

CCC, August 2007: Workshop: Using RIPE Routing Registry

Vesna Manojlovic, RIPE NCC

Summary

Analyse your ISPs routing policy, learn RPSL, take a look at the global BGP routing table, and see what's new about 32-bit AS numbers

Assumptions

- You are not working for an ISP / Local Internet Registry (LIR) [0], [1]
- You use (i)BGP in your networks
- You know how to use a “whois” Database [2]

Introduction

If you are not from an LIR, you are (for people from the RIPE NCC) an “End User”. That means that you can **not** directly **take part** in the “Routing Registry” (**RR**). But, since the RR is part of the “RIPE (whois) Database”, and that is a “Public Network Management Database”, you can still **use** the RR!

ISPs use RR to documenting routing policy, to publish it and achieve distributed communication of updates, consistent information between neighbors, and some kind of routing security. Policy is described on abstract level, using “Routing Policy Specification Language” (RPSL) [3], which overcomes incompatible hardware platforms. There are lot of tools for using, maintaining and analyzing RR data [5].

Looking-up routing policy

To see the routing policy of your ISP, you need to know which “Autonomous System” (AS) you belong to. To find the “AS Number” (ASN) for your IP address, you can use “RISwhois” [5], [6]:

```
# whois -h riswhois.ripe.net 212.212.212.212
```

The answer will show a (pseudo) “route” object, with the address range (prefix) to which a specific IP address belongs to in CIDR notation [4], and the “origin:” – the ASN that originated the prefix.

```
route:          212.212.0.0/16
origin:         AS4589
descr:         EASYNET Easynet Group Plc
lastupd-frst:  2007-05-29 07:34Z  1.2[...]1@rrc05
lastupd-last:  2007-08-06 04:20Z  23.[...]3.5@rrc00
seen-at:       rrc00,rrc01,rrc03,rrc04,... ,rrc14,rrc15
num-rispeers:  105
source:        RISWHOIS
```

Now you can look-up the ASN:

```
# whois -h whois.ripe.net AS4589
```

The answer will be something like this:

```
aut-num:       AS4589
as-name:       EASYNET
descr:         Easynet Group Plc
import:        from AS1 accept ANY
import:        from AS112 195.69.1.39 at 195.69.1.38 accept AS112
[...]# skipped 1000 lines
export:        to AS1 announce AS-EASYNET
export:        to AS112 195.69.1.39 at 195.69.1.38 announce AS-E
```

[...]# *skipped 1000 lines*

```
admin-c:      CL60-RIPE
tech-c:      EDEN-RIPE # Easynet Germany NOC
mnt-by:      EASYNET-MNT
source:      RIPE # Filtered
```

In order to understand what all the details in the import and export lines of this object mean, you need to know RPSL [3]. If you don't want that, you can use visual tools, that rely on RIS (later on).

Some AS Numbers with detailed policy: AS286 (KPN Eurorings), AS5400 (BT), AS1299 (Telia), AS8918 (Carrier1), AS3356 (Level3), AS3327 (Linxtelecom)

RtConfig

RtConfig is one of the tools from the IRRToolSet [5], which makes routing information more useful for network engineers, by providing automated router configuration. It reads policy from the RR, and generates parts of the router configuration file, in the format specific for several vendors: Cisco, Bay's BCC, Junos, Gated/RSd. It creates access list, route-map and AS path filters.

```
# RtConfig -h whois.ripe.net -p 43 -s RIPE -protocol ripe
-cisco_use_prefix_list
RtConfig> @RtConfig access_list filter AS4589

no ip prefix-list pl100
ip prefix-list pl100 permit 212.11.0.0/18
ip prefix-list pl100 permit 194.156.206.0/24
[...]# skipped 130 lines
ip prefix-list pl100 permit 155.91.88.0/24
ip prefix-list pl100 deny 0.0.0.0/0 le 32

RtConfig> @RtConfig import AS3000 10.3.15.2 AS2000 10.20.0.3

no ip prefix-list pl100
ip prefix-list pl100 permit 10.20.0.0/24
ip prefix-list pl100 permit 10.187.65.0/24
ip prefix-list pl100 deny 0.0.0.0/0 le 32
!
no route-map MyMap_2000_1
route-map MyMap_2000_1 permit 1
  match ip address prefix-list pl100
!
router bgp 3000
neighbor 10.20.0.3 route-map MyMap_2000_1 in
```

There are many options and possibilities for using RtConfig; see [1], [13], [17]

TEST (RIPE whois) Database

If you are not an ISP nor LIR, but still want to learn about usage of the RR and RPSL, you can use “TEST Database”: <http://www.ripe.net/db/support/testdb.html>

It is using the same server software and has the same client behavior as RIPE Database; however, it does not contain same information as operational DB. You can give to yourself as many AS numbers as you want, and create your own “aut-num” objects, and “route” objects for as much address space as you please!

