RIPE Database Operations Update

Shane Kerr, RIPE NCC
shane@ripe.net

Database Support

The RIPE DBM is the database support role at the RIPE NCC. It answers all e-mail to <ripe-dbm@ripe.net>. This includes requests for service, like recovering from lost passwords, or information, such as explaining the meaning of error replies from the Whois server.

RIPE DBM Changes

From the external point of view, the biggest recent change has been that the RIPE DBM mailbox now auto-replies to every e-mail that it receives. This is to let the sender know that the e-mail has arrived. It also lets the sender keep the same ticket number if they want to provide more related information before the RIPE DBM has had a chance to respond to the ticket.

A related change is that the <auto-dbm@ripe.net> address used to use <ripe-dbm@ripe.net> as the "From:" address when replying to database update e-mails. However, in order to prevent mail loops, this was changed. Details about this can be seen in this announcement:

http://www.ripe.net/ripe/mail-archives/db-wg/2004/msg00281.html

Full-Time RIPE DBM

Since 2003, the RIPE DBM function has been handled by 2 engineers from the Software Engineering Department for 4 days of the week, and 1 day by a RIPE NCC hostmaster. This will continue for the next month or two, after which time we expect to have a full-time RIPE DBM.

There are two reasons we have decided to have a full-time RIPE DBM:

1. Increasing load
   The RIPE DBM role will be receiving about 10% more tickets due to the rDNS migration (discussed in detail in the section on "Reverse DNS Changes", below). The RIPE DBM role will be handling an additional 10% more tickets due to the migration of the RIS support role. *Note: Table 1 does not include these tickets, or internal software support.*

2. More appropriate specialisation
   The staff currently providing RIPE DBM support, while trained software engineers, in most cases do not come from a background that specialises in customer support.

A software engineer will continue to provide support for software, and help build tools to handle user issues more efficiently.

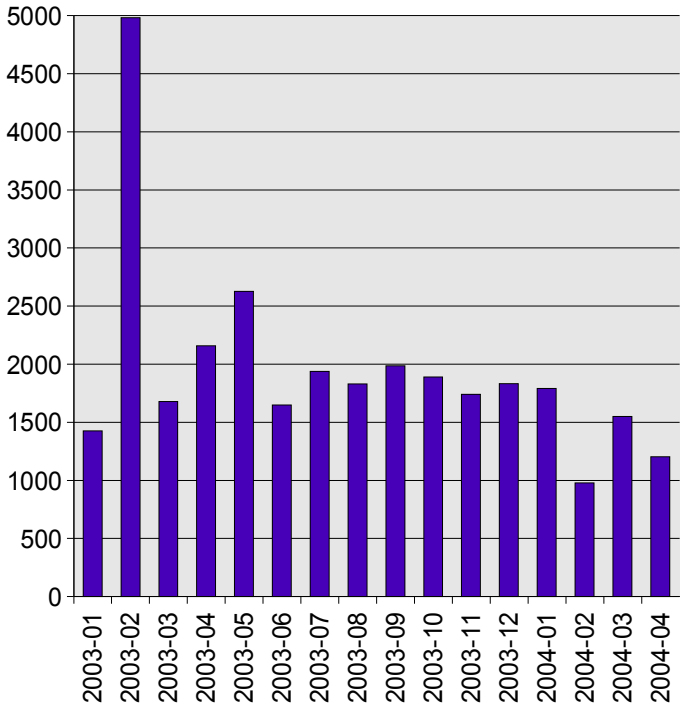

Table 1- RIPE DBM tickets per month

Database Statistics

The contents of the RIPE Database change as the information about the resources that it tracks is updated. For instance, new number resources are allocated, contact details are updated, and so on. You can use the Database Consistency and Statistics page to get full detailed information about the current contents of the database:

http://www.ripe.net/db/dbconstat/index.html

Database Contents

The RIPE Database contains about 49% INETNUM objects, 42% contact data (PERSON and ROLE objects), and 6% DOMAIN objects.

The database has had a steady growth in the number of objects since the RIPE 47 meeting, adding about 150000 new objects, for about a 7% increase.

2.15 million objects

150000 more

27 queries/second

2 queries/second *less*

7.5 updates/minute

3 times as many – due to clean-ups

Text 1- Database summary and change since RIPE 47

The INET6NUM objects had a 43% increase in the number of objects, with over 5000 now. These store IPv6 network information. This growth follows a 50% growth between the RIPE 46 and RIPE 47 meetings. Over 80% of these are assignments, with approximately equal numbers of /48 and /64 networks.

