

RIPE Whois Registration in 2005: What should be in Whois and Why?

Eva Ericsson Rabete, Leo Vegoda

Document ID: ripe-tbd

Date: April 2005

1) Introduction

This document discusses the use of the RIPE Whois database. It looks at what information is recorded in the database and why. Further, it asks whether it is appropriate to continue to use the RIPE Whois database today in the same way as it was used in the second half of the 1990s.

Eva Ericsson Rabete asked these questions at the Address Policy WG session at RIPE 49 in Manchester.¹ It was agreed that the issues should be raised on the address-policy-wg@ripe.net mailing list for further discussion.

2) EU Data Protection Principles

When assignments are made to private individuals, data protection law restricts what information can legally be published in a public database. EU Data Protection law sets down three requirements for processing personal data:

- The personal data must be collected and processed for specific and explicitly defined purposes.
- The purposes have to be legitimate.
- The processing has to be carried out in a careful manner and must be necessary for, and proportionate to, the legitimate purpose for which the processing is done.

The RIPE NCC is based in the Netherlands and must comply with Dutch Data Protection law, which is based on European Union (EU) Directive 95/46/EC. The EU consists of 25 countries, all of which must implement the EU directive in local law. A list of EU member countries is available from:

http://europa.eu.int/abc/governments/index_en.htm#members

The community needs to decide whether the legitimate purposes can be met with something more restrictive than the current Whois database, in order to meet the requirement for the data

¹ <http://www.ripe.net/ripe/meetings/ripe-49/presentations/ripe49-ap-whois-ipv6.pdf>

<http://www.ripe.net/ripe/meetings/ripe-49/webcast-files/ap-2.wmv>

processing to be proportionate. In most cases, the person whose data is stored and published in the Whois database needs to give their unambiguous consent.

3) What is Whois?

The Whois protocol was originally defined in RFC 812 and RFC 954. These RFCs defined policy for what information should be reported as well as the technical specification for the protocol. The current specification for the Whois protocol is RFC 3912. The major change since RFC 954 is that it removes all text not relevant to the on-the-wire protocol. The RIPE NCC has published IP address registration data in a Whois database since it began registering Internet number resources in the early 1990s.

4) What is registered?

IPv4, IPv6 and AS Number registrations are recorded in the RIPE Whois database. Anyone can query the database using the Whois protocol.

The main aim of registering Internet Number Resources in the public Whois database is to ensure uniqueness. Registration in the Whois database helps ensure that two different, but interconnected networks, do not accidentally use the same set of Internet Number Resources on their networks.

The public nature of the database and its information makes it clear who is using the resource. Network operators can use this information when configuring their networks to ensure that they do not advertise or route IP address space inappropriately.

5) Why is it there?

Contact information for the users of IPv4 networks² has traditionally been registered in the Whois database. Originally, this was to enable network operators to contact each other and troubleshoot connectivity problems. More recently, registration has become useful in the administration of the address space as a finite resource.

It is possible to query the RIPE Whois database for a range of different information. The most important database searches are for e-mail addresses, people's names, IP addresses and the various alphanumeric IDs (nic-hdls) given to people, companies and maintainers (change control protection mechanisms) in the database.

² Network prefixes shorter than /30

6) How is it used?

Routing Registry information is the information associated with the network routing policies of the Internet's various Autonomous Systems (ASs). This is published using **aut-num**, **route** and other database objects. Routing registry information is used by network operators to communicate with each other. It is not used by – and does not relate to – End Users.

IPv4 and IPv6 networks are registered in the Whois database as **inetnum** and **inet6num** objects. The data registered is used for a number of purposes, but there are no clear guidelines giving explicit guidance on the content or purpose of the registration and the permitted use of the data. The data registered includes details of the range in use, the user of the address space and contact information for the network operator. In some cases, this might be generic 'role' information; in others it could be the names, telephone numbers and e-mail addresses of the ISP's staff or customers.

7) Why do people register their networks?

Operators whose networks provide services or transit to third parties will often want those third parties to be able contact them. In cases where there is not a direct, contractual relationship it is useful to be able to access appropriate contact information. This is a key reason for the existence of the RIPE Whois database.

Stub sites whose networks do not provide services or transit to third parties, for instance enterprise and home networks, are less likely to want to be contacted by previously unknown third parties. In most cases, these network operators are unlikely to be able to diagnose or resolve the causes of network or other problems with their site. For this reason, there is little benefit in including these network operators' contact information in a Whois database.

8) Who uses data from the Whois database?

Operators often need to debug connectivity and reachability problems that cross intermediate networks. Being able to contact third party operators to report a problem can be useful.

End Users need to be able to report connectivity and other problems to third party network operators. It is likely that in many situations End Users will need a different set of contacts from those used by network operators.

RIRs currently use records in the Whois database to determine two things: how much of an LIR's existing address space has been sub-allocated or assigned and whether those sub-allocations or assignments are valid. The case is different for IPv4 and IPv6.

Even a very large IPv4 allocation is unlikely to generate a large number of individual assignment records. This is because ISPs can make large assignments holding the IP space for all their single-IP address customers. It is only networks larger than a /30 that need to be registered in the RIPE Whois Database.

Current policy requires Local Internet Registries (LIRs) to register all /48 IPv6 assignments in a database accessible to their Regional Internet Registry (RIR). Doing so allows the RIR to determine whether the LIR has made a sufficient number of /48 assignments to qualify for an additional IPv6 allocation. These database registrations might either be in the Whois database or an internal private database with limited access for the RIPE NCC. However, in most cases, End Users' residential connections will receive networks of the same size as large commercial enterprise sites, so neither company nor individual subscribers would be listed in a public Whois database. We should consider the fact that:

- stub sites are unlikely to be helpful when contacted by third parties.
- large ISPs might well make more than one million /48 assignments (a significantly larger number of assignments than would be made for IPv4 connections).

9) Who really needs access to Whois?

Law enforcement agencies, intellectual property owners and other organisations occasionally want to know who was using a particular IP address or network. In some cases they can get this information from Whois databases. In other cases they need to question the organisation running the access network for more specific information regarding the use of the addresses in which they are interested.

10) Country attribute

Both **inetnum** and **inet6num** database objects require publication of country information. More than one country can be specified for a single network.

It is not clear whether the country information is meant to be the country or countries where the IP space is in use, where the network connection is based or where the LIR is headquartered. The documentation for the objects is not much help. It just states, "Identifies the country". This is presumably because it was not clear to the RIPE NCC staff who produced the documentation, or

anyone else, what the 'country:' attribute is meant to signify. It is worth noting that because the purpose of the information is not well defined, it is not possible to rely on its accuracy. IP addresses do not have any national characteristics, so it is difficult to attempt to tie them to one or more countries with any meaning or relevance.