

A Fine Balance: Internet Number Resource Distribution and De-Centralisation

Internet Society

Executive summary

Internet number resources (IP addresses and AS numbers) are distributed to resource users through processes that have evolved over time. Although initially centralised, the processes of policy formation and resource distribution have more recently been devolved to regional organisations. This de-centralisation is a direct consequence of the expansion of the Internet to cover all regions of the globe and it serves a number of important functions. De-centralisation is not an end in itself however, and experience shows us that a careful balance and coordination are needed to ensure that the over-riding objectives of aggregation, conservation and registration continue to be met.

The commitment between the resource distributors and the resource users is bi-directional, and resource distribution is essentially an operational engineering function that requires careful co-ordination and consensus building to succeed. Network operators have very strong incentives to partner with operationally knowledgeable organisations when obtaining numbering resources and will choose not to interconnect with networks that disregard this reality. As a consequence, proposals either to further centralise or de-centralise the processes whereby Internet numbering resources are distributed should be given very careful consideration indeed, with maintenance of the fine balance that has served the community well to date uppermost in our minds.

Introduction

Debate about the appropriate models for the management of Internet-related resources continues to be a feature of the policy landscape. Indeed, given the centrality of the Internet to modern life, such debates are unlikely to cease. This short paper addresses one group of Internet resources, namely number resources, and specifically addresses the question of whether more, or less, centralisation of the number resource distribution function is desirable. Arguments for both more centralisation and more de-centralisation are regularly heard in inter-governmental, business and non-governmental fora and this paper is intended to provide some condensed guidance based on the history of Internet number resource distribution to date. The intended audience for this paper is anyone interested in how Internet number resource distribution is managed today, how the current arrangements came into being, and the important considerations to be borne in mind when contemplating amendments to current practice.

Distributing Internet numbering resources isn't a 'fire-and-forget' operation. It requires considerable associated administrative machinery and ongoing maintenance to ensure the continued smooth operation of the network. This maintenance is critically informed by the experiences of those operating the network themselves. For this reason, the bottom-up, inclusive and consensus-driven processes that today's resource distributors have developed are essential to the successful development and implementation of ongoing resource allocation procedures. Those processes have evolved over time and it is informative to briefly review the historical pathway that brought us to this point.¹

Internet number resource distribution, a brief history

Internet number resources are Internet Protocol (IP) addresses and Autonomous System (AS) numbers. For a detailed explanation of the role of IP addresses, how they differ from phone numbers, and many other important aspects of Internet number resource management, the reader is referred to Stephen M. Ryan et al.'s excellent analysis [1]. Suffice to say here that IP addresses are used to number Internet hosts and AS numbers are used to number the discrete networks that when combined constitute the global Internet.

Initially, distribution of Internet number resources was handled by one man, Jon Postel.

The assignment of numbers is also handled by Jon. If you are developing a protocol or application that will require the use of a link, socket, port, protocol, or network number please contact Jon to receive a number assignment.

RFC790 [2]

¹ Daniel Karrenberg et al. present this history in considerably more detail in [16].

Subsequently, in 1987, the responsibility for the assignment of IP numbers and ASNs was assumed by the Hostmaster at the DDN Network Information Center (NIC) [3]. In 1991, the NIC transitioned to Government Systems Inc., who subcontracted the work to Network Solutions, Inc. [4]. Even before this however, it was recognised that decentralisation of the number resource distribution function was desirable.

With the rapid escalation of the number of networks in the Internet and its concurrent internationalization, it is timely to consider further delegation of assignment and registration authority on an international basis. It is also essential to take into consideration that such identifiers, particularly network identifiers of class A and B type, will become an increasingly scarce commodity whose allocation must be handled with thoughtful care.

