

The Geography of Internet Addressing

Paul Wilson
Director General, APNIC

Version: 1, 15 April 2005

The ITU-T has proposed a new system of country-based IP address allocations which aims to satisfy a natural demand for self-determination by countries; however, the proposal also stands to realign the Internet's frontiers onto national boundaries, with consequences which are explored here.

Internet geography

As we've often heard, the term Internet originated with the concept of a network of networks, and a vision that many previously distinct computer networks could be linked together and act as one. The success of that early vision is clear - we do indeed see the Internet as a single entity, and we even speak of the Internet's architecture as if there was one designer who laid out a plan and supervised its construction. But despite all appearances, the Internet landscape is indeed made up of many separate networks, run by many independent operators and service providers; and it has a structure that has emerged and evolved over time, more like a geography than an architecture.

If the Internet landscape has a geography, it is a geography based not on physical countries and territories, but on the interconnected networks of which it is comprised. The essential character of the Internet, namely the ability to transmit traffic between any pair of connected points, relies not only on this interconnection of its component networks, but on the consistent operation of those networks according to common standards and policies. Of particular importance is the existence of a single common addressing and routing scheme, which allows new networks to connect into the Internet and immediately share traffic with all others. It must be understood, however, that this essential characteristic is not an assured outcome of the Internet itself; rather it is the result of administrative and operational systems that work specifically to preserve it.

This paper will explore these issues, particularly in light of recent proposals to introduce new mechanisms for IP address management, a prospect which could, over time, substantially alter both the geography of the Internet, and its essential characteristics as a single cohesive network.

Internet nations

To communicate across the Internet, we don't use phone numbers with country-code prefixes, but rather IP numbers (Internet addresses) with network address prefixes. The prefix of an IP address block is similar in function to that of an international phone number, except that the "nation" it identifies is an Internet network which can be of any size and of any physical extent - global, regional, national, or local. While the phone network currently uses some 220 prefixes, and must distinguish between these when routing phone calls between countries, a typical Internet router

currently has some 170,000 allocated address prefixes in its global routing table, and must consider all of these for every individual data packet that is routed between its “nations”.

Adding new country-code prefixes to the telephone network is an infrequent and highly regulated process, whose engineering impact is limited to the relatively small number of dedicated international switches. In the case of the Internet, new networks are established freely, in an environment of competition which features few specific regulations and no intrinsic alignment with national boundaries. The addition of each new network “nation” requires an engineering change that must ripple across the fabric of the entire Internet, and into every one of the hundreds of thousands of Internet routers that carry a complete global routing table.

The interconnections between Internet networks are extremely dynamic, and changes to the global routing tables track the ebb and flow of Internet markets and business relationships, traffic engineering adjustments, and automated network repair mechanisms. While the Internet routing system - which allows this level of dynamism - is highly automated through the use of sophisticated network protocols, it is not a system that can grow indefinitely without bounds.

Transparent borders

It seems that the Internet is like a world with many territories and many borders. This is true, but unlike the borders between countries, the borders between networks on the Internet are easily crossed indeed the very nature of the Internet requires that every point on the network is exposed not just to its neighbours, but to every other point. This is intrinsic to the Internet's flexibility and utility as a network, but as we know from our ongoing experience of network abuse (spam, hacking etc), it also has a downside - namely that the actions of one user can adversely affect many others. In a related way, the actions of an ISP or group of ISPs can and do affect all others on the Internet, either productively or adversely.

Act locally, impact globally

In particular, every new network on the Internet adds at least one IP address prefix to the global routing tables; any ISP can add additional prefixes, in small or large numbers. Since the routers which hold those tables represent the switchboards of ISP networks, they must adjust to changes rapidly and stably in order to continue to exchange traffic efficiently with other networks. A router which is holding a table which is too large for its memory capacity, or which is attempting to process dynamic changes at a rate higher than its processor's capacity, will certainly work more slowly than it should. This alone will generate disruption in immediate neighbouring networks. Worse, an overloaded Internet router may be forced to ignore routing updates or entire routes, effectively disrupting or preventing communications with either a few, many, or all other networks.

