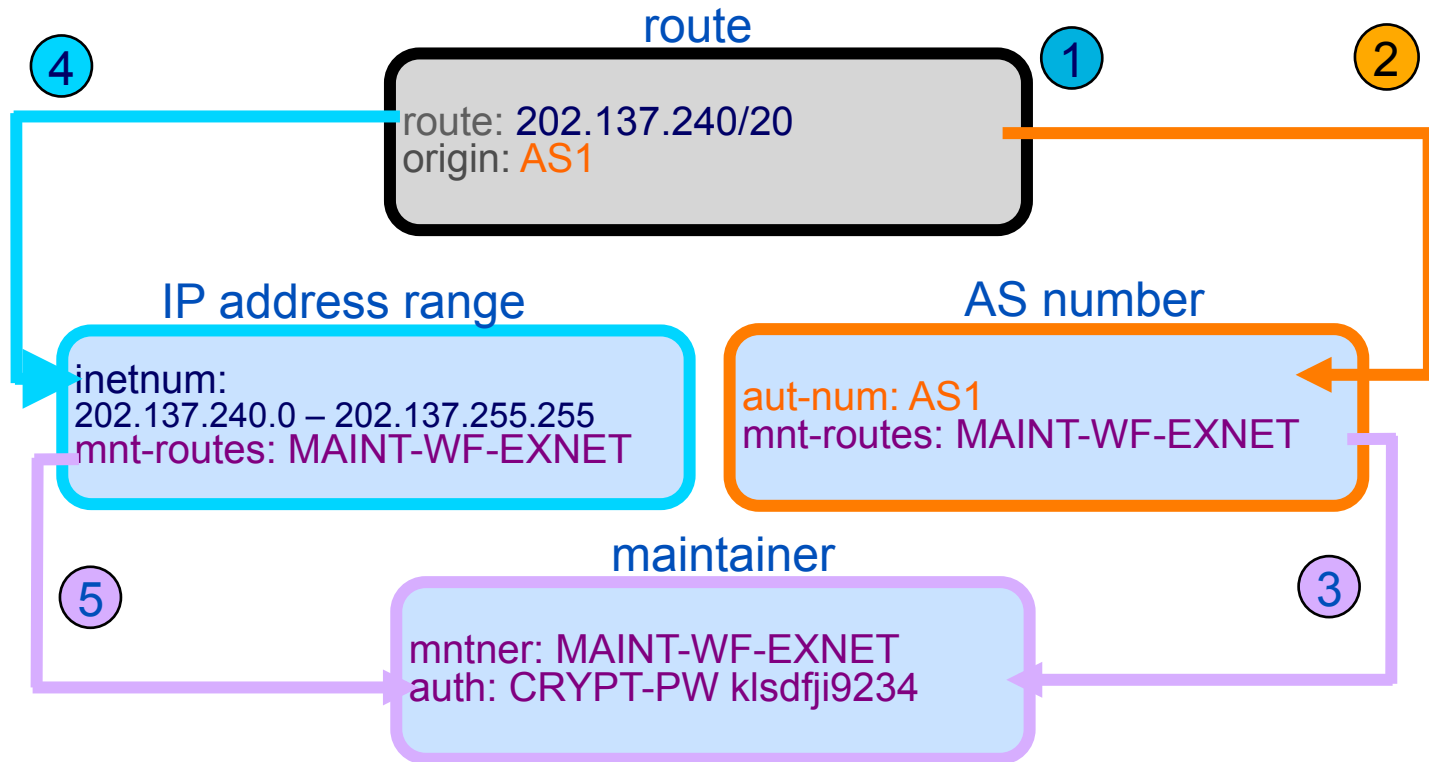
aut-num

inetnum

route

route

Creating Route Objects



1. Create route object and submit to APNIC RR database
2. DB checks aut-num obj corresponding to the ASN in route obj
3. Route obj creation must pass auth of mntner specified in aut-num *mnt-routes* attribute.
4. DB checks inetnum obj matching/encompassing IP range in route obj
5. Route obj creation must pass auth of mntner specified in inetnum *mnt-routes* attribute.

