IRR Tutorial

Taipei, Taiwan, 25 February 2003
APNIC 15 Open Policy Meeting

Objectives

- To provide an introduction to the APNIC Routing Registry
  - Explain basic concepts of the global RR
  - Outline the benefits of the APNIC Routing Registry
- NOT to:
  - Teach basic routing
  - Explain IP address policy and procedures
  - Provide advise on network configuration

Assumptions

- The audience
  - Knowledgeable about BGP routing
  - Familiar with basic APNIC database operations
  - Curious about Routing Registry usage

Overview of today

Two parts of today's tutorial:

- Morning: Background & theory
  - What is the IRR?
  - What does it look like?
  - How do I define a routing policy?
    - APNIC Trainers: Champika and Nurani
- Afternoon: Practical IRR usage
  - Complex routing policy
  - Generating BGP configuration using RtConfig
    - Andy Linton

Overview (Morning session)

- Introduction to the IRR
- APNIC Database
  - RPSL
  - updates, queries
  - Authentication
- Routing Policy
  - Why define routing policy?
  - Syntax, examples
  - simple "case studies"
- APNIC IRR

The Internet Routing Registry

What is an IRR?

- Global Routing Registry database
  - http://www.irr.net/
    - Uses RPSL
  - Established in 1995
- 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

Internet Routing Registry

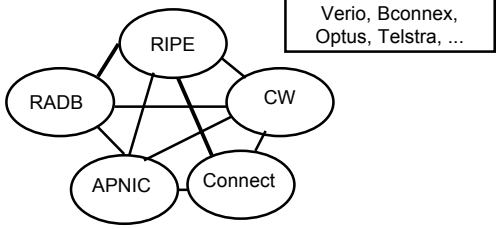

IRR = APNIC RR + RIPE DB + RADB + C&W + ARIN + …

List of IRRs

- ALLTEL
- ALTDB
- ANS
- AOLTW
- **APNIC**
- ARBOR
- ARCSTAR
- AREA151
- ARIN
- BCONNEX
- BELL
- CARYNET
- CCAIR
- CHTR
- CSAS
- CW
- DAKNET
- DERU
- DoDNIC
- EASYNET
- ENTERZONE
- EPOCH
- FASTVIBE
- FGC
- GT
- GTS
- GW
- HS
- I2
- KOREN
- KT
- LEVEL3
- LOOK
- NESTEGG
- NETRAIL
- Nyi.net
- OPENFACE
- OTTIX
- PANIX
- RADB
- REACH
- RGNET
- RIPE
- RISQ
- SAKURA
- SEMAPHORE
- SINET
- SOUNDINTERNET
- SPACELINK
- SPRINT
- TELSTRA
- US Data Authority
- VERIO
- WL2K
- WWNET

Why should I use the Internet Routing Registry?

- When peering
  - register your routes and filter your peers
- Some transit providers and big ISP's ask for this
- Useful for fixing problems
  - contact information
  - debug, configure and engineer Internet routing
    - RAToolSet developed by RIPE

Why should I use the Internet Routing Registry?

- Policy based routing
  - Allows different criteria as basis for routing decisions
  - Routing policy - description of the relationship between external BGP peers
    - Next level of abstraction: RPSL
- Ultimately: easier maintenance of routing configuration in big & complex networks

BGP Configuration from IRR

- Routing Policy specification Language (RPSL)
  - abstract, high-level policies
  - policies for each Autonomous System (AS)
- Internet Routing Registry
  - policies, routes and contact informatiom
  - benefit from the data and delegation of others
- RtConfig
  - RAToolSet
  - generate router configuration files
  - automates details and tedious aspects

Real-life Routing Registry examples

- Connect.au
  - `whois -h whois.ripe.net -s RADB -r -T aut-num AS2764`
- WIX (NZ), maintaining their own RR
- C&W, running private RR for their customers
- Some AS numbers with detailed policy:
  - `whois -h whois.apnic.net -r -T aut-num AS7474 (OPTUS)`
  - `whois -h whois.ripe.net -r -T aut-num AS4777 (APNIC)`
- Forums:
  - Routing SIG
    - http://www.apnic.net/meetings/archive/sigs/routing.html

APNIC Database

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 IRR

Integration of whois and RR

- Integrated APNIC whois database & Routing Registry
  - v3 & RPSL

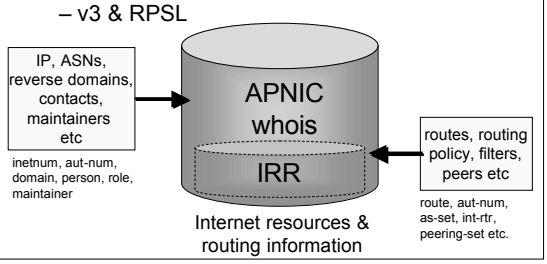

Object types

DB Intro

OBJECT	PURPOSE
person	contact persons
role	contact groups/roles
inetnum	IPv4 addresses
inet6num	IPv6 addresses
aut-num	Autonomous System number
domain	reverse domains
route	prefixes being announced
mntner	(maintainer) data protection

http://www.apnic.net/db/

Attributes & values

DB Intro

- 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

Object templates

DB Intro

To obtain template structure*, use :
whois -t <object type>

```
% whois -h whois.apnic.net -t person

person:   [mandatory]  [single]     [primary/look-up key]
address:  [mandatory]  [multiple]   [ ]
country:  [optional]   [single]     [ ]
phone:    [mandatory]  [multiple]   [ ]
fax-no:   [optional]   [multiple]   [ ]
e-mail:   [mandatory]  [multiple]   [look-up key]
nic-hdl:  [mandatory]  [single]     [primary/look-up key]
remarks:  [optional]   [multiple]   [ ]
notify:   [optional]   [multiple]   [inverse key]
mnt-by:   [mandatory]  [multiple]   [inverse key]
changed:  [mandatory]  [multiple]   [ ]
source:   [mandatory]  [single]     [ ]
```

