

European Internet Registry

Policies and Procedures

*S. Dolderer
D. Karrenberg
M. Kuehne
M. Norris
C. Orange
W. Woeber
J. Zsako*

*Document ID: ripe-136
Obsoletes: ripe-104, ripe-105*

ABSTRACT

The distribution of IP address space follows the hierarchical scheme described in [Gerich93a]. For Europe and parts of the surrounding area address space is allocated by IANA to the RIPE NCC which acts as a regional Internet registry. Address space is allocated by the RIPE NCC to local Internet Registries (IR), who assign it to end users. In this document, we describe the policies and procedures associated with address space management that must be followed by local IRs. Moreover, we present a number of services available to local IRs to simplify the tasks associated with address space management.

1. Scope

This document describes the European Internet registry system for the distribution of globally unique Internet address space and its operation. Particularly it describes the rules and guidelines governing the distribution of this address space. The rules set forth in this document are binding for all address space allocated and assigned via the RIPE NCC.

This document does not describe private Internet address space and multicast address space. This document does not describe local additions to the European guidelines. While providing an overview about the global Internet registry system this document does not describe allocation and assignment rules used by other regional registries.

1.1. Overview

The main body of this document comprises eight sections, with content as follows.

Section 2 (Internet Address Space and the Internet Registry System) defines different types of IP address space and their purposes. It explains the goals used in assigning such addresses and outlines the hierarchical nature of the Internet Registry system used to achieve these goals. The important distinction between Provider Aggregatable and Provider Independent address space is also covered.

Section 3 (Address Space Assignment Procedures) describes the procedures to be followed by European IP registries when assigning IP addresses to users. The importance of documentation is stressed, while the various elements of information required are explained in detail. Next, the criteria and standards of evaluation are dealt with. Finally, the actual assignment of address space, of various kinds, is described, as are the accompanying steps which a registry must take.

Section 4 (Rules and Guidelines for Allocations) explains how the RIPE NCC allocates IP address space to registries in an efficient and equitable manner and how the status and nature of such allocations are made publicly available in the RIPE database.

Section 5 (DNS and Reverse Address Mapping) documents the role of the RIPE NCC in providing reverse delegation, and explains how registries can manage subsidiary reverse delegation of assigned address space.

We conclude with a glossary in which the key terms used in this document are defined.

2. Internet Address Space and the Internet Registry System

2.1. Types of IP Addresses

IP addresses for the purposes of this document are 32-bit binary numbers used as addresses in the IPv4 protocols. There are three main types of IP addresses

Public Addresses

The public IP addresses make up the Internet address space. They are assigned to be globally unique according to the goals described in Section 2.2. The main purpose of this address space is to allow communication using IPv4 over the Internet. A secondary purpose is to allow communication using IPv4 over interconnected private internets. One can currently distinguish two kinds of public addresses: provider independent (PI) and provider aggregatable (PA) addresses; see Section 2.4 for more details.

Private Addresses

Some address ranges have been set aside for the operation of private networks using IP. Anyone can use these addresses in their private networks without any registration or coordination. Hosts using these addresses can not be reached from the Internet. For a thorough description of private address space, please refer to [96a].

Special and Reserved Addresses

There are a number of address ranges reserved for applications like multicasting. These are described elsewhere (cf [Deering89a]) and are beyond the scope of this document.

2.2. Goals of Public Address Space Distribution

In the remainder of this document, we are primarily concerned with the management of public Internet address space, as defined in the previous section. Every assignment of Internet addresses must *guarantee* that the following restriction is met.

Uniqueness

Each public Internet address worldwide must be unique.

This is an absolute requirement which guarantees that every host on the Internet can be uniquely identified.

In addition to the uniqueness requirement, public Internet address space assignments should be made with the following three goals in mind.

Aggregation

The distribution of public Internet addresses in a hierarchical manner, permitting the aggregation of routing information. This is necessary to ensure proper operation of Internet routing. This goal could also be called *Routability*.

Conservation

The fair distribution of public Internet address space according to the operational needs of the end users operating networks using this address space. In order to maximize the lifetime of the public Internet address space resource, addresses must be distributed according to need, and stockpiling must be prevented.

Registration

The provision of a public registry documenting address space allocation and assignment. This is necessary to ensure uniqueness and to provide information for Internet trouble shooting at all levels.

It is in the interest of the Internet community as a whole that these goals are pursued. It is worth noting that "Conservation" and "Aggregation" are often conflicting goals, and therefore that each assignment must be evaluated carefully. Moreover, the above goals may occasionally be in conflict with the interests of individual end users or Internet service providers. Careful analysis and judgement are necessary in each individual case to find an appropriate compromise. The rules and guidelines in this document are intended to help Internet registries and end users in their search for good compromises.

2.3. The Internet Registry System

The Internet Registry system has been established to achieve the goals stated in Section 2.2. It consists of hierarchically organized *Internet Registries* (IRs). Address space is typically assigned to end users by *local* IRs. The address space assigned is taken from that allocated to the local IR by the regional IR. *End users* are those organizations operating networks in which the address space is used. The address space may, however, be requested by a consultant (*requester*) acting on behalf of the end user. Local IRs are typically operated by *Internet Service Providers* (ISPs). Local IRs hold *allocations* of address space for *assignment* to end users. Assigned address space is actually used to operate networks, whereas allocated address space is held by local IRs for future assignments to end users. To achieve both the conservation and aggregation goals, only IRs can hold allocations of address space.

IANA

The Internet Assigned Numbers Authority has authority over all number spaces used in the Internet. This includes IP address space. IANA allocates public Internet address space to regional IRs according to their established needs.

Regional IRs

Regional IRs operate in large geopolitical regions such as continents. To date, three Regional IRs have been established, namely the InterNIC serving North America, the AP-NIC serving the Asian Pacific region, and the RIPE NCC serving Europe and surrounding areas. Since these do not cover all geographical areas, regional IRs also serve areas around their core service areas. The number of regional IRs is expected to remain small.

Regional IRs are established under the Authority of IANA. This requires consensus within the Internet community of the region. In particular, the ISPs in the region under consideration should be involved in the process. The duties of a regional IR include the coordination and representation of the local IRs in its region.

Local IRs

Local IRs are established under the authority of a regional IR. Local IRs are typically operated by ISPs and serve the customers of those ISPs as well as the customers of smaller ISPs who are connected to the rest of the Internet through the larger ISP. Other organizations such as large international Enterprises can also operate local IRs.

Much of this document is concerned with the responsibility of the local IR in the assignment process. In some cases, the local IR assigning the address space is not run by the ISP that will provide connectivity. It is important to note that maintenance of the administrative information regarding the assigned address space is the responsibility of the IR that makes the assignment, and not of the ISP providing the connectivity. Furthermore, only IRs can hold address allocations.

End-Users

Strictly speaking end users are not part of the IR system. They do, however, play an important role with respect to the goals defined above. In order to achieve the conservation goal, for example, end users should plan their networks to use a minimum amount of address space. They must document their addressing and deployment plans to the IR and furnish any additional information required by the IR for making assignment decisions. To achieve the aggregation goal, an end user should choose an appropriate local IR. End users should be aware that changing ISPs may require replacing addresses in their networks. Finally end users must provide and update registration data for the address space assigned to them.

Requesters

In addition to these key players in the Internet Registry System, there are often consultants who setup and manage networks for end users. The consultants may be the people actually submitting a request for address space to an IR on behalf of an end user. We refer to the person making the request for an end user as a requester, whether that person is employed by the organization, or is simply acting on behalf of the organization with respect to the address space request.

For Europe, the Internet Registry System hierarchy consists of the following entities (from the top down): IANA, the RIPE NCC, and local IRs.

2.4. Provider Independent vs Provider Aggregatable Addresses

Provider Aggregatable Address Space

Local IRs operated by Internet service providers are allocated Provider Aggregatable (PA) address space which they assign to their end users. This is done in such a way that routing information for many end users of an ISP can be aggregated on the borders of the provider's routing domain. This keeps the number of routes and state changes in the interdomain routing

system (between providers) at an acceptable level. The cost of propagating a relatively small number of aggregated routes is much lower than that of propagating each end user's individual routes throughout the entire interdomain routing system.

If an end user changes service providers, their PA address space will have to be replaced. As a consequence, all hosts and routers at the end user's organization will have to be reconfigured. The end user will need to obtain a new address space assignment, and return the previously assigned address space. To ensure the address space is properly returned, a clear, preferably contractual, understanding is needed between the local IR and the end user. The agreement should state that the assignment of the address space becomes invalid when the provider no longer provides Internet connectivity to the end user or shortly thereafter.

The goal of this arrangement is to minimize the load on the interdomain routing system. If the end user continued to use PA address space obtained from their previous service provider when connecting to another service provider, their routing information could not be aggregated and would have to be propagated separately throughout the whole interdomain routing system.