RFC1174 [5]

A later document enabling this recommendation also noted that, 'The demand for network numbers has grown significantly within the last two years and as a result the allocation of network numbers must be approached in a more systematic fashion ... The major reason to distribute the registration function is that the Internet serves a more diverse global population than it did at its inception. This means that registries which are located in distinct geographic areas may be better able to serve the local community in terms of language and local customs.' [6]

This document also set out the initial criteria for organisations desirous of qualification as Internet number resource distributors, or 'regional registries':

It is important that the regional registry is unbiased and widely recognized by network providers and subscribers within the geographic region. It is also important that there is just a single regional registry per geographical region at this level to provide for efficient and fair sub-allocation of the address space. To be selected as a distributed regional registry an organization should meet the following criteria:

a) networking authorities within the geographic area legitimize the organization

b) the organization is well-established and has legitimacy outside of the registry function

c) the organization will commit appropriate resources to provide stable, timely, and reliable service to the geographic region

d) the commitment to allocate IP numbers according to the guidelines established by the IANA and the [Internet Registry (IR)] IR

e) the commitment to coordinate with the IR to establish qualifications and strategies for sub-allocations of the regional allocation.

Today there are five Regional Internet Registries (RIRs) with responsibility for Internet number resource distribution within their service region. They are as follows:

- AfriNIC – Africa
- APNIC – Asia-Pacific
- ARIN – North America & Caribbean (part)
- LACNIC – Central & South America & Caribbean (part)
- RIPE-NCC – Europe, Middle East & Central Asia

These five organisations co-ordinate their activities through the Number Resource Organisation (NRO) which, through an MoU with ICANN, fulfils the role of the Address Supporting Organisation (ASO) in advising ICANN on number resource matters and providing global policy for the allocation of number resources [7]. The Internet Assigned Numbers Authority (IANA) coordinates the global IP and AS number space, and allocates blocks of that space to the five RIRs in accordance with applicable policies ratified by the ICANN Board [8]. The IANA function is performed under the terms of a contract between the United States government and ICANN that is in effect until September 30, 2011. IANA is directed in matters regarding the technical parameters of numbering resources by the IETF [9]. These relationships are illustrated in Figure 1.

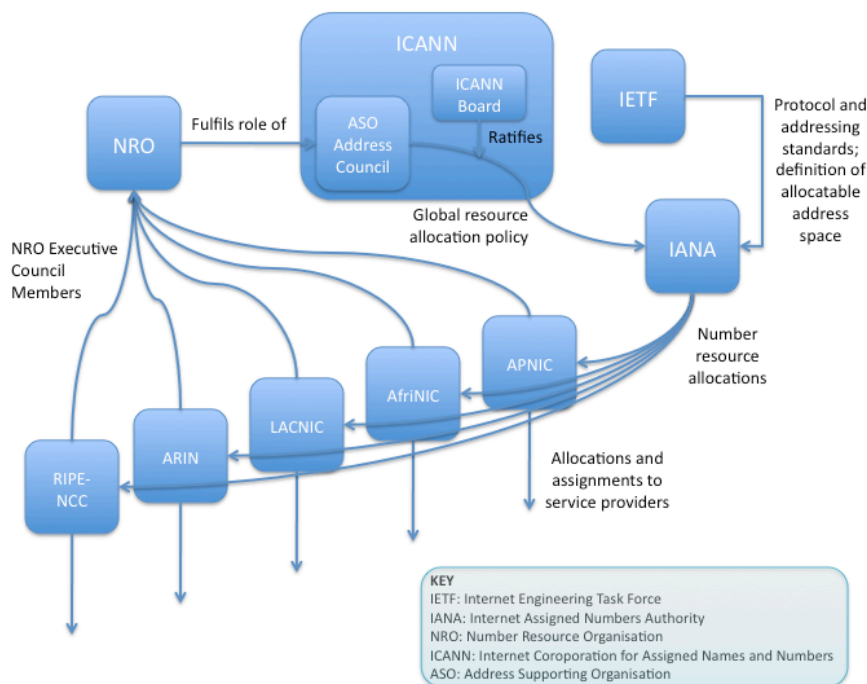


Figure 1 Number resource distribution organisational relationships

The brief history outlined above illustrates how the Internet number resource distribution function has evolved over time. We can now say that Internet number resource distribution is essentially an engineering function co-ordinated between network operators and other stakeholders under consensus agreements [10].