*Recognised by the RIPE whois client/server

Inter-related objects
- Internet Resources

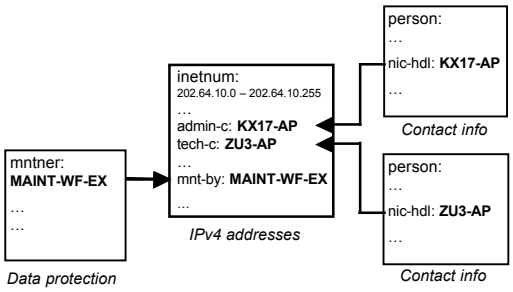

Inter-related objects
- Routing information

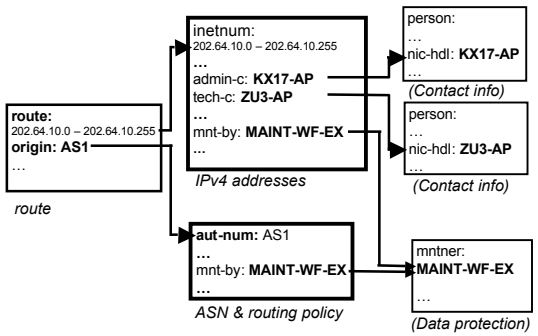

RPSL

Routing Policy Specification Language

RPSL

- Powerful language
  - RPSL is more expressive
    - Policies can be expressed at the AS level
- Policies can be detailed
  - router configurations

- Use secure, encrypted submission paths

Features of RPSL

- Support the exchange of complex routing policy information between ISPs in a secure and openly agreed manner

  - ISPs can configure filters for their boarder routers, or check router configurations against routing policies

Why RPSL ?

- RPSL also includes many non-routing related concepts and data structures
  - general address management services such as whois db
- RPSL is capable of more functionality
  - (than previous ripe-181 code)
- RPSL is of most value in attributes that has more complex structures
  - *AS-Set, Route* and related objects

Objects in RPSL

- Object Names
  - Can have '-' or '_' inside
    - APNIC-HM_Maintainer
  - Can have digits, first character alphabetic
  - Last character must be a letter or a digit
  - Reserved names, Reserved prefixes
- Attributes and values
  - Value of an attribute has a type
    - <object-name>, <as-number>, <ipv4-address>, <address-prefix> etc.

Objects in RPSL

- Line continuation possible
  - Space, tab, +
- Comments
  - Begin with #, can be anywhere inside an object, cannot start at the beginning of a line (column 0)
- Some attributes are mandatory
  - *Mnt-by* is *mandatory in all objects*
- The order of attributes is flexible
- Empty attributes not allowed, not removed
- Object ends at blank line (\n\n)

Objects in RPSL - Example

```
person:     Ky Xander
address:    ExampleNet Service Provider
address:    2 Pandora St Boxville
address:    Wallis and Futuna Islands
country:    WF
phone:      +680-368-0844 # day time
fax-no:     +680-367-1797
e-mail:     kxander@example.com
remarks:    Example object
            for training
nic-hdl:    KX17-AP
mnt-by:     MAINT-WF-EX
changed:    kxander@example.com 20020731
source:     APNIC
```

Comment

Line continuation

Mntner object example

– Specifies authentication information

```
mntner:       MAINT-2DAY-NZ
descr:        2day.com
descr:        Auckland
admin-c:      PM5-NZ
tech-c:       JA39
upd-to:       peter@2day.com
upd-to:       jabley@automagic.org
auth:         CRYPT-PW X/8b6sCRno6LQ
mnt-by:       MAINT-2DAY-NZ
referral-by:  APNIC-HM
changed:      hm-changed@apnic.net 20021105
source:       APNIC
```

Inetnum object example

– Specifies IP allocations & assignments

```
inetnum:      202.36.0.0 - 202.37.255.255
netname:      NZGATE-NZ
descr:        NZ Gate National Service Provider
descr:        Administered by Telecom New Zealand Ltd
descr:        New Zeland
country:      NZ
admin-c:      DBK1-AP
tech-c:       KS61-AP
tech-c:       KS61-AP
remarks:      service provider
notify:       dbmon@apnic.net
mnt-by:       APNIC-HM
changed:      ARRON@WAIKATO.AC.NZ 19950612
changed:      hostmaster@apnic.net 20011004
changed:      hm-change@apnic.net 20020722
status:       ALLOCATED PORTABLE
source:       APNIC
```

Route Object

- Each interAS route originated by an AS is specified using a route object
  - "route:" attribute is the address prefix of the route
  - "origin:" attribute is the AS number of the AS that originates the route into the interAS routing system
  - "route:" and "origin:" attribute pair is the primary key