Provider Independent Address Space

In contrast to PA address space, PI address space can remain assigned to its user as long as the criteria for the original assignment are met. The duration of the assignment is independent of the use of a particular provider's services. The apparent advantage of PI address space is that a user's hosts and routers need not be reconfigured if the user decides to change service providers. However, PI addresses are expensive to route because no use can be made of aggregation. All early Internet address space assignments were provider independent. Many assignments made by local IRs are also formally provider independent due to a lack of prior agreements between ISP and the end user that the assignment will be terminated when the service is.

Validity of assignment

Assignments of any kind of address space are valid as long as the original criteria on which the assignment was based are still valid. If an assignment is made for a specific purpose and the purpose no longer exists, then the assignment is no longer valid. If an assignment is based on information that turns out to be invalid so is the assignment.

3. Address Space Assignment Procedures

3.1. Introduction

In this section, we describe the procedures to be followed by local IRs when assigning address space to their users. We start with a description of the information to be gathered from the user. The purpose of the information gathering is twofold. First, the information is required to make address assignment decisions, with respect to the aggregation and conservation goals. Second, the information is required for registration purposes.

We go on to describe how this information should be evaluated to make appropriate assignments, and introduce additional considerations that may be essential in the assignment decision. Finally we specify the procedures to be followed in the assignment process.

Before going into the factors in the assignment process, we start with some general background information and policies that determine the information to be gathered, and the procedures to be followed.

Address space is assigned by IRs to end users who use it to operate the specific networks described in an address space request. IRs guarantee that no other end user will be assigned the same address space during the validity of the assignment. An assignment is valid as long as the criteria on which it is based remain valid.

In accordance with the conservation goal, end users are not permitted to reserve address space. Evaluation of IP address space requests must be based on the documentation provided for the following 24 months, as specified in the addressing plan and the current address space usage described in the next section. The amount of address space assigned must be justified by this documentation. This means that address space assigned in the past should be used to meet the current request if possible. Once an organisation has used its assigned address space, it can request additional address space based on an updated estimate of growth in its network.

3.2. Documentation

To make appropriate assignment decisions, information must be gathered about the organisation, addressing requirements, network infrastructure, current address space usage, and future plans of the end user requesting address space. Some information is essential to the assignment process, and is

formally required by the IR's. Other information is very helpful in making assignment decisions, but is not strictly required. The Local IR must assure that the required information is complete before proceeding with the request.

For gathering the required information, the RIPE NCC provides a set of forms and a set of instructions to fill them in. Although use of the forms provided (or a local-language equivalent) is strongly encouraged, each local IR can obtain and manage this information as it considers appropriate. Requests requiring evaluation by the NCC must, however, be submitted on a current version of the "European IP Address Space Request Form" (e.g. ripe-137: [Caslav96a]).

The information gathered in the assignment process must be maintained permanently by the IR making the assignment, and must be made available to the regional registry immediately upon request. The IR is responsible for protecting the end user's privacy. Aside from the data specified in Section 3.2.1.5, which is published in the registry database, the information gathered must be kept in strict confidence. The IR is not authorised to provide the information to anyone not representing IANA or a regional registry, unless explicitly requested to do so by the end-user.

In the subsections that follow, we outline the specific data to be gathered and the reasons for doing so.

3.2.1. Required Information

The following set of information must be provided with every request for an address assignment. The data is essential both to properly assigning addresses and to maintaining a global overview of assignments. With the exception of the information specified in Section 3.2.1.2, all information refers to the currently requested address space.

3.2.1.1. Overview of Organisation

To properly assess the user's address space requirements, it is essential to understand the structure of the organisation to which the addresses are being assigned, and which part of the organisation will make use of the addresses.

Consider the following situation. A bicycle manufacturer based in Belgium has a variety of departments. Some, such as the *Front Fork* and *Derailer* departments, specialise in specific bike parts. Others, such as the *Sales* and *Development* departments are more general by nature. In such a company, the departments Sales, Development, and Manufacturing may fall directly under

the top management, whereas the subdepartments *Derailer*, *Chain*, *Pedal*, and *Front Fork* fall under the Manufacturing department. If someone submits a request for address space, we must know which part of the organisation will make use of the assigned addresses. Suppose, for example, the *Manufacturing* department is assigned address space for use by all bike parts subdepartments. If shortly thereafter, the *Chain* department requests address space it is important that we know an assignment has already been made to the organisation to meet the *Chain* department's needs. A similar situation may occur if the *Sales* department has groups of representatives in several countries. It is essential to know if addresses being requested by the central office will be used in Antwerp or in Madrid. We want to prevent assignments being made for the same subnet by two different parts of the organisation. In the case of a distributed sales department, this must also be known to assure a proper assignment with respect to aggregation.

The person responsible for making the assignment can only be aware of this situation if an overview of the organisation, and the requester's role therein is known. It is therefore important that a brief overview of the organisational structure be provided. This should include details of the parent company, subsidiaries and contact persons.

In the case of our bicycle manufacturer, one would expect someone representing the *Chain* department to produce general information about the structure of the organisation in Belgium, and contact persons for the *Manufacturing*, *Sales*, and *Development* departments. We would not expect the same person to present information on the structure within the *Sales* department, such as who manages the office in Rome.

Clearly, the assignment process is greatly simplified if an organisation coordinates its address space management, and if all requests are made by a single body representing the entire organisation.

Contact Persons

To facilitate handling the request, contact information is required for the person making the request and for someone at the organisation to which the address space will be assigned. The information should be entered on the *Requester* and *User* contact templates, respectively. These templates contain *name*, *organisation*, *country*, *phone*, *fax-no*, and *e-mail* fields. In each template, the appropriate person's name should be specified in full. The organisation refers to that in which this person works, and the country refers to that in which the person's office is located. The telephone and fax numbers should include the country prefixes in the form *+code*, and if the person can be reached by e-mail from the Internet, the address should be specified.

The contact person information is only collected to facilitate the address space request. It may or may not include data for persons that will later be entered in the RIPE database.

3.2.1.2. Current Assignment Space Usage

To meet the conservation goal in address space assignments, one must have information regarding address space assignments made to the user in the past before new address space can be assigned. A detailed description of how the address space is currently being used is required. Using this information, we can prevent assigning new address space, where already assigned addresses can be employed to meet the user's needs.

Each set of addresses already assigned to the organisation must therefore be reported. The current use of these addresses must be documented in a table similar to that below. An entry must be included for each physically separate subnet in the user's network. Subnets are considered to be physically separate if there is an IP router between them. Each row in the table below contains an entry for a subnet in the end user's organisation.

Prefix	Subnet Mask	Size	Addresses Used			Description
			Current	1yr	2yr	
193.1.193.0	255.255.255.192	64	28	34	50	Derailer
193.1.193.64	255.255.255.224	32	10	12	15	Chain
193.1.193.96	255.255.255.224	32	8	13	17	Front Fork
193.1.193.128		128				Unused
193.1.194.0	255.255.255.0	256	132	170	210	Frame
193.1.195.0	255.255.254.0	512	317	350	380	Assembly
		1024	495	579	672	Totals

Each entry in the table above is made up of the following fields which specify the current and projected use of the address space in the subnet. The *Description* field is used to specify a short but semantically meaningful description of the role of the subnet in the user's organisation. In our example, we have descriptions corresponding to various bike parts. Together with the size information, this provides significant insight as to the network structure in the organisation.

The number of network interfaces currently used in the subnet, along with the number expected to be needed in the coming one and two years must also be specified. These numbers are to be entered in the *Current*, *1yr*, and *2yr*

fields of the subnet entry, and include the number of network interfaces to be used, such as those for hosts, routers, gateways, terminal concentrators, printers, and any other machines requiring one or more network interfaces.

The *Size* field is used to specify the size of the subnet, which determines the maximum number of network interfaces that can be incorporated in the subnet. It must be a power of two, and of course should be greater than or equal to the number specified in the *2yr* field. If it is smaller, this may be the motivation for the address request, or it may be a mistake in the requester's planning.

The *Subnet Mask* field is used to specify just that, and finally, in the *Prefix* field, the position in the assigned address space at which the addresses for this subnet start is specified.

As in the example, entries should be made in the table for assigned address space which is currently not used.

3.2.1.3. Request Overview

The network overview is used to obtain a quick idea about the scale of the request. This information allows the IR processing the request to gain immediate insight as to the nature of the assignment request. The exact information to be gathered is:

Size of Request To give the IR an immediate indication of the scale of the request, the total number of Internet addresses being requested must be specified under *request-size* on the network overview form. If the *request-size* is 512, the user specifies a need for that number of Internet addresses. Prior to the use of CIDR, the user would have asked for two Class C networks. Because classless addressing is now used, the size of the request may be less than 256 or fall between the class borders (e.g. 32, 288, 384).

Addresses to be Used To obtain an overview of the structure of the requester's network, one must know how many Internet addresses will actually be used at different points in time. This corresponds to the number of interfaces to the network, which often will be slightly higher than the number of hosts.

In the network overview form, the fields *addresses-immediate*, *addresses-year-1*, and *addresses-year-2* are used to specify how many of the requested network addresses will be used immediately following the assignment, within 12 months, and within 24 months, respectively.

