APNIC Training

Internet Routing Registry (IRR)

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)

Assumptions

- The audience
  - Knowledgeable about Routing
  - Curious about Internet Routing Registry usage (IRR)
  - But not yet familiar with Routing Policy Specification Language (RPSL) and IRR

Overview

- What is IRR?
- Why use an IRR?
- APNIC database and the IRR
- Using the Routing Registry
- Using RPSL in practice
- Benefit of using 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, ...

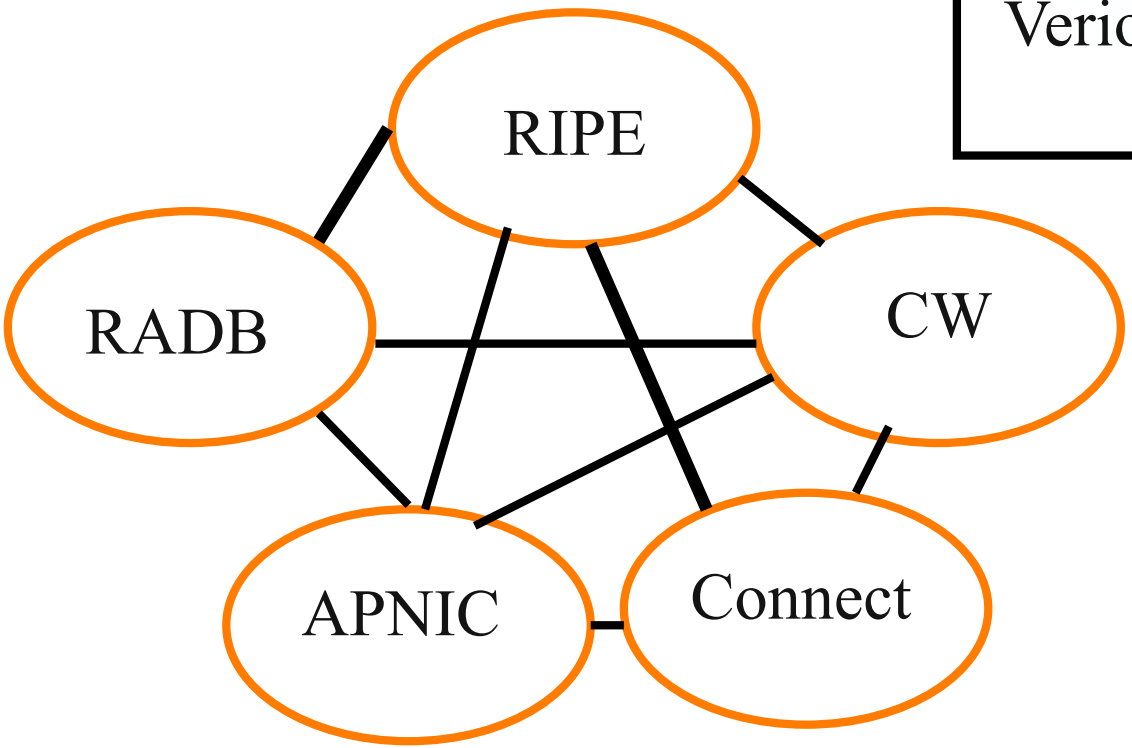

IRR = APNIC RR + RIPE DB + RADB + C&W + ARIN + …

Routing Registry Objects

- Route, aut-num, inet-rtr, peering-set, AS-set, rtr-set, filter-set
  - Each object has its own purpose
  - Together express routing policies
- More details covered later

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

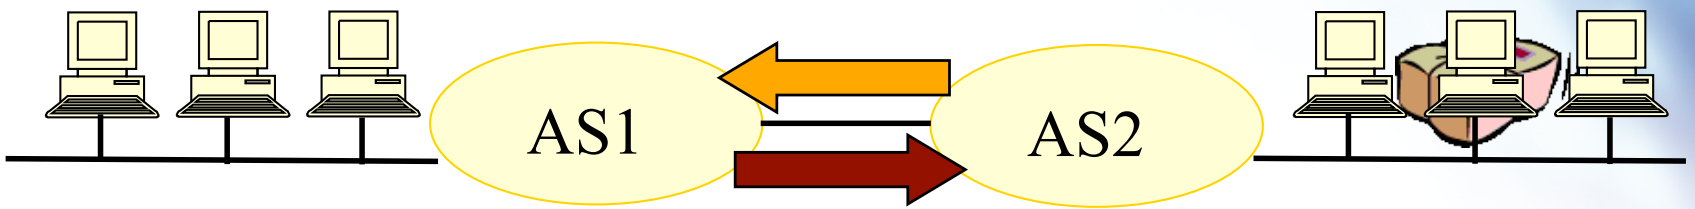

NET1 NET2

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.)

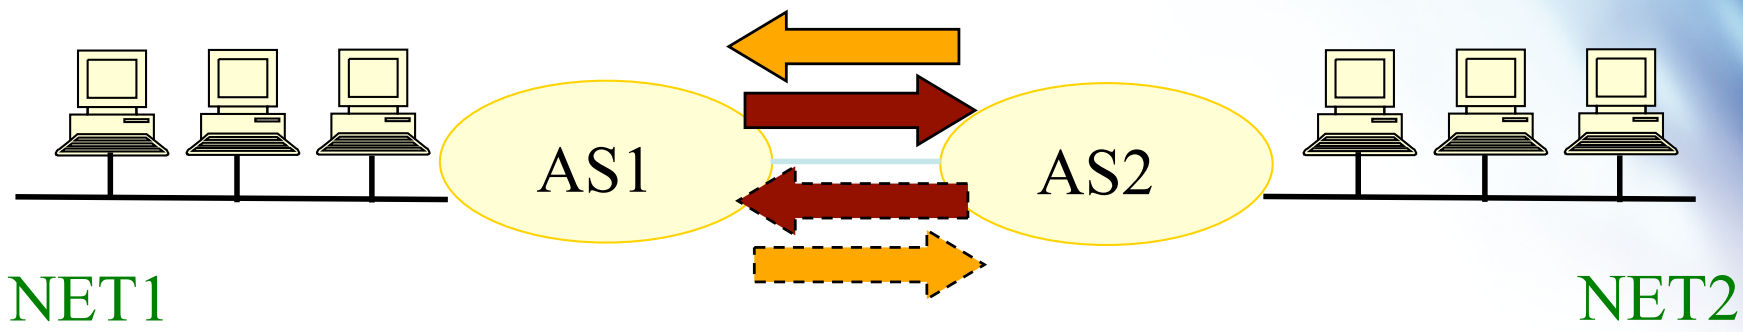

In order for traffic to flow towards from NET1 to NET2:

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
- Relevant RFCs
  - Routing Policy Specification Language
  - Routing Policy System Security
  - Using RPSL in Practice