Attributes:

- source: TEST
- referral-by: TEST-DBM-MNT
- nic-hdl: xxxx-TEST
- TEST-NCC-HM-MNT to maintain ALLOCATED inentnums

Interfaces:

- whois -h test-whois.ripe.net
- mailto: <test-dbm@ripe.net>
- <http://www.ripe.net/webupdates-test>
- <http://www.ripe.net/db/syncupdates/syncupdate-test-minimal.html>

Running your own RR

If you want to document your internal routing policy in RPSL; or, if you want to have a private RR for your local IXP (where you have BGP between your peers, but you are all using private AS numbers: reserved AS numbers (RFC-1930): 64512 – 65535) -- you can do it. Here are the steps:

1. Download, install and set-up server SW
 - RIPE DB Server SW: <http://www.ripe.net/db/cvs-bugzilla.html>
 - IRRd: <http://www.ird.net/>
2. Register your RR with the IRR
 - <http://www.ird.net/docs/list.html>
3. Possibly: get the mirroring agreement with the RIPE DB
 - <http://www.ripe.net/db/mirrored.html>
4. Give your customers access to your RR
 - Read-only? / With privileges to update objects?

Private routing registries:

- Savvis (ex C&W) - rr.savvis.net (AS3561)
- whois.connect.com.au (CCAIR) (AS2764)
- wix.net.nz - Wellington Internet eXchange

Routing Information Service (RIS)

The RR contains the theoretical, documented, intended routing policy. The reality is often different then the plans, or from the documentation. To collect information about the real-world BGP routing, RIPE NCC has a project called Routing Information Service (RIS) [5]: Remote Route Collectors (RRC) at multiple Internet eXchange Points (IXP) that peer with more then 500 peers, listen to and collect the route announcements, and store them for at least 3 months. Data can be searched through command line, looking-glass, web-interfaces, weekly reports, various visual tools, and notification systems.

RISwhois that maps prefix to originating AS number(s), by looking it up in BGP data collected by RIS, rather than in RR. It also gives quick & summarised view of prefix visibility, by listing the RRCs in which the prefix was seen. <http://www.ris.ripe.net/cgi-bin/riswhois.cgi>

BGPlay is a Java application which displays animated graphs of the routing activity of a certain prefix within a specified time interval. Its graphical nature makes it easy to understand how BGP updates affect the routing of a specific prefix. (See also iBGPlay [7]) <http://www.ris.ripe.net/bgplay/>

MyASn service notifies network operators when a prefix (of your choice) is announced with an incorrect AS path. An AS path is seen as incorrect when it does not match with a regular expression that you have set up. The system can send alerts by e-mail or to your own syslog server. Alarms can be divided in different groups, and have many other options to configure.

<http://www.ris.ripe.net/myasn.html>

Routing Registry Consistency Check Project (RRCC) project helps identify, correct and keep the routing data in the RIPE Database up-to-date, by comparing it to RIS data. <http://www.ripe.net/rrcc/>

Routing security

RR offers basic routing security by allowing operators to check who is registered as a user of a certain ASN (since the “aut-num” object can be created only by RIPE NCC, on behalf a client that was either represented by the LIR, or they are an LIR). Routable prefixes are then registered as “route” objects, which can be created only by the “owners”/maintainers of both ASN and the IP address space (“inetnum”). If ISPs generate access-list filters based on the information from the RR, they can prevent mistakes which happen due to mis-configuration, when some operators leak other people’s ranges, thus black-holing them.

Bogons is a name for un-allocated and/or reserved address ranges (also called **martians** [8]), that should not be routed on the public Internet. Problem is that people are sometimes leaking bogons due to mis-configuration, or some (other) people use bogon space as source for spamming, DDoS, probes...

To stop sending or receiving bogon routes, you can either:

- maintain your bogon filters manually
- apply “Secure BGP Template” [9]
- peer with the bogon route server [10]
- in RPSL / RR: add “AND NOT fltr-bogons” to all your import and export filter rules

Outdated “bogon” filters: when new /8 block is allocated to RIR / LIR, it is sometimes unreachable from networks with stale filters. RIPE NCC has a “debogonising” project [11] where you can test if your filters are up-to-date, or, if you are an owner of new address space, where will your network **not** be visible.

You can also find what are all the ranges used by the RIPE-region LIRs:

<https://www.ripe.net/ripe/docs/smallest-alloc-sizes.html>

Or, you can follow the mailing lists where RIRs announce new /8 blocks e.g. RIPE routing-WG or NANOG [12].

32-bit AS numbers

The 16-bit AS numbers are running out (there are 65000 of them, and we have used about 40000 already. So, the new, improved, larger version was introduced: 32-bit (RFC-4893). . Right now, and until end 2008, you can get (through your LIR) either, but you will get a 16 bit ASN by default, 32 bit on request. In 2009, it will be the other way around. After 1/1/2010 RIPE NCC will always give a 32 bit AS Number.