Internet number resource distribution requires careful stewardship

As described, the evolution of the present system was driven by a number of factors, namely:

- growth in the number of networks and consequent demand for IP address space;
- internationalisation of the internetwork leading to an increasingly diverse population of users, and;
- increasing scarcity of addresses leading to a need for careful allocation.

Building on these historical drivers, the three primary goals of the number resource distribution function are as follows:

- *Conservation*: to ensure efficient use of a finite resource and to avoid service instabilities due to market distortions (such as stockpiling or other forms of manipulation);
- *Aggregation*: to assist in maintenance of Internet routing tables at a manageable size, by supporting techniques to ensure continued operational stability of the Internet, and;
- *Registration*: to provide a public registry documenting address space allocations and assignments, necessary to ensure uniqueness and provide information for Internet troubleshooting at all levels [11].

Within the global framework, each regional community has always developed its own specific policies and procedures as it has always been recognized that these goals can be in conflict with each other and with the interests of individuals and organizations. It is also recognized that legitimate regional interests could justify varying approaches in balancing these conflicts. Nevertheless, the process by which IP addresses are provisioned to resource users can be stated generically as that presented in Figure 2 below.

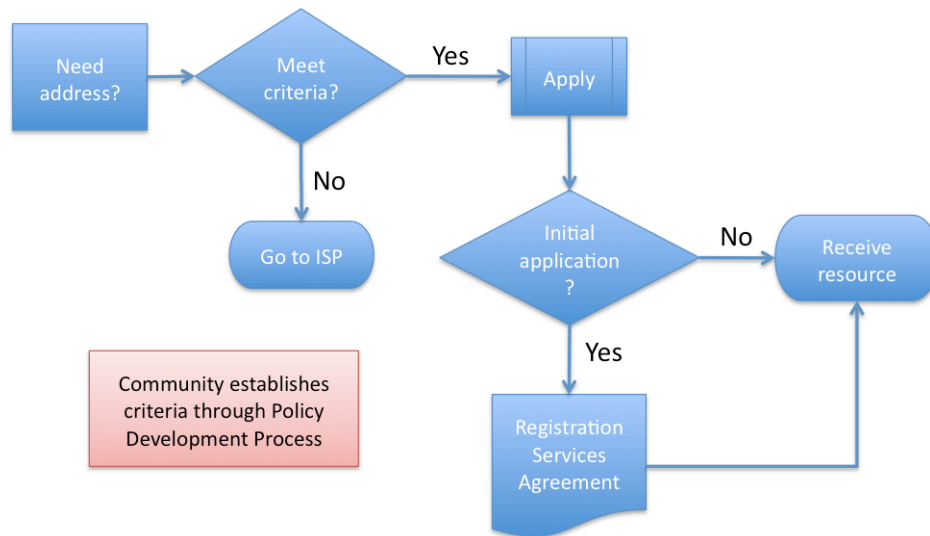


Figure 2 Generic IP address provisioning process²

However, whereas the specific approaches may differ across the regions, all regional distributors operate on a basic principle of open, transparent, consensus-based decision-making, following self-regulatory practices that exist elsewhere in the Internet and other industries. The Internet is a collaborative network-of-networks and, as such, necessitates these consensus-based processes in order to obtain maximal community support for the institutions and policies that bind them together. Furthermore, the resource registries all maintain not-for-profit cost-recovery systems and organizational structures that seek to include all interested stakeholders. In the same way that ‘given enough eyeballs, all bugs are shallow’ [12], the inclusive, bottom-up approach to policy formation adopted by the RIRs maximises the likelihood that adopted policies meet the needs of the community in the widest sense, and simultaneously minimises the likelihood of conflict given that we understand a priori the potential for conflict between the competing goals of the number resource distribution function.

The activities and services of each of the registries are defined, performed, discussed, and evaluated in open fora, whose participants are ultimately responsible for decision-making. Ongoing discussions are carried out on the public mailing lists of each registry, which are open to both the registry constituents and the broader community. This has the combined effect of ensuring that everyone’s views can be heard, and their opinions will be considered. It simultaneously ensures that the policy decisions that are made are widely supported by a significant cross-section of the community. This is vital to ensure relevant policies that effect real change on the network.