The current system for IP address management is concerned, therefore, not only with fair distribution of addresses, but also with maintenance of IP address routability, for without the capability to route an address, the address is useless. There are a number of ways in which the address management system assists and maintains the Internet's routing system. Firstly, IP address distribution is "provider-based", meaning that addresses are allocated to the discrete IP networks which comprise the Internet, and which are able to maintain the aggregation of those address blocks. In addition, address management policies specifically aim to limit the addition of new routes to levels which are sustainable with current routing technology. They stipulate for instance that except under special circumstances, networks below a certain size cannot receive their own address

prefix allocations (instead, such networks are required to join an existing network and receive address space from that provider, coexisting within a single global routing entry). Policies also stipulate that ISPs should limit their fragmentation of address blocks, and limit their announcement of more specific address prefixes to the global routing tables. Such measures are generally effective in ensuring reasonable stability of today's Internet infrastructure, but it is important to understand that such policies are themselves dynamic and can be adapted as necessary to the changing Internet environment.

At an operational level, ISPs typically manage routing table growth by configuring their routers to ignore certain classes of prefixes (such as those for very small networks), and thereby maintain efficient operations. However, in a scenario where the number of routes to larger networks increases dramatically, for instance through mismanagement of address space by registries or ISPs, many providers would have to implement far stricter policies. These measures would inevitably result in loss of connectivity between some existing networks, but if implemented widely, they would result in widespread loss of global Internet connectivity, particularly affecting smaller and more remote networks and users (those networks that are unable to employ the latest high capacity router technologies, and who are perhaps less likely to represent commercial priorities for larger providers). If we ever reach a point of routing crisis in the Internet, it will be the smaller and more isolated networks which first experience the impact of selective isolation.

Experiments in geographical address management

In the early days of IP address management, until some time in the early 1990s, it was commonly assumed that the Internet's geography would follow that of the physical world. In some cases, large address blocks were set aside for entire countries, and in some of these cases, organisations were formed within those countries to manage that address space (often these were called NICs or Network Information Centres). Early examples of these were JPNIC in Japan and AUNIC in Australia, and by the mid-nineties, several national NICs were formed.

At the same time, the ongoing growth of the Internet was forcing other changes in our approach to address management. The increasing workload experienced by the InterNIC, the global address registry, combined with the need for more careful address management, prompted a call to regionalise the address management task. By 1993, new Regional Internet address Registries (RIRs) had been formed in Europe and the Asia Pacific. The growth of transnational ISPs meant that many larger players lost interest in national registries, so that by the late nineties few new national NICs were being formed, while some were even disbanded.

Regional Internet address registries

Since their establishment, the RIRs have become the sole mechanism for distribution of IP address space to their users, namely ISPs and network providers, throughout the world. Today, 5 RIRs are in operation: AfriNIC serving the African continent, APNIC for the Asia Pacific, ARIN for North America, LACNIC for Latin America, and RIPE NCC for Europe. All of these operate as independent and neutral non-profit organisations, based on an industry self-regulatory model in which open and transparent, bottom-up processes are used to consider the inputs of all stakeholders in the formulation of address management policies.

National IP address management – the APNIC experience

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. At the same time, several of these existing organisations, most notably JPNIC, supported and contributed greatly to the establishment of APNIC.

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.

What about IPv6?

It is important to note that for the purposes of this discussion, the IPv4 and IPv6 addressing systems behave identically. There is no solution offered by IPv6 to the issue of fragmentation or routing table growth, so it is to be expected that routing tables in an IPv6 Internet would be of a similar size to today's tables. On the other hand, the much larger size of IPv6 address space appears to provide the great danger of an explosion in routing table sizes, particularly if allocation mechanisms are introduced which conflict with today’s measures for the control of table sizes.