Queries

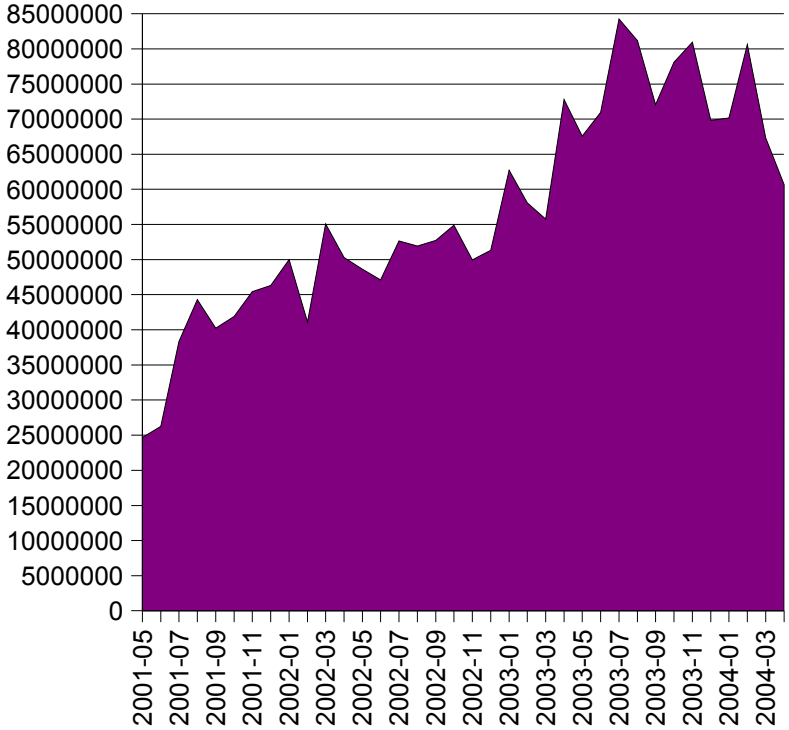

Table 2- Whois Queries Received per Month

One of the main tasks of the server is to answer Whois queries. The rate peaked in 2003 at over 30 queries per second, and is currently almost at this level. About 20% of the queries actually come from the web query interface at:

http://www.ripe.net/perl/whois

These queries come from about 100000 unique IP addresses each day.

The Database can handle a wide variety of query types. However, about 60% of the queries are simple IP lookups, such as "193.0.1.17". The next most common query type is a domain query, most of which get forwarded to another Whois server. This is done because several ccTLD registries have the "refer:" attribute set in their DOMAIN objects, which was intended to help ease the transition when ccTLD registries moved their DNS and related information out of the RIPE Database and into their own databases. All other query types combined account for about 20% of the queries received.