References, links, more information:

[0] RIPE NCC is a Regional Internet Registry that distributes IP resources (IPv4, IPv6 & ASN) to Local Internet Registries, in Europe, Middle East and parts of Central Asia., and does other co-ordination work (RIPE (whois) Database), reverse DNS + DNSSEC, K-root nameserver, ENUM, outreach, statistics...). It is an independent, not-for-profit, membership association.

RIPE is an open community of network operators that decides on policies and procedures for technical running of the Internet in “Europe”.

<http://www.ripe.net/>

[1] If you **are** working for an ISP / Local Internet Registry (LIR), come to the “**Routing Registry Hands-on Course**”: ripe.net/training/rr/outline.html

[2] Using RIPE Database: where to learn more:

- <https://e-learning.ripe.net>
 - Creating and Updating Basic Database Objects
 - Using PGP Keys & X509 Certificates in mntner Objects
- RIPE Database Users Manual: Getting Started
 - <http://www.ripe.net/ripe/docs/ripe-402.html>
- Whois Queries Reference Card
 - <http://www.ripe.net/db/support/db-refcard.pdf>

[3] RPSL:

- RFC-2622: Routing Policy Specification Language
- RFC-2725: Routing Policy System Security
- RFC-2650: Using RPSL in Practice

[4] CIDR – Classless Inter-Domain Routing

- RFC-1518: An Architecture for IP Address Allocation with CIDR
- RFC-4632: The Internet Address Assignment and Aggregation Plan
- <http://www.ripe.net/info/info-services/addressing.html>
- The cidr report: <http://www.cidr-report.org/>

[5] Tools:

- RIS: <http://www.ripe.net/ris/> * *RIPE NCC*
- IRRToolSet – including RtConfig * *ISC*
 - <http://www.isc.org/index.I?/sw/IRRToolSet>
- Routers configuration / BGP filters generator * *Marco d'Itri*
 - based on the popular Template::Toolkit perl package
 - <http://www.linux.it/~md/software/rpsltool-1.0.tgz>
- IRR Power Tool * *Richard Steenbergen*
 - a PHP and CVS-based SW for using IRR data to manage prefix-lists
 - <http://sourceforge.net/projects/irrpt/>

[6] RIPE whois command line client:

<ftp://ftp.ripe.net/tools/ripe-whois-latest.tar.gz>

[7] iBGPlay

“BGPlay” for ISP, to have a view of internal network, and to follow the paths of outgoing traffic. The “route collector” is within your ASN, and gathers info from your multiple routers.

<http://www.dia.uniroma3.it/~compunet/www/view/tool.php?id=ibgplay>

Contact: Maurizio Pizzonia, pizzonia @ dia.uniroma3.it

[8] Martians

- RFC-1918 <ftp://ftp.ripe.net/rfc/rfc1918.txt>
- <http://www.isi.edu/~bmannig/dsua.html>

[9] Secure BGP Template

<http://www.cymru.com/Documents/secure-bgp-template.html>

[10] Bogon route server, Team Cymru, AS65333, community 65333:888

www.cymru.com

[11] De-bogonising project

www.ripe.net/ripe/draft-documents/deboganising-draft.html

[12] Routing discussion fora:

- <http://www.ripe.net/ripe/wg/routing/>
- <http://www.nanog.org/>

[13] RtConfig documentation & examples:

<http://www.isc.org/index.pl?sw/IRRToolSet/documentation/RtConfig.php>
<http://www.ripe.net/training/rr/material/examples/>

[14] BGP Communities

- Article about communities:
 - cisco.com/warp/public/759/ipj_6-2/ipj_6-2_bgp_communities.html
- RFC-1997: BGP Communities Attribute
- RFC-1998: An application of the BGP Community Attribute in Multi-home Routing

[15] BGP Update Report

<http://bgpupdates.potaroo.net/>

[16] Netkit

<http://www.netkit.org>

[17] More presentations, seminars, workshops:

LIR Basics and Routing Registry, Related Services and News

European Peering Forum, Frankfurt , DE, November 2006

<http://www.ripe.net/info/ncc/presentations/RIPE-NCC-LIR-and-RR-seminar.pdf>

Using Routing Registry and Related Tools for Configuring Routers

APRICOT, Kyoto, Japan, February 2005

http://apricot.net/apricot2005/slides/C5-4_2.pdf

An Introduction to the Routing Registry

LINX Meeting, London, UK , November 2004

<http://www.ripe.net/info/ncc/presentations/linx-rr-200411/page1.htm>

Benefits of using RIPE Routing Registry and Related RIPE NCC Tools Telfor, Belgrade, CS , November 2004

<http://www.ripe.net/info/ncc/presentations/rr-telfor-20041124/page1.htm>

RIPE Routing Registry and Related RIPE NCC Tools

Ficix Workshop, Helsinki, FI , June 2004

<http://www.ripe.net/info/ncc/presentations/ficix-200410/page1.htm>