RIPE

Support for Classless Internet Addresses

in the RIPE Database

Tony Bates
Daniel Karrenberg
Marten Terpstra

Document ID: ripe-121

ABSTRACT

This paper describes the necessary changes to the RIPE database schema and software in order to support classless internet addresses.

1. Introduction

The features described in this document will be usable in the RIPE database at a time specified in [9]. Please refer to this document for more details.

Classless Inter-Domain Routing (also known as supernetting) defines a method of address allocation in the Internet, and a mechanism to reduce routing table size in Internet routers.

The internet (IPv4) address is a 32-bit value split into two parts called network part and host part to provide hierarchical routing[1]. The host part of an address is used for local routing, the network part for global routing. The boundary between the network and host parts was originally defined by run-length encoding in three classes called A (8/24 bit network/host), B (16/16) and C (24/8). This type of internet address will herein after be referred to as classful address.

The proliferation of local area networks made it necessary to introduce more structure into local routing. This was achieved by "subnetting"[2], a technique which divides the host-part of an address into subnet and host parts. This division is visible only to the routers connected to the collection of networks designated by the network part of the address. These routers can use the subnet part for local routing between subnets. Addresses with subnets are still classful as the division between the network part and the local part(s) is still determined by static run-length encoding.

Classless addresses differ from the classful addresses in that the division between the network part and the host part is no longer determined by run-length encoding but by additional information carried in the routing protocols. The additional information is often represented as and referred to as the address mask. This is a 32-bit value where 1-bits represent the address bits in the network part of the address. Classless addressing allows address space to be allocated in almost totally arbitrary and thus more suitable sized pieces.

Furthermore the division need not be the same everywhere in the Internet routing system. As one moves away from the local environment addresses can be aggregated into more global units forming a routing hierarchy:

LAN segment < building/department < enterprise < service provider

However this can only be achieved if addresses are allocated according to that hierarchy, contiguous and on bit boundaries. The aggregate addresses can then be perceived by routers as a single address, and thus reduce the size of the routing tables. For details of these see[3,4]. In CIDR terms the classless internet address if often known as an "IP prefix":

> An IP prefix is a 32-bit value and an indication of the leftmost contiguous significant bits within this address representing the network part.

The RIPE Database

The RIPE Database[5] stores information about address space allocated by the RIPE NCC[6] and routing policies of European Internet service providers[7,8]. The RIPE database schema can currently only represent classful addresses. With the introduction of CIDR, the representation of internet addresses in the RIPE database should be extended to deal with classless addresses. This document details the current representation of internet addresses in the RIPE database, and proposes extensions in terms of representation for classless addressing. This document does not deal with the consequences for the database query mechanisms. It also does not deal with other aspects of the database schema.

2. Current Address Representation

Internet addresses are traditionally represented as "dotted quads". These consist of four decimal numbers in the range of 0-255, each representing 8 bits of the address starting with the first bit of the network part. This nicely fits with the 8-bit granularity of classful addresses.

The current database deals only with classful representation (i.e. standard Class A, B and C style addresses). The addresses can be represented in two ways: **classful net** and **classful range**.

2.1. Classful Net

The *classful net* representation is a dotted quad with zeroes in the parts representing the host part of the network address. These represent the whole address range corresponding to the network address.

```
representation   range of addresses covered

192.1.1.0        192.1.1.0 - 192.1.1.255
```

2.2. Classful Range

The *classful range* representation is two dotted quads separated by "blank dash blank" (" - "). Both dotted quads represent classful nets, i.e. the host part of the respective net contains all zeroes. This represents the address range corresponding to the networks represented by the first and second dotted quads as well as all addresses in the interval between them.

```
   representation      range of addresses covered

192.1.1.0 - 192.1.2.0  192.1.1.0 - 192.1.2.255
```

3. Classless Address Representations

Unfortunately, there are a number of representations for classless addresses in use. While this potentially aids the user in understanding classless addressing, it makes it more difficult to use a single representation. The following choices have to be made w.r.t. the RIPE database:

- Which representations to accept for the different addresses that appear in the schema.

- Which representation to use when presenting the information in response to a query.

We will present the common representations in turn followed by a discussion on how they are used in the RIPE DB for submission and presentation of classless addresses.

3.1. Prefix Length

This representation is a dotted quad followed by a slash and the decimal length in bits of the prefix. This is used in the CIDR[4] and BGP-4[9] documents as well as in popular router software.

```
representation   range of addresses covered

192.1.1.0/24     192.1.1.0 - 192.1.1.255
192.1.128.0/17   192.1.128.0 - 192.1.255.255
```

If we look at these examples in terms of the mask depicted by the length we seeing the following:

```
                   10         20         30
           +----------|----------|----------|--+
           |12345678 12345678 12345678 12345678|
           |===================================|
      24 = |11111111 111111111 1111111 00000000|
      17 = |11111111 111111111 1000000 00000000|
           +-----------------------------------+
```

Table 1: mask length as network bits in 32 bit IP address

Table 1 shows which bits form the network part (represented by "1"s) for the two example lengths of ‘network prefix’ (24 and 17) given.