Using the Routing Registry

IRRToolSet

- Set of tools developed for using the Internet Routing Registry (IRR)
- Work with Internet routing policies
 - These policies are stored in IRR in the Routing Policy Specification Language (RPSL)
- The goal of the IRRToolSet is to make routing information more convenient and useful for network engineers
 - Tools for automated router configuration,
 - Routing policy analysis
 - On-going maintenance etc.

IRRToolSet

- Now maintained by ISC:
 - <http://www.isc.org/software/irrtoolset>
 - Download: <ftp://ftp.isc.org/isc/IRRToolSet/>
 - Installation needs: lex, yacc and C++ compiler

Use of RPSL - RtConfig

- RtConfig
 - part of IRRToolSet
- Reads policy from IRR (aut-num, route & -set objects) and generates router configuration
 - vendor specific:
 - Cisco, Bay's BCC, Juniper's Junos and Gated/RSd
 - Creates route-map and AS path filters
 - Can also create ingress / egress filters
 - (documentation says Cisco only)

Why use IRR and RtConfig?

- Benefits of RtConfig
 - Avoid filter errors (typos)
 - Expertise encoded in the tools that generate the policy rather than engineer configuring peering session
 - Filters consistent with documented policy
 - (need to get policy correct though)

Using RPSL in practice

Overview

- Review examples of routing policies expression
 - Peering policies
 - Filtering policies
 - Backup connection
 - Multihoming policies

RPSL - review

- Purpose of RPSL
 - Allows specification of your routing configuration in the public IRR
 - Allows you to check “Consistency” of policies and announcements
 - Gives opportunities to consider the policies and configuration of others

Address Prefix Range Operator

Operator	Meanings
\wedge_-	Exclusive more specifics of the address prefix: E.g. $128.9.0.0/16^{\wedge_-}$ contains all more specifics of $128.9.0.0/16$ excluding $128.9.0.0/16$
\wedge_+	Inclusive more specific of the address prefix: E.g. $5.0.0.0/8^{\wedge_+}$ contains all more specifics of $5.0.0.0/8$ including $5.0.0.0/8$

Address Prefix Operator (cont.)

Operator	Meanings
n	<p>n = integer, stands for all the length “n” specifics of the address prefix: E.g. 30.0.0.0/8¹⁶ contains all the more specifics of 30.0.0.0/8 which are length of 16 such as 30.9.0.0/16</p>
$^{n-m}$	<p>m = integer, stands for all the length “n” to length “m” specifics of the address prefix: E.g. 30.0.0.0/8²⁴⁻³² contains all the more specifics of 30.0.0.0/8 which are length of 24 to 32 such as 30.9.9.96/28</p>

AS-path regular expressions

- Regular expressions
 - A context-independent syntax that can represent a wide variety of character sets and character set orderings
 - These character sets are interpreted according to the current The Open Group Base Specifications (IEEE)
- Can be used as a policy filter by enclosing the expression in “<” and “>”.

Filter List- Regular Expression

- Like Unix regular expressions
 - . Match one character
 - * Match any number of preceding expression
 - + Match at least one of preceding expression
 - ^ Beginning of line
 - \$ End of line
 - \ Escape a regular expression character
 - _ Beginning, end, white-space, brace
 - | Or
 - () Brackets to contain expression
 - [] Brackets to contain number ranges

