The RIPE Database Requirements Task Force

Status: Draft Report
Date: July 2021

Task Force Members:

- Nick Hilliard
- James Kennedy (co-Chair)
- Shane Kerr (Vice Chair)
- Peter Koch
- Sara Marcolla
- Bijal Sanghani (Chair)

RIPE NCC Staff Support

- Boris Duval
- Edward Shryane
- Maria Stafyla

Contents

1. What is the RIPE Database?

Since August 1992, the RIPE Database has served as the authoritative registry of Internet number resources and related information within the RIPE NCC service region (Europe, Middle East, parts of Central Asia). The RIPE Database was built to facilitate coordination between network operators across this region. It also provides information about Internet number resources distributed prior to the current Regional Internet Registry (RIR) system.

RIPE Database information was originally provided on a voluntary basis. This changed when the RIPE NCC began allocating IP addresses and established data collection processes as part of its role as an RIR.

There were several iterations of the RIPE Database as its structure and functionality evolved. New features and objects were added over the years, making the database information richer and more complex. Some data was also migrated out of the RIPE Database in the past, such as information about ccTLD domain names.

In 1995, the Internet Routing Registry (IRR) was created. This is a subset of the RIPE Database that provides routing information. Its purpose is to ensure the stability and consistency of the Internet-wide routing system by sharing information between network operators.

More recently, the RPKI Database was created to offer verifiable proof of holdership of resource registrations by an RIR.

2. The difference between the RIPE Database and the RIPE Registry

The RIPE community has tasked the RIPE NCC to maintain a repository of all allocated Internet number resources in its service region. This information is stored in the RIPE Registry and the RIPE Database.

The RIPE Registry is maintained by the RIPE NCC and contains all data, private and public, about resources and resource holders in its service region. The RIPE Database provides a public view of some RIPE Registry data. The information disclosed in the RIPE Database aims to facilitate cooperation and coordination between network operators and other stakeholders for a variety of operational tasks, including troubleshooting and preventing outages.

The RIPE Database is maintained by both the RIPE NCC and resource holders. Usage of the database is covered by the RIPE Database Terms and Conditions.

The RIPE NCC is responsible for allocating resources to its members as well as avoiding discrepancies between the RIPE Registry and the RIPE Database.

Resource holders are responsible for updating information regarding their resource usage in the RIPE Database.

Figure 1: Public and non-public data in the RIPE Registry and the RIPE Database(s)