Route object example

- Each interAS route originated by an autonomous system

```
route:          202.37.240.0/23
descr:          route originating from 2day.com
origin:         AS17914
mnt-by:         MAINT-2DAY-NZ
changed:        jabley@automagic.org 20021220
source:         APNIC
notify:         noc@2day.com
```

Aut-num object

- Autonomous System Number (ASN) object
  - Value of the "aut-num:" attribute is the AS number of the AS described by this object
  - routing policies of the AS are specified using the "import:", "export:" and "default:" attributes

Aut-num object example

– Describes an Autonomous System

```
aut-num:   AS17914
as-name:   ASN-2DAY-NZ-AP
descr:     2Day Internet Limited
country:   NZ
import:    from AS17914:AS-TRANSIT action pref=100; accept ANY
import:    from AS17914:AS-PEERS action pref=120; accept PeerAS
export:    to AS17914:AS-TRANSIT announce AS17914:AS-CUSTOMERS
export:    to AS17914:AS-PEERS announce AS17914:AS-CUSTOMERS
admin-c:   PM5-NZ
tech-c:    JA39
remarks:   2day.com peers at the Auckland Peering Exchange
mnt-by:    MAINT-2DAY-NZ
changed:   jabley@automagic.org 20021104
source:    APNIC
```

routing policy

As-set object

- Defines a set of **aut-num** objects
  - Previously as- macro
- The "as-set:" attribute defines the name of the set
  - The "members:" attribute lists the members of the set
    - list of AS numbers or other as-set names

AS-set & aut-num objects

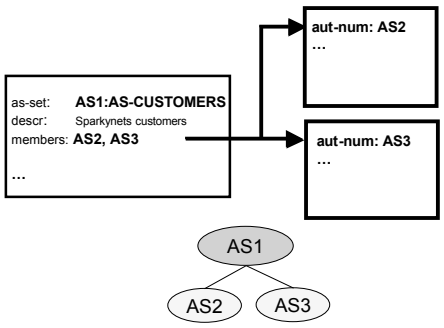

As-set object example

– defines a set of **aut-num** objects

```
as-set:       AS17914:AS-CUSTOMERS
descr:        ASes which obtain transit from 2day.com
members:      AS17914, AS9327
admin-c:      PM5-NZ
tech-c:       JA39
notify:       peter@2day.com
notify:       jabley@automagic.org
mnt-by:       MAINT-2DAY-NZ
changed:      jabley@automagic.org 20021104
source:       APNIC
```

Related IRR objects

- Route-set object
  - Defines a set of routes that can be represented by **route** objects or by address prefixes
    - Previously community
- Inet-rtr object
  - Routers are specified using the *inet-rtr* object
- rtr-set
  - Defines a set of routers specified by inet-rtr names, ipv4_addresses or other rtr-set names
- Peering-set
  - Defines a set of peerings that are listed in its "peering:" attributes
- filter-set object
  - Defines a set of routes that are matched by its filter

Queries and Updates

Basic Database Queries

- Unix
  - whois –h whois.apnic.net <lookup key>
- Web interface
  - http://www.apnic.net/apnic-bin/whois2.pl
- Look-up keys
  - usually the object name
  - Check template for look-up keys

Advanced db 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
- d enables use of flags for reverse domains
- r turn off recursive lookups

Database Query - inetnum

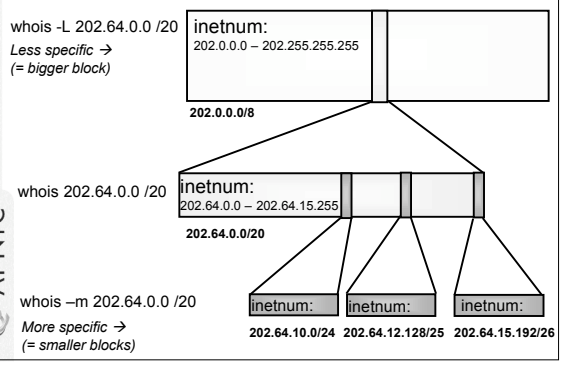

Database Query - Inetnum

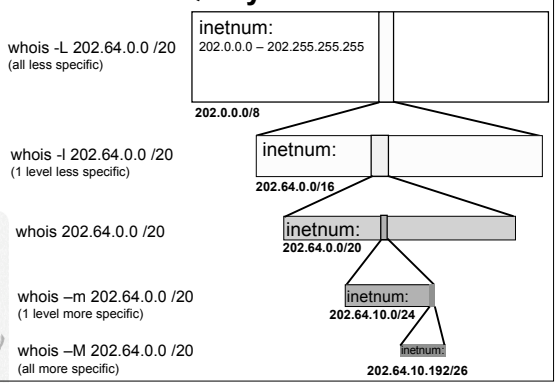

Database Query - Inetnum

'-M' will find all assignments in a range in the database

```
% whois -M 202.64.0.0/20

inetnum:    202.64.10.0 - 202.64.10.255
netname:    SILNET-AP
descr:      Satyam Infoway Pvt.Ltd.,
.....
inetnum:    202.64.12.128 - 202.64.12.255
netname:    SOFTCOMNET
descr:      SOFTCOM LAN (Internet)IP.
.....
inetnum:    202.64.15.192 - 202.64.15.255
descr:      SILNET
descr:      Satyam Infoway's Chennai LAN
.....
```