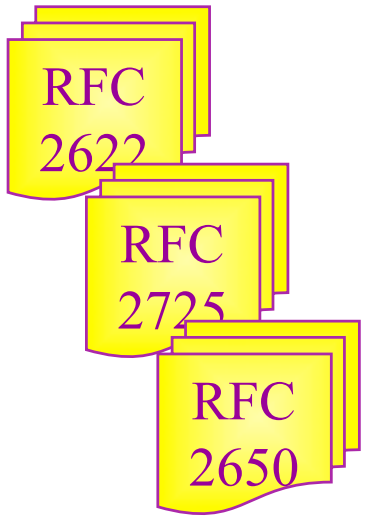

Routing Policy - Examples

Basic concept

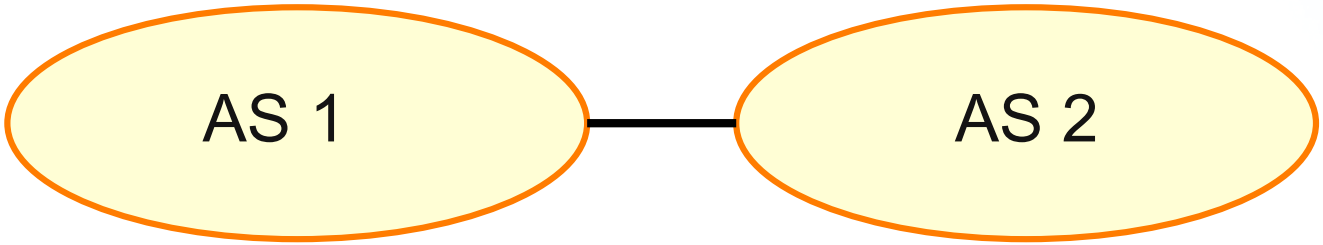

"action pref" - the lower the value, the preferred the route

```
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

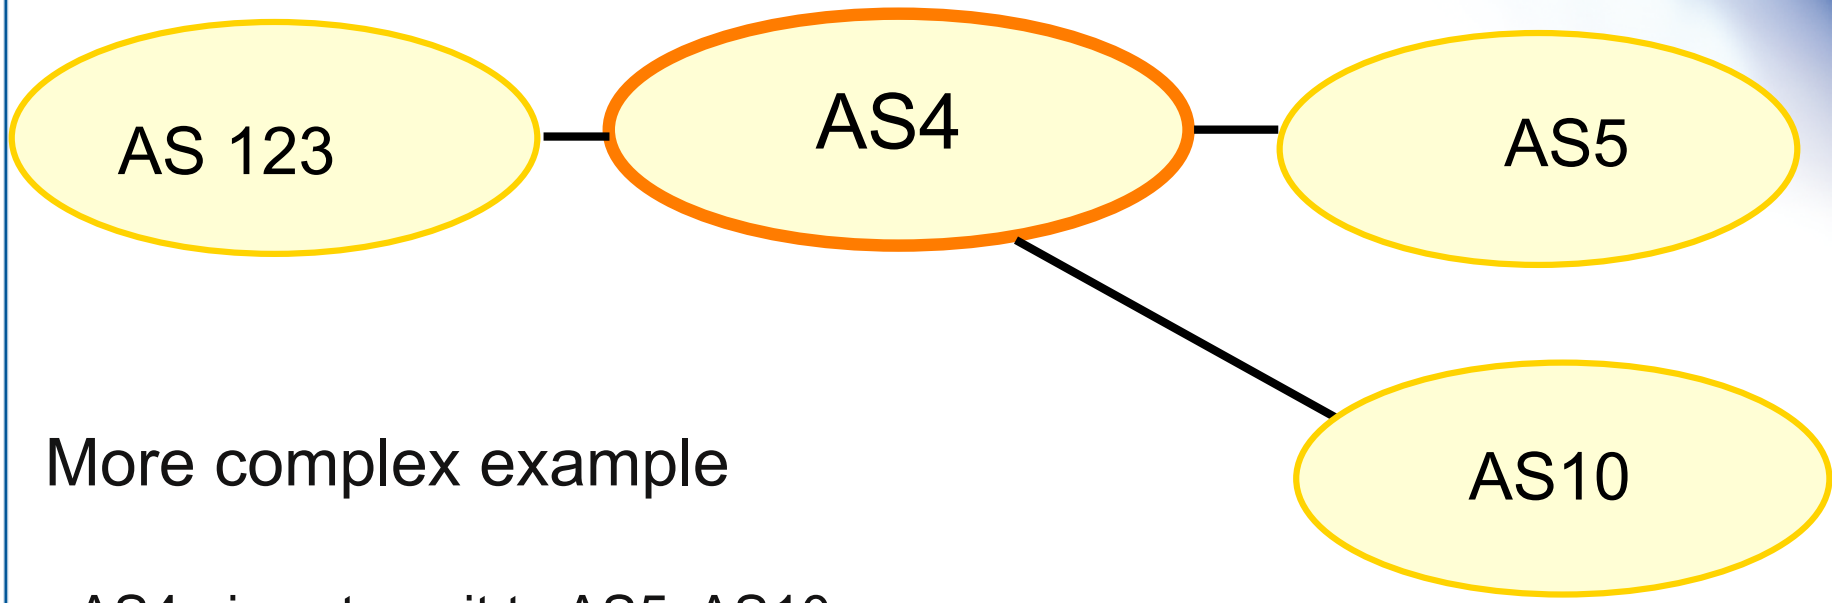

More complex example

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

Routing Policy - Examples

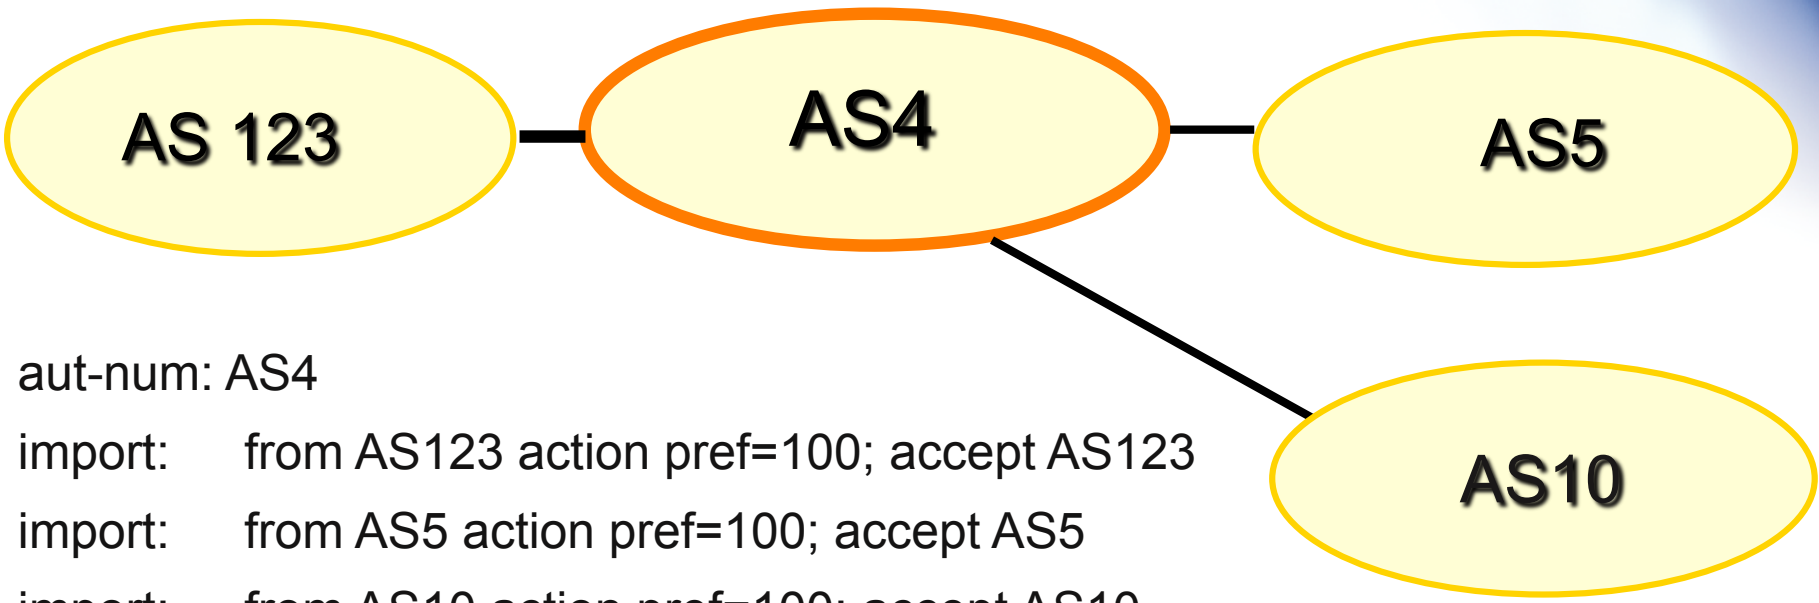

aut-num: AS4
import: from AS123 action pref=100; accept AS123
import: from AS5 action pref=100; accept AS5
import: from AS10 action pref=100; accept AS10
export: to AS123 announce AS4
export: to AS5 announce AS4 AS10
export: to AS10 announce AS4 AS5

Not a path

Routing Policy - Examples

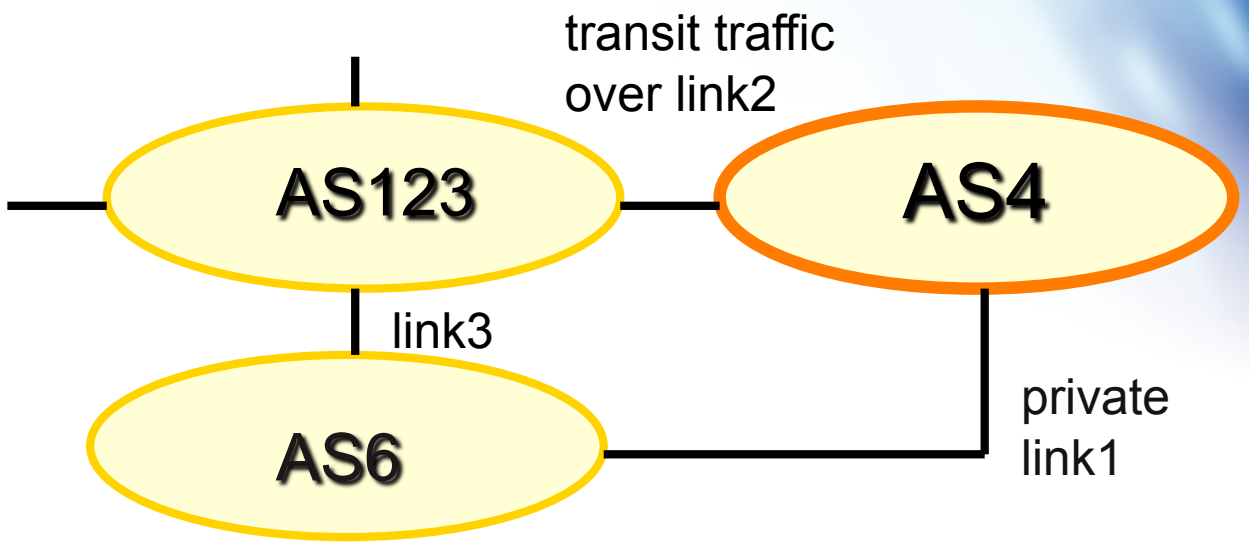

More complex example

- 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

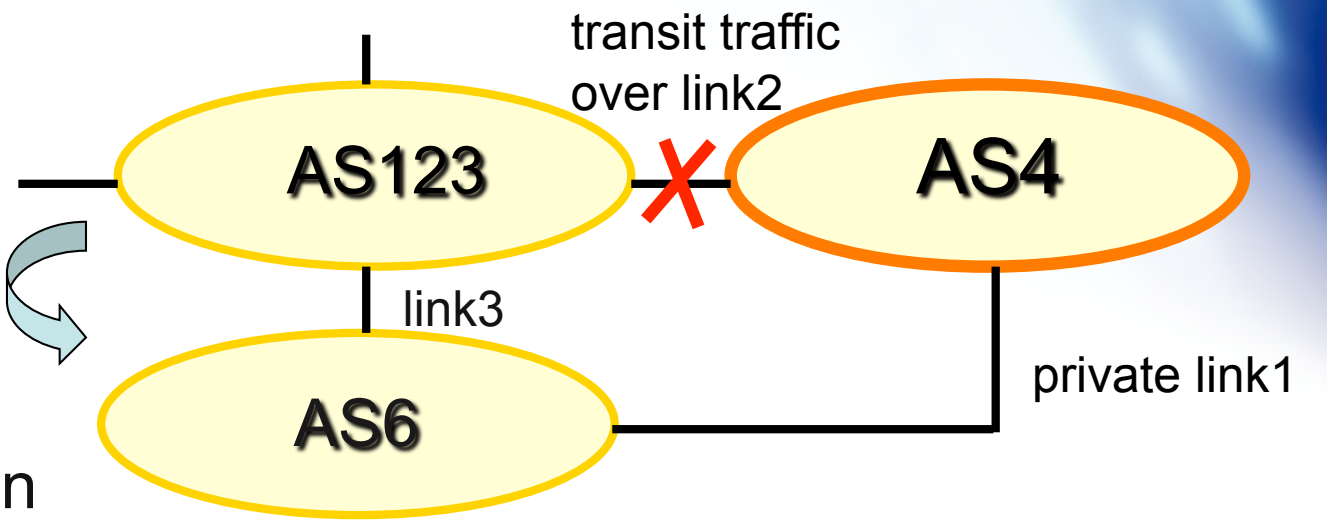

AS representation

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

Why use an IRR?