Number of Subnets In practice, the end user will want to employ the requested addresses in one or more subnets in an organisation. The number of physically separate subnets in which the requested addresses will be used is an important factor in making correct assignments. Together with the number of addresses to be used, this provides a global picture of the requester's envisioned network infrastructure. In the network overview form, the fields *subnets-immediate*, *subnets-year-1*, and *subnets-year-2* are used to specify the number of subnets in the requester's network plan to be implemented immediately, within 12 months and within 24 months, respectively.

Internet Connection Prior to assigning address space, it is essential to know if the end user requesting IP addresses is already connected to the Internet. If so, then the selection of appropriate address space for this user may depend on which provider(s) currently supplies connectivity. If the user is not connected, but is planning to be, this should also be taken into account. This information is essential if the conservation and aggregation goals of the public address space distribution are to be met. The current and planned connectivity of the user is specified in the *inet-connect* field of the network overview form.

Country Finally, the country or countries in which the addresses will be used must be specified using the ISO 3166 two letter code. The *country-net* field of the network overview form is reserved for this purpose. If the ISO 3166 code is not known, the full name of the country should be specified.

Private Address Space Using private addresses helps to meet the conservation goal. For this reason, users should always be informed that private addresses might be a viable option. In particular, private address space can be employed if not all hosts require network layer access to the Internet. Although users are not required to use private address space even if it would satisfy their needs, it is important that they have considered the possibility. The *private-considered* field in the network overview form should be checked after the requester has indicated whether it is applicable for the user's network.

Request Refused If a user's organisation has had an assignment request refused in the past, then it is useful to know when and by which IR. Whatever the case, it is useful to know if a request has been refused, and why. This information should be specified in the *request-refused* field in the network overview form.

PI Requested If provider independent address space is requested by the user, special steps will have to be taken by the local IR processing the request. The *PI-requested* field in the network overview form should be checked if

this is a request for PI address space.

3.2.1.4. Addressing Plan

To assess the suitability of assigning the requested address space, an *addressing plan* is required. This provides detailed information on the projected use of the requested address space. Like the *current address space usage*, the addressing plan is a table in which every subnet is specified.

With few exceptions, the entries in the following table are the same as those in the table of current address space usage.

Relative Prefix#	Subnet Mask	Size	Addresses Used			Description
			Immediate	1yr	2yr	
0.0.0.0	255.255.255.192	64	8	16	30	Systems Group
0.0.0.64	255.255.255.224	32	17	22	26	Engineering
0.0.0.96	255.255.255.224	32	12	17	20	Manufacturing
0.0.0.128	255.255.255.224	32	10	15	20	Sales
0.0.0.160	255.255.255.240	16	5	9	12	Management
0.0.0.176	255.255.255.240	16	7	8	9	Finance
		192	59	87	117	Totals

The number of network interfaces immediately required in the subnet, along with the projected need for the coming 12 and 24 months must be specified. These numbers are to be entered in the *Immediate*, *1yr*, and *2yr* fields of the subnet entry.

In the *Relative Prefix* field, we specify the relative position in the assigned address space at which the addresses for this subnet will start. The relative position of the first subnet is always 0.0.0.0. For each subsequent subnet, the start position is selected to allow for the total number of hosts in the *Size* fields of the subnets which precede it.

To conserve address space, the start positions of the subnets should be selected to minimise padding in the address space. In the example above, we arrange the rows in decreasing order of the *Size* field. This scheme can be applied in general to prevent wasting address space between subnets. The size of every valid request for address space will be the sum of sizes of the subnets specified in the addressing plan.

Current evaluation criteria assume that addressing is classless. This means that all possible prefixes of any length can be used. If there are technical restrictions preventing the use of certain address ranges or the choice of optimal subnet sizes, these restrictions need to be explicitly documented. Documentation needs to include the precise nature of the restriction, the make, model and version of the hardware or software causing the restriction, and its precise location in the network.

3.2.1.5. Database Information

For registration purposes, information is required about the organisation needing address space. Information is also required about the persons involved in the request and administration of the addresses. Some of the information may be redundant because the same person can play multiple roles. However, every role can be filled by someone different, so all information must be supplied in full. The data specified below is to be gathered by the local IR handling the assignment, and will be stored in the registry database, at which point it becomes publicly accessible.

Organisation Some information about the organisation that will be using the addresses must be supplied for maintenance of the RIPE database. The *Network Template* is supplied for this purpose.

To help identify this assignment in the RIPE database, a short, but semantically meaningful name must be entered in the *netname* field. A short description of the organisation that will use the assigned addresses is needed. The information is specified using one or more *descr* fields in the *Network Template*. If, for example, the assigned addresses will be used by the Department of Neural Surgery at Catatonic State University, then the department and university names may be specified in two *descr* fields. The ISO 3166 country code should be specified in the *country* field. The full country name can be used if the code is not known.

The *admin-c* and *tech-c* fields are used to specify the IR handle (*NIC handle*) for the administrative and technical contact persons, respectively. The administrative person specified in the *admin-c* field must be physically located at the site of the network. The technical person specified in the *tech-c* field may be a network support person located on site, but could also be a consultant that maintains the network for the organisation. In both cases, more than one person can be specified. The use of NIC handles to specify each contact person is required, as it assures each person has a unique entry in the database. If the person doesn't have an entry in the database, a unique NIC handle can be acquired upon request.

Personal Data For every person involved in an assignment request, we need a full set of personal data. This data can only be omitted if up to date information for the given person is already stored in the RIPE database. If new data is provided for a person with an entry in the database, it will be viewed as an update upon submission, and overwrite the current person data. Otherwise, the following set of data must be specified in the *Person Template*. The person's name should be specified in full in the *person* field. The full postal address is specified using multiple *address* fields. The international telephone number which can be used to reach the person at work must be entered in the *phone* field, and the fax number should be entered in the *fax-no* field. Of course, the NIC handle for this person must be entered in the *nic-hdl* field to uniquely identify this person in the database.

Submission Information

In both the Network Template and Person Template, space is reserved to identify the person submitting these entries to the registry database. The submitter's e-mail address must be specified in the *changed* field together with the date the template is submitted.

Similarly, the *source* field is used to specify the registry database where the requester information can be found after an assignment is made. In this case it will be RIPE, as the requester information for this assignment will be stored in the RIPE database.

3.2.2. Additional Information

In the assessment of an assignment request, the additional information described below is always useful. In some cases, IR's may require this data be provided as part of the evaluation process.

Deployment Plan Suppose we are dealing with a new corporation that wants to have access to the Internet, and estimates an immediate need for 10,000 addresses. In such cases, a *deployment plan* may be requested from the user. The plan should include a list of events which will lead to the use of the requested addresses, along with the dates that the events will occur. This can be used to determine how realistic the user is being, and if suitable to phase the assignment process according to the user's plans.

Topological Map The old saying "a picture is worth a thousand words" certainly holds in the case of networks. If a *topological map* of the current and planned network infrastructure in the requesting organisation can be acquired, it can provide insight on the network structure. Such maps are often

available, and are quite useful when combined with the addressing plan and current address space usage.

Special Circumstances Sometimes, due to the use of old systems or special purpose hardware, the user is unable to make use of assignments based on classless addressing. If this is the case, information should be gathered from the user as to the *specific hardware or software* which presents a problem. Moreover, it is useful to know how long the user will be using the hardware or software which presents a problem.

Verification Information In working with a user who hasn't had substantial network experience, it is sometimes hard to determine whether the user's request is based on a realistic plan. It can therefore be useful to request information which might indicate the degree to which the user understands network planning and management. First, one may ask how accurate the user thinks the estimations in the addressing plan are, and how they have been derived. The corresponding name space plans provide another indication of how well considered the user's plans are.

3.3. Evaluation

Having collected the above information, we must now determine a proper assignment with respect to the conservation and aggregation goals stated in Section 1. Every request requires an individual evaluation process that takes current assignment guidelines into account.

Given the above documentation, one must determine whether IP addresses should be assigned, and if so, how many and of what type. In the process, it is essential that IR's work to prevent the stockpiling of address space. The use of classless addressing will contribute to the conservation of address space. Meanwhile, to enable proper routing, one must make strategic decisions with regard to aggregation. These concerns motivate the evaluation process outlined in this section.

Evaluation Steps

1. Current address space usage One should start by comparing the current address space usage provided by the requester with other information available to the IR. After verifying the current address space usage, one should check to see if the requested IP addresses can be taken from those already assigned to the user.

2. Network Overview Next, the size of the request, specified in the *Network Overview* should be compared with the number of addresses to be used immediately, and within two years of the time the request is made. Here we

evaluate the utilisation rates, that is address space requested in relation to that to be used. Unless there are special circumstances, immediate utilisation should be at least 25% of the assigned address space, and the utilisation rate one year later should be at least 50%.

3. Private Address Space If *private address space* might be suitable for this network, it must be established that the user is aware of this option and has decided against it. Moreover users should be aware that they will have more address space at their disposal if they use private address space.