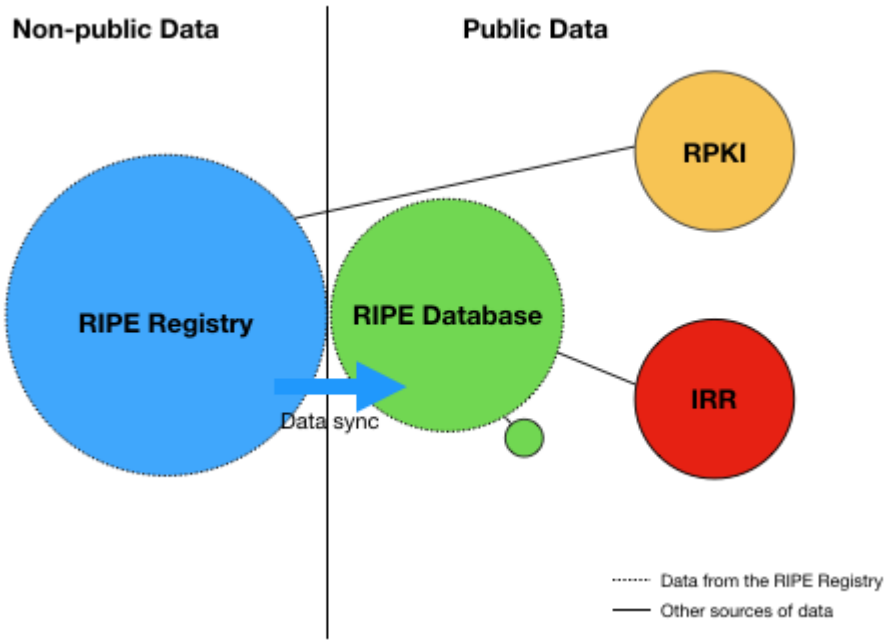

3. Why are we reviewing the RIPE Database functionality now?

The RIPE Database provides essential information to members of the RIPE community, which helps them to keep networks and the Internet running in their region. Many stakeholders depend on the accuracy and availability of the data stored in the database to do their job properly, especially regarding cyber security. Some database users, such as ISPs or IXPs, have been part of the RIPE community for years, while others are relatively new, such as Law Enforcement Agencies (LEAs) or regulators. These user groups have different needs and expectations regarding the database which is creating friction within the community. Changing privacy requirements and the question of whether personal data is needed in the RIPE Database are also ongoing discussions inside the RIPE community. Although there was consensus on this topic in the past, this is less obvious today.

While the RIPE Database Working Group and the RIPE NCC are able to solve a lot of the operational issues, a high-level approach was needed to establish a general consensus about the functionality of the RIPE Database. The RIPE Database Requirements Task Force (DBTF) was formed to tackle this challenge and provides here a list of high-level requirements and recommendations that attempt to resolve ongoing and possible future issues regarding the functionality of the RIPE Database and the data it contains.

The task force did its best to anticipate the community's needs by taking a holistic approach for each requirement and steering away from technicalities. Similarly, the implementation of these requirements and recommendations are not addressed in this document. This could be carried out by another task force or the relevant working groups.

4. Data management principles

Over the years the need for certain information to be available in the RIPE Database has changed. While new attributes have been added and increased the amount of data stored in the database, there has not been a thorough cleanup to remove what is no longer relevant or required. The task force identified four data management principles that should guide how users add and update data in the RIPE Database.

4.1 Data accuracy

The data added to the database should be accurate to ensure uniqueness of Internet number resources and to provide reliable registration information to all parties involved in network operations. For example, contact details or information about a specific assignment should be accurate to facilitate the contact and identification of the organisation holding the assignment.

4.2 Data consistency

The data stored in the database should be consistent across all objects according to RIPE Policies and other requirements. The same data should be added to each object depending on what type of data is needed to avoid data disparity and for data completeness. For example, this principle applies to assignments where users should ideally enter the same level of data.

4.3 Data minimisation

The data stored in the RIPE Database should be adequate, relevant and limited to what is necessary to fulfil the purposes of the RIPE Database. Therefore, the use of personal data should be strictly limited to what's necessary.

For example, this applies to ROLE objects where a generic email address related to a role should be provided (e.g. support@example.com) instead of a personal email address (e.g. johndoe@example.com).

With regards to the contact details inserted for operational purposes, see section 5.

4.4 Data security

To ensure data security, every object in the RIPE Database must be protected. This is currently done using a so-called MNTNER (maintainer) object. It serves as a "lock" to protect another object that you control. A MNTNER object can hold one or more authentication methods that can be used to unlock and modify the objects it protects.

5. Purposes, Requirements and Recommendations

Purposes
The following section lists the purposes, requirements and recommendations established by the task force.

To produce its requirements, the task force looked at four purposes of the RIPE Database:

- Providing authoritative and accurate registration of Internet number resources
- Provisioning of the Reverse Domain Name System (rDNS)
- Publishing routing policies by network operators (RIPE IRR)
- Facilitating Internet operations and coordination

Even though this document is based around the four purposes mentioned above, the task force is aware that there is a fifth purpose that should be taken into consideration:

- Enabling scientific research into network operations and topology

Requirements and other considerations
For each purpose, the task force evaluated a list of requirements and other considerations and listed their current status and rationale.

Recommendations
Based on this information and following the data management principles established earlier, the task force made recommendations for each requirement and other consideration.

5.1 Purpose: Providing authoritative and accurate registration of Internet number resources

The need to maintain an accurate public record of Internet number resource holders is common to all Regional Internet Registries (RIRs). This is outlined in RFC 7020:
“A core requirement of the Internet Numbers Registry System is to maintain a registry of allocations to ensure uniqueness and to provide accurate registration information of those allocations in order to meet a variety of operational requirements.”

The results of the user survey conducted by this task force in January 2020 confirm that having access to trustworthy and accurate information is one of the most valued aspects of the RIPE Database for users.