Information to share

- Routes and AS objects give an abstract specification of the policy of an AS
  - Provides device independent view of routing policy
  - Neighbouring ASes can lookup, verify and understand the other party's policy
  - Provides a clear picture where this AS fits into the Internet

Information to share (cont.)

- Information – if every AS registers its policy and routes….
  - a global view of routing policy could be mapped
    - This global picture has the ability to improve the integrity of global Internet routing
  - Provides LIR/ISP with a mechanism to find all possible paths between any two points in the Internet
- Provides a high level of abstraction

Network Planning

- Network planning
  - Simulation
    - Changes in polices can be simulated first by changing the registry but not the routers
      - To understand effects of policy changes to the existing networks
      - To make better network planning
      - To make it easier to adjust policies to maximise the performance of the network
  - Route filtering
    - Peering networks
    - A provider and its customer

Router configuration and Network troubleshooting

- Router configuration
  - By using IRRToolSet
    - Extract information from IRR to create a router readable configuration file
    - Vendor independent
    - Verification of Internet routing and Protect against inaccurate routing info distribution
- Network troubleshooting
  - Easier to locate routing problems outside your network

Asia Pacific Network Information Centre

APNIC

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

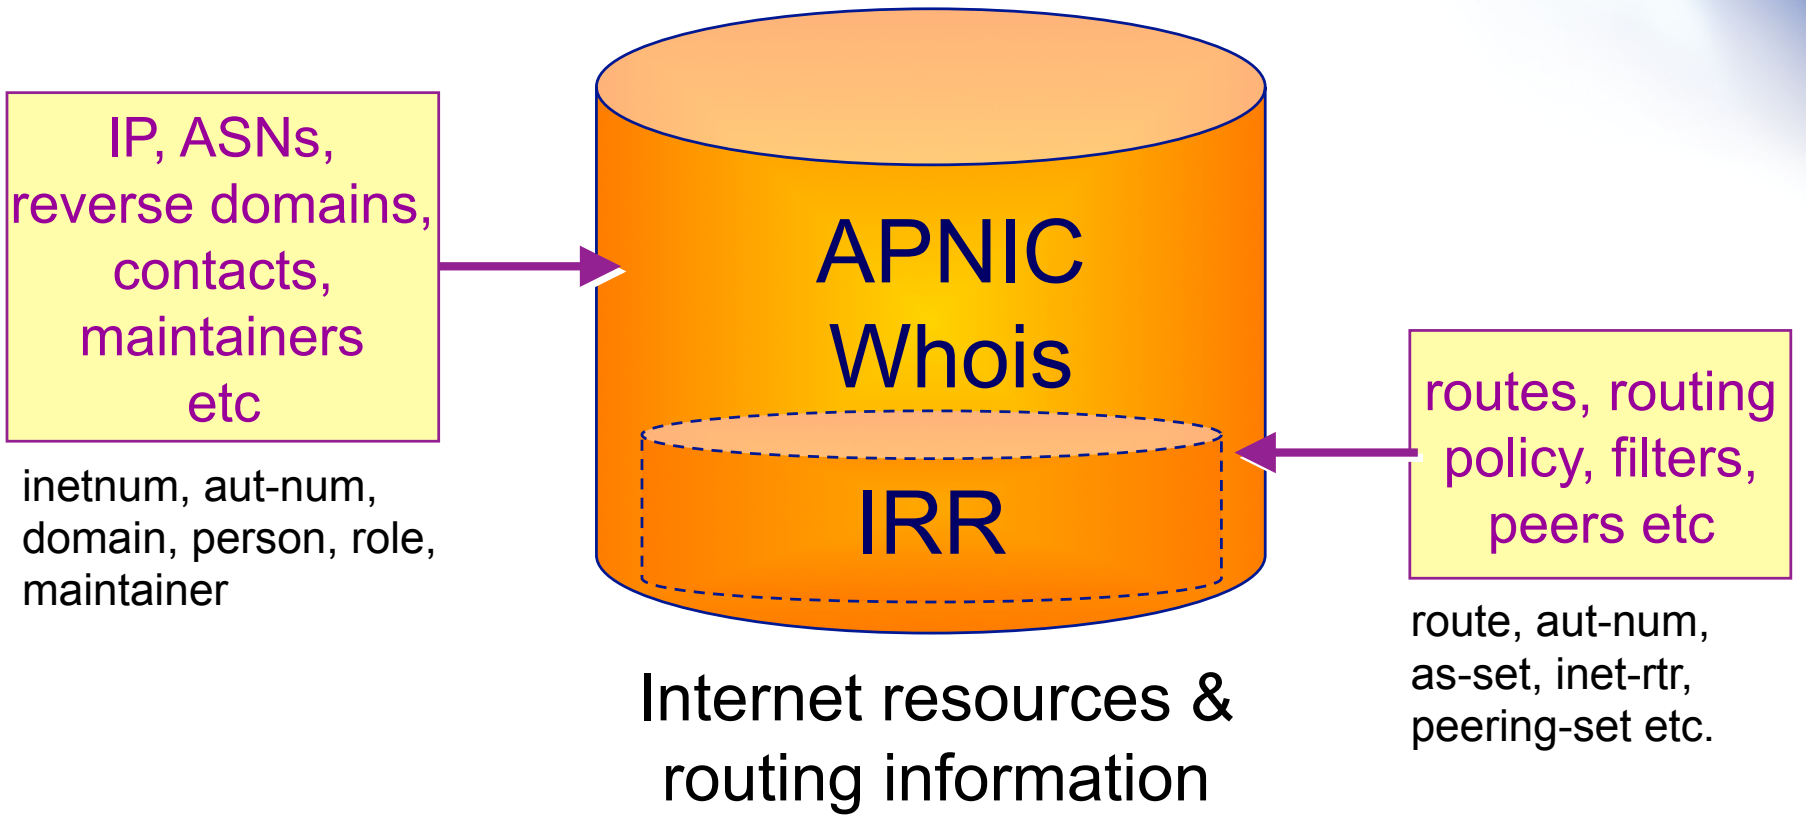

IRR Objects

- 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

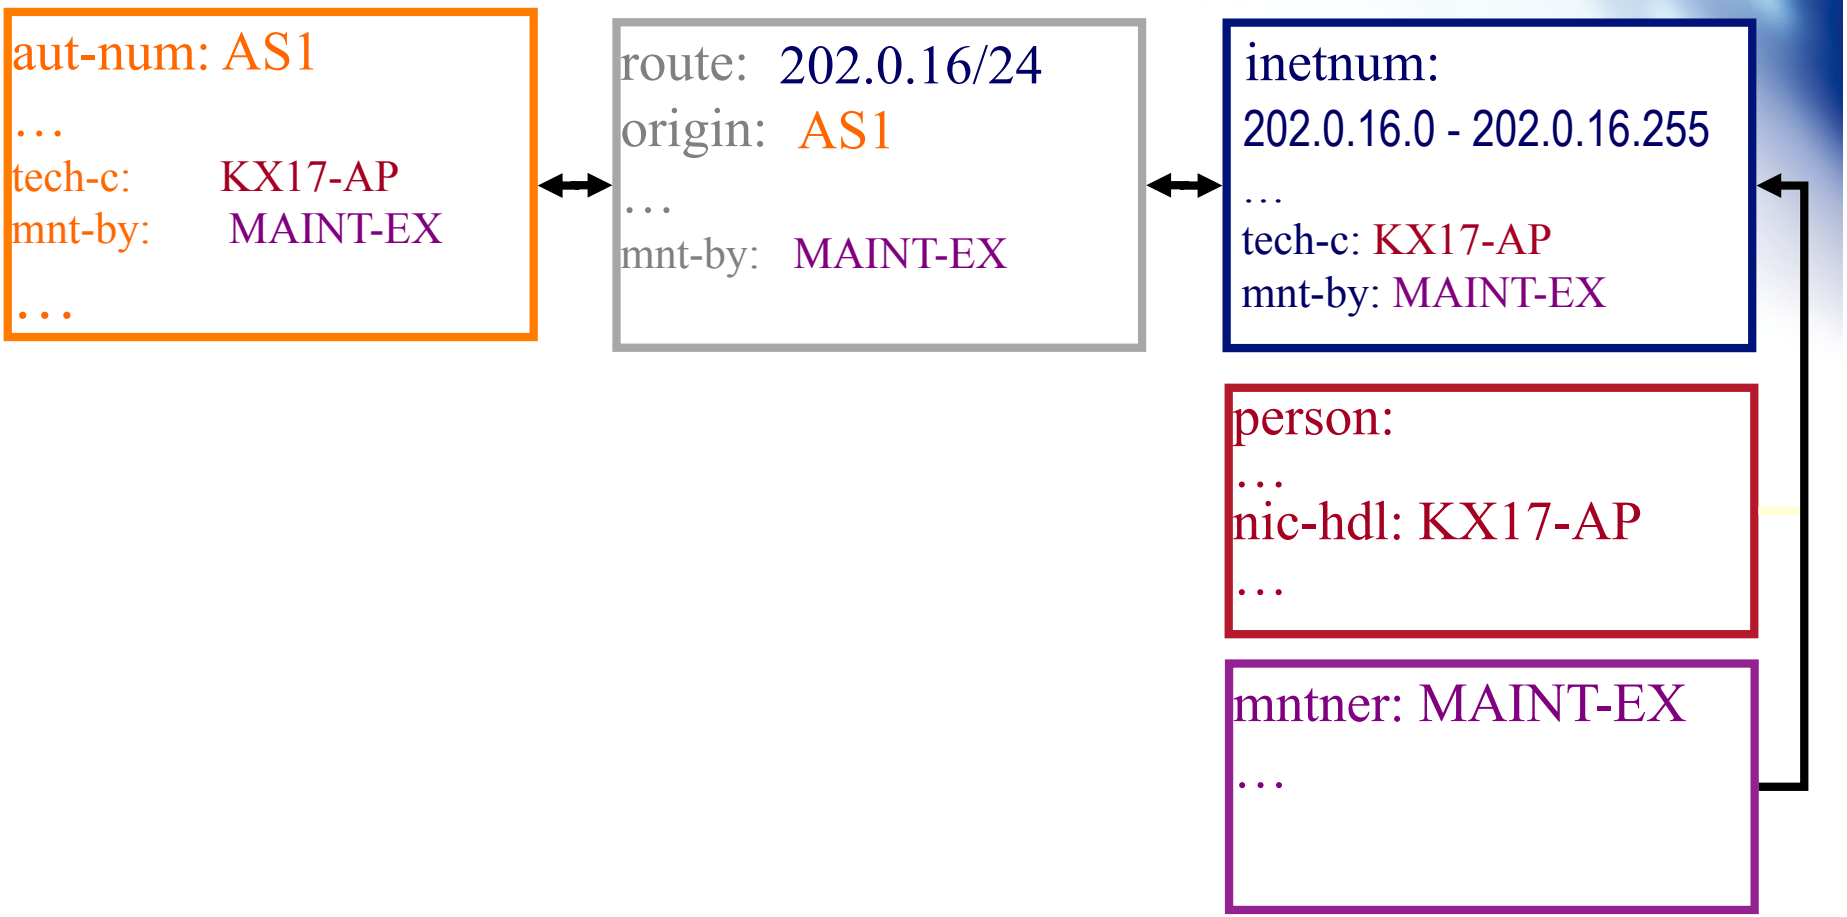

Inter-related IRR Objects

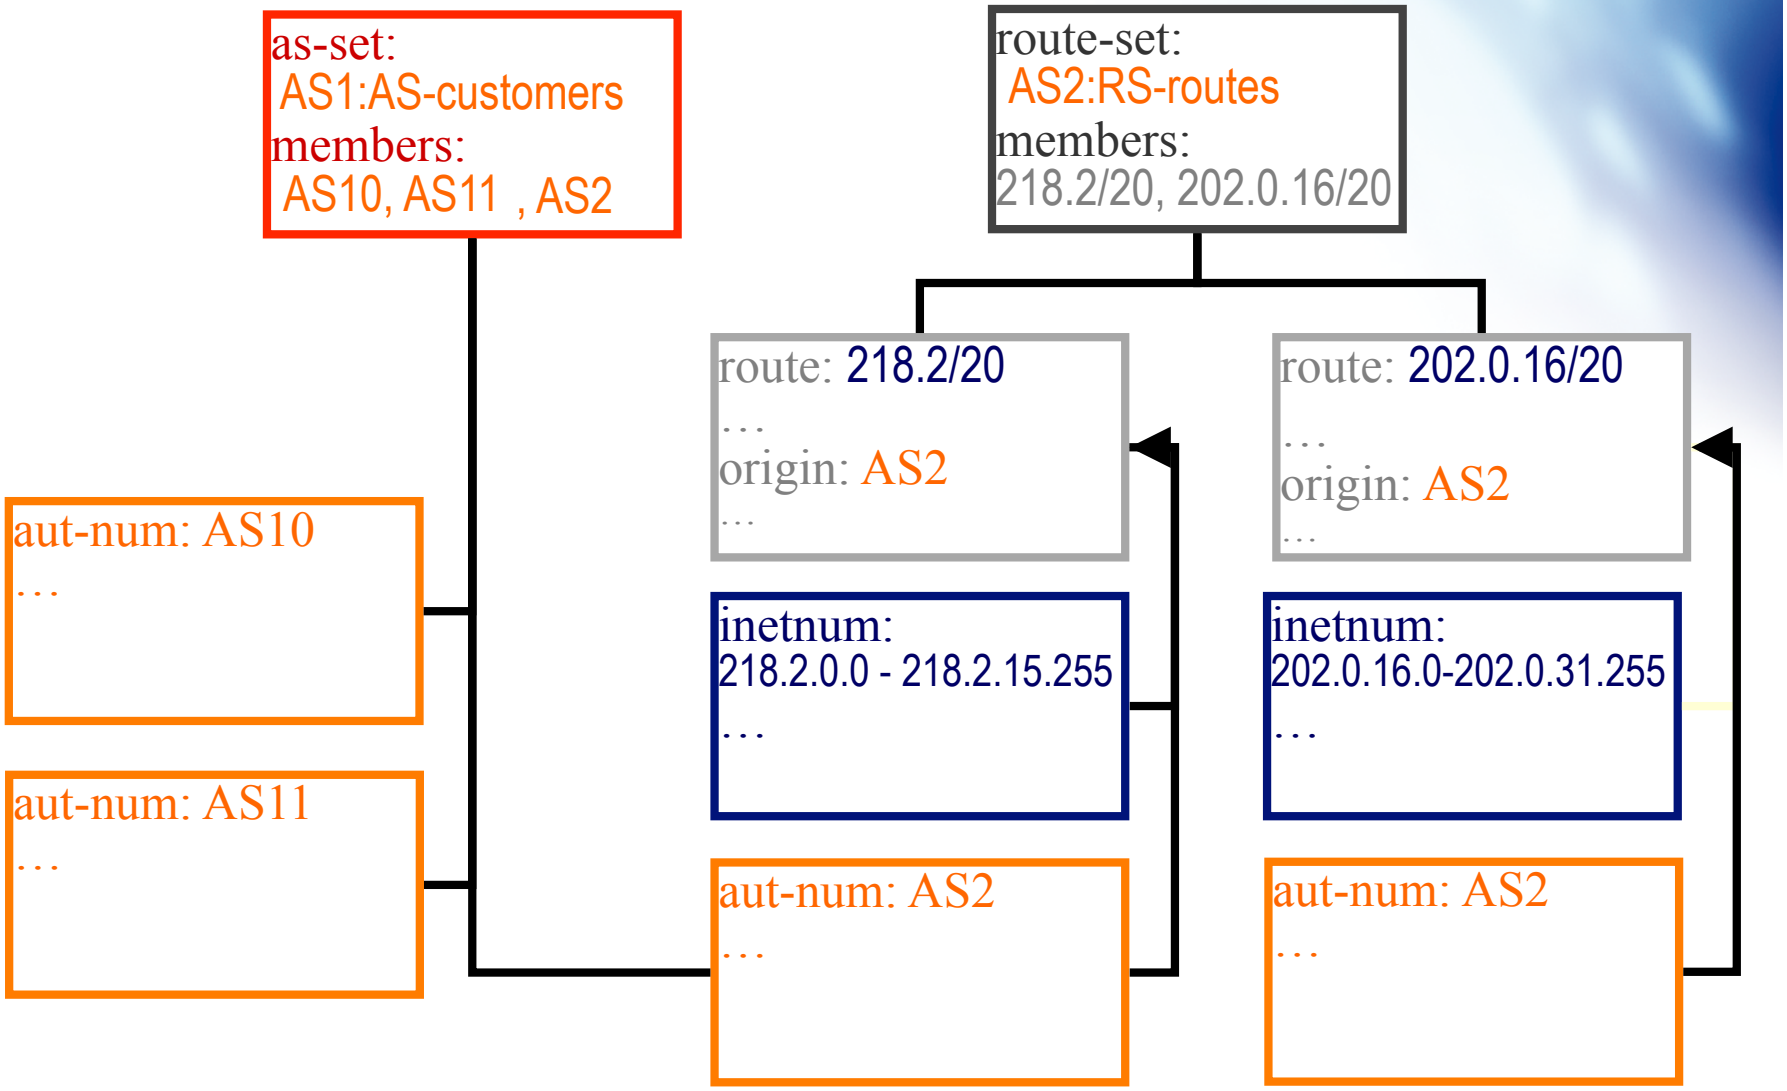

Hierarchical Authorisation

- **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:

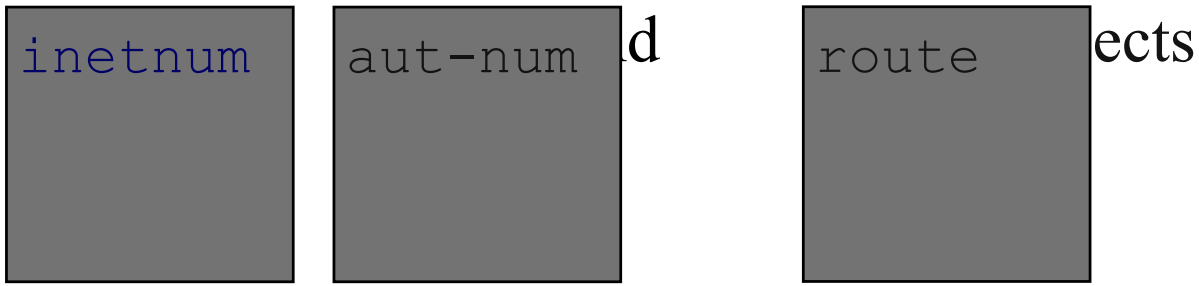

Authorisation Mechanism