The staffs of the individual registries also participate actively in other Internet conferences and organizations and, importantly, each registry has a strong tradition of participating in the public activities of the others.

² Based on a slide in <http://www.nro.net/documents/presentations/RIR-GAC-Vancouver1.ppt>

De-centralisation: A fine balance

This last point is key. The current system works precisely because of the close co-ordination between the regional registries that meet regularly as peer institutions operating collectively to provide the global Internet number resource distribution function. Historically, as described above, it has been desirable to de-centralise the resource distribution function *to some extent* for several reasons:

- to improve scalability;
- to bring the function closer to the resource users;
- to ease the establishment of appropriate funding structures, and;
- to obtain greater support from the local community.

These important goals have motivated the establishment of the five regional registries that we have today. This history teaches us that centralisation of the resource distribution function in a single global organisation is undesirable and has several negative consequences. De-centralisation has been shown to be necessary. However, there are limits to this approach, and the history of National Internet Registries (NIRs) in the Asia-Pacific region is informative here.

At the time of APNIC's establishment, in 1993, several National NICs were established or emerging and these were incorporated into the initiative through the confederation or NIR membership structure. The benefits of this structure included service to local ISPs in the local language and timezone, and integration of additional services relevant to the local community.

Unfortunately, as time went on, the NIR structure of APNIC became problematic in certain respects. Each NIR received its own allocations, which they were able to manage according to local policies, but these policies could not be easily coordinated. This resulted in a situation in which IP address blocks became fragmented, with adverse impacts on ISPs and on the global Internet. After some years of operating in this mode, problems had increased to the extent that APNIC suspended the admission of new NIRs (in 1998).

Some years later (since 2002), new APNIC NIRs are being established again, but with certain specific conditions, which address the previous problems. First, an NIR is committed to follow regional and global policies, in order to avoid incompatible policies, which could conflict with those of other countries or networks. Further, in order to reduce fragmentation of address space, which also has global impact, an NIR does not receive its own block of addresses. The NIR is able to process and approve IP allocations, but those allocations are taken from the APNIC pool rather than from a separate national pool. This "shared address pool" model of regional address space management was introduced with the consensus of the APNIC community including the NIRs themselves, and is critical to the efficacy of APNIC's NIR system.

Source: Paul Wilson, *The Geography of Internet Addressing*, [13]

The point of this brief history is to show that, while there are good reasons for de-centralising the number resource distribution function on a regional basis, further de-centralisation has been shown to be counter-productive. Indeed, the existing regionalised system works well precisely because the registries maintain close links with each other and seek to co-operate, as peers, on matters of global importance.

Risks of further de-centralisation

Excessive de-centralisation of the number resource distribution function carries with it a number of risks.³ As the APNIC experience illustrates, policy confusion

³ Geoff Huston describes these risks in some detail in [17].

and dilution is inevitable if the number of organisations with responsibility for setting address policy is increased. This de-centralisation is also misaligned with widespread business models whereby many network operators have trans-national concerns – requiring them to deal with a multiplicity of resource distributors will increase costs for them, which will in turn lead to increased prices for consumers.

The interdomain routing system, a truly fundamental piece of the Internet’s machinery, crucially depends upon the ability to aggregate smaller networks into larger, and consequently fewer, routing announcements. This is what makes it possible to scale the network up to global proportions. Concerns are already being expressed about the scalability of today’s Internet routing system in the face of mounting levels of de-aggregation [14]. Limitless de-centralisation of the resource allocation function is in conflict with the need to maximise aggregation of announced routing prefixes. Increased strain on the routing system implies additional costs for operators, but also a less stable and less resilient network. Figure 3 illustrates the increased number of entries in the core routing table required when a single ISP is forced to obtain its addressing resources from multiple de-centralised sources. The existing RIR system has this effect today for ISPs that require address space for a global network footprint, but that isn’t the norm. Further de-centralisation of the resource distribution function would result in this cause of de-aggregation becoming much more widespread.

Even where registries reserve space in advance for ISP business growth (thereby allowing an ISP to return to the registrar for more space and to have a new block allocated that is contiguous with their existing allocation), this can only prevent de-aggregation within blocks allocated from a specific address registry (e.g. B in the figure) and cannot prevent de-aggregation across multiple address registries.