Database Query - Recursion

Recursion is enabled by default

```
% whois  203.113.0.0/19

inetnum:     203.113.0.0 - 203.113.31.255
netname:     TOTNET-AP
descr:       Telephone Organization of THAILAND(TOT)
descr:       Telephone and IP Network Service Provider
descr:       State Enterprise Thailand Government
country:     TH
admin-c:     NM18-AP
tech-c:      RC80-AP
   ..... .
person:      Nopparat Maythaveekulchai
address:     YTEL-1234 Office
address:     Telephone Organization of THAILAND(TOT)
   ..... .
person:      Rungsun Channarukul
address:     YTEL-1234 OfficeP
address:     Telephone Organization of THAILAND(TOT)
```

Database Query – Turn off recursion

- No nic-handle lookup

```
% whois -r  203.113.0.0/19

inetnum:      203.113.0.0 - 203.113.31.255
netname:      TOTNET-AP
descr:        Telephone Organization of THAILAND(TOT)
descr:        Telephone and IP Network Service Provider
descr:        State Enterprise Thailand Government
country:      TH
admin-c:      NM18-AP
tech-c:       RC80-AP
mnt-by:       APNIC-HM
mnt-lower:    MAINT-TH-SS163-AP
changed:      hostmaster@apnic.net 19990922
source:       APNIC
```

Database Query - Inverse

Inverse lookup with '-i '

```
% whois -i person DK26-AP

inetnum:      202.101.128.0 - 202.101.159.255
netname:      CHINANET-FJ
descr:        chinanet fujian province network
country:      CN
admin-c:      DK26-AP
……

domain:       128.103.202.in-addr.arpa
descr:        in-addr.arpa zone for 128.103.202.in-addr.arpa
admin-c:      DK26-AP
……
aut-num:      AS4811
as-name:      CHINANET-CORE-WAN-EAST
descr:        CHINANET core WAN EAST
descr:        connect to AT&T,OPTUS
country:      CN
admin-c:      DK26-AP
……
person:       Dongmei Kou
address:      A12,Xin-Jie-Kou-Wai Street,
address:      Beijing,100088
country:      CN
phone:        +86-10-62370437
nic-hdl:      DK26-AP
```

Database Update Process

– Email requests to <auto-dbm@apnic.net>
– Each request contains an object template

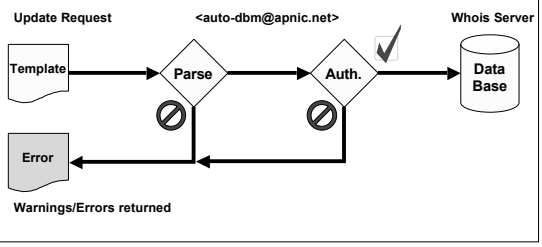

Database Update Process

- Update transactions
  - Creating, Modifying, Deleting objects
- Updates are submitted by email
  - E-mail to: <auto-dbm@apnic.net>
- Email message contains template representing new or updated object

Database Update - Web Interface

- Creates a template through the web form
  - Template will be sent to you by email
  - This should be forwarded to:

    <auto-dbm@apnic.net>
  - Common error
    - to reply to the email
      - (Adds extra character in front of each line)

http://www.apnic.net/services/whois_guide.html

Authentication

Data Protection

- mnt-by attribute refers to mntner object
  - Checked at every update
  - Include: password: <clear_text_pwd>
  - Or: sign the message with private key
- auth:
  - CRYPT-PW, PGPKEY-<id>
  - Key-cert object
    - Possible to use PGP & GPG
  - http://www.gnupg.org

Authorisation Mechanism

inetnum:	202.137.181.0 – 202.137.185.255
netname:	EXAMPLENET-WF
descr:	ExampleNet Service Provider
…….	
mnt-by:	**MAINT-WF-EX**

mntner:	**MAINT-WF-EX**
descr:	Maintainer for ExampleNet Service Provider
country:	WF
admin-c:	ZU3-AP
tech-c:	KX17-AP
upd-to:	kxander@example.com
mnt-nfy:	kxander@example.com
auth:	CRYPT-PW apHJ9zF3o
mnt-by:	MAINT-WF-EX
changed:	kxander@example.com 20020731
source:	APNIC

Mnt-by & Mnt-lower

- ‘mnt-by’ attribute
  - Can be used to protect any object
  - Changes to protected object must satisfy authentication rules of ‘mntner’ object.
- ‘mnt-lower’ attribute
  - Also references mntner object
  - Hierarchical authorisation for inetnum & domain objects
  - The creation of child objects must satisfy this mntner
  - Protects against unauthorised updates to an allocated range - highly recommended!