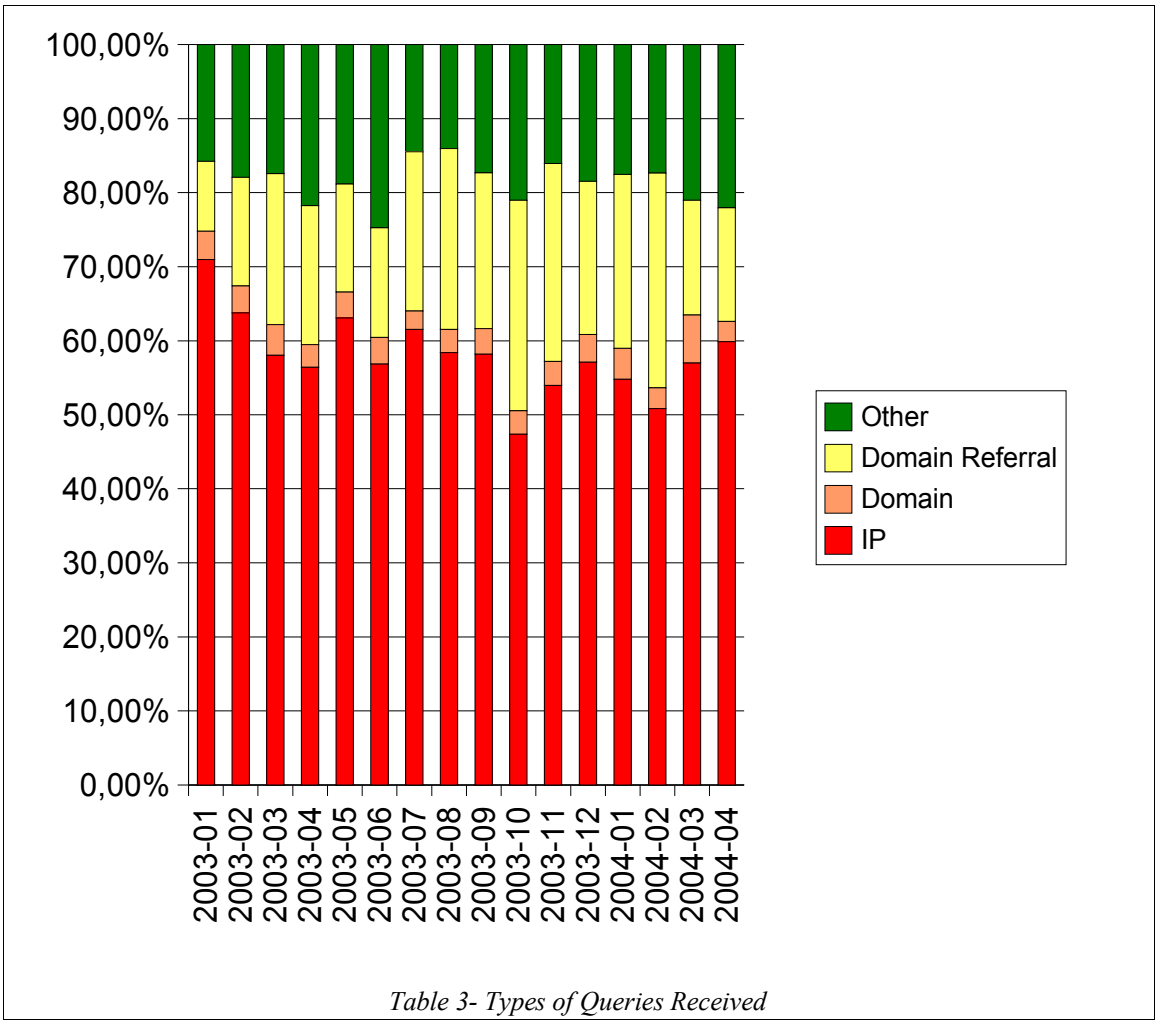

Table 3- Types of Queries Received

X.509 Support Added

As part of the Improved Secure Communication System for RIPE NCC Members, X.509 support was added to the RIPE Database.

More information about X.509 authentication can be found here:

http://www.ripe.net/db/x509.html

Database Schema Changes

The KEY-CERT class was updated to allow users to represent their X.509 certificates, like this:

```
key-cert:   X509-42
method:     X509
owner:      /C=NL/O=RIPE NCC/OU=Members/CN=eu.ripe-ncc.shane/Email=shane@ripe.net
fingerpr:   D5:92:29:08:F8:AB:75:5F:42:F5:A8:5F:A3:8D:08:2E
certif:     -----BEGIN CERTIFICATE-----
certif:     MIID/DCCA2WgAwIBAgICAIQwDQYJKoZIhvcNAQEEBQAwcTELMAkGA1UEBhMCRVUx
certif:     EDAOBgNVBAgTB0hvbGxhbmQxEDAOBgNVBAoTB25jY0RFTU8xHTAbBgNVBAMTFFNv
  .
  .
  .
certif:     hZNmF5c/d0gauqvL+egb+3V+Zg+sJTzHMVKQLF1ybWgJjU75Pi+mO7BG0zsQ13pT
certif:     YxuZCR2W15nwt7zLiHtmfw==
certif:     -----END CERTIFICATE-----
remarks:    Sample Key Certificate
notify:     ripe-dbm@ripe.net
mnt-by:     RIPE-DBM-MNT
changed:    ripe-dbm@ripe.net 20040101
source:     RIPE
```

The “auth:” attribute was updated so that the KEY-CERT objects that contain X.509 certificates can be referenced.

Interface Changes

E-mail support has been extended to allow S/MIME messages, which is the standard for using X.509 authentication in e-mail.

Client-side certificates are recognised by both the webupdates and syncupdates interfaces when SSL is used.