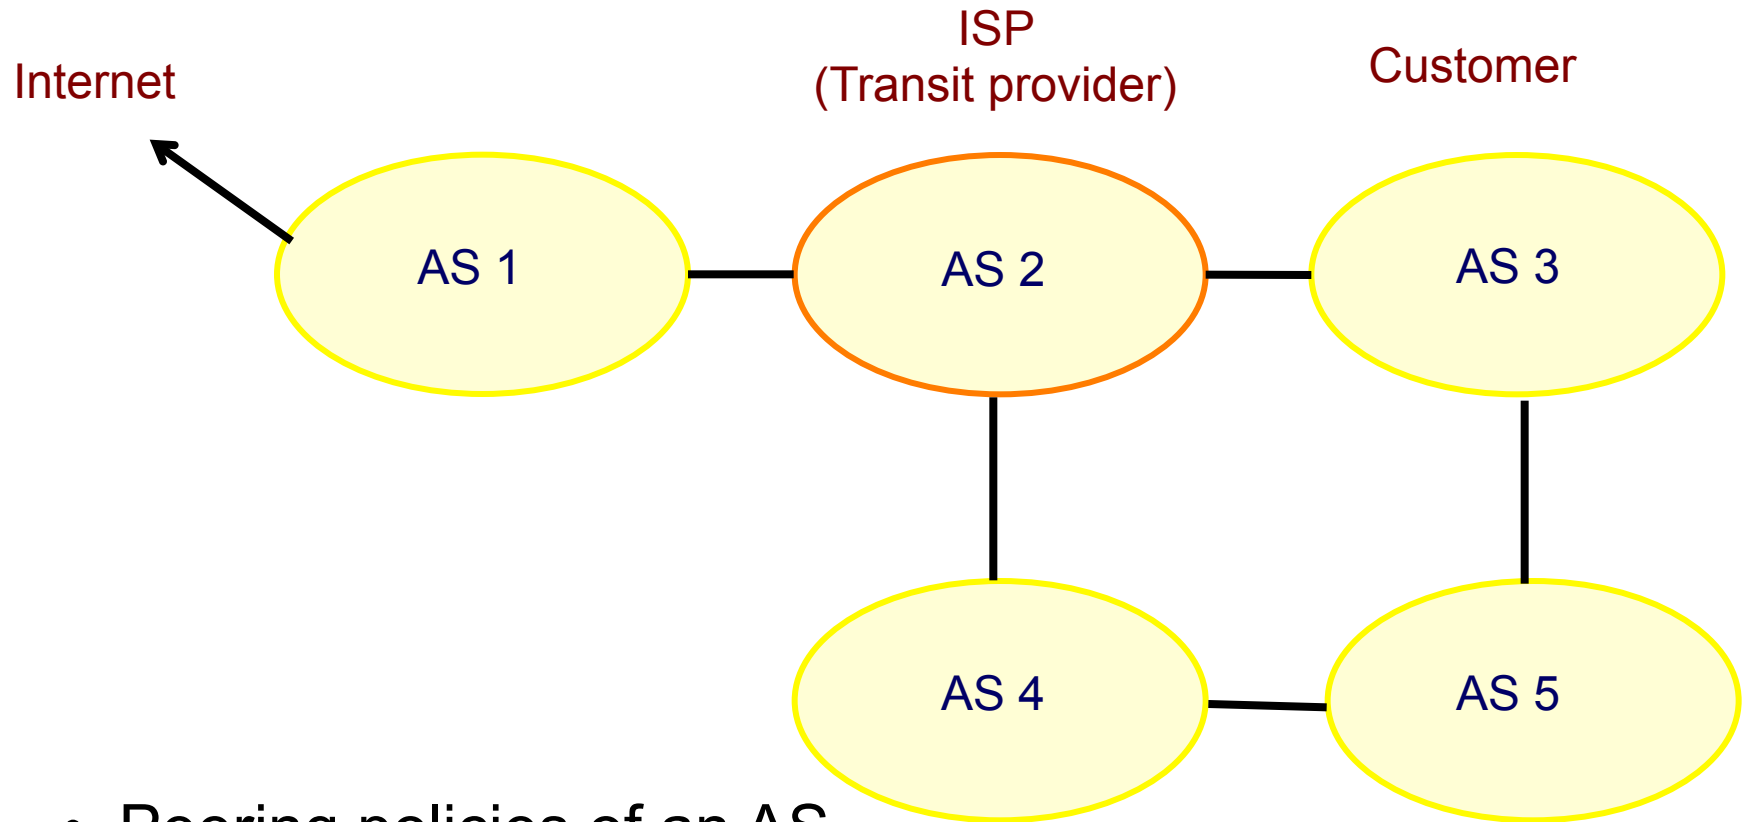
AS-path Regular Expression

Operator	Meanings
<AS3>	Route whose AS-path contains AS3
<^AS1>	Routes whose AS-path starts with AS1
<AS2\$>	Routes whose AS-path end with AS2
<^AS1 AS2 AS3\$>	Routes whose AS-path is exactly "1 2 3"
<^AS1 . * AS2\$>	AS-path starts with AS1 and ends in AS2 with any number ASN in between
<^AS3+\$>	AS-path starts with AS3 and ends in AS3 and AS3 is the first member of the path and AS3 occurs one or more times in the path and no other AS can be present in the path after AS3

AS-path Regular Expression (cont.)