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

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

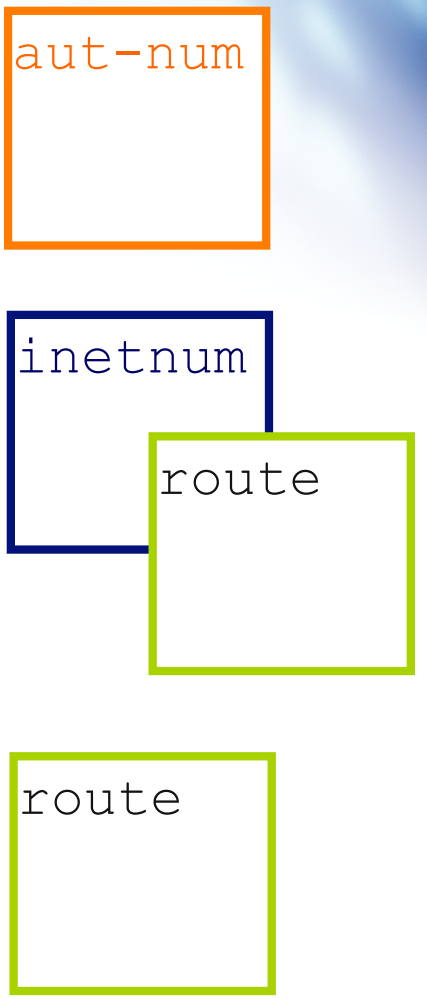

Creating Route Objects

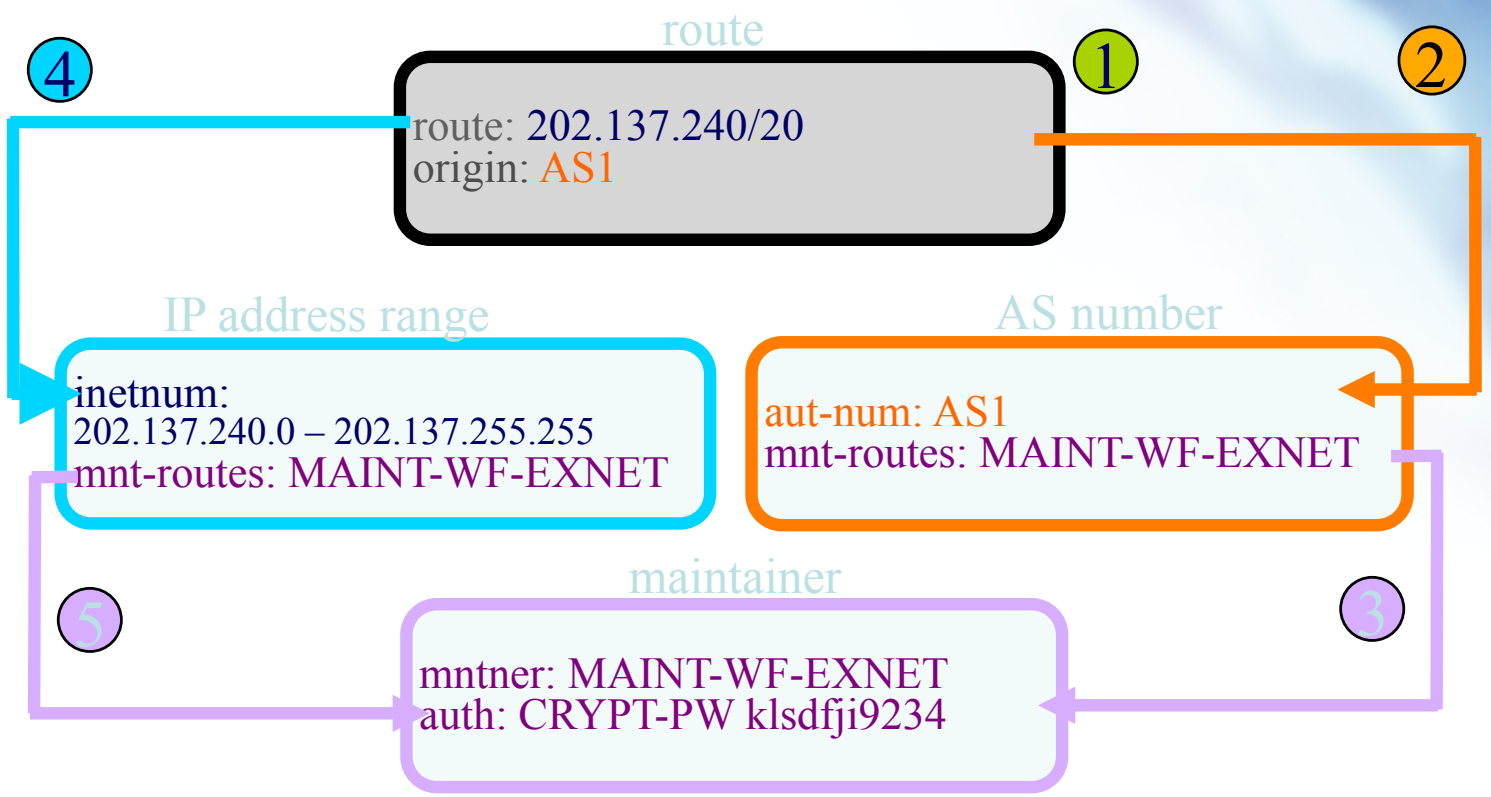

1. Create route object and submit to APNIC RR database
2. DB checks aut-num obj corresponding to the ASN in route obj
3. Route obj creation must pass auth of mntner specified in aut-num *mnt-routes* attribute.
4. DB checks inetnum obj matching/encompassing IP range in route obj
5. Route obj creation 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://irrtoolset.isc.org
  - Download: ftp://ftp.isc.org/isc/IRRToolSet/
    - Installation needs: lex, yacc and C++ compiler

Use of RPSL - RtConfig

- RtConfig v4
  - 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

Common Peering Policies

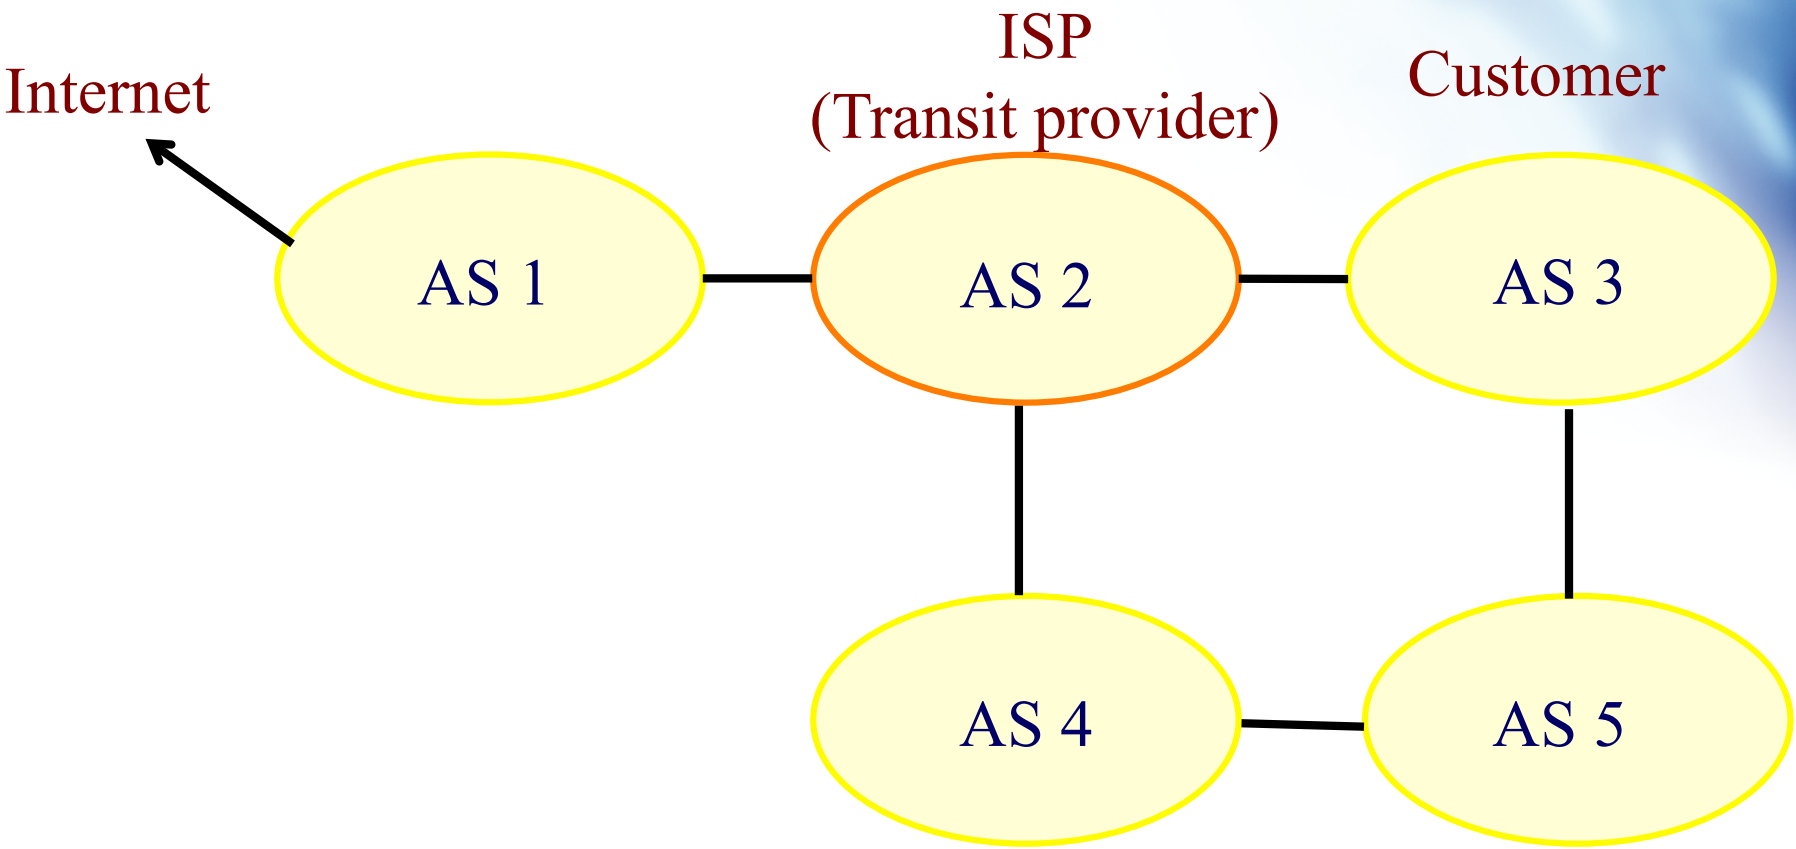

- Peering policies of an AS
  - Registered in an aut-num object

Common Peering Policies

- Policy for AS3 in the AS2 aut-num object