Organisation Object Added

A new type of object, the ORGANISATION object, is now available in the RIPE Database. Full information can be found at:

http://www.ripe.net/db/organisation.html

The RIPE Whois Database stores three main types of contact information: PERSON, ROLE, and ORGANISATION objects. The PERSON and ROLE objects provide a way to find people responsible for operations or usage of the resources represented in the RIPE Whois Database (IP blocks, Autonomous Systems, and domain names). However, these do not provide an easy way of mapping resources to a particular organisation. The ORGANISATION object fulfils this need.

Object Details

A sample organisation object:

```
organisation: ORG-RBI1-RIPE
org-name:     Ruritania Banking Interchange
org-type:     NON-REGISTRY
address:      1 High Street
address:      Polarcity
address:      Northern Nowhere
phone:        +31 20 5354444
e-mail:       bit-bucket@ripe.net
admin-c:      HOHO15-RIPE
tech-c:       HOHO15-RIPE
ref-nfy:      bit-bucket@ripe.net
mnt-ref:      RURITANIA-MNT
mnt-by:       RIPE-NCC-HM-MNT
changed:      ripe-dbm@ripe.net 20040419
source:       RIPE
```

All object types may refer to an object, by adding the "org:" attribute. The maintainer specified in the "mnt-ref:" of the organisation object must authorise the reference (`RURITANIA-MNT` in the above object). This is to prevent people from referencing organisation objects that they have no relationship with.

Queries

The organisations can be looked up by handle, such as `ORG-RBI1-RIPE`, or by name, such as "Ruritania Banking Interchange". Also, you can look up all of the objects that reference a given organisation via an inverse query. For example:

```
whois -r -i org ORG-SANT1-RIPE
```

Will return all objects that have `ORG-SANT1-RIPE` in their "org:" attributes.

By default, any organisation objects referenced by another object will be returned with that object, the same as person and role objects are.

IANA, RIR, and LIR Objects

The IANA and RIRs have organisation objects in the RIPE Whois Database that are maintained by the RIPE NCC. These are used to mark appropriate resources, such as INETNUM objects for /8 allocations from the IANA to the RIPE NCC.

LIRs have organisation objects that are created and maintained for them by the RIPE NCC. The ORGANISATION object for an LIR is created in the RIPE Whois Database when the organisation becomes an LIR.

INETNUM, INET6NUM, and AS-BLOCK objects were updated to include references to the appropriate organisation.

LIRs can update parts of their organisation objects through the LIR Portal:

https://lirportal.ripe.net/

Reverse DNS Changes

Some changes have been made to the way that reverse DNS domains (`in-addr.arpa` and `ip6.arpa`) are handled, including the database support for reverse domain objects. The best place for current reverse DNS information is:

http://www.ripe.net/reverse/

Old Way

The reverse data was kept in the RIPE Database as DOMAIN objects, and also in DNS zone files, maintained separately.

When users wanted to update their reverse DNS information, they would send an e-mail to <auto-inaddr@ripe.net>. A program would then verify the request is valid, and then update both the RIPE Database and the zone files.

There were several problems with this set up:

- LIRs had to use a separate interface for the maintenance of DNS.
- LIRs would update the RIPE Database directly, causing inconsistencies between the view from Whois and DNS.
- The database update software developed more rapidly than the DNS update, making features like web updates and X.509 authentication unavailable to DNS administrators.
- Full automation was impractical, requiring human intervention and the related delays.
- The policy required significant additional work from LIRs.

New Way

The reverse data is kept in the RIPE Database as DOMAIN objects. DNS zone files are periodically built from the contents of the RIPE Database.

When users want to update their reverse DNS information, they send an e-mail to <auto-dbm@ripe.net>. The database update program will verify the DNS information and update the RIPE Database. The information will appear in the DNS after a short delay.

The policy constraints have been reduced. Previously only space assigned to end users could be reverse delegated in DNS. This caused administrative burden to LIRs, as every time space was assigned, reverse DNS had to be set up for that space. Now reverse DNS can be set up for an entire allocation.

There are other advantages, in addition to solving the problems above:

- The introduction of the "mnt-domains:" attribute is allows separate the DNS to be administered by a different set of people to those that maintain INETNUM or INET6NUM objects.
- The rules for domain object names have been made more strict, preventing accidents (for example, `666.193.in-addr.arpa`).
- Deployment of DNS Security Extensions (DNSSEC) requires a method for the exchange of public keys. Using the Whois Database as the authoritative source for zone file creation enables the use of the Whois Database authorisation mechanisms including the LIR Portal, PGP keys and X.509 certificates, for DNSSSEC public key exchanges.