Operator	Meanings
<AS3 AS4>	Routes whose AS-path is with AS3 or AS4
<AS3 AS4>	Routes whose AS-path with AS3 followed by AS4

Common Peering Policies



- Peering policies of an AS
 - Registered in an aut-num object

Common Peering Policies

- Policy for AS3 in the AS2 aut-num object

```
aut-num:      AS2
as-name:      SAMPLE-NET
dsdescr:      Sample AS
import:       from AS1 accept ANY
import:       from AS3 accept <^AS3+$>
export:       to AS3 announce ANY
export:       to AS1 announce AS2 AS3
admin-c:      TP1-AP
tech-c:       TP2-AP
mtn-by:       MAINT-SAMPLE-AP
changed:      sample@sample.net
```

ISP Customer – Transit Provider Policies

- Policy for AS3 and AS4 in the AS2 aut-num object

```
aut-num:      AS2
import:       from AS1 accept ANY
import:       from AS3 accept <^AS3+$>
import:       from AS4 accept <^AS4+$>
export:       to AS3 announce ANY
export:       to AS4 announce ANY
export:       to AS1 announce AS2 AS3 AS4
```

AS-set Object

- Describe the customers of AS2

```
as-set:      AS2:AS-CUSTOMERS
members:    AS3 AS4
changed:    sample@sample.net
source:     APNIC
```

Aut-num Object referring as-set Object

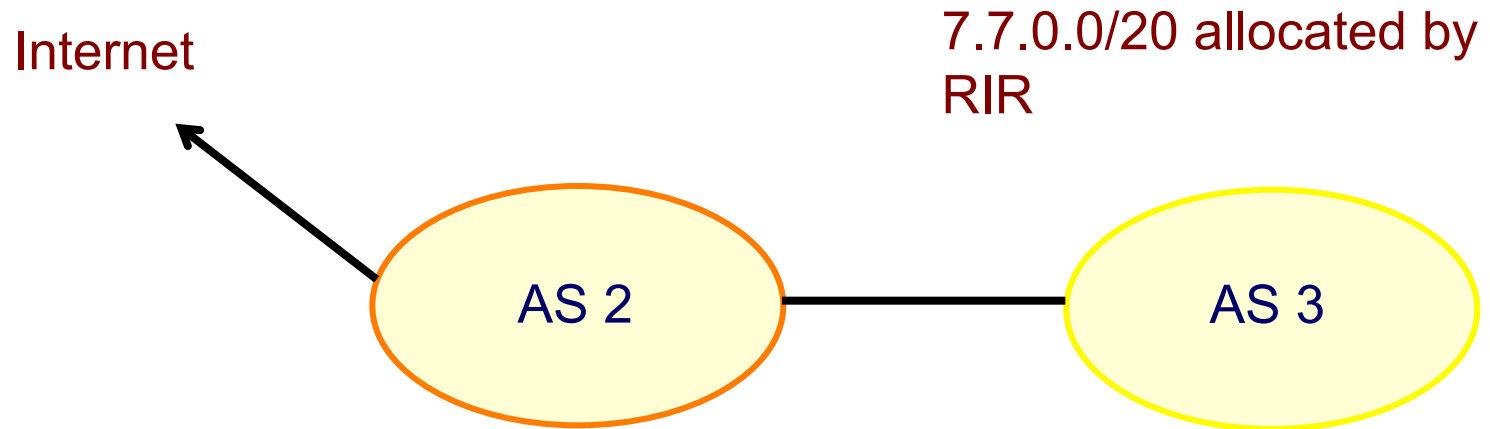
```
aut-num: AS2
import: from AS1 accept ANY
import: from AS2:AS-CUSTOMERS accept
<^AS2:AS-CUSTOMERS+$>
export: to AS2:AS-CUSTOMERS announce ANY
export: to AS1 announce AS2 AS2:AS-CUSTOMERS
```

```
aut-num: AS1
import: from AS2 accept <^AS2+AS2:AS-CUSTOMERS+$>
export: .....
```

Express Filtering Policy

- To limit the routes one accepts from a peer
 - To prevent the improper use of unassigned address space
 - To prevent malicious use of another organisation's address space

Filtering Policy



AS3 wants to announce part or all of 7.7.0.0/20 on the global Internet.