Authentication/Authorisation

– APNIC allocation to member
  • Created and maintained by APNIC

```
Inetnum:     203.146.96.0 - 203.146.127.255
netname:     LOXINFO-TH
descr:       Loxley Information Company Ltd.
Descr:       304 Suapah Rd, Promprab,Bangkok
country:     TH
admin-c:     KS32-AP
tech-c:      CT2-AP
① mnt-by:    APNIC-HM
② mnt-lower: LOXINFO-IS
changed:     hostmaster@apnic.net 19990714
status:      ALLOCATED PORTABLE
source:      APNIC
```

1. Only APNIC can change this object
2. Only Loxinfo can create assignments within this allocation

Authentication/Authorisation

– Member assignment to customer
  • Created and maintained by APNIC member

```
Inetnum:     203.146.113.64 - 203.146.113.127
netname:     SCC-TH
descr:       Sukhothai Commercial College
Country:     TH
admin-c:     SI10-AP
tech-c:      VP5-AP
mnt-by:      LOXINFO-IS
changed:     voraluck@loxinfo.co.th 19990930
status:      ASSIGNED NON-PORTABLE
source:      APNIC
```

Only LOXINFO-IS can change this object

Hierarchical Authorisation

- **mnt-routes: <your_mntner>**
  – authenticates ***creation*** of more specific route objects
  – Only one level below
  – In: aut-num, inetnum, route objects,

Creating route objects

- Must pass multiple authentication (logical AND):
  - Originating ASN
    - The mntner in the mnt-routes is checked
    - Then in the mnt-lower
    - Then in the mnt-by
  - 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

Creation of route objects

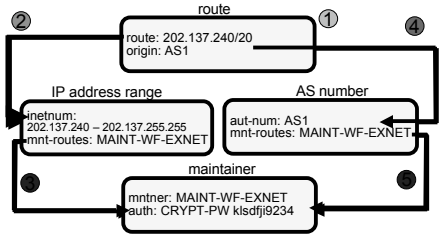

1. Create route object and submit to APNIC RR
2. The db checks inetnum obj matching/encompassing address range in route obj
3. Route obj creation must pass auth of mntner specified in inetnum *mnt-routes* attribute.
4. The db checks aut-num obj corresponding to the ASN in route obj
5. Route obj creation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

Questions?

Routing Policy

What is a Routing Policy?

- Exchange of routing information between Autonomous Systems

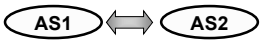

- Usually policies are not configured for each network separately
  - Configured for groups of networks

Why define a Routing Policy?

- Documentation
- Consistency across your AS
  - routers / implementations
- Scalability
- Provides routing security
  - Can peer originate the route?
  - Can peer act as transit for the route?

How define a Routing Policy?

- Who are my BGP neighbours?
  - (customers/ peers/ upstreams)
- What routes are:
  - Originated by each neighbour?
  - Imported from each neighbour?
  - Exported to each neighbour?
  - Preferred when multiple routes exist?
  - How are they treated (modified routing parameters?)

Defining the Routing Policy

- Routing and packet flows

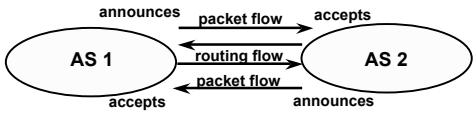

For AS1 and AS2 networks to communicate
- AS1 must announce to AS2
- AS2 must accept from AS1
- AS2 must announce to AS1
- AS1 must accept from AS2

Defining the Routing Policy

Basic concept

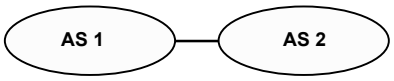

"action pref" - the lower the value, the preferred the route

```
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
```

Defining the Routing Policy

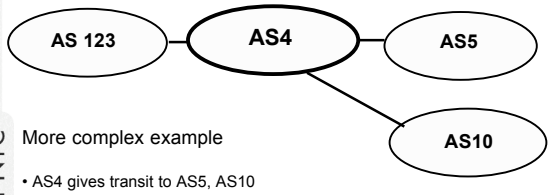

More complex example

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

Defining the Routing Policy

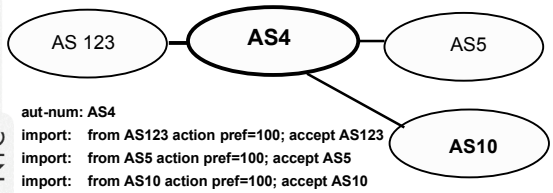


```
aut-num: AS4
import:  from AS123 action pref=100; accept AS123
import:  from AS5 action pref=100; accept AS5
import:  from AS10 action pref=100; accept AS10
export:  to AS123  announce  AS4
export:  to AS5 announce AS4 AS10
export:  to AS10 announce AS4 AS5
```