Clean-up

Because the old way allowed inconsistencies, a clean-up was necessary to make the database match the contents of the zone files. This involved creating, modifying, and deleting domain objects, and sending users e-mail about the changes.

NONE Authentication Scheme Deprecated

The RIPE Database will no longer accept updates using the NONE authentication scheme.

NONE was intended to be used consciously, as a notification facility or as a means to tag objects. This authentication scheme was deprecated because it is likely that in many cases NONE is used simply because it is easy.

An announcement with full details was sent to the Database Working Group, but not to a larger audience, because of the potential security concerns. Full details about the process, as well as instructions for users affected by the change, can be found here:

http://www.ripe.net/db/none-deprecation-042004.html

Maintainer Modification

The procedure used by the RIPE NCC to update maintainers was straightforward. Any "auth: NONE" attributes were removed. If that was the only authentication scheme, then a password was generated, and an MD5-PW "auth:" attribute was added. Any such password generated was e-mailed to the contact(s) of the maintainers.

RIPE-NCC-NONE-MNT

A maintainer with NONE authentication, `RIPE-NCC-NONE-MNT`, was added to objects without any maintainer when the database was converted from RIPE-181 format to RPSL format in April 2001. The main use of this maintainer was for INETNUM objects. There were approximately 60000 such objects, making it impractical to create new maintainers for all of them.

Instead, a special maintainer was used to replace these, `RIPE-NCC-LOCKED-MNT`, which has an authentication only available to the RIPE NCC. A URL was sent to the contacts on the INETNUM or other objects referencing this maintainer, which will let them generate a new maintainer or assign another existing maintainer to the object.

Routing Policy

Another use of the `RIPE-NCC-NONE-MNT` has been to allow the creation of objects representing routing policy for resources not allocated or assigned by the RIPE NCC. This is done by using "mnt-routes: RIPE-NCC-NONE-MNT" or "mnt-lower: RIPE-NCC-NONE-MNT" as appropriate. A new maintainer object, `RIPE-NCC-RPSL-MNT`, was created for these cases, with a well-known password, published in the object.

Database-Related Changes to the LIR Portal

The LIR Portal has had several new or changed features that are database-related.

Organisation Object Editor

There are records in the RIPE Database that are protected by RIPE NCC maintainers. For instance, INETNUM objects for ALLOCATED space are maintained by RIPE-NCC-HM-MNT. This is because the RIPE NCC is responsible for these records. However, some information in INETNUM objects for ALLOCATED space, such as the administrative or technical contacts, may be modified by LIRs with the LIR Portal Allocation Editor.

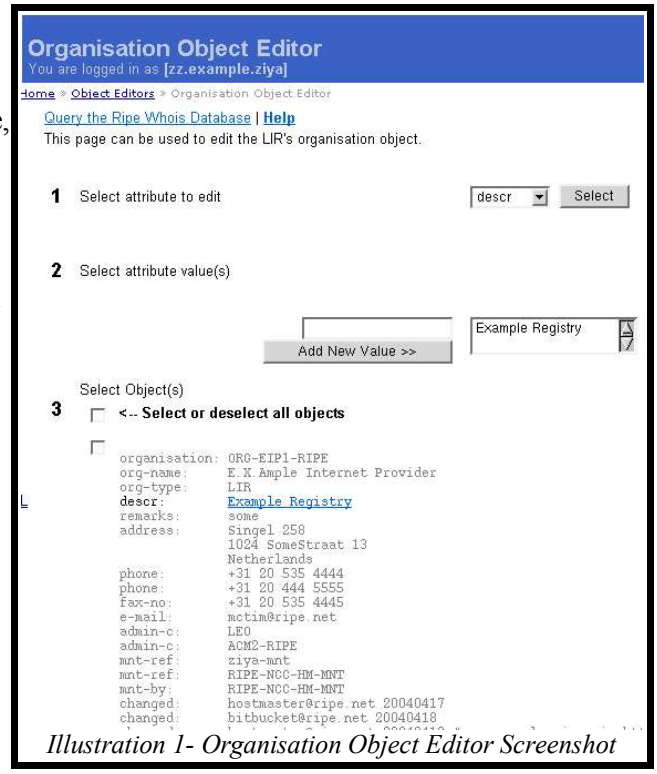

Illustration 1- Organisation Object Editor Screenshot

When the organisation object type was added, the LIR Portal software was updated so that changes to LIR information are automatically reflected in the corresponding organisation object. However, there is some information that is only meaningful in a database context, such as the "mnt-ref:" attribute. It is possible to edit this with the organisation object editor, which works like the allocation editor.