4. Very Small Enterprises (VSE's) An end user with a small number of hosts (currently ≤ 32) is referred to as a very small enterprise (VSE) regardless of the size of the organisation. Address space for VSEs should be assigned in a classless way. As with all address space requests, care should be taken to avoid assigning more address space than is required. VSEs that do not intend to connect to the Internet should not be assigned public address space but rather should be advised to use private address space. This is especially the case for VSEs that request PI space so they can easily arrange connectivity at a later date. These enterprises should be advised that for VSEs in general, the effort required to renumber at a later date is minimal.

5. Addressing Plan In evaluating the addressing plan, one should first check that the totals for the number of addresses to be used immediately, in one year, and in two years, correspond to those specified in the request overview. The validity of the network masks should then be checked to see if they are consistent with the size of the subnet. Sometimes address space can be saved by using different subnet masks than specified by the user. If so, the user should be requested to resubmit an addressing plan with a more appropriate use of network masks.

In general, there should not be a large gap between the number of addresses requested for a subnet (size) and the number which will be used. This holds even if the requester argues that network administration will be greatly simplified by an addressing scheme with lots of padding.

6. Additional Information If a deployment plan has been provided, the addressing plan should be reviewed to see if the two correspond. Likewise, one should inspect the topology map if it is available to see if it agrees with the addressing plan. Any information gathered which can be used to check the validity of the current address space usage and addressing plans should be employed.

7. Address Reservations Assignments must be based solely on realistic expectations as specified in the addressing plan and the current address space usage. End users are not permitted to reserve addresses based on long term

plans, because it fragments the address space. Such reservations are generally fruitless because they turn out to be unnecessary or insufficient for the user's needs.

8. Static Dialup Due to constraints on the available free pool of IPv4 address space, the use of static IP address assignments (e.g., one address per customer) for dial-up users is strongly discouraged. While it is understood that the use of static addressing may ease some aspects of administration, the current rate of consumption of the remaining unassigned IPv4 address space does not permit the assignment of addresses for administrative ease. Organizations considering the use of static IP address assignment are expected to investigate and implement dynamic assignment technologies whenever possible. If allocations for this purpose are indeed made, special allocation and verification procedures apply. Please contact the RIPE NCC for details.

9. Virtual Hosts Sometimes a single host is assigned more than one IP address on the same interface and physical network. Often this is used to circumvent shortcomings in higher level protocols such as HTTP. Large scale assignments for this purpose are discouraged for the reasons mentioned in the paragraph above. If allocations or assignments for this purpose are indeed made, special allocation and verification procedures apply. Please contact the RIPE NCC for details.

3.4. Assignment Procedures

We now describe the specific procedures to be followed in assigning address space. In the following, we assume that the required information has been gathered, and evaluated as outlined in the previous subsections. The procedures described in this subsection lead to the assignment of specific address space for the request under consideration.

3.4.1. Assignments within Allocations

Once an IR has assured that the address space request merits the assignment of some amount of address space, the actual set of addresses to be assigned must be selected. If the addresses are to be assigned from a range of address space allocated to the local IR making the assignment, then care must be taken to prevent fragmentation of the allocated space. Specifically, each set of address space assigned should start on a CIDR (bit) boundary. This means the start address for each range of addresses to be assigned must be divisible by the size of the range. This helps to achieve the aggregation goal in address space assignments.

Suppose a request can be satisfied with either a number of small chunks of

address space or with a single large one. For example, if 384 addresses are sufficient to satisfy a request, but no more than 256 will be used in a single physical subnet, then the user can be assigned a /24 and a /25 rather than a /23, which results in saving a /25 for another user. In accordance with the conservation goal, local IRs are encouraged to assign multiple ranges of addresses in such cases, rather than a single large range. Of course the effort to do so should increase as the amount of address space that can be saved in doing so increases.

Local IRs are always welcome to request advice from the RIPE NCC in making assignment decisions. In some cases, however, an assignment must be approved by the RIPE NCC even if it is made from a block of address space allocated to the local IR making the assignment. In particular, if the size of the assignment is larger than the local IR is permitted to make, it must be approved by the RIPE NCC in advance. The size of assignments a local IR is permitted to make without prior approval is determined by the local IR's *assignment window*, discussed in Section 3.6.

If the addresses are to be assigned from a range which has not allocated to the local IR, the selection will be made by the IR to which the address space has been allocated. In general, this will be the regional registry.

3.4.2. PA vs PI Space

The criteria used to determine the amount of address space and the registration requirements are identical for PA and PI address space. For example, regardless of whether the assignment is for PA or PI address space, an assignment with a prefix longer than /24 can be made if the request can be satisfied with less than 256 addresses.

Local IRs must make it clear to the user which type of address space is assigned. Clear contractual arrangements which specify the *duration* of the address space assignment are mandatory for every assignment of PA address space. Although not strictly required, it is strongly recommended that contractual arrangements are made when assigning PI space as well.

With respect to aggregation, the clear advantages of assigning PA space mandate that IRs recommend its use whenever possible. End users should therefore be advised to use PA space if it appears to be sufficient for their situation.

The consequences of using PA or PI space must always be made clear to the end user. In particular, to be sure that end users understand the consequences of using PA space, the assigning IR must give each user requesting PA space a warning similar to the following:

Assignment of this address space is valid as long as the criteria for the original assignment are still met and only for the duration of the service agreement between yourself and ISP XXXX. ISP XXXX has the right to re-assign the address space to another user upon termination of the agreement or following an agreed period thereafter. If the address space assignment becomes invalid, you may have to re-configure the addresses of all equipment using this address space. The reconfiguration is only necessary if you continue to require global uniqueness of the addresses for that equipment. It is important to realise that some Internet services do not require globally unique addresses. For example, services that can be accessed through a NAT (network address translator) or through an application layer gateway/firewall don't require the machines which access them to have globally unique addresses.

Of course, the consequences of using PI space must also be made clear to the end user. This is accomplished by giving each user requesting PI space a warning similar to the following:

Assignment of this address space is valid as long as the criteria for the original assignment are met. Note that the assignment of PI address space does not imply that it will be routable on any part of the Internet. Users may have to pay a premium for routing of PI addresses (as opposed to PA addresses). Eventually, it may become impossible to get relatively small amounts of PI space routed on most of the Internet. We strongly suggest you contact any prospective service providers for information regarding the possibility and pricing of service when using PI addresses.

The type of assigned address space must be registered in the *status* attribute of the *inetnum* object in the RIPE database by the assigning IR. The possible values of this attribute are

ASSIGNED PA

This is used for PA address space that has been assigned to an end user.

ASSIGNED PI

This is used for PI address space that has been assigned to an end user.

Every new address space assignment must be marked as PA or PI in the RIPE database. Moreover, to improve registration of old assignments, IRs are

encouraged to mark past assignments in the registry database with PA or PI as appropriate. Any assigned address space without an explicit type in the status attribute is assumed to be PI space. Priority should therefore be given to marking PA space as such.

In general, local IRs are provided with ranges of PA space from which they can make assignments. If an end user requests address space of a type which an IR does not assign (typically PI), the IR must refer the end user to a registry which can fulfill the request. If a local IR wants to assist one of its customers in obtaining an assignment of address space for which it does not hold an allocation, it should support an IR that does provide this service. This includes aiding the end user in the preparation of a properly documented request and in furnishing background information to the assigning IR as required. The local IR can of course refer the user to a local IR which is able to make the assignment.

Local IRs which do not normally assign large amounts of a given type of address space (again typically PI) need not hold an allocation to handle address space requests. The address space can be acquired from the RIPE NCC as needed. In general, such assignments are not aggregatable.

3.5. Replacing IP Addresses

Much of the address space assigned in the past is aggregated in practice but *formally* is not of type PA. Formally, address space is not of type PA unless there are contractual agreements regarding the assignment's purpose and term of validity. It is therefore recommended that IRs ask end users to release non-PA address space upon termination of service. Similarly, if an end user moves to a new IR to obtain Internet services, the new IR should encourage the user to release any non-PA address space they hold, and to request a new assignment (a process with which the new IR should be more than happy to help). To minimise the use of non-PA space in the future, IRs should work to make contractual arrangements to formally convert aggregated address space to PA address space.

While the procedures for numbering and renumbering hosts in an IP network are becoming less troublesome, asking or forcing end users to renumber is sometimes problematic. The renumbering process usually requires a considerable amount of time and effort both on the part of the end users and on the part of the ISPs and local IRs involved. In some cases, there is a clear obligation to replace address space assignments, and local IRs should be prepared to support their customers in the process. A more general and very important case is the (voluntary) replacement of PI address space which for historical reasons has been randomly assigned and cannot be aggregated with other PA assignments. Such replacements can play a key role in containing the growth

of routing tables, and thus for the maintainability of the Internet as a whole. Because the renumbering process is nontrivial, the Internet Registry System as a whole must support the process as far as possible.