```
aut-num:      AS2
as-name:      SAMPLE-NET
dsescr:       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:       CW89-AP
tech-c:        CW89-AP
mtn-by:        MAINT-SAMPLE-AP
changed:      sample@sample.net
```

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

Source: www.cisco.com

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

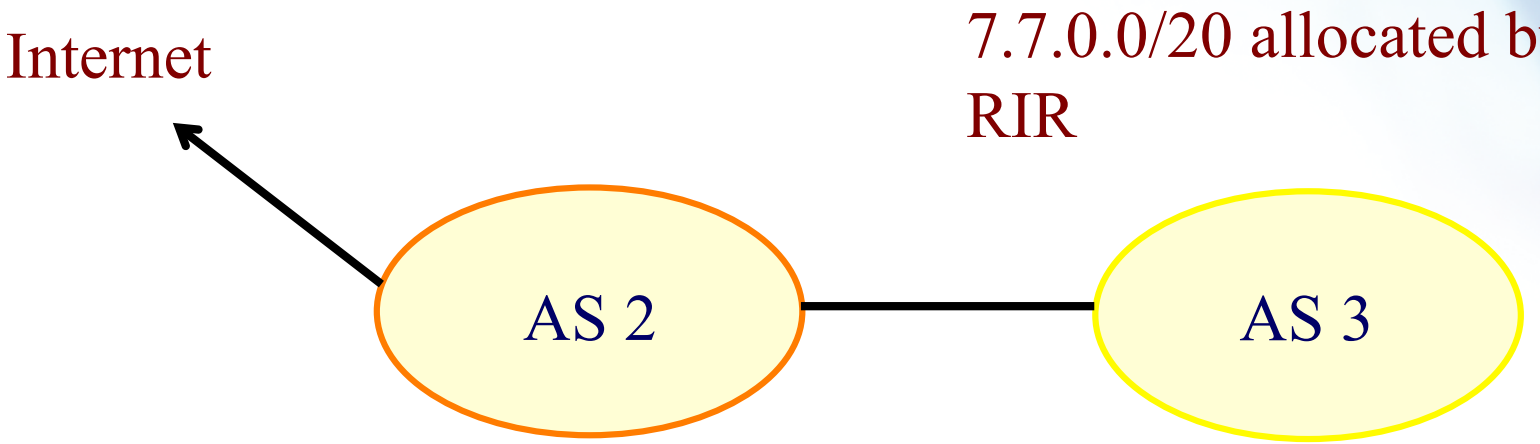

AS3 wants to announce part or all of 7.7.0.0/20 on the global Internet.

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

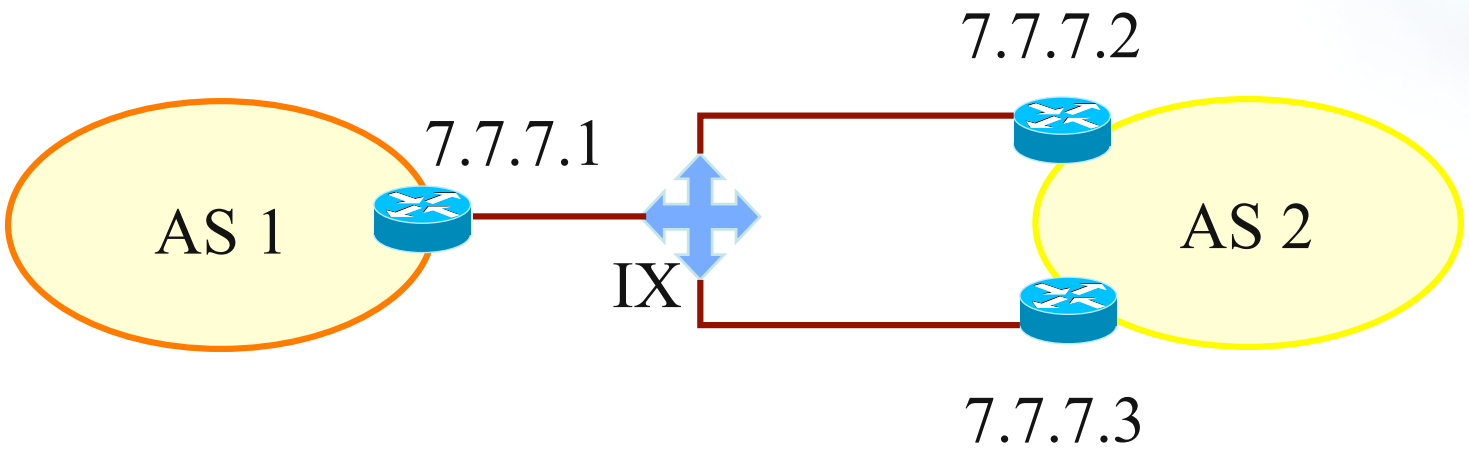

How to define AS1's routing policy by specifying its boundary router?

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+$>
```