https://lirportal.ripe.net/lirportal/liruser/resource_editor/organisation.html

Maintainer Creation

The creation of new MNTNER (pronounced "maintainer") objects in the RIPE Database is restricted. Even with the deprecation of the NONE authentication scheme, there are still many objects in the Database that are unprotected, for example PERSON objects. This means that any user with a maintainer object can "take over" these objects, simply by adding their maintainer to them. In order to limit this risk, some relationship between the RIPE NCC and the users for maintainer objects was considered necessary.

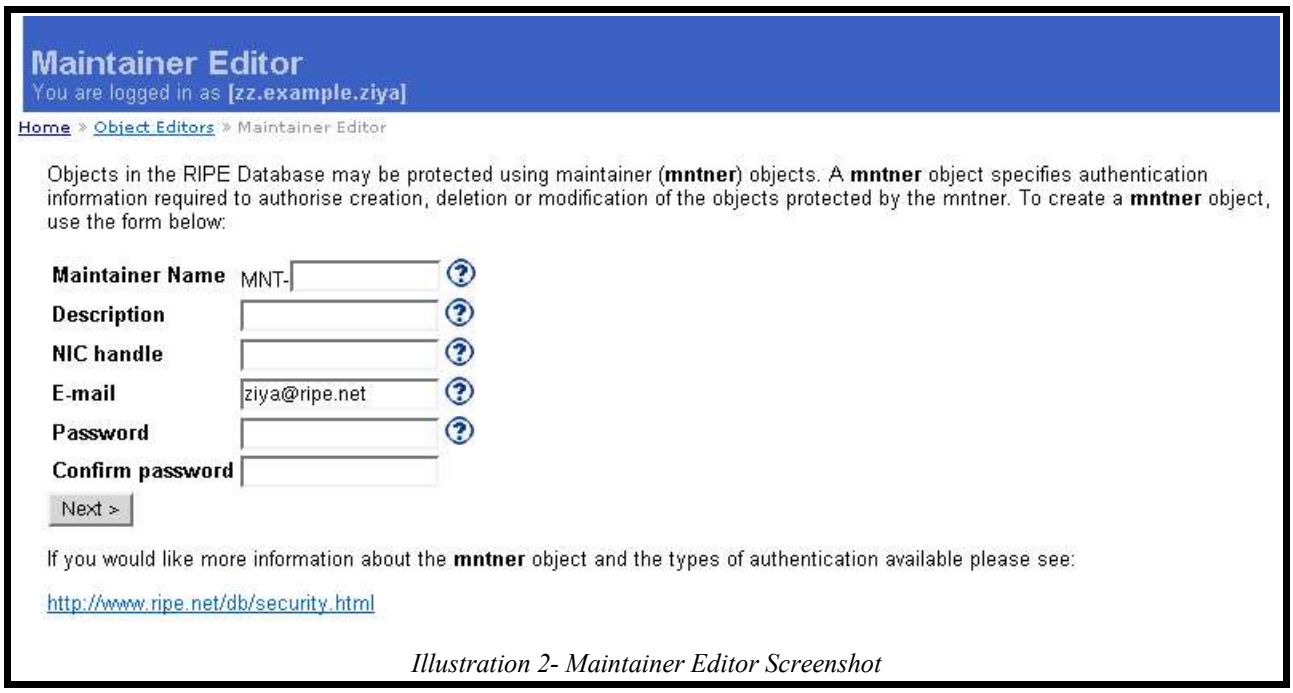

Illustration 2- Maintainer Editor Screenshot

For users of the LIR Portal, the RIPE NCC has a contractual relationship with the users – which is a very strong relationship – so there is no reason to restrict creation of maintainer objects for these users.

https://lirportal.ripe.net/lirportal/liruser/resource_editor/create-mnt.html

Forms Changed

Some changes were necessary to the forms on the LIR Portal to provide additional information for the database, due to the reverse DNS changes and the introduction of the organisation object. Also, some information was removed from the LIR Portal forms to simplify them. Details may be found in the announcement:

http://www.ripe.net/rs/newforms-042004.html

Other Database Changes

There were a number of relatively minor changes to the database that were also made. Details may be found in various announcements:

http://www.ripe.net/ripe/mail-archives/db-wg/2004/msg00282.html
http://www.ripe.net/ripe/mail-archives/db-wg/2003/msg00035.html