5.1.1 Requirement: Baseline requirements for registration information of Internet number resources

Current status and rationale:
The current baseline requirements for registration information in the RIPE Database are:

- Full legal name of resource holder
- Full postal address of resource holder
- Contact information for administrative and technical matters (admin-c, tech-c)

Please note that the legal name information is maintained by the RIPE NCC and that the postal address and admin-c/tech-c contact information is maintained by the resource holders.

The task force analysed three policy documents relating to the registration of Internet number resources (IPv4, IPv6, AS Numbers) to understand the level of information required to register resources in the RIPE Database.

The main takeaways from this analysis:

- All three RIPE policy documents require assignments and allocations to be registered in the RIPE Database.
- The reasons stated for the registration requirement are ensuring uniqueness of IP addresses and supporting network operations.
- Only the IPv4 policy provides guidance on the level of registration information required.
- Privacy considerations are included in the IPv4 and IPv6 policies. The IPv4 policy provides more details on how End User resources should be registered in the RIPE Database.
- The AS Numbers policy doesn't mention how detailed AS Number registrations should be.

As of 2020, there is a fourth baseline requirement needed to register information in the RIPE Database: "Country Code of where the resource holder is legally based". This came from Numbered Working Item 10 (NWI-10) and was still being implemented at the time of writing.

Recommendation:
The task force believes that the following baseline requirements are sufficient to ensure uniqueness and to provide accurate registration as defined in RFC 7020:

- Full legal name of resource holder
- Contact information for administrative and technical matters
- Country Code of where the resource holder is legally based

The task force didn't recognise the full postal address of resource holders as a baseline requirement for registration information of Internet number resources. Therefore, the task force recommends that information about the postal address should be made optional and not compulsory. In the long term, the task force recommends taking this information out of the database. If the community accepts this recommendation the relevant supporting documents should be updated accordingly.

5.1.2 Requirement: IPv4 PA assignments

Current status and rationale:
IPv4 policies require all PA assignments ("status: ASSIGNED PA" INETNUM objects) to be registered in the RIPE Database, A core reason for registration of IPv4 PA assignments was to justify an LIR's need for additional IPv4 allocated address space. However, since the RIPE NCC ran out of IPv4 in 2019, this policy has been rendered obsolete.

Also, some resource holders register more information than needed (e.g. create PA assignments for individual IP addresses), while many others don't make any PA assignments at all (see example below).

Figure 2: Over- and under-assigning examples as of May 2021

Over assigning	Under assigning
/32 PA assignments: 530,995 (out of a total of 4,206,427)	PA allocations without any child PA assignments: 16,232
420,518 assignments held by 13 LIRs with more than 10,000 /32 ASSIGNED PA	793 allocations held by 12 LIRs with 50 or more "empty" PA allocations each
Few LIRs registered the bulk of the tiny assignments	9,986 LIRs hold PA allocations containing no PA assignments

Recommendations:
The task force recommends that as resource holders have full responsibility over the registration of their IPv4 PA assignments(s), they are free to make assignments or not. If the community accepts this recommendation, the relevant RIPE Policies should be updated accordingly and documenting IPv4 PA assignment(s) will stop being a policy requirement.

Please note that the task force does NOT recommend that these assignments be deleted but that resource holders can choose to document this information in the RIPE Database.

However, if a resource holder wants to sub-allocate or partition part of their IPv4 resources to another entity, the task force strongly recommends documenting this sub-allocation or assignment in the RIPE Database.

Following the data consistency principle, the task force also recommends resource registration requirements to be applied consistently to all Internet number resources, regardless of their type or status.

To ensure that the information published in the RIPE Database is correctly updated by resource holders, the task force recommends that the RIPE NCC continue to use ARCs (Assisted Registry Checks) to verify this data.

5.1.3 Other consideration: Using the RIPE Database as an IPAM solution

Current status and rationale:
The RIPE Database is also used for other operational functions that are not directly related to its core purposes. From its January 2020 user survey, the task force identified the use of the RIPE Database as an enterprise IPAM (IP Address Management) solution.