During the period in which end users migrate individual services or parts of their networks to the new address space, complications may arise. In many cases, they may need to be connected to more than one ISP for the duration of the transition period, and sometimes addresses from both the old range(s) and the new might have to be managed and used in parallel. With the goals of aggregation and conservation in mind, as well as to minimise duplicate logistics, this period should be kept as short as possible.

IP Address Space Replacement Procedures:

In general, address space should be replaced on a one-to-one basis. An assignment of PA space to replace previously assigned PI space can be made if the original assignment criteria are still met and if the address space to be replaced is currently used for the end user's network.

Only if a large percentage of the original assignment is not in use (50% or more than 4096 addresses) will an end user be required to submit the usual documentation to the new registry. This part of the request is then treated like any other request for assignment of additional addresses.

The address space to be replaced (the individual address ranges and the total size) must be properly documented with the standard IP address space assignment request forms. For address space that was allocated by local IRs established within the framework of the RIPE NCC, a copy of the documentation is forwarded to the registry or registries that assigned the address space being replaced. Before assigning the new address space, an agreement (preferably contractual) should be reached regarding the maximum period within which the previously assigned addresses will be returned to the original registry or to the regional registry for eventual reassignment.

Whenever a large amount of addresses are to be replaced (the equivalent of a /20 or more) the regional IR must be informed about the intended replacement and the agreements on the maximum period of parallel assignments. In complex cases, the regional IR may decide to provide guidance in the process of managing the address space replacement.

In general a period of 3 months should be allowed for the end user to complete the transition to the new addresses. For exceptional cases, where the end user requests to keep both assignments for more than 6 months, approval should be obtained for the proposed time frame from the RIPE NCC.

For those addresses that have not been assigned by a local IR, or which were assigned by an IR that has since closed, the regional IR will act in lieu of the original registry.

3.5.1. Multihomed Users

An end user may have reason to obtain connectivity through more than one service provider. If so, it may be necessary to obtain address space assignments from more than one IR to support different parts of the user's network. In general, there is no problem with users acquiring address space and service from more than one IR. Their networks are then referred to as *multihomed*.

Because users can be multihomed, IRs must be especially careful in reviewing address space requests, and the corresponding *current address space usage* described in Section 3.2.1.2. One must be sure that users are not acquiring multiple assignments for the same purpose from different IRs. Moreover, one must check that a similar address space request has not been refused by another IR for some valid reason.

3.5.2. Update the RIPE Database

Registration of objects pertaining to an assignment in the RIPE database makes it possible to track the use of address space, facilitate operational contacts, and facilitate studies of address allocation. These activities are essential to effective maintenance of the Internet.

Submission of objects to the database is the final and required step in making an assignment. This step makes the assignment definite, and makes public information regarding the assignment available to anyone seeking it. Care should therefore be taken to assure the correctness of the assignment and of all relevant data prior to completing this step.

The information collected about the user's organisation in the *Network Template* (see Section 3.2.1.5) is used to construct an *inetnum* database object. The range of addresses assigned to the user is also entered in the *inetnum* object, and is specified in *dotted quad notation*. For example, if an organisation is assigned a /22 address set for 1024 network addresses, the range will be something like: *193.1.192.0 - 193.1.195.255*. And, as stated above, the *status* field of the *inetnum* object is used to specify whether the assigned addresses are PI or PA.

Unless up-to-date objects are already available in the RIPE database, in addition to the *inetnum* object, a *person* object must be submitted for each person specified in the *tech-c* and *admin-c* fields of the *inetnum* object.

Assuming the assigning IR has properly stored information gathered in the assignment process for future reference, submission of the objects described above completes the assignment process. The IR can now provide the end user with the assigned address range and any other data relevant to its use.

3.6. Assignment Window

It is essential that local IR staff follow the procedures for address space assignments and apply the evaluation criteria used to determine assignment sizes as discussed above. The procedures are straightforward. The evaluation criteria however, can be difficult to apply to new situations.

To assure the conservation, aggregation, and registration goals are met, we must be sure the assignment criteria and procedures are properly applied. In general, this means that local IRs with little or no experience should receive maximal support in the assignment process, whereas local IRs with more experience should be allowed to make most assignments without consulting the RIPE NCC. Large assignments always require prior approval because of their impact on the available address space.

To achieve this variation in support level, each local IR is given an *assignment window* by the RIPE NCC. The assignment window is the maximum number of addresses that can be assigned by the local IR to an end user without prior approval by the RIPE NCC. This is expressed in /nn notation. Therefore, a local IR with an assignment window of /23 is allowed to assign up to and including 512 addresses to an end user without prior approval of the RIPE NCC. As the local IR staff gain experience, the window size is increased.

Every assignment of address space which is larger than the local IR's assignment window is formally invalid until explicit approval is acquired from the RIPE NCC. Without this approval, the address space can not be used as it may result in a failure to meet the uniqueness requirement for Internet addresses at a later date.

The assignment window is not only applied to individual assignments, but to multiple assignments to the same end user in a 12 month period. If a local IR makes more than one assignment to an organisation in any 12 month period, the total amount of address space assigned may not exceed the local IR's assignment window. Additional address space can only be assigned to that organisation after approval by the RIPE NCC.

In general the assignment window for new registries is 0. This means that every assignment requires prior approval by the RIPE NCC before becoming effective.

As the local IR staff become familiar with assignment procedures, the assignment window can be raised. In general, it will first be raised to /24. At this point, the local IR staff can make assignments for up to and including 256 addresses without prior approval from the RIPE NCC. If the RIPE NCC receives a request to raise the assignment window for a local IR, it will be evaluated based on the proficiency of the local IR staff. This is determined based on whether assignment documentation presented to the RIPE NCC is correctly completed, whether good judgement is shown in the evaluation of address space requests, whether past assignments have been properly registered, and on the experience of the local IR with handling larger assignments. Currently, the maximum size of the assignment window for any local IR is a /19. Therefore, every assignment involving more than 8192 addresses requires the approval of the RIPE NCC.

An established local IR is responsible to train new staff to handle address space assignments according to the policies and procedures described in this document. Sometimes, due to time constraints on experienced registry staff the training is not performed sufficiently, and new staff members of a well established local IR may make errors both in judgement and procedure before they are properly trained to make assignments. If such errors are noticed by the RIPE NCC, the local IR will be notified, and if it happens repeatedly, the assignment window for the local IR may be decreased to prevent the new staff members from making erroneous assignments involving large amounts of address space. The assignment window can again be increased based on the criteria stated above.

4. Rules and Guidelines for Allocations

In Section 3, we described the procedures for handling requests for address space, including the process used to determine which addresses should be assigned to an end user. In most cases, the address space will be selected from a range that has been *allocated* to the local IR for this purpose. Of course, before a local IR can select addresses to fulfill a request, it must have a range of address space to choose from. If the local IR does not have sufficient address space of the type to be assigned, then a request for an (additional) allocation can be submitted to the RIPE NCC.

To meet this need, a range of addresses is made available to a local IR by the RIPE NCC. Such an address range is referred to as an *allocation*. In Europe, the RIPE NCC is the only IR permitted to allocate address space. A local IR is not allowed to re-allocate part of its allocation to any other organisation. Moreover, without prior approval by the NCC, local IRs are not permitted to exchange allocated address space among themselves.

In some cases, a local IR makes address space assignments for the customers of another IP service provider which itself does not operate a local IR. The local IR is of course responsible for all assignments from its allocation, even if there is an agreed involvement of staff from the other IP service provider. Whereas staff of the other IP service provider can and should be involved in processing the end user's request, local agreements about shared responsibility in the registration process are not recognised by the regional registry and are strongly discouraged.

If at some point, an IP service provider decides to establish a local IR rather than using an existing local IR to obtain address space, then all subsequent assignments it makes should be from address space allocated to it from the RIPE NCC. Furthermore, any address space used by the ISP's customers which was acquired from a transit provider's allocations should be returned to the transit provider as soon as possible, and new assignments should be made to the end users from the ISP's allocated address space.

In the following subsections, we describe how a local IR can obtain an allocation and how allocated address space should be managed.

4.1. The Slow Start Mechanism

To prevent allocating large blocks of address space that won't be assigned, the RIPE NCC has introduced the concept of a *slow start* for allocations. The idea is to allocate address space to local IRs at the rate it will be assigned. The minimum size of an individual address space allocation is /19 (8192 addresses), and the maximum size is /16 (65536 addresses). The size of an allocation to a particular local IR is based solely on the rate that the IR

has assigned previously allocated address space to end users.

The slow start mechanism implements a consistent and fair policy for every local IR with respect to allocations. Although the mechanism is similar to the assignment window mechanism described in Section 3.6, the policy it implements is different. The size of further allocations depends exclusively on the assignment rate of the local IR concerned while the assignment window depends on the proficiency of the registry staff in evaluating requests and processing assignments.

4.2. First Allocation

When a new local IR starts up, it has no address space allocated to it. The first allocation will be made automatically by the RIPE NCC, generally upon receipt of the first assignment request from the local IR. Because there is no information about the rate at which a new IR will make address assignments, the size of the first allocation is always a /19 (8092 addresses).