Describing simple backup connections: AS1

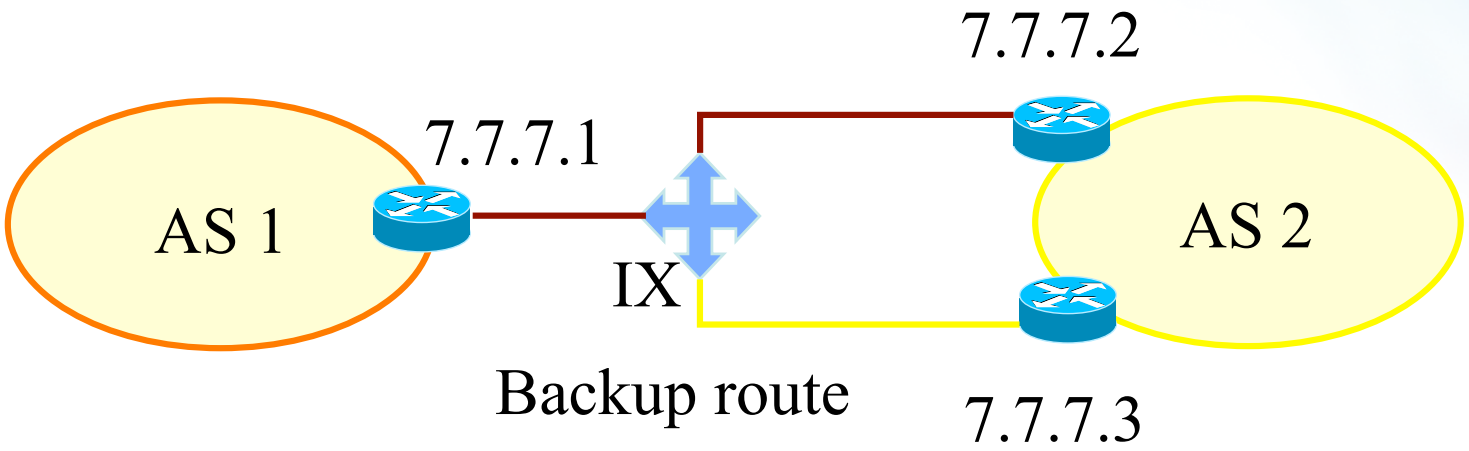

How to define AS1's routing policy of its backup route?

⇨ Use preference

Describing simple backup connections: AS1 (cont.)

```
aut-num:   AS1
import:    from AS2 7.7.7.2 at 7.7.7.1 action pref=10;
           from AS2 7.7.7.3 at 7.7.7.1 action pref=20;
           accept <^AS2+$>
```

Use of pref

- pref is opposite to local-pref
- Smaller values are preferred over larger values

Describing simple backup connections: AS2

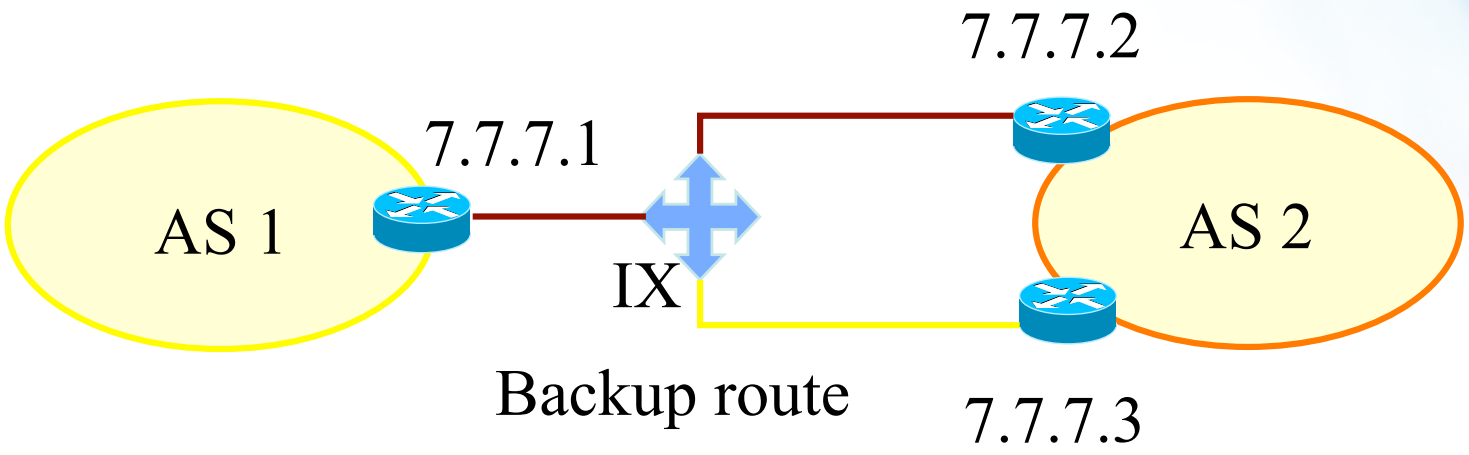

How to define AS2's routing policy of AS1's backup route?

⇨ multi exit discriminator metric (med) can be used

Describing simple backup connections: AS2 (cont.)

```
aut-num:   AS2
export:    to AS1 7.7.7.1 at 7.7.7.2 action med=10;
           to AS1 7.7.7.1 at 7.7.7.3 action med=20;
           announce <^AS2+$>
```

Use of med

- Suitable for load balancing including backups

Multihome Routing Policy

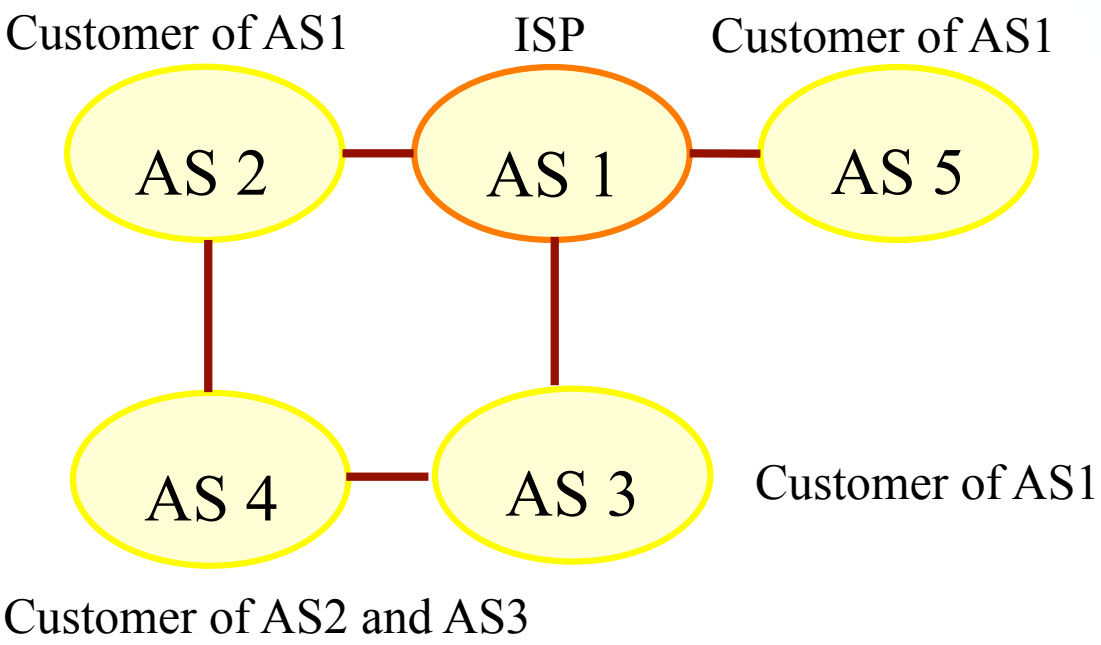

AS1's base policy

- 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

prop-079: Abuse-c field for APNIC Whois Database

- There is no consistent way to provide details of where to send abuse reports in the APNIC Whois Database
- Abuse report usually sent to admin-c or tech-c (even though sometime they are not responsible handle this report)
- If whois contacts are not valid APNIC forward invalid contact report to private contact database to update invalid contacts

prop-079: Abuse-c field for APNIC Whois Database

- Make it mandatory to include a reference to an IRT (Incidence Response Team) object in inetnum, inet6num and aut-num objects
- Existing allocation/assignment record need to add it if they would like to update the record
- New allocation/assignment need to add it at the time to allocation/assignment (HM will do that)

prop-079: Abuse-c field for APNIC Whois Database

- All spam/abuse report need to send to IRT object listed contact
- Another policy will ensure that APNIC whois DB object will be updated regularly
- APNIC will focus more training on IRT object in fururе

APNIC RR service scope

- Support
  - APNIC Helpdesk support

<helpdesk@apnic.net>

- Training
  - IRR Training

- Mirroring
  - APNIC mirrors IRRs within Asia Pacific and major IRRs outside of the region.

Summary

- APNIC RR integrated in APNIC Whois DB
  - whois.apnic.net
  - <auto-dbm@apnic.net>
- IRR benefits
  - 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?

