Routing Registry Consistency Check

Joachim Schmitz
AOL Inc

Engin Gündüz
Shane Kerr
Andrei Robachevsky
João Luis Silva Damas

RIPE NCC

Document ID: ripe-201
Date: 4 December 2001

ABSTRACT

This document describes a project to crosscheck European Internet routing against data registered in the RIPE Routing Registry. This project was proposed as a new activity by the RIPE Routing Working Group. The document gives an overview of the project, discusses the implementation, and mentions some points that have to be addressed in future design documents. This document is intended to solicit input from interested parties.

Table of Contents

---

1. Introduction

For quite some time, the RIPE Database has been part of a global system known as the Internet Routing Registry (IRR). It consists of a set of online databases that offer the possibility to store routing policy information that may be used in network verifications or configurations. Other databases in the IRR include the RADB, ARIN, and a growing number of databases operated by Internet Service Providers (ISPs).

Practical applications of the IRR have ranged from simple registration of Internet resources, such as autonomous system (AS) numbers, to registration of routing policies, and complete routing configurations.

However, the IRR is not generally considered to form an essential or even adequate tool for network operations. Sometimes, it has been viewed as a somewhat obscure system that was difficult to use and its potential benefits have not become available to a wider audience.

Over time, the degree of usage and attention to data maintenance has varied among users resulting in a situation where the quality of data stored is mixed.

This project aims at correcting that situation by establishing a stronger link between routing data as seen on the Internet and as stored in the RIPE Database. It is also a goal of this project to address the causes that led to the current unsatisfactory situation by providing tools and information that enhance the usefulness of the IRR to ISPs.

2 Motivation

As previously described, the IRR stores Internet routing information that can be used by ISPs to configure and debug their connectivity. This information is stored in the RIPE Database as **route**, **aut-num**, **inet-rtr**, **route-set**, **as-set**, and other RPSL objects containing information such as the intended AS origin of an announced prefix, peering information between different ASes, etc. The quality of these data varies widely, while it is supposed to be a representation of what is actually seen in the Internet Routing system.

During the initial stages of RIPE Database development, two major drawbacks existed. On one hand, the limitations of the RIPE-181 representation of routing policies became obvious over time, resulting from growing demands and evolving routing protocols. To a significant extent, this has been amended by the introduction of RPSL[1].

On the other hand, mechanisms for data protection within the Database were not strong, and unauthorised (usually accidental) modifications of data occurred. In addition, there was no authorisation mechanism for registration of data in the Database and any user could essentially store any information in the Database (e.g. a new origin for a prefix already registered in another **route** object).

Nowadays, the available protection mechanisms are strong and users can ensure that their data may only be modified by themselves. Regarding authorisation, the RPSL version of the RIPE Database (version 3.0) implements the Routing Policy System Security (RPSS) specification [2], which enables security mechanisms for object creation. Thus a major concern regarding reliability of data has been removed.

As a result of the two factors described above, and due to the highly dynamic nature of the Internet, a significant amount of data in the Routing Registry is incorrect, out of date, or simply missing. With a sound basis in place from which to start, this situation may now be remedied.

It is worthwhile to note that existence of incorrect data in the IRR does not necessarily make it unusable in its entirety. For example, ISPs that ask their customers to register in the IRR can still configure their routers from the IRR because the subset of data that is of interest for this purpose would be complete and accurate, even if other data in the IRR might not be.

While improvements to the Database system will help prevent erroneous registrations in the future, there is a need to correct the inaccurate data currently stored in the Database. Moreover, development of support mechanisms is required to enable ISPs to maintain data and keep it up to date as the Internet evolves.

The RIPE Routing Working Group identified these needs and requested that the RIPE NCC addressed them. This document describes the corresponding RIPE NCC project.

3. Goals

From the previous considerations, the main goals of the project can be outlined as follows:

- To track, report on, and increase the accuracy of IRR data
- To track and report IP allocation usage in the Internet
- To develop tools to improve user interaction with the IRR; in particular, to increase ease of use of the IRR and to benefit from its resources
- To disseminate tool and user information

The following sections discuss how to attain these goals.

3.1 Types of Routing Data in the IRR

The current IRR exhibits two types of data

- IP routes
- routing policies

These are tied to the following objects in the Database

- **route** objects, describing which AS originates each prefix
- **aut-num** objects for routing policies
- **inet-rtr**, **as-set**, **route-set**, and other RPSL objects supporting the description of routing policies.

3.2 Types of Routing Data on the Internet

In general, routing data on the Internet is distributed via BGP announcements between peering partners. A full set of Internet routes is available by analyzing the full set of BGP announcements. The set will vary depending on where in the Internet the BGP feed is collected.

BGP announcements contain

- the route prefix and its length;
- the AS path under which the announcement is visible and its originating ASes;
- additional BGP attributes that propagate through the routing system.

Any of these data items are available by establishing appropriate BGP peerings. The main source of live routing data for this project will be the Routing Information Service (RIS) project [3] being carried out by the RIPE NCC.

3.3 Comparing Registry and Internet Data

The two sources of data (IRR and live Internet routing data) have quite different formats. Tools that can understand both sources and make comparisons will be developed.

3.3.1 Route Advertisements

A route, which is the combination of network addresses with an originating AS, may be:

- recorded in the RIPE Database and advertised on the Internet (which is the optimum case)
- recorded in the RIPE Database but not advertised on the Internet, and
- advertised on the Internet but not recorded in the RIPE Database.

