Internet address and naming systems: Global harmonization or territoriality?

Mohsen Souissi

Abstract:
Underpinning the operation of the Internet's infrastructure are address and naming systems. The multiparty management of high-level Internet resources (IP addresses, the DNS root zone) has evolved technically, operationally and politically. This analysis of the evolution of the management processes for IP addresses and domain names draws attention to a combination of two approaches to managing these two categories of resources: the centrality of authority and the "territoriality" of the beneficiaries (and of late-comers). Despite the tensions surrounding the complex worldwide governance of these resources and the often polarized debates about a "centralized" versus a "federated" management, actors in the field have often, out of pragmatism, cleared a path between these two extreme positions for the purpose of coping with technical, operational and economic limitations.

The address and naming systems — Internet Protocol (IP) routing and the Domain Name System (DNS) — are essential to the operation of the Internet's infrastructure.[1] After a brief history of the initial procedures for assigning high-level Internet resources (IP address blocks as well as the DNS root zone and top-level domain names), this article describes the evolution of the techniques, procedures and policies having to do with the management of these resources — complex and often tense owing to the presence of diverse stakeholders and their sometimes diverging interests. When relevant, light will be shed on the Internet's geographical dimension with regard to the management of IP and DNS resources. In some cases, the pertinence of this dimension (whether decreed by rules or technical procedures, or introduced and imposed *de facto* by factors related to operations or business) will be discussed.

At the origin…

…Of IP addresses

According to the Internet Protocol (ISI/USC 1981), which was worked out at the end of the 1970s, a node in an Internet network has to have a single, unique IP address (IPv4) of a fixed length (4 bytes). This solution could satisfy a theoretical need for a little more than four billion addresses. An IP address is made up of two successive parts, the one identifying the network and the other a node in the network. An IP address is said to be both an identification and a location.

[1] This article has been translated from French by Noal Mellott (Omaha Beach, France). The translation into English has, with the editor's approval, completed a few bibliographical references. All websites were consulted in January 2021.

At the start, these addresses were assigned and managed hierarchically to the clear advantage of the pioneers of the Internet. Under IPv4, these first-comers benefitted from address blocks of a much more suitable size than did newcomers.[2] Since the Internet was born in the United States under the exclusive authority of the Department of Defense (DoD) but out of a program that involved American universities, the large majority of these initial beneficiaries were, logically enough, American.

So, we can say that geographical considerations were taken into account when allocating IP network numbers to the pioneers of the Internet for nearly ten years. Only the Information Sciences Institute (ISI) of the University of Southern California (USC) (in fact John Postel at the start) was, under the contract with the DoD, allowed to allocate these resources. This centralization in what was called the Internet Assigned Numbers Authority (IANA) operated in combination with the geographical dimension related to the location of beneficiaries. To illustrate the gradual extension of the Internet's geographical scope to other countries, in particular in western Europe, we need but observe the changes, over the years, under the heading "Assigned network numbers".[3] However the geographical scope did not have a wide reach till the early 1990s, when uses of the Internet exploded owing to the success of the World Wide Web. The IP addressing system, it was then feared, would soon be saturated. It was becoming more difficult to effectively manage the distribution of blocks of IP addresses through a central authority.

At this point, regional actors emerged, such as the RIPE Network Coordination Center (RIPE NCC, which, roughly speaking, covers Europe and western Asia), the core of what would later become the regional Internet registries (RIR).

…Of domain names

A few years after being launched, the size of the Internet was still small enough to continue manually mapping the names of the connected nodes and their IP addresses. This was done through the computer file hosts (named HOSTS.TXT), which the Stanford Research Institute (SRI) managed using very little automation. This mapping was done in close cooperation with John Postel's ISI team, which managed the IP addresses on IANA's list.[4] The HOSTS.TXT file had to be regularly downloaded to each Internet node so that the changes made would be taken into account. These updates started occurring more and more frequently — an engineering and operational challenge that forced Postel's team to look for a fitter solution.

This was the context in which Paul Mockapetris (1983a, 1983b) designed the DNS. The Network Working Group, which would later become the Internet Engineering Task Force, formally adopted this proposal as two Request For Comments in November 1983: RFC 882 and RFC 883 (later obsoleted by RFC 1034 and RFC 1035). The IETF went on to become the worldwide organization for setting the standards for Internet protocols.