AS2 wants to be certain that it only accepts announcements from AS3 for address space that has been properly allocated to AS3.

Aut-num Object with Filtering Policy

```
aut-num:      AS2
import:      from AS3 accept { 7.7.0.0/20^20-24 }
.....
```

For an ISP with a growing or changing customer base, this mechanism will not scale well.

Route-set object can be used.

Route-set

```
route-set: AS2:RS-ROUTES:AS3
members: 7.7.0.0/20^20-24
changed: sample@sample.net
source: APNIC
```

Specifies the set of routes that will be accepted from a given customer

Set names are constructed hierarchically:

AS2 : RS-ROUTES : AS3



indicates whose sets these are



indicates peer AS

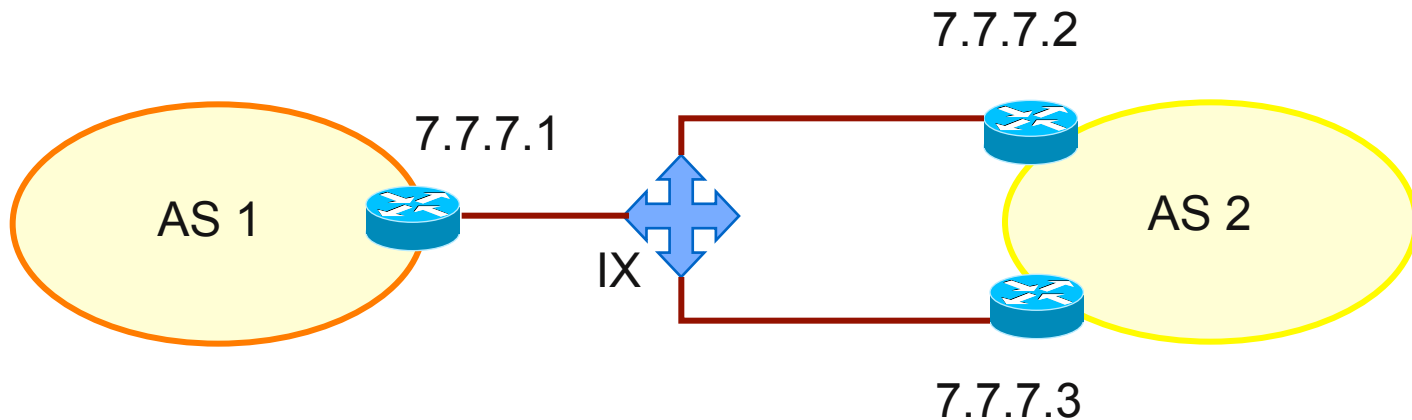
Filter configuration using route-set – AS2

```
import:    from AS1 accept ANY
import:    from AS3 accept AS2:RS-ROUTES:AS3
import:    from AS4 accept AS2:RS-ROUTES:AS4
export:    to AS2:AS-CUSTOMERS announce ANY
export:    to AS1 announce AS2 AS2:AS-CUSTOMERS
```

RPSL allows the peer's AS number to be replaced by the keyword PeerAS

```
import:    from AS2:AS-CUSTOMERS accept
           AS2:RS-ROUTES:PeerAS
```

Including interfaces in peering definitions: AS1



How to define AS1's routing policy by specifying its boundary router?

Including interfaces in peering definitions: AS1 (cont.)