Even though this usage is tolerated, using the RIPE Database as an authoritative source for the organisation of number resources assignments is generating information that is not

strictly necessary to fulfil the purposes of the RIPE Database. For this reason, it goes against the data minimisation principle defined earlier.

Recommendation:
The task force recommends limiting and discouraging the use of the RIPE Database as an enterprise IPAM solution.

5.1.4 Requirement: Historical data and personal data filtering

Current status and rationale:
Since 2013, the RIPE Database has stored historical data, as requested by the RIPE community.

This includes:

- Every time an object is updated, the previous version is saved. A standard query will return the most recent version. Old versions are available by using the history query flags.
- If an object is deleted and re-created, a query will return only the most recent version. Deleted objects are not returned in historical query results.
- Objects that are supposed to contain personal data are excluded from historical queries.

Measures are in place to minimise the exposure of personal data, and objects that are meant to contain personal data are filtered out from queries. However, personal data might still be returned in other attributes.

From the community feedback the task force gathered, historical data seems to be used for both operational and research purposes. Operational usage includes fraud investigations and troubleshooting (e.g. explaining outdated configurations). Research usage includes IP measurements and understanding traffic patterns. While it would be possible to list all of the data necessary for operational usage, research usage widely depends on the research itself and can include many data types.

Recommendations:
The task force recognises historical data as a requirement of the RIPE Database, however access to this data should be limited to what's necessary for the most common type of use cases. Regarding research usage, the task force recommends that the RIPE NCC grants access to a wider set of historical data to researchers who need it on a case-by-case basis and according to specific criteria. These criteria should be discussed and defined by the RIPE community.

There is no easy way to track the chain of ownership for address blocks that have been split or merged. The community should consider adding this functionality to historical data.

5.2 Purpose: Provisioning of Reverse Domain Name System (rDNS)

DNS reverse mapping is a DNS-based service that maps IP addresses back to domain names. The reverse DNS tree is structured to follow the address “hierarchy” for both IPv4 (on octet boundaries) and IPv6 (on nibble boundaries). There is no formalised DNS mapping service for AS Numbers.

Since DNS reverse mapping is closely tied to the address space, delegations usually go to the party registered as the holder for that space. Providing DNS reverse mapping management functions (which do not include DNS name service itself) can be seen as a genuine function of both an RIR and an LIR. The RIPE Database is used as a provisioning and documentation tool for reverse DNS for IP addresses under RIPE NCC management. This enables the use of the core address registry for provisioning authorisation purposes (reverse mapping follows INETNUM and INET6NUM).

There are operational procedures, including technical checks, that guide the operation of the reverse DNS by the RIPE NCC. These have been developed and maintained under guidance from the DNS and Database Working Groups. Other (non-DNS specific) general rules apply to the objects used for provisioning reverse DNS to the database.

The task force didn’t find the need for any requirements or recommendations attached to this purpose and therefore recommend maintaining the current status quo.

5.3 Purpose: Publishing routing policies by network operators (RIPE IRR)

5.3.1 Requirement: Routing information

Current status and rationale:
The RIPE Routing Registry is a subset of the RIPE Database which holds information about routing on the Internet. Since the RIPE Database is authoritative for both IP addresses and AS Numbers which have been allocated or assigned by the RIPE NCC, it provides a natural way to publish authoritative information about how Internet number resources are routed on the Internet.

The current RIPE Database authorisation model prevents users from creating or modifying routing policy information for number resources that they do not maintain. Historically lax permissions have left the routing database with a legacy of non-authoritative content.

Recommendations:

- The RIPE Database will provide routing information for:
    - Internet number resources delegated by the RIPE NCC.
    - Internet number resources which fall under the terms of the "RIPE NCC Services to Legacy Internet Resource Holders" policy.
    - Other Internet number resources which already have routing information in the RIPE Database.

- Routing information is maintained by the holders of these resources.
- The holders of these resources will be authenticated by the RIPE NCC or by the holders of parent resources, and only the holders will be authorised to manage routing information for the resources that they hold.
- Routing information for resources delegated to holders that have not been authenticated by the RIPE NCC should be labelled as non-authoritative. This should apply to both non-RIPE NCC resources and legacy resources with no formal relationship with the RIPE NCC.
- The RIPE community should aim to create policies to delete stale and inaccurate non-authoritative routing information.
- It should not be possible to add new routing information to the RIPE Database for address resources delegated by other Regional Internet Registries.