3.2. Network and Mask

This representation is based on the subnet mask representation. It is a dotted quad representing the address followed by whitespace and a dotted quad representing the mask covering the prefix bits.

```
     representation         range of addresses covered

192.1.1.0 255.255.255.0     192.1.1.0 - 192.1.1.255
192.1.128.0 255.255.128.0   192.1.128.0 - 192.1.255.255
```

3.3. Classless Range

This representation is a dotted quad followed by "blank hyphen greater-than blank" and another dotted quad. This representation is specific to the RIPE database. This represents the range defined by the classless addresses represented by the first and second dotted quad well as all addresses between them. The difference to the *classful range* representation is that the host part of the high end of the range is not assumed to be all ones.

It should be noted that this representation is the only representation in which one can specify a range that is not necessarily bit-aligned. Although this should be avoided in light of CIDR, it does make this representation the most flexible of the three.

representation	range of addresses covered
192.1.1.0 -> 192.1.1.255	192.1.1.0 - 192.1.1.255
192.1.128.0 -> 192.1.255.255	192.1.128.0 - 192.1.255.255

The separator between the begin address and end address has changed to clearly indicate the different semantics.

representation	range of addresses covered	hosts
192.1.1.0 - 192.1.2.0	192.1.1.0 - 192.1.2.255	512
192.1.1.0 -> 192.1.2.0	192.1.1.0 - 192.1.2.0	257

4. New representation of IP addresses in the RIPE database

The choice of representation of IP addresses in the RIPE database is depending on the context. The first part concerns the object in the database that deals with address assignment information, the *inetnum* object, the second concerns the newly proposed object that contains routing information, the *route* object [8].

4.1. Representation in the "inetnum" object

With the newly proposed *route* object, the *inetnum* object will only contain assignment information. Because of the large number of entries currently in the database using the classful net and classful range representation, these two representation will remain valid. New to the inetnum will be the classless range representation.

Representations accepted

To help registries when sending in assignment information, the network and mask notation and the prefix/length notation will be accepted by the software in addition to the three notations mentioned above, *but* they will be rewritten by the software to the classless range notation. The database output will **only** contain the classful net and range notation and the classless range notation. The table below shows some rewrite examples.

representation	rewritten to
192.87.45.0/24	192.87.45.0 -> 192.87.45.255
192.1.128.0 255.255.128.0	192.1.128.0 -> 192.1.255.255

Please note that **non-contiguous** subnet masks will **not** be allowed.

It should also be noted that the classless range representation is the only representation that supports subnets of the existing classful style networks. While these subnets can be expressed in prefix/length and network and mask representation, they will always be rewritten to a classless range.

4.2. Representation in the "route" object

The *route* object will contain routing information for IP address space. Because many vendors have implemented the prefix/length notation, the route object shall only contain this representation.

Representations accepted

All other representations, classful and classless, will be accepted by the database software, *but* will be rewritten to the prefix/length representation *if possible*.

In cases where there is no one to one mapping between the representation sent in and the prefix/length representation, an error will be generated, and the object will be refused. These cases happen when ranges cannot be represented by a common mask, i.e. do not start and end on a common bit boundary.

This means that objects that would require to be split into multiple prefix/length representations will not be accepted. The table below shows some rewrite examples.

representation	rewritten to
192.87.45.0	192.87.45.0/24
192.1.128.0 - 192.1.255.0	192.1.128.0/17
192.1.128.0 -> 192.1.255.255	192.1.128.0/17
192.1.128.0 255.255.128.0	192.1.128.0/17
192.87.45.0 - 192.87.46.0	ERROR
192.1.128.0 -> 192.1.130.255	ERROR

Please note that also here **non-contiguous** subnet masks will **not** be allowed.

5. References

[1] J. Postel, "Internet Protocol", RFC 791, January 1981.

[2] J. Mogul, "Internet subnets", RFC 917, January 1984.

[3] Fuller, V., Li, T., Yu, J., Varadhan, K., "Supernetting: an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

[4] Y. Rekhter, T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.

[5] R. Blokzijl, "RIPE Databases", ripe-013, August 1990.

[6] D. Karrenberg, M. Terpstra, "IP Address Space Assignment Procedures", ripe-104, December 1993.

[7] T. Bates, J-M. Jouanigot, D. Karrenberg, P. Lothberg, M. Terpstra, "Representation of IP Routing Policies in the RIPE Database", ripe-081, February 1993.

[8] T. Bates, Gerich, E., Joncheray, L., Joanigot, J-M, Karrenberg, D., Terpstra, M, Yu, J., "Representation of IP Routing Policies in the RIPE Database", ripe-181, October 1994.

[9] T. Bates,D. Karrenberg, M. Terpstra, "RIPE Database Transition Plan", ripe-123, October 1994.