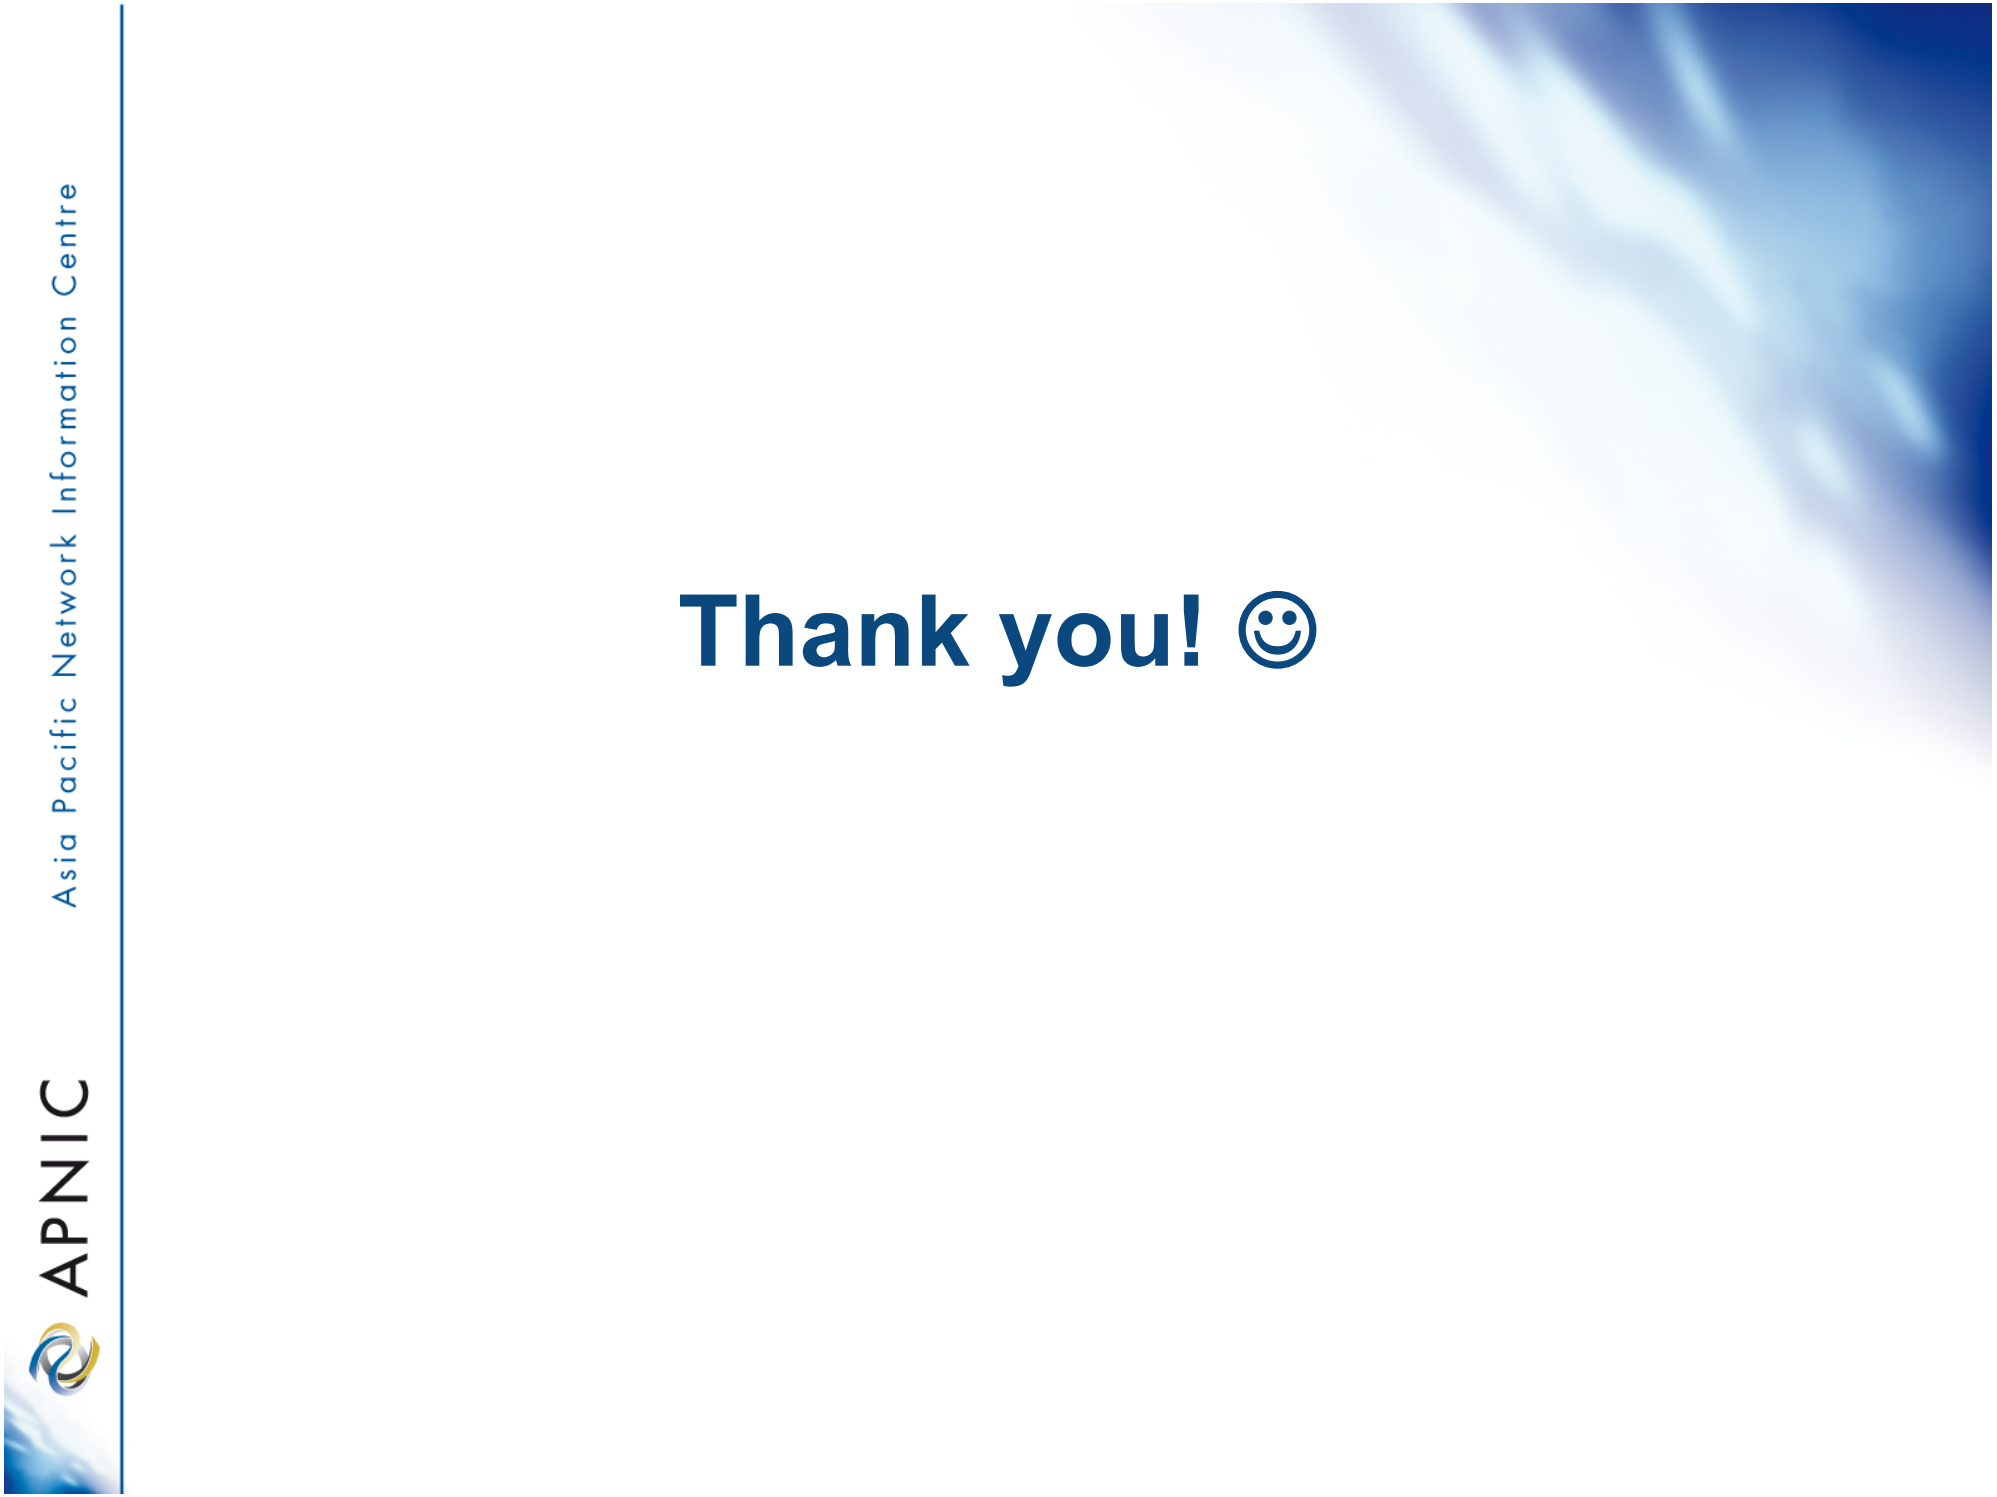

Thank you! ☺