[2] At the start, the IPv4 addresses were divided into five classes from A to E. A, B and C were to be used for "unicast" addresses: 2^{24} for class A, 2^{16} for B and 2^{8} for C. The class A prefixes, which were of a very limited number (fewer than 128), had to be allocated to the American pioneers, while the Europeans managed to obtain one or more class B prefixes. When many of the IPv4 addresses had been allocated and routing tables were becoming congested, the Internet Engineering Task Force (IETF) had to adopt emergency measures, such as Classless Inter-Domain Routing (CIDR) for "abolishing" classes, and the reservation of private IP address ranges (RFC 1918) (in contrast with public ones) once Network Address Translation (NAT) was introduced.

[3] In the various versions of the document "Assigned Numbers" published and managed by the ISI, which hosted "IANA functions". This document underlaid a series of "request for comments" (RFCs), which were released in one or two versions per year from 1981 till 1986. The series can be found by following the links "obsoleted by" on https://tools.ietf.org/html/rfc790.

[4] The SRI, which later became SRI International, was the first entity to have set up a network information center (NIC), a phrase that would be used for registries of domain names or IP addresses.

The DNS was a much more robust and scalable solution thanks to three key properties: it was hierarchical, distributed and redundant. Its author had not at all imagined that his invention would be such a success and last so long. Indeed, it is still in effect. The top level in this hierarchy is the DNS root zone, based on top-level domains (TLDs). The initial list (POSTEL & REYNOLDS 1984) of top-level domains included the so-called "general purpose domains" (gTLDs: .com, .edu, .gov, .mil and .org in addition to the very special .arpa).[5] There were also country-code top-level domains (ccTLDs), which have two letters (alpha-2 as required by standard ISO 3166).

Once again, we observe that the allocation of domain names has, since the start, followed an approach that combines a central authority with the geographical location of beneficiaries.

A dose of pluralism in the management of Internet resources

Toward the end of the 1990s, prior to the Internet's phenomenal, worldwide success, the National Telecommunications and Information Administration (NTIA) in the US Department of Commerce proposed a dose of pluralism for improving the centralized management of Internet resources (top-level names and addresses).[6] For the sake of efficiency, it was deemed wise to consult other stakeholders, namely private and international stakeholders who increasingly wanted to be involved in the process of management. In this context was born in 1998 the Internet Corporation for Assigned Names and Numbers (ICANN), a nonprofit organization under Californian law. Its first assignment was to manage IANA, in particular, high-level Internet resources such as IPv4 and IPv6 address blocks, AS number blocks, the DNS root zone and TLDs. ICANN also had the assignment of organizing, in 2000, the allocation of a new set of generic TLDs (among them: .biz, .info and .museum).

The advent of ICANN opened a new era of multistakeholder governance. The variety of participants had two facets: profiles (private and public sectors, governments, users, etc.) and geographical representativeness. This trend in Internet governance has gone through several cycles as a function of: the assessments made of ICANN's operations by the US administration, the results of negotiations between stakeholders and, above all, the fallout from geopolitical events, like the quake set off by the Snowden affair in the digital realm in 2013. Several stakeholders — most of them not American — profited from this affair to stomp on the accelerator for moving toward independence of the Internet's governing bodies from US authorities. These changes involved endless debates and negotiations, in particular during ICANN meetings or the annual Internet Governance Forum (IGF). One issue in these negotiations was the creation in 2016 of the Public Technical Identifiers (PTI), an ICANN subsidiary.[7] This new organization would take over IANA's duties.

These debates were often polarized. For the one side, the changes made did not provide a solution since the Internet's "real boss" was still the US Department of Commerce, which had the power to make all major decisions (*e.g.*, whether or not to validate technical proposals or to end contracts with ICANN at any time and reassume its control over the Internet). This side was calling for a clean slate in order to build a totally new Internet that would be multistakeholder from the start and freed from the "yoke" of the United States. In contrast, the other side advocated a pragmatic approach. Though not satisfied with the changes made (most of them taking place very slowly), it tried to be positive and favored a broader approach to this making of continuous (incremental) improvements. The aim was to effectively take account of the opinions of stakeholders who did not necessarily have sufficient representation on governing bodies. For some

[5] Initially planned to be temporary, .arpa ended up with a lasting function in Internet management as the "Address and Routing Parameter Area" with the subdomains in-addr.arpa and ip6.arpa (for mapping to domain names addresses from IPv4 and IPv6 respectively) and e164.arpa (for mapping to URIs in the context of ENUM).