5.3.2 Requirement: Maintaining accurate routing origin information

Current status and rationale:
Routing information in the RIPE Database falls into two broad categories:

- Routing origin information, which documents associations between address blocks and AS Numbers.
- Information about routing relationships between different AS Numbers.

Both types of information are useful, depending on the context.

In the RIPE Database, routing information is stored using certain Routing Policy Specification Language (RPSL). Routing origin information is also stored in the RPKI Database. There is no direct link between these two databases, although both support authorisation from address.

Even if both address and ASN are needed for routing origination, only the address holder's authorisation is currently used.

Recommendations:
Maintaining accurate routing origin (address prefix and autonomous system number) information is a requirement of the RIPE Database:

- Routing origin information is published via ROUTE: / ROUTE6: objects in the RIPE Database.
- ROAs are created in RPKI to represent routing origin information.

5.3.3 Other Consideration: Routing Policy Specification Language (RPSL)

Current status and rationale:
The RIPE Database implements the RPSL base standard (RFC 2622), and RPSLng (RFC 4012), which are based on the key-value format specified in ripe-181. Although tools are readily available to help operators validate routing policy, in particular relating to prefix set management, there are few known deployments of other components of RPSL in production

networks, and no known complete software implementations of either server-side or client-side RPSL.

RPSL is an old standard that never fully matured and has not been maintained over the decades.

Recommendation:
RPSL is *not* a requirement for the RIPE Database. As such, the RIPE Routing working group should look at formally deprecating RPSL, with the cooperation of the RIPE Database Working Group. The specific recommendation is not to adopt a new syntax such as XML or JSON, but rather to consider what routing information is useful to operators and design a way of representing that.

Until RPSL is re-evaluated, the RIPE Database must continue to support it.

5.3.4 Requirement: RPKI Database

Current status and rationale:
The Resource Public Key Infrastructure (RPKI) allows digital certificates to be associated to number resources, thereby providing resource holders with proof of holdership. Currently, the RIPE NCC RPKI database stores Route Origin Authorisation (ROA) information, but it is capable of storing other information.

Each LIR operates its own Certificate Authority (CA) or CA hierarchy, which is signed by the RIPE NCC's CA. The RIPE NCC acts as a root CA for the RPKI, and provides the option to host CA services for each LIR, and also the option to delegate this authority to a CA operated by the LIR.

Currently the portion of RPKI operated by the RIPE NCC is separate from the RIPE Database, and is managed as a RIPE NCC service.

Recommendation:
The task force recommends that RIPE NCC members and the RIPE community consider whether the RPKI Database should be treated as a community resource (like the RIPE Database) with policies and rules set by the community, or continue to be treated as a RIPE NCC service.

5.4 Purpose: Facilitating Internet operations and coordination

Since its inception, the RIPE Database has helped to foster communication between stakeholders, quickly becoming one of the main sources of information to help troubleshoot and develop networks in the RIPE region. The contact information available in the database is historically provided on a best-effort basis by its users.

The RIPE Database should facilitate communication and cooperation among stakeholders for the following reasons:

- Operational issues such as measuring or troubleshooting networks
- Handling abuse cases, supporting the handling of cyber incidents, as well as supporting LEA investigations

5.4.1 Requirement: Operational Contact Information (PERSON and ROLE Objects)

Current status and rationale:
The PERSON object provides information about a real person. The original intention was that this should only be used for contacts responsible for technical or administrative issues relating to Internet resources registered in the RIPE Database.

However, there have been growing concerns around the rising number of PERSON objects in the RIPE Database. As of May 2021 there were 1.92 million PERSON objects in the RIPE Database, 13,000 of which have been locked by the RIPE NCC and 57,000 are unreferenced. Most of these objects are referenced from assignments.

Along with GDPR compliance risks, it is difficult for maintainers to keep high volumes of individual PERSON objects up-to-date, which subsequently impacts the data accuracy goal of the database. This also goes against the data management principle of data minimisation. It is also important to note that maintainers also have to comply with the terms and conditions and legal requirements of the RIPE Database. The requirement for contact information in the RIPE Database could be sufficiently fulfilled by using ROLE objects and generic/group email addresses.