Not a path

Defining the Routing Policy

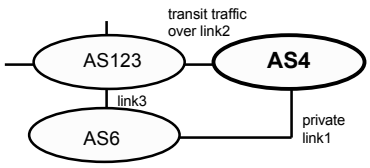

More complex example

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

Defining the Routing Policy

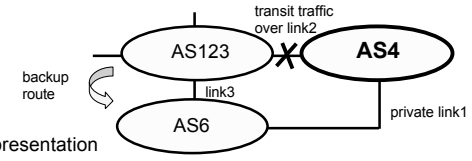

AS representation

```
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

Experimental setup: AS relations

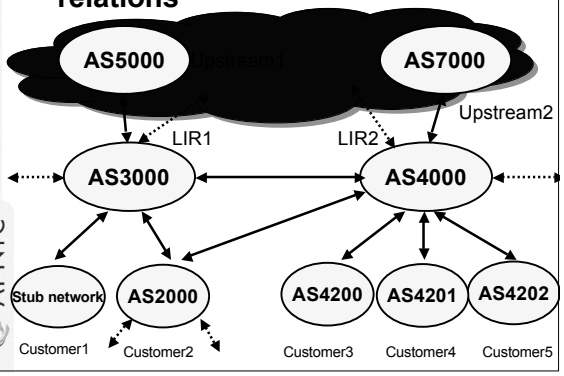

AS relations, including allocations & assignments

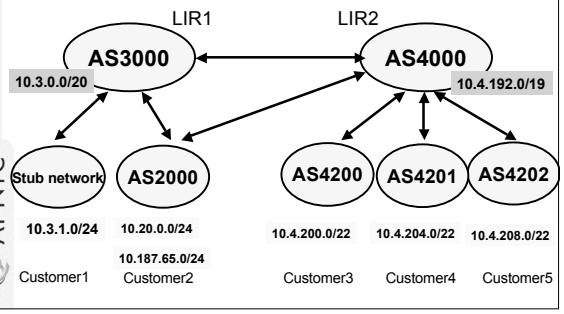

Case studies, overview

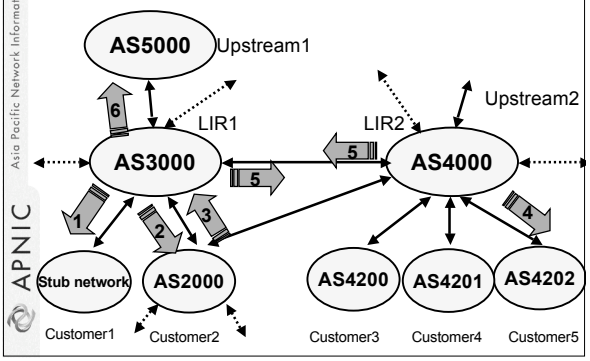

Case1: Static end-user set-up

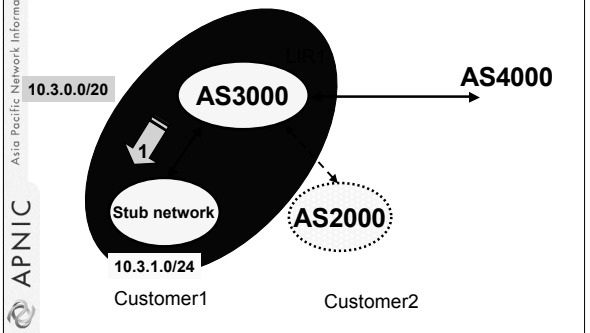

Case 1: Static route importation – aut-num

```
aut-num: AS3000
import:  protocol STATIC into BGP4
            from AS3000
              accept {10.3.1.0/24}
export: to AS4000 announce AS3000
[…]
```

Case 2: Multi-homed customer - provider set-up

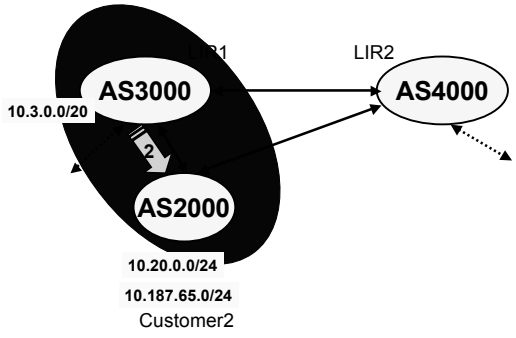

Case 2: BGP customers, - provider aut-num

```
aut-num: AS3000
import: from AS2000
                accept AS2000
export: to AS2000 announce ANY
[…]
```

- The simplest policy is strict customer/provider relationship
  - Customer sends its routes to provider
  - Customer accepts everything the provider sends

Case 3: Multi-homed customer - customer set-up

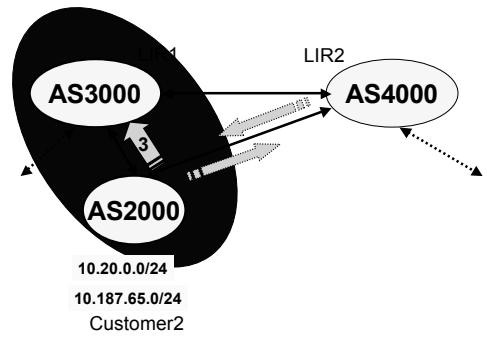

Case 3.1: Not Full Multihoming - customer aut-num

- DB objects:

```
aut-num: AS2000
import:  from AS3000 accept ANY
export:  to AS3000 announce AS2000
import:  from AS4000 accept AS4000
export:  to AS4000 announce AS2000
[…]
```

```
route: 10.20.0.0/24
origin: AS2000
[…]
```

```
route: 10.187.65.0/24
origin: AS2000
[…]
```

Case 3.2: Full Multihoming - customer aut-num

- Introducing policy, setting the “pref” value
  - lower the “pref”, the preferred the route