Remember that the amount of space allocated does not determine the size of assignments a local IR can make. As discussed at the end of Section 3, a new local IR must have every assignment approved by the RIPE NCC until its assignment window is increased.

4.3. Further Allocations

A local IR can obtain additional allocations as required. A request should be submitted to the RIPE NCC when the currently allocated address space is nearly used up, or if a single assignment will require a larger set of addresses than can be satisfied with the allocated address space. To obtain a new allocation, a local IR should submit a request to the RIPE NCC which includes a complete list of the assignments made from all of their allocations. The RIPE NCC will check this information, compare it with assignments registered in the database and may request further information to determine the need for a new allocation. Additional address space will be allocated after the information supplied with the request has been verified, and a new allocation has been deemed necessary.

Unfortunately, there is a tradeoff between the aggregation and conservation goals in making allocations. To further aggregation, the RIPE NCC aims to allocate contiguous address ranges to a local IR. However, because conservation of address space must also be taken into account, this is not always possible. A new allocation to a registry can therefore not be expected to be contiguous with past allocations. While the RIPE NCC always aims to further the aggregation goal, and therefore to allocate contiguous space, *the RIPE policy is that under no circumstances are multiple allocations made to the*

same local IR guaranteed to be contiguous.

4.4. Allocation Registration

Allocations are registered in the RIPE Database by the NCC using *inetnum* objects. Information about the local IR, which is similar to that for an end user receiving an assignment is stored together with the range of allocated address space and its type. The range of addresses is stored in the *inetnum* field in dotted quad notation, and the type is stored in the *status* field and can have one of the following values:

ALLOCATED PA

This address space has been allocated to an IR, and all assignments made from it are provider aggregatable. This is the most common allocation type for local IRs.

ALLOCATED PI

This address space has been allocated to an IR, and all assignments made from it are provider independent.

ALLOCATED UNSPECIFIED

This address space has been allocated to an IR, and assignments made from it may be either provider aggregatable or provider independent. *This type of allocation is obsolete, and will not be applied to future allocations. Some older allocations have been used for both PA and PI assignments, and as such receive this allocation type.*

These objects are maintained by the RIPE NCC. When hierarchical authorisation is implemented, authorisation can be used for the creation of *inetnum* objects "under" the allocation objects.

5. DNS and Reverse Address Mapping

Applications such as *ftp*, *e-mail*, and *telnet* allow users to specify an Internet host as a *domain name*, such as "ns.ripe.net". With this notation, the full name of a host is determined in a hierarchical fashion. In this example, *net* is one of the top level domains in the system, *ripe* is the name of a subdomain defined in the *net* domain, and *ns* is the name of a host in the "ripe.net" domain. This hierarchy is similar to the UNIX file system, and it enables unique naming of hosts on the Internet.

Before an application can send IP packets to a given destination, however, its IP address must be determined. The Domain Name System (DNS) is a distributed hierarchical directory service which makes it possible to obtain the IP address for a host given its symbolic name specified in the domain name notation described above. The inverse procedure which produces the domain name from the IP address is called *reverse address mapping*, and is the focus of this section.

We begin with a brief introduction to the DNS because reverse address mapping is simply a specific application thereof. Detailed information on the DNS can be found in [Albitz94a]. In this section, we aim to provide a sufficient basis to understand the procedures involved in making reverse address mapping possible for address space allocated by the RIPE NCC.

5.1. Forward Name Mapping

The DNS is a client/server system. If data is properly administered on the domain name servers, then every public IP address in use has exactly one domain name associated with it. The IP address which corresponds to a given domain name can be extracted with a *resolver*, which works as follows. If a machine needs the IP address for a host identified by its symbolic name, and it cannot be obtained locally, the resolver is used, first to inspect the domain name, and then to determine what name server should be able to provide the IP address that corresponds to it. The resolver then sends a request to the appropriate name server which either returns the required IP address, or the address of a server that should be able to provide more details. If the latter, the resolver repeats its request to the new server. This continues until the IP address is found (or the host is reported to be unknown). This procedure is called forward *mapping* or *resolution*.

Of course, before a host can be identified with the DNS, it must be registered with a server. The responsibility for name registration is hierarchical. The administration of a subset of the DNS name space is *delegated* to a representative of an organisation by the party which holds responsibility for the name space it falls under. In the example above, the administration of the top level domain *net* is performed by the InterNIC. The InterNIC delegated the

subdomain *ripe* to a representative of the RIPE NCC, who chose to use the name *ns* to refer to one of the hosts in its network. After the name *ns* has been properly specified in the name server for *ripe.net*, the domain name *ns.ripe.net* can be used to find the Internet host with IP number 193.0.0.193.

The suffix of each domain name corresponds to a *top level domain* (TLD) in DNS, and authority to administer it is delegated by IANA. Generally, this will be an ISO 3166 two letter country code such as "nl" for The Netherlands, "fr" for France, or "us" for the United States. These codes have been delegated by IANA as country specific TLDs. (The only exception is the domain ".uk" which has been delegated to Great Britain; "gb" is in fact the ISO code.) The administration of each country specific TLD is delegated to a representative residing in the country. If responsibility for a country specific TLD has yet to be delegated, then a resident can request permission from IANA to manage it. Responsibility for the TLD will be delegated to that person if the guidelines specified in [Postel94a] are agreed to and if no objections are made within some short period after the possible delegation is announced.

When the DNS was first implemented, a small number of "generic" three letter codes were defined as TLDs. These domains are administered by the InterNIC and are still in wide spread use within the US. Historically, organisations have selected TLDs based on their primary business. For example academic institutions usually have names that end in "edu", military organisations in "mil", and companies in "com".

Delegation policies are up to the party responsible for the administration of the domain from which delegations are made. For example, policies regarding delegation of second level domain names ending in "edu" are determined by the InterNIC. Delegation policies for third level domain names, however, are determined by the body to which the corresponding second level domain name has been delegated. For example, a representative of Catatonic State University may be responsible for the delegation of subdomains which fall under "cat.edu". In general, the delegation policies applied by DNS domain administrators are expected to remain within the guidelines outlined in [Postel94a].

In mapping a domain name to an IP address, the name servers administered by those responsible for the associated domains must provide the information sufficient to resolve it. Suppose a request is received for the IP address of a host named "bite.dog.cat.edu". Because the InterNIC is responsible for all delegations in the TLD "edu", the request can first be passed to InterNIC's name server. If the domain "cat.edu" has been delegated to the Catatonic State University name server, the InterNIC's name server will probably pass the request to the university's name server, which in turn might pass it on to

the appropriate department's name server. If all name server files are in order, the department's name server should provide the IP address for the domain name in question.

This is a simplified model of how name resolution occurs. It ignores caching and other alternatives that are used to optimise the DNS. It does, however, give a realistic view of which parties are responsible to provide which information in the resolution process.

5.2. Reverse Address Delegation and Mapping

Just as it is necessary to obtain the IP number for a host with a given domain name, it is often necessary to do the reverse, that is to obtain the domain name associated with a specific IP address. Authorisation checks used by some Internet applications some diagnostic services need the fully qualified domain name associated with an address, for example. Given an IP address, the procedure used to obtain the domain name associated with it is called *reverse mapping*.

In the DNS, a pseudo domain called "in-addr.arpa" (a historical abbreviation for "inverse addresses in the Arpanet") has been defined, to make this possible. Delegations in this domain are made by IRs, because they allocate and assign address space. For example, the RIPE NCC has been delegated the domain "193.in-addr.arpa", because it is responsible for allocations and assignments in 193/8 (among others). The RIPE NCC delegates authority for names within the domain "193.in-addr.arpa", after the corresponding address space has been allocated and assigned.

Given the IP address 193.3.20.100 in dotted quad notation, suppose its domain name is required. First, a pseudo domain name "100.20.3.193.in-addr.arpa" is generated by reversing the order of the address components and adding the suffix ".in-addr.arpa". This name is then used to find the domain name corresponding to the IP address with reverse mapping. Once the name has been generated in the pseudo domain, the reverse mapping mechanism is technically equivalent to the forward mapping mechanism.

Although the mechanisms used for forward and reverse mapping are equivalent, authority of the domain hierarchies is different. In particular, while delegation in the generic and country specific TLDs follows the organisational structure described in the previous section, delegation in the pseudo domain "in-addr.arpa" involves those responsible for the allocation and assignment of the corresponding address space.

The term *reverse delegation* refers to the delegation of IP address names in the pseudo domain "in-addr.arpa".

For example, the inverse domain name "193.in-addr.arpa" has been reverse delegated to the RIPE NCC, which is therefore responsible to supply information which can lead to domain names corresponding with assigned IP addresses in the 193/8 range. It is important to note that reverse delegation of address names in the pseudo domain does not occur automatically either upon allocation or upon assignment of address space. Rather, for all allocations and assignments from the address space managed by the RIPE NCC, *reverse delegation only occurs in response to an explicit request submitted to the RIPE NCC*. This is of course a prerequisite if reverse mapping is to be used for IP address to domain name translations.