The latter two describe mismatches that may be detected by comparing the data sets. These mismatches may be caused by errors in aggregation (e.g. recording 193.0.0.0/22 in the Database when advertising 193.0.0.0/21), errors in origin (e.g. an advertisement from an old provider), simple omissions, or other reasons. Initially, no attempt will be made to classify errors. They will just be recorded.

3.3.2 RPSL Policy and AS Paths

Routing data collected in the Routing Information System (RIS) includes AS path information. It can be used to verify whether the policies recorded in corresponding aut-num objects allow the AS path seen in the live routing data. This means that missing "import:" and "export:" attributes in the **aut-num** objects may be detected.

It is not possible to detect unnecessary "import:" and "export:" attributes, due to the limitations of the RIS data imposed by the nature of BGP. Some routes will not get advertised to any of the route collectors, but still be used on the Internet. This may occur, for example, if a slow or expensive link is used for backup purposes that are activated only in case the regular connection breaks. Another example is upstream route aggregation, which removes longer prefixes from the routing table.

3.3.3 Allocation Registry Checks

By including data from the IP allocation registry, it will be possible to track address space usage and detect announcements of space that should not be advertised on the Internet. This is not limited to the use of private address space or AS numbers, but also reveals usage of unallocated AS numbers and IP addresses.

4 Reporting and Tool Development

Information produced by this project will be made available at the RIPE NCC web site and/or through e-mail subscription. This will include the following:

4.1 General IRR Consistency Report

The report will present general statistical data regarding the size of the IRR, amount of inconsistencies of different types, and delegated address space usage.

4.2 Specific IRR Consistency Report

This type of report will present detailed information about IRR inconsistencies for specific portions of the IRR. A specific portion may be defined by the user based on:

- AS number or **as-set**, AS information: In this case, the report will list deviations in AS numbers (missing or extra) and AS paths from registered routing policies.
- AS number or **as-set**, route information: In this case, the report will contain a list of inconsistent routes originated by the specified AS or ASes.
- IP address space as defined by **inetnum**, **route**, or **route-set** objects. In this case, the report will contain a list of inconsistent routes concerning this address space.
- **mntner** object name, which is referenced by "mnt-routes:" attribute in **aut-num**, **inetnum**, or **route** objects. In this case, the report will contain a list of inconsistent routes related to address space or autonomous systems covered by this attribute.

- **mntner** object name, which is referenced by "mnt-lower:" attribute in **as-block**, **aut-num**, **inetnum**, or **route** objects. In this case, the report will contain a list of inconsistent routes related to address space or autonomous systems covered by this attribute.
- **mntner** object name, which is referenced by "mnt-by:" attribute in **aut-num**, **inetnum**, or **route** objects. In this case, the report will contain a list of inconsistent routes related to address space or autonomous systems specified by these objects.

Two types of the reports may be requested: periodic report and monitoring report sent when an inconsistency appears.

4.3 IRR Correction Wizard

In the event that registrations in the IRR require updating, any user today must set up the corresponding objects by hand, which often is cumbersome and subject to human error. To improve ease of use on the IRR, the project will introduce an IRR correction wizard. This tool will take the routing policy as found on the Internet and provide the user with a set of corresponding objects which can be applied to update the IRR. This includes additions and deletions. The user must verify correctness of the information presented, add the corresponding authorisation information, and forward the objects in a message to the Database update mechanism (<auto-dbm@ripe.net>). The "correction wizard" is a supplementary tool to be used along with reports produced in section 4.2.

4.4 Router Configuration Checker

It will be possible to identify some potential problems in the parts of a router configuration built from the IRR using tools like RtConfig. Normally, the automatic router configuration generators generate only parts of the configuration using IRRs, rather than a complete configuration. Similarly, this tool will check subsets of configurations and identify the statements that are inconsistent with the real state of the Internet. It can be used as router configuration debug help.

4.5 Other Tools

Tools developed in the framework of this project will not be limited to the ones described above. Feedback from the community may identify demand for other tools.

5 Side-Note on Data Privacy

Due to business confidentiality or other considerations, several service providers do not want to register their routes or policies in a public Routing Registry. However, it should be noted that in order to reach a certain IP address, routing information regarding this IP address must be public. If it were not public, this IP address would not be reachable on the Internet.

As a matter of fact, any site that wants to be reachable on the Internet must have or be part of one active public route in the Internet routing table. If a service provider chooses to route certain networks only on limited parts of the Internet (if not using private address space anyway; also in conjunction with NAT) there is no immediate need to feed these data into a public Routing Registry. As an example, we may consider a private peering between two service providers where the corresponding routes are not propagated outside their own autonomous systems.

It is obvious that private agreements on routes and policies do not need to be registered in the public IRR because they do not show up on the Internet.

Consequently, private or special agreements that are a result of mutual consent among the parties involved do not need to show up in the public IRR. As soon as coordination becomes impractical due to mere size, or if a guaranteed authoritative reference is needed, the public IRR comes in.

6 Timeline

Work on this project started during the year 2000. Initial results of this project are available at http://www.ripe.net/ripencc/pub-services/db/rrcc.
An initial prototype of this system was presented at the RIPE 38 Meeting

In April 2001, the RIPE Database migrated to version 3.0, which supports RPSL. Having the new Database in production will enable faster development of this project and more convenient use for ISPs.

References

[1] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg and M. Terpstra, "Routing Policy Specification Language (RPSL)", RFC 2622, June 1999.

[2] C. Villamizar, C. Alaettinoglu, D. Meyer and S. Murphy, "Routing Policy System Security", RFC 2725, December 1999.

[3] A. Antony, H. Uijterwaal. "Routing Information Service. Design Note", RIPE-200, October 1999.