```
aut-num: AS2001
import:       from AS3000 action pref=50; accept ANY
export:       to AS3000 announce AS2001
import:       from AS4000 action pref=100; accept ANY
export:       to AS4000 announce AS2001
```

Questions?

- More cases to be discussed in the afternoon.
  - How to create BGP config

Internet Routing Registry

Why use the Internet Routing Registry ?

- Storing the routing policies
  - Consistent configuration between BGP peers (peers & customers & upstreams)
- When peering
  - Register your routes and filter your peers
  - Some transit providers and big ISPs ask for this
  - Useful for debugging problems
    - Compare reality versus policy
    - Contact info

Why use the Internet Routing Registry ?

- Automatically configure backbone routers
  - Expertise encoded in the tools that generate the policy rather than engineer configuring peering session
  - Enables router configuration (rtconfig)
- Shows routing policy between any two ASes (prpath)

Why use the Internet Routing Registry ?

- Enables creation of **aut-num** based on router conf (aoe)
- Providing defense against bogus routes and unintentional routing leaks
- Perform network planning

APNIC RR Specific Benefits

- Easier maintenance
  - Use one set of maintainer and person objects
    - to manage both Internet resources and routing information.
- Integrated resource & routing management
  - Ensures IP & ASN are within APNIC resource ranges.
  - Registered resource holder has control over routing objects
    - ensured through mnt-by, mnt-lower and mnt-routes Reduced costs
  - The APNIC Routing Registry free to APNIC members.

RtConfig

- Part of the IRRToolSet
- Generates router configuration based on the RR
  - Cisco, Bay's BCC, Juniper's Junos and Gated/RSd
- Creates route-map and AS path filters
- Can also create ingress / egress filters

Practical Usage of the RR

Potential Practical Problems

- Policy can easily get very complex and result in even more complex router configuration
- Line limit on cisco AS path filters
  - need to be careful when using as-set
- Nervous about configuring routers from public data?
  - Compare this with anti-virus SW updates!

Next steps

- Tasks for your own AS:
  - Create person and maintainer objects
    - Set up PGP authentication
  - Create aut-num objects for each AS
  - Identify IP prefixes associated with each AS
    - Create route objects in the database
  - Create as-set objects where policy is common
  - Either in the APNIC RR
    - Or in your own routing registry database

References

- RFC 2622 *"Routing Policy Specification Language (RPSL)"*
- RFC 2650 *"Using RPSL in Practice"*
- RFC 2725 *"Routing Policy System Security"*
- APNIC Routing Registry Guide
  - http://www.apnic.net/services/apnic-rr-guide.html
- IRRToolSet
  - http://www.ripe.net/ripencc/pub-services/db/irrtoolset/index.html

Questions?

Summary

- The Internet Routing Registry
- APNIC Database
  - RPSL
  - Queries and updates
  - Authentication
- Routing Policy
  - Case studies
- Routing Registry Benefits

Appendix

Object Templates in RPSL

Mntner object template

```
mntner:         [mandatory] [single]   [primary/look-up key]
descr:          [mandatory] [multiple] [ ]
admin-c:        [mandatory] [multiple] [inverse key]
tech-c:         [optional]  [multiple] [inverse key]
upd-to:         [mandatory] [multiple] [inverse key]
mnt-nfy:        [optional]  [multiple] [inverse key]
auth:           [mandatory] [multiple] [ ]
remarks:        [optional]  [multiple] [ ]
notify:         [optional]  [multiple] [inverse key]
mnt-by:         [mandatory] [multiple] [inverse key]
auth-override:  [optional]  [single]   [ ]
referral-by:    [mandatory] [single]   [inverse key]
changed:        [mandatory] [multiple]   [ ]
source:         [mandatory] [single]     [ ]
```

Inetnum object template

```
inetnum:      [mandatory]  [single]     [primary/look-up key]
netname:      [mandatory]  [single]     [lookup key]
descr:        [mandatory]  [multiple]   [ ]
country:      [mandatory]  [multiple]   [ ]
admin-c:      [mandatory]  [multiple]   [inverse key]
tech-c:       [mandatory]  [multiple]   [inverse key]
rev-srv:      [optional]   [multiple]   [inverse key]
status:       [generated]  [single]     [ ]
remarks:      [optional]   [multiple]   [ ]
notify:       [optional]   [multiple]   [inverse key]
mnt-by:       [mandatory]  [multiple]   [inverse key]
mnt-lower:    [optional]   [multiple]   [inverse key]
mnt-routes:   [optional]   [single]     [inverse key]
changed:      [mandatory]  [multiple]   [ ]
source:       [mandatory]  [single]     [ ]
```

Route object template

```
route:          [mandatory] [single]     [primary/look-up key]
descr:          [mandatory] [multiple]   [ ]
country:        [optional]  [single]     [ ]
origin:         [mandatory] [single]     [primary/inverse key]
holes:          [optional]  [multiple]   [ ]
member-of:      [optional]  [multiple]   [ ]
inject:         [optional]  [multiple]   [ ]
aggr-mtd:       [optional]  [single]     [ ]
aggr-bndry:     [optional]  [single]     [ ]
export-comps:   [optional]  [single]     [ ]
components:     [optional]  [single]     [ ]
remarks:        [optional]  [multiple]   [ ]
cross-mnt:      [optional]  [multiple]   [inverse key]
cross-nfy:      [optional]  [multiple]   [inverse key]
notify:         [optional]  [multiple]   [inverse key]
mnt-lower:      [optional]  [multiple]   [inverse key]
mnt-routes:     [optional]  [multiple]   [inverse key]
mnt-by:         [mandatory] [multiple]   [inverse key]
changed:        [mandatory] [multiple]   [ ]
source:         [mandatory] [single]     [ ]
```

Aut-num object template