[6] https://www.ntia.doc.gov/legacy/ntiahome/domainname/domainname130.htm

[7] https://pti.icann.org/

advocates of this approach, it was necessary to be wary both of those (the United States) whom they already knew but whose actions had been relatively "correct" (*i.e.*, without any flagrant abuse of authority) and of those who were not familiar (or not enough) but who might slow down management even more and be backed by dictatorial governments that wanted to censor free speech on line and turn the Internet into a tool for state control over people.

Global versus regional…

…In the management of IP addresses:

Parallel to ICANN's growing power, regional Internet registries (RIR) were being set up, five of them with the creation AFRINIC in 2005. These five RIRs correspond to the continents (cf. Figure 1). To summarize: an RIR receives an allocation from IANA of IP address blocks and, on the basis of the prefixes, allocates subprefixes to local Internet registries (LIRs, typically network operators in the same region). For instance, French network operators receive their allotments of IP prefixes from RIPE NCC. Some regions (*e.g.*, APNIC) might even have national Internet registries (NIRs).

Figure 1: The five regional Internet registries (RIR):
AFRINIC (African Network Information Center), APNIC (Asia-Pacific Network Information Center), ARIN (American Registry for Internet Numbers), LACNIC (Latin America and Caribbean Network Information Center) and RIPE Network Coordination Center (RIPE NCC).
Source: https://en.wikipedia.org/wiki/Regional_Internet_registry

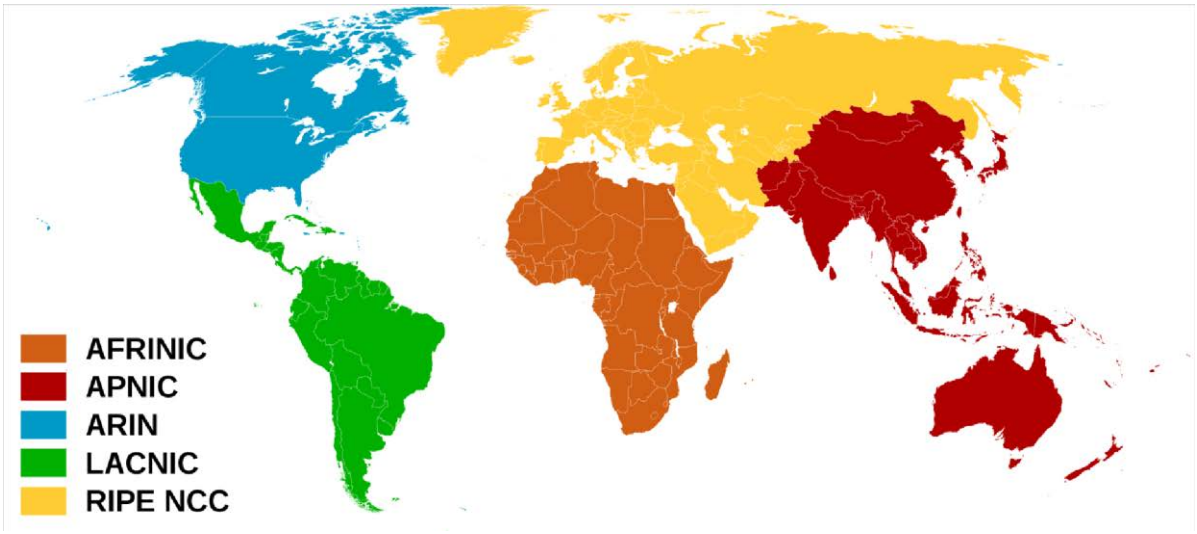

Le duties of these geographical RIRs have been increasingly formalized at three different levels:

- INTERACTIONS WITH THEIR REGIONAL INTERNET COMMUNITIES (mainly network operators and service-providers). This is how most RIRs have shored up their legitimacy. Most decisions about technical and operational changes are, it is worthwhile pointing out, made at this level. The intent might be to bring pressure to bear on ICANN or IANA. The RIR mantra has often amounted to: "We are applying technical procedures drafted and approved by our conformity members in a bottom-up, transparent process".

● INTERACTIONS WITH OTHER RIRS. Within their margin of autonomy, RIRs often consult each other to draft harmonized regional policies that take account of specific characteristics in each region.

● INTERACTIONS WITH ICANN/IANA. By forming a coalition called the Number Resource Organization (NRO), the RIRs have set up a formal organization that represents them as a group in the debates and negotiations engaged with ICANN and, too, with organizations such as the ITU Telecommunication Standardization Sector.