As described above, pseudo domain names are generated in terms of dotted quad notation for IP addresses. This requires that reverse delegation take place on octet boundaries. Suppose the RIPE NCC allocates a /17 to a local IR named Eye-Pea, for example, 193.1.0/17. Then no reverse delegation of the name "1.193.in-addr.arpa" will be made to Eye-Pea, because it is only responsible for a subset of the address space corresponding to the inverse domain "1.193.in-addr.arpa". The RIPE NCC therefore remains responsible for the inverse domain name "1.193.in-addr.arpa" and all reverse delegations that fall under it.

On the other hand, suppose a /16 allocation is made to a local IR called Aye-Queue, for example 193.2/16. Then, the *zone* "2.193.in-addr.arpa" can be reverse delegated to Aye-Queue upon request because that IR is responsible for all assignments in the address range 193.2.0.0 - 193.2.255.255. Subsequently, Aye-Queue will be expected to provide pointers to reverse domain name information for addresses in the range 193.2.0.0 - 193.2.255.255. Note that if the request is granted, Aye-Queue is said to have *authority* over the "2.193.in-addr.arpa" zone.

Following the procedures specified in Section 3 of this document, Aye-Queue may then assign a /24, for example 193.2.40.0 - 193.2.40.255 to some organisation called Organiser. Subsequently Aye-Queue can make a reverse delegation for "40.2.193.in-addr.arpa" so that requests for domain names associated with addresses in the range 193.2.40.0 - 193.2.40.255 will be forwarded to Organiser.

Note that with the classless scheme, both address space allocations and assignments may take place on non-octet boundaries, whereas reverse delegations must occur on octet boundaries because the the reverse domain name is specified in terms of dotted quad notation for the IP address. This means that allocations and assignments made on non octet CIDR boundaries, a slightly different delegation strategy is required, that will be explained in this section. The basic system, however, remains unchanged.

The RIPE NCC together with the local IRs act together to assure that reverse delegation is correctly performed. This allows reverse mapping to be used to find the domain names corresponding to IP addresses from the range managed by the RIPE NCC. The role of both parties is covered in the following subsections.

5.3. Local IRs and Reverse Delegations

If a local IR obtains reverse delegations for the address space it assigns, it is able to efficiently provide expected services, namely IP number to domain name mapping, for the end users it services. In this section, we describe how reverse delegations can be obtained.

We start with a description of the responsibilities which accompany authority over inverse address domain name zones. We then discuss the proper distribution and maintenance of the reverse address database when CIDR address space allocations and assignments are made. The specific procedures used to obtain reverse delegations are described. Finally we consider issues relevant to local IRs regarding PA versus PI address space assignments.

5.3.1. Responsibilities

Prior to the delegation of domain name zones (e.g. "cat.edu"), the person or organisation to whom authority over the zone is delegated agrees to provide some key services necessary to support domain names extending from the zone. Similar agreements are of course necessary for reverse delegations if DNS is to provide accurate mappings from IP addresses to domain names.

When a reverse domain zone is delegated to a local IR, care should of course be taken in the proper construction of the DNS configuration files for the zone. Known pitfalls and some useful tips for avoiding them can be found in [Beertema93a].

For each zone, a secondary server must be set up to improve the reliability of the database under adverse conditions. To increase the probability that the secondary server can be reached if the primary server becomes unavailable, the secondary server is required to be on a subnet physically separated from the primary server. For reverse delegations corresponding to /16 allocations to local IRs, an additional secondary server is provided by the RIPE NCC. This does *not* replace the required secondary, but rather provides extra reliability for these substantial delegations. It is customary for local IRs and other organisations managing reverse domain names to provide secondary services for one another.

As is required for top level domain name servers, both the primary and

secondary reverse domain name servers must be directly reachable from the Internet.

If a local IR is given authority over a reverse domain name zone, it is responsible for subsequent reverse delegations in that zone. This means the local IR must assure that an organisation to which authority is delegated satisfies the requirements described in this section for its zone. In Section 5.4, we describe the services provided by the RIPE NCC to assure proper working of the reverse domain name system. Local IRs should use this as a reference for the services they provide to end users to whom they make reverse delegations.

5.3.2. CIDR and Reverse Delegations

As mentioned above, making allocations and assignments on CIDR boundaries, rather than on traditional class (octet) boundaries, requires a slight variation on the reverse delegation scheme. Basically, if an allocation or an assignment is made on a nonoctet boundary, authority over the corresponding reverse domain zone must not be delegated, but must be maintained by the IR that makes the address space allocation or assignment.

5.3.2.1. Allocations and Reverse Delegations

If an allocation smaller than a /16 is made to a local IR, such as the 193.1.0/17 allocation to Eye-Pea in our example, then authority for an immediate subdomain of 193.in-addr.arpa cannot be granted, because a subset of the corresponding address space may be allocated to another local IR.

For any single allocation under /16 in the 193/8 address range, the RIPE NCC will maintain authority for the immediate subdomain of 193.in-addr.arpa, and reverse delegations will be made upon request if preceded by corresponding address space assignments. This of course holds for reverse delegations corresponding to any /8 address space allocations managed by the RIPE NCC.

If at some point, the remainder of the block (in this example 193.1.128/17) happens to be allocated to Eye-Pea, a request (accompanied by a *domain* database object) can be submitted for a reverse delegation of 1.193.in-addr.arpa. When a reverse delegation is made to a local IR, the RIPE NCC will forward all the relevant data and the responsibility for any reverse delegated zones to the local IR. In this example, if 1.193.in-addr.arpa is delegated to Eye-Pea, all past and future delegations made from that domain fall under the authority of Eye-Pea.

Suppose, on the other hand that a /16 has been allocated to the local IR in the

first place, such as the 193.2/16 allocated to Aye-Queue in the example above. Then Aye-Queue (e.g. the local IR receiving the allocation) may immediately request authority for a subdomain of 193.in-addr.arpa, in this case 2.193.in-addr.arpa. Because Aye-Queue is responsible for all address space corresponding to the reverse domain name 2.193.in-addr.arpa, the reverse delegation can be granted.

5.3.2.2. Assignments and Reverse Delegations

With respect to reverse delegations, we can distinguish two address space assignment categories, namely those assignments that involve an integral number of /24's, and those that do not. We begin with the latter.

We first consider an assignment made by a registry holding a full /16 allocation. Continuing with our example, suppose that Aye-Queue has been allocated 193.2/16 and has a reverse delegation for 2.193.in-addr.arpa. Aye-Queue might assign the 64 addresses in 193.2.0/26 to an end user, say Use-IQ. Following the reasoning applied for the /17 allocation to Eye-Pea above, Use-IQ cannot obtain a reverse delegation for 0.2.193.in-addr.arpa, because some of the corresponding address space may be assigned to other end users.

To address this problem, Aye-Queue can essentially delegate 0.2.193.in-addr.arpa to itself, and maintain the IP address to domain name information for Use-IQ and any other end users to whom the corresponding address space is assigned. Such a delegation to the same organisation is of course not necessary, but it can be helpful in easing the administration of multiple domains. When a local IR makes a reverse delegation to itself for address space it assigns, it is recommended, but not strictly required that a domain object be submitted to the RIPE database to register the reverse delegation.

Suppose a similar assignment is made by Eye-Pea from the 193.1.0/17 address space allocated to it. If Eye-Pea assigns 193.1.0.0/26 to an end user, say Use-IP, the problem arises again. Moreover, because Eye-Pea does not have authority for 1.193.in-addr.arpa, it cannot delegate 0.1.193.in-addr.arpa to itself. Rather Eye-Pea can receive a reverse delegation for 0.1.193.in-addr.arpa upon request, after at least one assignment has been made from the corresponding /24 address space. The procedures to obtain a reverse delegation are outlined in Sections 5.3.3 and 5.5 below.

Of course, Use-IQ could have been assigned an integral number of /24's. For example, suppose it was assigned 768 addresses in the range 193.2.0.0-193.2.2.255. Then Aye-Queue can make reverse delegations of {0,1,2}.2.193.in-addr.arpa to Use-IQ. The procedures Aye-Queue should follow in making reverse delegations and the services it should provide to its end users are described in Sections 5.4 and 5.5.

Suppose now that Eye-Pea assigns the 256 addresses in the range 193.1.0.0 - 193.1.0.255 to Use-IP. Then Eye-Pea need not manage the reverse domain 0.1.193.in-addr.arpa, because Use-IP is the only end-user with addresses assigned from the corresponding address space. In this case, Eye-Pea should support the end user in requesting a reverse delegation (using the *inetnum* object) from the RIPE NCC, and provide secondary database and other services as necessary. See the next section for more information.

In summary, after an assignment which is smaller than /24 has been made. a local IR should obtain a reverse delegation for the reverse name corresponding to the assigned address space. It should maintain the reverse mapping entries to reflect IP address to domain name information provided by the end user. More information on management of inverse mappings when allocations and assignments are made on non-octet CIDR boundaries is available in [Eidnes96a].