```
aut-num:        [mandatory]  [single]     [primary/look-
up key]
as-name:        [mandatory]  [single]     [ ]
descr:          [mandatory]  [multiple]   [ ]
country:        [optional]   [single]     [ ]
member-of:      [optional]   [multiple]   [ ]
import:         [optional]   [multiple]   [ ]
export:         [optional]   [multiple]   [ ]
default:        [optional]   [multiple]   [ ]
remarks:        [optional]   [multiple]   [ ]
admin-c:        [mandatory]  [multiple]   [inverse key]
tech-c:         [mandatory]  [multiple]   [inverse key]
cross-mnt:      [optional]   [multiple]   [inverse key]
cross-nfy:      [optional]   [multiple]   [inverse key]
notify:         [optional]   [multiple]   [inverse key]
mnt-lower:      [optional]   [multiple]   [inverse key]
mnt-routes:     [optional]   [multiple]   [inverse key]
mnt-by:         [mandatory]  [multiple]   [inverse key]
changed:        [mandatory]  [multiple]   [ ]
source:         [mandatory]  [single]     [ ]
```

As-set object template

```
as-set:         [mandatory]  [single]     [primary/look-up key]
descr:          [mandatory]  [multiple]   [ ]
country:        [optional]   [single]     [ ]
members:        [optional]   [multiple]   [ ]
mbrs-by-ref:    [optional]   [multiple]   [inverse key]
remarks:        [optional]   [multiple]   [ ]
tech-c:         [mandatory]  [multiple]   [inverse key]
admin-c:        [mandatory]  [multiple]   [inverse key]
notify:         [optional]   [multiple]   [inverse key]
mnt-by:         [mandatory]  [multiple]   [inverse key]
changed:        [mandatory]  [multiple]   [ ]
source:         [mandatory]  [single]     [ ]
```

Route-set object template

```
route-set:    [mandatory]  [single]     [primary/look-up key]
descr:        [mandatory]  [multiple]   [ ]
members:      [optional]   [multiple]   [ ]
mbrs-by-ref:  [optional]   [multiple]   [inverse key]
remarks:      [optional]   [multiple]   [ ]
tech-c:       [mandatory]  [multiple]   [inverse key]
admin-c:      [mandatory]  [multiple]   [inverse key]
notify:       [optional]   [multiple]   [inverse key]
mnt-by:       [mandatory]  [multiple]   [inverse key]
changed:      [mandatory]  [multiple]   [ ]
source:       [mandatory]  [single]     [ ]
```

Inet-rtr object template

```
inet-rtr:     [mandatory]  [single]     [primary/look-up key]
descr:        [mandatory]  [multiple]   [ ]
alias:        [optional]   [multiple]   [ ]
local-as:     [mandatory]  [single]     [inverse key]
ifaddr:       [mandatory]  [multiple]   [lookup key]
peer:         [optional]   [multiple]   [ ]
member-of:    [optional]   [multiple]   [inverse key]
remarks:      [optional]   [multiple]   [ ]
admin-c:      [mandatory]  [multiple]   [inverse key]
tech-c:       [mandatory]  [multiple]   [inverse key]
notify:       [optional]   [multiple]   [inverse key]
mnt-by:       [mandatory]  [multiple]   [inverse key]
changed:      [mandatory]  [multiple]   [ ]
source:       [mandatory]  [single]     [ ]
```

Peering-set object template

```
peering-set:  [mandatory]  [single]     [primary/look-up key]
descr:        [mandatory]  [multiple]   [ ]
peering:      [mandatory]  [multiple]   [ ]
remarks:      [optional]   [multiple]   [ ]
tech-c:       [mandatory]  [multiple]   [inverse key]
admin-c:      [mandatory]  [multiple]   [inverse key]
notify:       [optional]   [multiple]   [inverse key]
mnt-by:       [mandatory]  [multiple]   [inverse key]
changed:      [mandatory]  [multiple]   [ ]
source:       [mandatory]  [single]     [ ]
```

Filter-set object template

```
filter-set:     [mandatory]  [single]     [primary/look-
up key]
descr:          [mandatory]  [multiple]   [ ]
filter:         [mandatory]  [single]     [ ]
remarks:        [optional]   [multiple]   [ ]
tech-c:         [mandatory]  [multiple]   [inverse key]
admin-c:        [mandatory]  [multiple]   [inverse key]
notify:         [optional]   [multiple]   [inverse key]
mnt-by:         [mandatory]  [multiple]   [inverse key]
changed:        [mandatory]  [multiple]   [ ]
source:         [mandatory]  [single]     [ ]
```

Rtr-set object template

```
rtr-set:        [mandatory]  [single]    [primary/look-up key]
descr:          [mandatory]  [multiple]
members:        [optional]   [multiple]
mbrs-by-ref:    [optional]   [multiple]
remarks:        [optional]   [multiple]
tech-c:         [mandatory]  [multiple]  [inverse key]
admin-c:        [mandatory]  [multiple]  [inverse key]
notify:         [optional]   [multiple]  [inverse key]
mnt-by:         [mandatory]  [multiple]  [inverse key]
changed:        [mandatory]  [multiple]
source:         [mandatory]  [single]
```