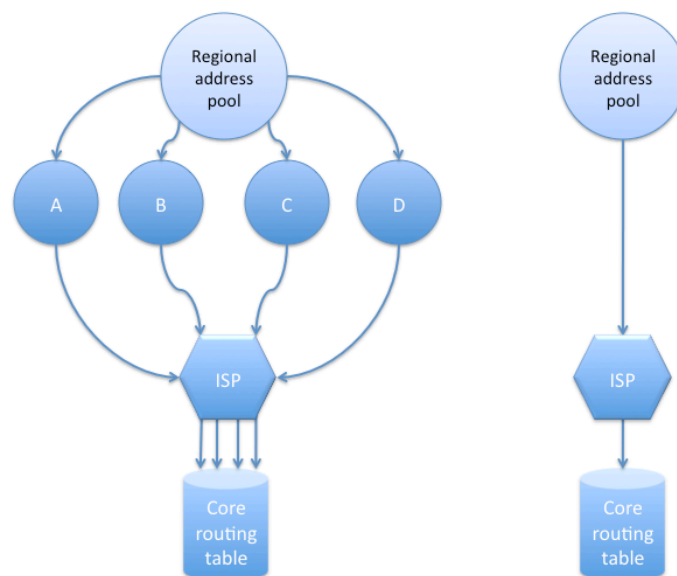


Figure 3 De-centralisation of resource distribution increases pressure on the routing function. On the left we show that where an ISP is required to obtain address space from four separate registries (A, B, C, and D) it results in four de-aggregated entries in the core routing table. On the right we show that allocation from a larger regional address pool prevents this de-aggregation.

De-centralisation of resource distribution also increases the likelihood that it will not be possible for the local resource registry to allocate contiguous space to repeat applicants as the registry may have exhausted its own allocation from the regional pool and will therefore only have discontinuous space available. This is true for IPv4 allocations today, and may also be a concern for IPv6 allocations as IPv6 deployment evolves. The point is illustrated in Figure 4, which also shows that resource distribution is a matter of ongoing relationships between resource consumers and resource distributors over time. It isn't a 'fire-and-forget' process.

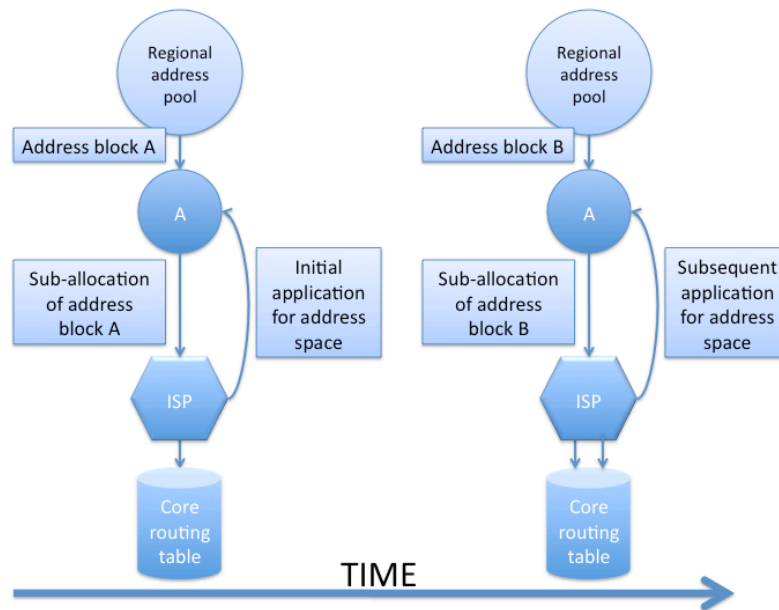


Figure 4 De-centralisation causes de-aggregation over time. On the left we show an ISP making an initial application for address space from the local registry. Over time, the registry exhausts its allocation from the regional pool and is allocated a new block (B). When the ISP makes a subsequent application it is unable to obtain space contiguous with the original allocation. This results in two de-aggregated entries in the core routing table.