CIDR Notation for INETNUM Creation Supported

We now allow the use of CIDR notation when creating INETNUM objects. The CIDR string is replaced by the expanded range notation before the object is processed. This feature is *only* permitted when creating an object. For example, if the following object is submitted to the RIPE Database:

```
inetnum: 1.2.0.0/16
```

Before any further processing is done it will be converted to:

```
inetnum: 1.2.0.0 - 1.2.255.255
```

Prefix Range Lists for “mnt-routes:” Implemented

The "mnt-routes:" attribute syntax has been extended to allow prefix range lists. This will enable people to specify which maintainer has to authorise the creation of the specific routes. For example:

```
mnt-routes: MY-MNT { 20.34.0.0/1617-18, 20.34.0.0/16^- }
mnt-routes: NOT-MY-MNT { ANY }
```

Support for the previous syntax is unchanged. The extended syntax complies with RFC2725 which can be found at:

ftp://ftp.ripe.net/rfc/rfc2725.txt

“cross-nfy:” and “cross-mnt:” Attributes Deprecated

The database previously had a cross notification mechanism. This mechanism used "cross-nfy:" and "cross-mnt:" attributes of ROUTE and AUT-NUM objects to notify particular maintainers and contacts when an overlapping ROUTE object had been added or removed.

In May 2003, this mechanism was deprecated, but there were still approximately 180 AUT-NUM and 620 ROUTE objects with these attributes. These were modified in February 2004 to remove the attributes, and the contacts notified.

Creation of Overlapping INETNUM Objects Prevented

It used to be possible to create INETNUM objects with overlapping begin and end IP addresses. For example:

```
inetnum:  10.0.2.0 - 10.0.3.255
inetnum:  10.0.3.0 - 10.0.4.255
```

The database will no longer allow this sort of behaviour. Nested INETNUM objects may still be created.

Future Plans

"abuse-mailbox:" Attribute

The Database Working Group has proposed the addition of the "abuse-mailbox:" attribute to INETNUM, INET6NUM, PERSON, ROLE, and MAINTAINER objects. It would look like this:

```
abuse-mailbox:    abuse@example.com
```

The full proposal may be found here:

> http://www.ripe.net/ripe/mail-archives/db-wg/2004/msg00286.html

The implementation is straightforward, so if there is consensus the RIPE NCC will proceed with putting this into production.

RPSLng

A revised RPSLng draft was published on 2004-04-28, which addresses all known issues with the proposed IETF RFC:

> http://www.ietf.org/internet-drafts/draft-blunk-rpslng-04.txt

The implementation has been completed for some time, based on one of the previous drafts, so once the draft becomes an RFC the RIPE NCC will put the RPSLng-enabled database server into production.

CRISP

The Cross-Registry Internet Service Protocol (CRISP) is a protocol designed to act as a query-only protocol for finding registry information: domain names, IP addresses, AS numbers, and so on. It is being developed at the IETF, and may function in many ways like an improved WHOIS protocol. You can find information about the IETF Crisp Working group here:

> http://www.ietf.org/html.charters/crisp-charter.html

Drafts defining core protocol, domain type, and a specific transport mechanism have been submitted as RFCs. The RIPE NCC, along with the other RIRs, have agreed to provide final editing to the number requirements and type. It is expected that this work will occur in the next several months.

Route Creation Authorisation from More Specific Addresses

There are cases when people would like to create a route object that aggregates networks, for example:

```
inetnum: 12.0.0.0 - 12.0.0.255
inetnum: 12.0.1.0 - 12.0.1.255
```

To authorise the creation of a single ROUTE object for `12.0.0.0/23`. A detailed proposal was sent to the Database Working Group:

http://www.ripe.net/ripe/mail-archives/db-wg/2004/msg00185.html

However, the rules are relatively complicated, and the number of cases where this is useful are limited.

The RIPE NCC proposes to *not* implement this change.

“proposal: haiku”
Was the e-mail subject,
For the working group.

Is it serious?
Or just a clever joke?
Only Robert knows.

The mail thread is long.
Another proposal comes.
Ideas abound.

The rain falls gently.
Well, at least in Amsterdam.
It always rains here.

http://www.ripe.net/ripe/mail-archives/db-wg/2004/msg00223.html