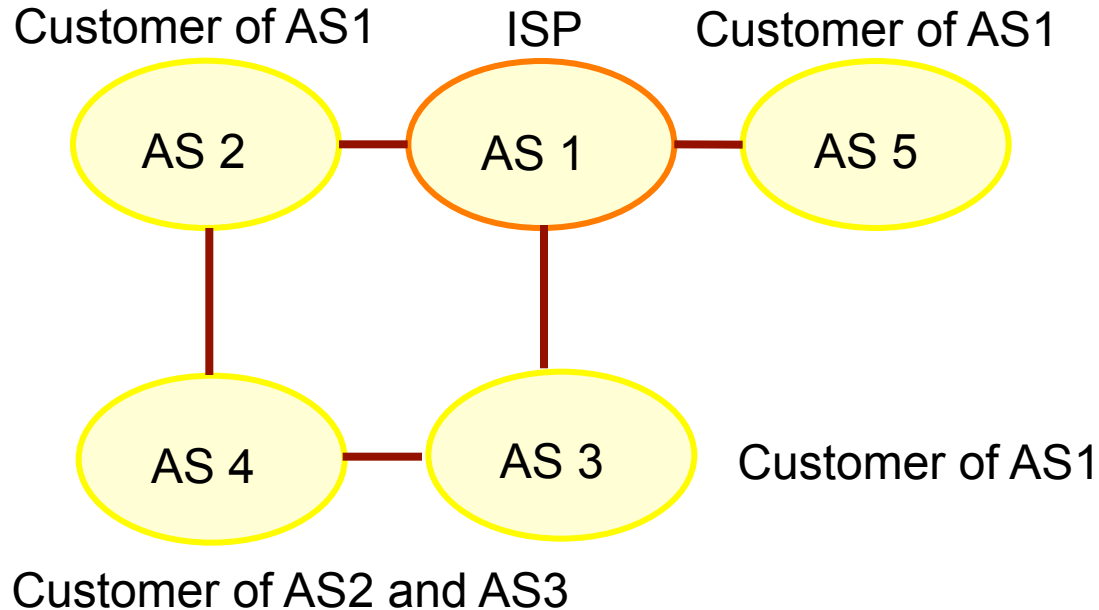
```
aut-num: AS1
import: from AS2 at 7.7.7.1 accept <^AS2+$>
```

AS1 may want to choose to accept:

- only those announcements from router 7.7.7.2
- discard those announcements from router 7.7.7.3

```
aut-num: AS1
import: from AS2 7.7.7.2 at 7.7.7.1 accept <^AS2+$>
```

Multihome Routing Policy



AS1's base policy

- Only accepts routes from customers that are originated by the customer or by the customer's customers

Multihome Routing Policies (cont.)

```
aut-num: AS1
import:  from AS2 accept (AS2 or AS4) AND
<^AS2+AS4*$>
import:  from AS3 accept (AS3 or AS4) AND
<^AS3+AS4*$>
import:  from AS5 accept AS5 AND <^AS5+$>
```

Benefit of using IRR

Using the Routing Registry



Costs

- Requires some initial planning
- Takes some time to define & register policy
- Need to maintain data in RR

Benefits

- You have a clear idea of your routing policy
- Consistent config over the whole network
- Less manual maintenance in the long run

APNIC RR service scope

- Routing Queries
 - Regular whois clients
 - APNIC whois web interface
 - Special purpose programs such as IRRToolSet
- Routing Registration and Maintenance
 - Similar to registration of Internet resources

Summary

- APNIC RR integrated in APNIC Whois DB
 - Facilitates network troubleshooting
 - Generation of router configuration
 - Provides global view of routing
- APNIC RR benefits
 - Single maintainer (& person obj) for all objects
 - APNIC asserts resources for a registered route
 - Part of the APNIC member service!

Questions?

Need any help?

Member Services Helpdesk

- One point of contact for all member enquiries
- Online chat services

Helpdesk hours

9:00 am - 9:00 pm (AU EST, UTC + 10 hrs)

ph: +61 7 3858 3188 fax: 61 7 3858 3199



- *More personalised service*
 - Range of languages:
Bahasa Indonesia, Bengali, Cantonese, English, Hindi, Mandarin, Thai, etc.
- *Faster response and resolution of queries*
 - IP resource applications, status of requests, obtaining help in completing application forms, membership enquiries, billing issues & database enquiries