With the introduction of IPv6, IANA and the RIRs agreed to extend the procedures in effect for IPv4 while introducing improvements. The latter mainly came out of feedback from the management of IPv4 resources and had to do with specific aspects of IPv6. Although this amounted to an evolution instead of a revolution, a fresh start could thus be made in a "healthier" and "fairer" context for the introduction of IPv6. For IPv4, the "lucky" pioneers had received the lion's share of addresses while all the others — the much more numerous "latecomers" — had to be satisfied with the "crumbs". Given the announced abundance of IPv6 space, the new resources were, we can conclude, fairly distributed. However this remark cannot be extended to the rollout and operation of Internet infrastructures. For the latter, some pundits have talked about a wider, persistent digital divide.

So, from the end of the 1990s till the first years of the new century, IANA, the RIRs and their Internet communities were working on joint policies for the distribution (with technical support from the IETF) of IPv6 prefixes from the top (IANA) down to network operators (LIRs). However the procedures ensuing from these policies, once implemented, did not suit everyone. Some stakeholders saw them as being "bureaucratic", inefficient, unadapted to "real" needs, or even counterproductive (since they would not help to adopt IPv6). The increasing number of loopholes, here and there, made a review of these policies necessary. The arguments used in favor of this review often hinged on specifically regional factors or on the need to be "pragmatic" or "flexible" — arguments that the other RIRs would later adopt in hope of a new phase of harmonization.

To illustrate these cycles of harmonization/exceptions/reharmonization, let us take the example of "provider independent" (PI) IP prefixes. An organization with a PI prefix may use it for numbering its network in full independence from its Internet service-provider, whose role is restricted to connecting clients and routing their traffic. The majority of PI prefixes in IPv4 dated back to the period around 1990 — at a time when the RIRs did not yet exist and, above all, the allocation of IP address blocks was not yet done under contract. In the main, the organizations lucky enough to have IPv4 IP prefixes could keep them "for life" or else sell them speculatively on a thriving market given the growing scarcity of IPv4 address space.[8] Owing to this painful experience with IPv4 PI prefixes, some network operators were relieved when, at the end of the 1990s, they learned that the global IPv6 policy would make a break with this way of allotting PI prefixes. According to the IETF's first publications, this policy was to see to it that only LIRs would allot the IPv6 prefixes under the "provider-aggregatable address space", which had become the rule for IPv4.[9] Given the operational and technical difficulties (in particular the absence of a multihoming standard in IPv6), the American Registry for Internet Numbers (ARIN) became the first regional organization to loosen its policy in order to authorize the direct allotment of IPv6 PI prefixes to end sites.[10] Forced to respond to similar needs in their regional communities, the other RIRs had to loosen their policies too.

[8] https://www.ripe.net/manage-ips-and-asns/resource-transfers-and-mergers/brokers

[9] https://tools.ietf.org/html/rfc2073. The IETF's recommendations and opinions about allocating IPv6 prefixes had, in turn, to change to such a point that the IETF finally decided to obsolete its previous RFCs and replace them with new ones (3587 and 6177). The new RFCs tended toward the conclusion that none of all this was the IETF's business but, instead, a matter for IANA, the RIRs, LIRs and clients to take up. These were the parties who had to reach an agreement with each other.

[10] https://www.arin.net/vault/policy/proposals/2005_1.html

…In the management of domain names and the DNS

Let us briefly look backwards at the first "general purpose domains" (.com, .edu, .gov, .mil, and .org), which would eventually be called "generic top-level domains" (gTLDs). The .com and .org domains were soon extended for registration outside the United States. As much can be said about .net, which had been adopted a little later. However the three other TLDs (.edu, .gov and .mil) were reserved for organizations in the United States.[11] These three gTLDs were the precursors of what would later be called "sponsored top-level domains".

The TLDs with rules of eligibility based on a geographical criterion are the country and geographical TLDs (respectively, ccTLDs and geoTLDs). The latter emerged in 2009 with the creation of .cat and .asia. The latest cycle for creating gTLDs started in 2012. It has produced a new set of TLDs for cities (*e.g.*, .paris), regions (*e.g.*, .Bzh for Brittany) and continents (*e.g.*, .africa). John Postel decided to refer the question of the eligibility for a ccTLD to standards that another international organization (the International Organization for Standardization, ISO) was to make and manage. This well-advised decision proved its mettle as new nation-states formed while others dissolved (the .yu for Yugoslavia making way for the new ccTLDs: .rs for Serbia and me. for Montenegro).