An often-repeated mantra of resource allocation is that obtaining IP address resources does not guarantee routability on the network. The Internet is fundamentally consensus-based and co-operative in nature, which means that de-centralised resource distribution policy regimes that were viewed to be harming the good of the inter-network may very well find that their constituents were no longer able to obtain global routability. Or, in other words, without community support for the policies and processes adopted by the resource distributors, the resource itself becomes worthless.

Considerable effort is currently underway to address the lack of security for the Internet's interdomain routing infrastructure. The proposals currently being developed place considerable responsibility on the resource distribution function to operate reliable trust anchors. Using public-key cryptography, it is possible to provide a cryptographically verifiable source attesting to *who* is authorized to originate reachability information for *which* blocks of address space on the Internet. Similar to the way in which security is provided for some web transactions today, these trust anchors require very sophisticated and robust processes and IT machinery to deliver the reliability and security benefits required. Further de-centralisation is in conflict with the need to operate this

system reliably and effectively and is likely to erode the potential for improved Internet security in the medium term. Concerns for a secure, stable and reliable network are best addressed by continuing to provide for a well-coordinated, consensus-based resource distribution model. Indeed the introduction of any resource distribution function that was not fundamentally part of the existing global collaboration would severely undermine efforts to improve Internet security.

The open, bottom-up, consensus-driven processes that have served the Internet community so well to date are also put at risk if de-centralisation of resource allocation policy control is made to organisations with a restricted view of who is suitable to operate a network, and at what cost. This could seriously harm the networks and the entrepreneurs that built them in many developing regions of the global economy, for example [15].

Conclusions

Distributing Internet numbering resources isn't a 'fire-and-forget' operation. It requires considerable associated administrative machinery and ongoing maintenance to ensure the continued smooth operation of the network. This maintenance is critically informed by the experiences of those operating the network themselves. For this reason, the bottom-up, inclusive and consensus-driven processes that today's distributors have developed are essential to the successful development and implementation of ongoing resource allocation procedures.

The commitment between the resource distributors and the resource users is bi-directional, and resource distribution is essentially an operational engineering function that requires careful co-ordination and consensus building to succeed. Network operators have very strong incentives to partner with operationally knowledgeable organisations when obtaining numbering resources and will choose not to interconnect with networks that disregard this reality. As a consequence, proposals either to further centralise or de-centralise the processes whereby Internet numbering resources are distributed should be given very careful consideration indeed, with maintenance of the fine balance that has served the community well to date uppermost in our minds.

References

1. Stephen M. Ryan, et al., Legal and Policy Aspects of Internet Number Resources, <http://www.chtlj.org/authors/ryan>
2. RFC790, Assigned Numbers
3. RFC1020, Internet Numbers
4. RFC1261, Transition of NIC Services
5. RFC1174, IAB Recommended Policy on Distributing Internet Identifier Assignment and IAB Recommended Policy Change to Internet "Connected" Status

6. RFC1366, Guidelines for Management of IP Address Space
7. About the NRO, <http://www.nro.net/about/index.html>
8. Global Addressing Policies, <http://www.icann.org/en/general/global-addressing-policies.html>
9. IANA Number Resources, <http://www.iana.org/numbers/>
10. NRO response to ITU comments on the management of Internet Protocol (IP) addresses, <http://www.nro.org/documents/nro17.html>
11. RFC2050, Internet Registry IP Allocation Guidelines
12. Eric S. Raymond (1999), *The Cathedral & the Bazaar*
13. Paul Wilson, The Geography of Internet Addressing, http://www.circleid.com/posts/the_geography_of_internet_addressing
14. RFC4984, Report from the IAB Workshop on Routing and Addressing
15. Internet Governance Reform: Pitting African Internet Communities Against National Governments?, <http://www.cipesa.org/131>
16. Daniel Karrenberg, et al., Development of the Regional Internet Registry System, http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_4-4/regional_internet_registries.html
17. Geoff Huston, Possible responses to the ITU-T proposal regarding IPv6 address distribution, <http://www.ripe.net/info/internet-management/possible-responses.html>