Recommendation
The task force recommends to promote ROLE objects instead of PERSON objects but still make it possible for users to create PERSON objects if/where necessary. However, the task force is aware that users could also add personal data to ROLE objects. This is why stricter checks and measures should be implemented to prevent users from involuntarily entering personal data in both object types. This will also allow users to progressively move away from PERSON objects.

Implementation details should be discussed in the RIPE Database Working Group in collaboration with the RIPE NCC.

5.4.2. Other consideration: Publishing the legal address of resource holders

Current status and rationale
The task force evaluated if the RIPE NCC should publish the legal address of resource holders in the RIPE Database, i.e. all direct recipients of resources from the RIPE NCC. This information is already stored in the RIPE Registry but not available in the RIPE Database. The RIPE NCC does not provide any confidential or private information to LEAs without a court order or other legally enforceable order or request under Dutch law. This includes the legal address of resource holders.

It's also important to note that since the implementation of NWI-10, a new attribute was added to the ORGANISATION object with the Country Code for the country in which the RIPE NCC Member or End User is legally based.

The recommendation of publishing the legal address of resource holders was supported by LEAs and CSIRTs but received low support from the wider network operators community who raised relevant concerns about disclosing this data.

The arguments against this recommendation were:

- Distinguishing natural persons from legal persons is complicated as some operators use their home address as their business address.
- This data is not useful for day-to-day network operations.
- Implementing this recommendation would create a lot of overhead for the RIPE NCC.
- The name and legal country of the RIPE NCC Member or End User will be available in the RIPE Database once NWI-10 has been fully implemented.

The arguments in favour of this recommendation were:

- It supports LEAs investigations in fraud and abuse cases by saving time searching for the legal address in cases where suspected criminal activity has taken place and to issue a subpoena.
- It will spare LEAs from spending time in a long-winded cross-border translation process to get this information.
- Smaller LEAs will profit from this as they might not have the resources to get this information from somewhere else.
- The privacy concerns can be addressed by differentiating legal persons from natural persons.

Recommendation

After weighing the pros and cons and listening to community feedback, the task force decided not to go ahead with this recommendation as there was no clear consensus. However, the task force recognises the LEAs need to access this information in a timely manner to be able to quickly respond to criminal activity on the Internet. Therefore, the task force's recommendation for legal address is that the community explore alternative solutions. The task force recommends that this work is carried out by the relevant working groups.

6. Terminology

Accuracy: In this document, the term "accuracy" refers to "registration accuracy" as defined in RFC 7020: "A core requirement of the Internet Numbers Registry System is to maintain a registry of allocations to ensure uniqueness and to provide accurate registration information of those allocations in order to meet a variety of operational requirements. Uniqueness ensures that IP addresses and AS numbers are not allocated to more than one party at the same time."

Assisted Registry Checks: The Assisted Registry Check (ARC) is the name for the RIPE NCC's "audit" and "additional allocation audit" activities. During the ARC review, the RIPE NCC performs a variety of consistency checks to assess the quality of LIRs' registry data.

Internet Number Resources: IPv4 addresses, IPv6 Addresses and Autonomous System Numbers.

Registration: The documentation of Internet number resources within the RIPE NCC Service Region.

Resource Holder: An organisation or individual that has been allocated Internet number resources in the RIPE NCC service region. This includes RIPE NCC Members and End Users.

RPKI: The Resource Public Key Infrastructure (RPKI) allows Local Internet Registries (LIRs) to request a digital certificate listing the Internet number resources they hold. It offers verifiable proof of holdership of registration of resources by a Regional Internet Registry (RIR).

7. Relevant Policies and Documents

- The RIPE Registry
- The RIPE Database Terms and Conditions
- The Internet Numbers Registry System
- IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region (Section 4.0 and 6.2)
- IPv6 Address Allocation and Assignment Policies for the RIPE NCC Service Region (Section 3.3, 5.3 and 5.5.)
- Autonomous System (AS) Number Assignment Policies (Section 6.0)