5.3.3. Obtaining a Reverse Delegation

We now describe the procedures to be followed by a local IR, requesting a reverse delegation from the RIPE NCC. As can be concluded from the previous section, a local IR can obtain authority for a /16 following its allocation, or for a /24 after one or more assignments have been made from it. The two procedures are very similar. However, in this section we describe only the steps to be taken for a /16 reverse delegation, while the procedures for a /24 are described in Section 5.5. To acquire authority over the reverse domain name zones corresponding with the address space involved, the following steps should be taken.

1. Configure the DNS configuration files for the zone for which the delegation is requested. Following the findings in [Beertema93a], for a direct subdomain of the domains under RIPE NCC authority (193.in-addr.arpa, etc.), the following settings are strongly recommended:

Refresh = 8 hours (28800 seconds)
Retry = 2 hours (7200 seconds)
Expire = 7 days (604800 seconds)
Time To Live = 1 day (86400 seconds)

We recommend a lower retry level if connectivity is unstable.

2. Install the configuration files on the primary and secondary servers.
3. Gather information required for a RIPE database *domain object*. The name of the domain requested, for example "1.193.in-addr.arpa" must be entered in the *domain* field. A description of the organisation

to maintain the zone under consideration should be included in one or more *descr* fields. The *admin-c*, *tech-c*, and *zone-c* fields should specify the persons who are responsible for administration at the site of the primary server, technical support for the site, and authority over the zone, respectively. The *nserver* fields, of which there must be at least three, specify the domain names of the primary, secondary, and RIPE NCC reverse name servers. Finally, as for all RIPE database objects, there is a *changed* and *source* field used to specify the e-mail address of the person making the request and the database source "RIPE", respectively.

4. Submit a request for a reverse delegation to <auto-inaddr@ripe.net>. This step implies the requirements described in Section 5.3.1 have been satisfied. The domain object described above should be included in the request, as well as zone file entries for the zone above the one requested. For example if a reverse delegation is requested for 1.193.in-addr.arpa, the relevant zone file entries should be included for 193.in-addr.arpa, whereas if a reverse delegation is requested for 2.2.193.in-addr.arpa, the zone file entries should be included for 2.193.in-addr.arpa. (This is one request to <auto-inaddr@ripe.net> for which the X_NCC_RegID field may be omitted, but the omission will result in the a low priority for the request.)

As described below, the RIPE NCC will test the proper working of the primary and secondary servers. If the local IR making the request has been allocated the address space corresponding to the reverse domain name zone for which the request has been submitted, and the servers are running properly, the request for delegation will be granted. In the next section, the services provided by the RIPE NCC are described.

5.3.4. Side Effects for PA/PI Assignments

End users have a right to reverse mapping services. An end user holding non-PA address space from a zone that has been reverse delegated to one service provider is permitted to keep the address space, and obtain connectivity services from another provider. Because the address space falls in the reverse delegation zone of the initial local IR, that IR is required to continue to provide reverse mapping services for the address space assigned to the end user. Moreover, the local IR has to provide this service under the same conditions it applies to its other end users (e.g. extremely high fees for this service are unacceptable - unless they are applied to all end users.)

For PA addresses, contractual agreements confine the provision of reverse mapping services directly to the time period for which the assignment is valid. Clearly this is one reason why converting non-PA address space to PA

address space is in the best interests of the assigning local IRs and their customers.

5.4. The RIPE NCC Services for Reverse Delegation

Because the RIPE NCC allocates address space to local IRs from 193/8 and other /8's allocated to it by IANA, it is responsible for reverse delegations that correspond to the address space. Requests to resolve reverse mappings for address space assigned from that allocated by the RIPE NCC should therefore be sent to the RIPE NCC. If a reverse mapping is required for an address, and one is not sure which regional IR the address originated from, a request can be sent to the RIPE NCC; if the address space originated from one of the other regional registries, its contact details will be returned. Of course, one cannot perform reverse mappings for IP addresses that have not either been allocated (/16) or assigned and registered (/24).

Upon receiving a request for a reverse delegation of an immediate subdomain of 193.in-addr.arpa or 194.in-addr.arpa, the RIPE NCC will first check if all of the corresponding address space has been allocated to the requesting IR. If the request is for a /24, then a check will be made as to whether some or all of the address space has been assigned. If the required allocation (assignment) has been made, the following services will be performed:

1. Access to the primary and secondary servers specified in the domain object will be tested.
2. Data for any already registered reverse zones in the corresponding address space will be forwarded to the requesting organisation, for inclusion in the newly defined reverse zone files. (If the reverse delegation is made, further responsibility for past delegations is transferred to the requesting organisation.)
3. The reverse domain name server will be tested using the information provided in the request.

For a /16 reverse delegation, the next two tasks are also performed:

4. If the reverse delegation request is to be granted, the RIPE NCC will set up secondary server for the reverse domain on ns.ripe.net, and notify the local IR.
5. Once the reverse delegation has been made, requests made to the RIPE NCC for reverse delegations for address space corresponding to the delegated zone will be forwarded to the mailbox specified in the *SOA RR* field of the reverse zone configuration file, and to the person specified in the *zone-c* field of the domain object.

All requests for reverse delegations at the RIPE NCC are now being handled by an automatic procedure. Zone information for the request described above will be checked automatically upon receipt in the <auto-inaddr@ripe.net> mailbox. If the zone is set up properly, the request will be evaluated by a RIPE NCC staff member, and if everything is in order, the request will be granted.

Reverse delegations allow a local IR to administer reverse mapping services for IP addresses assigned to its end users. The end users can then be sure they can be identified quickly and easily from hosts on the network having only access to the IP address. Because of the distributed nature of the database, its proper working depends on the correct administration of all zones. On some rare occasions, an organisation managing a delegated zone proves unable to correctly perform the required services. If repeated complaints are made regarding the administration of a zone delegated by the RIPE NCC, the RIPE NCC may revoke the delegation, as a final service to support efficient and correct reverse mapping.

5.5. Making Reverse Delegations to End Users

Up to this point, we have been concerned with the procedures surrounding the reverse delegation of a zone to a local IR. Because the reliability of the data distributed with the DNS increases as the distance to the data source decreases, reverse delegations of a zone can also be made to end users for each /24 they are assigned. In this context a /22 assignment is simply a multiple /24 assignment, for which multiple reverse delegations should be requested.

A local IR should always support end users requesting reverse delegations corresponding to address space (one or more /24's) which has been assigned to them from address space allocated to the local IR. The end user must be made aware of the means to acquire a reverse delegation and the responsibilities that go with it.

Basically, the same criteria hold in the case of reverse delegations to end users as hold for local IRs. The end user requesting authority for a particular zone must agree to fulfill the responsibilities described in Section 5.3.1. There is a slight variation of the procedures described in Section 5.3.3. While the end user is responsible for the reverse delegation and therefore for the procedures surrounding it, local IRs traditionally support end users in obtaining and in maintaining reverse delegations for their address space. For example, it is common for the assigning local IR to provide a secondary server for the reverse delegation.

If a local IR such as Eye-Pea has an allocation for a /19, /18, or a /17 address

range, then the reverse delegation must be made by the RIPE NCC rather than the local IR. In this case, an *inetnum* database object (rather than a domain object) should be submitted to the RIPE database and with the request sent to <auto-inaddr@ripe.net>.

On the other hand, if a local IR such as Aye-Queue has a /16 allocation, it may make the reverse delegation itself, but is encouraged to submit an *inetnum* object to the RIPE database to register the reverse delegation.

In both cases the local IR is expected to perform services similar to those performed by the RIPE NCC to assure the requested delegation is appropriate and properly administered. For example, the assigned address space must correspond to the zone requested, and the primary and secondary servers must be tested to assure that they are reachable and properly configured.

Acknowledgements

The authors would like to thank the staff of the RIPE NCC for reviewing multiple versions of this document.

References

- 96a. Y. Rekhter , R. Moskowitz, D. Karrenberg, G. de Groot , and E. Lear, *Address Allocation for Private Internets*, RFC 1918 (02/1996).
- Albitz94a.
Paul Albitz and Cricket Liu, *DNS and BIND*, O'Reilly & Associates, Inc., Sebastopol, CA (1994).
- Beertema93a.
P. Beertema, *Common DNS Data File Configuration Errors*, RFC 1537 (10/1993).
- Caslav96a.
P. Caslav, M. Kuehne, and C. Orange, *European IP Address Space Request Form*, ripe-137 (05/1996).
- Deering89a.
S. Deering, *Host Extensions for IP Multicasting*, RFC 1112 (08/1989).
- Eidnes96a.
H. Eidnes and G. de Groot, *Classless in-addr.arpa delegation*, Internet Draft (05/1996).
- Gerich93a.
E. Gerich, *Guidelines for Management of IP Address Space*, RFC 1466 (05/1993).
- Postel94a.
J. Postel, *Domain Name System Structure and Delegation*, RFC 1591 (03/1994).