The ITU proposal for national allocations

The recent ITU proposal that countries should receive and manage separate IPv6 allocations carries a certain risk in this respect. Apart from imposing a potential cost and obligation on every country to establish an agency to manage this resource, certain technical risks would be created which have global implications. The possibility of even a small number of different IP address policy regimes, let alone the potential for some 200 different policy regimes, could certainly produce negative effects not before seen on the Internet. Excessive consumption and subdivision of address space under such policies could result in very large numbers of additional address prefixes within the IPv6 routing tables, which would need to be carried by every ISP on the Internet. Carriage of such routes would impose performance and cost impacts that many ISPs could not afford, while address space which is dropped from routing tables is effectively unreachable by some or all of the Internet, generating an obvious impact by selectively isolating network users from each other.

One response to this problem of excessive fragmentation in the routing space could be to contemplate further national regulatory intervention. A country may need to establish not only a management system for address space, but also support specific shared infrastructure for carriage and management of Internet traffic at the national level (for instance by way of national Internet gateways and aggregation points), as well as inter-provider settlement schemes which have been difficult if not impossible to establish within the Internet context. Another possible outcome is the prospect of a gradual degrading of the Internet as a single cohesive global network, into one in which specific agreements are made by every pair of networks that wishes to exchange traffic. While such a system works in the scale of the global telephone network with some hundreds of providers, it must be remembered that a full set of bilateral agreements among the tens of thousands network operators would require hundreds of millions of such agreements. Clearly this is not a universal solution, and it is a more likely outcome that smaller network providers will be driven out of the market by a small set of larger multi-national providers.

It is clear that addressing systems lie at the very heart of networks, and that there is a close relationship between the address system, the services that a network can offer, and the nature of the business structures that support the deployment and operation of the network. Placing an inappropriate or badly attuned address system into an existing network model risks not only disruption and burgeoning cost overheads, but ultimately the destruction of the cost value of the network and its very reason for existence. The substantial cost and potential risks of such changes must surely be well justified by the real benefits that are offered.

Conclusion

The structure of today's Internet is a geography of independent networks around the world, with transparent borders allowing traffic to flow freely between any pair of locations. While there are cases of inequality in terms of inter-network arrangements for funding certain network connections, there is an overall equality implied by the ability of all networks, once connected, to exchange traffic as peers.

Many have claimed that the Internet's new paradigms will force a restructure of society, even threaten the nation-state. This is proving far from correct, but there are certainly aspects of the Internet which do not sit well with the traditional view of world geography. This in itself does create challenges, however, in responding to these we must take approaches which recognize the nature of the Internet we have today, and ensure that essential characteristics are preserved. If our approach is wrong, the end result could be a new form of digital divide, in which the erstwhile global uniformity of the Internet is shared only by privileged countries and companies, while others are left in a dramatically poorer situation.

Further Reading

1. For more discussion of the technical aspects of Internet address management and aggregation, see Geoff Huston's paper IP Addressing Schemes - A Comparison of Geographic and Provider-based IP Address Schemes at:

<http://www.potaroo.net/ispcol/2004-12>

2. The ITU-T proposal on IP address management is contained within Houlin Zhaos paper ITU and Internet Governance at:

<http://www.itu.int/ITU-T/tsb-director/itut-wsis>

3. The Regional Internet Registries have websites available as follows:

AfriNIC: <http://www.afrinic.net>

APNIC: <http://www.apnic.net>

ARIN: <http://www.arin.net>

LACNIC: <http://www.lacnic.net>

RIPE NCC: <http://www.ripe.net>

Thanks

Thanks to Geoff Huston for his valuable contributions to this article.

The Author

Paul Wilson has served as Director-General of APNIC since August 1998. Previously, he was Chief Executive Officer of Pegasus Networks, the first private ISP in Australia (est 1989). Over this period he also worked as a consultant to the United Nations and other international agencies on Internet projects in many countries. Since 1994, he has worked with the International Development Research Centre (IDRC) on its Pan-Asia Networking (PAN) Programme, supporting Internet development projects in Mongolia, Vietnam, Cambodia, Maldives, Nepal, Bhutan, PNG, and China.

Please send any comments or questions to the author, by email to: dg@apnic.net