The allotment of domain names follows a pattern close to what has been described for the RIRs. IANA records the TLD in a registry of domain names, which, in turn, delegates derived domain names (which also have to be registered) to eligible parties.

As a global regulatory authority, ICANN has formalized *ex post* (often through a contract) its relation with the gTLD registries, but it has no legal basis for exercising any authority over the ccTLD registries. A mere exchange of letters sufficed for the ccTLD registries, which existed prior to ICANN, to keep their powers for managing domain names and setting, as they saw fit, the geographical criteria of eligibility, which vary from country to country.

While remaining autonomous, most ccTLD registries have chosen to form regional associations for exchanges with peers about the problems they have to settle (benchmarks, cooperation, good practices, etc.). One of the more dynamic associations of this sort was founded in 1998: the Council of European National Top-Level Domain Registries (CENTR), which now has more than fifty members.

…In the DNS root zone?

Contrary to the ccTLDs, the DNS root zone has undergone major changes since the days when John Postel managed nearly everything by himself. With the coming of ICANN, the duties of IANA and its interfaces had to be more formally defined, especially with regard to the TLDs (on account of the creation of new gTLDs).

Faced with the gradual development of new forms of technology to be rolled out in the root zone, ICANN has often invoked its concern for security, stability and resilience to justify the multiplication (sometimes excessive) of consultations and tests prior to making changes in this zone. Among these new forms of technology: the Domain Name System Security Extensions (DNSSEC), "anycast" for addressing and routing, or internationalized domain names (IDNs), not to mention the problem of a massive delegation to thousands of new gTLDs.[12]

[11] There were a few exceptions, such as Cso.edu: https://www.sciencespo.fr/cso/.

[12] It should be pointed out that, in this context, the TLDs or RIRs often pushed ICANN to make necessary, even urgent changes by showing that there was, on their part, no longer any reason for waiting or them putting off. This was the case for IPv6 and then for DNSSEC in the root zone.

Figure 2: IPv6 enabled global sites
Source: https://root-servers.org/

As of 01/23/2021 10:40 a.m., the root server system consists of 1368 instances operated by the 12 independent root server operators.

The change with the most positive impact on resilience was for the operators of root servers to gradually adopt anycast, a trend set off by Internet Systems Consortium (which operates F-Root) following the denial-of-service attack against root name servers in 2002 (ISC 2003). There were thirteen of these servers named from A to M: ten of them located in the United States, while K was located in the United Kingdom, I in Sweden and M in Japan. The gradual rollout of hundreds of new instances of anycast on the five continents (Figure 2) put an end to the myth that DNS root servers were an American monopoly. Nevertheless, dozens of "instances on each continent" does not necessarily mean that this distribution is geographically balanced. By looking more closely at this map, we notice that geographical coverage is far from uniform. Despite the "*1368 instances operated by the 12 independent root server operators*" according to the caption, we are forced to admit that these instances are distributed there where the traffic is heaviest, in particular at Internet exchange points (IXPs).

References

ISC [Internet Systems Consortium] (2003) "Hierarchical Anycast for global service distribution" (Newmarket, NH: Internet Systems Consortium) available at https://www.isc.org/pubs/tn/isc-tn-2003-1.html.

ISI/USC [Information Sciences Institute, University of Southern California] (1981) "Internet Protocol: DARPA Internet Program Protocol Specification", September, RFC 791 prepared for Defense Advanced Research Projects Agency (DARPA, Arlington, VA), available via https://tools.ietf.org/html/rfc791.

NETWORK WORKING GROUP (1996) "RFC 1918: Address allocation for private internets", *RFC Editor*, February, available at https://www.rfc-editor.org/info/rfc1918.

MOCKAPETRIS P. (1983a) "RFC 882: Domain names: Concepts and facilities", *RFC Editor*, November, available at https://www.rfc-editor.org/info/rfc882 [obsoleted by RFC 1034].

MOCKAPETRIS P. (1983b) "RFC 883: Domain names: Implementation specification", *RFC Editor*, November, available at https://www.rfc-editor.org/info/rfc883 [obsoleted by RFC 1035].

POSTEL J. & REYNOLDS J. (1984) "RFC 920: Domain requirements", *RFC Editor*, October, available at https://www.rfc-editor.org/info/rfc920.