NRO response to LS 5 - Country Internet Registry (CIR) model

Overview

This document responds to ITU IPv6 Group Liaison Statement, LS 5 - Country Internet Registry (CIR) model. This Liaison Statement refers to document TD 3 'Concerns on Ipv4 Address Policy with regard to Ipv6 Deployment'—hereafter "NAv6 paper"—a study commissioned by the ITU and performed by research students at NAv6, a research group in Malaysia, in which the authors propose "the creation of a parallel structure to the RIRs [Regional Internet Registries] for the allocation and distribution of IPv6 addresses".

The five Regional Internet Registries who form the Number Resource Organization (NRO)[1] would like to thank the Chairman of ITU IPv6 Group, Mr Mohammed Al Khamis, for the opportunity to share their views on the NAv6 paper with the members of the group.

The NAv6 paper proposes a Country Internet Registry (CIR) model that, it claims:

- Together with the RIR model, can potentially value add to the creation of a more fairly balanced IPv6 address allocation model (p. 25)
- Provides an alternative for users and a possible appeal process (p. 30)
- Facilitates in bridging the digital divide among developed and developing nations by more efficiently handling the management of IPv6 addresses (p. 13)
- Brings down the charges currently incurred by the applicants (p. 30)
- Does not introduce any additional costs to the ISPs and those managing the Internet (p. 16)

Furthermore, the NAv6 study suggests the ITU perform a role in distributing IPv6 address blocks to CIRs. These CIRs then would allocate the addresses based on locally defined policies.

Keeping down costs, bridging the digital divide, and providing an appropriate range of alternatives to meet the needs of networks, are goals shared by all Internet stakeholders, including AfriNIC, APNIC, ARIN, LACNIC, and the RIPE NCC. However, the NAv6 paper does not demonstrate a causal connection between its proposed system and the claimed benefits, let alone, in most cases, any possible relationship between the two.

Below you will find further elaboration of the views on the NAv6 paper from the five RIRs who together form the NRO.

[1] The five RIRs are AfriNIC, APNIC, ARIN, LACNIC, and the RIPE NCC. Four of the RIRs who form the NRO are ITU sector members: AfriNIC, APNIC, ARIN, and the RIPE NCC.

Part 1: Today's Internet registry system

This section provides a background to today's system of Internet address management—including Regional Internet Registries (RIRs)—as it has been developed and refined by the Internet community over the past (nearly) two decades. We stress firstly that this system has supported the development of the Internet through the entire period of its commercialization and successful development as a global critical infrastructure.

RIR Criteria: ICP-2

"ICP-2: Criteria for Establishment of New Regional Internet Registries"[2], based on the earlier Address Supporting Organization (ASO) 1990 "Memorandum of Understanding (MoU)"[3], is the institutional basis that governs how ICANN can recognize new RIRs. Amongst the criteria for recognition of RIRs is the following:

> "Each region should be served by a single RIR, established under one management and in one location. The establishment of multiple RIRs in one region is likely to lead to:
>
> - fragmentation of address space allocated to the region;
> - difficulty for co-ordination and co-operation between the RIRs;
> - confusion for the community within the region.
>
> The internal administrative or membership structure of an RIR must also not be such as to cause any of these effects."

Some of the other criteria in ICP-2 include that the prospective RIR:

- Be an independent, not-for-profit and open membership association
- Serve a region not served by any other RIR
- Be technically capable of providing the required allocation and registration services to the community in its region
- Have clearly documented procedures for the development of regional and global resource management policies. These procedures must be open and transparent, be accessible to all interested parties, and ensure fair representation of all constituencies within the region

Needs-based delegations

Needs-based delegation of IP addresses has been practiced since the inception of IP address allocations, has been refined under the RIR system, and permeates the entire addressing system, both in relation to IANA to RIR allocations as well as RIR to network allocations. The community that uses addresses developed this model as the fairest system of distribution; networks that need addresses can easily qualify for addresses, while the risk of hoarding is reduced. The needs-based addressing model has

[2] 2001, http://www.icann.org/en/icp/icp-2.htm

[3] 1999, Section 9, http://aso.icann.org/documents/memorandum-of-understanding-1999/

underpinned the growth of the Internet. Definitions of what is a qualifying "need" have been adjusted by community consensus via the open and transparent Policy Development Processes (PDPs) of the RIRs over time to meet the needs of a changing Internet environment and business models. The open, community-based policy development process provides a forum for consideration of the tradeoffs between address availability, registration requirements, and the routing implications for each Internet address policy proposal.

Policy development process

The development of policies for the management of Internet number resources—and in particular, the specific rules by which IP addresses are distributed, registered and administered)—is undertaken on a regional basis, in regular open forum meetings which are open to all interested participants. These regional "Policy development processes" are formally defined, and constitute a key self-regulatory aspect of the management of Internet number resources.

In all regions, policy is developed according to the following common principles:

- **Accessible and Open.** The process is open to everyone. Participation is not contingent upon membership or any other status. The process is conducted in public policy forums and on public electronic mail lists;
- **Transparent**. All electronic mail is archived and is publicly available to anyone. The minutes of all public policy forums and the meetings of the RIR Executive Boards are publicly available to anyone;
- **Documented**. All policies are formally documented. All procedures used to implement Internet Number Resource management policy are formally documented. These documents are publicly available to anyone at no charge;
- **"Bottom-up".** Policies are developed to reflect:
    - The evolving needs of the operators and users of Internet services;
    - Changes in technology.

Part 2: An analysis of the NAv6 document

IPv6 network management needs

The NAv6 paper states that:

> *"The CIR's will allocate IPv6 addresses prudently to the LIR/ISP and end-user sites based on their needs and requirements. The address prefix allocation and assignment range is as given in section 3.1.2"* (p. 22).

The paper also includes the following table:

> *"Table D-1. IPv6 Address delegation recommendations*" (p. 48).

The proposed IPv6 address block sizes given in Table D-1 would constitute a return to the "classful" addressing system that was used in IPv4 prior to the introduction of CIDR (Classless Inter Domain Routing)[4] and the RIR system. It is precisely this approach to address management that resulted in the rapid early depletion of the IPv4 address pools. The outcome would be that subsequent assignments to end users would then be fragmented and contrary to the primary goal the NAv6 paper states that it seeks to preserve: aggregation (p. 23).

Allocation requirements

The NAv6 paper proposes a reservation system for CIRs:

> *"The allocation scheme would facilitate to allocate the IPv6 address space more efficiently and fairly among nations. The address allocation scheme can use the following criteria,*
>
> *a) The size of population*
> *b) Growth rate in terms of utilization of the IPv6 address space*
> c) *Business and organizational growth"* (p. 10).

The paper uses the term "allocation" for these reservations. It is important to note that each block given to a CIR cannot be considered to be "allocated". "Allocation" occurs when an RIR provides addresses to a network. The pool that an RIR or CIR holds, prior to delegating those resources to a network, cannot be considered "allocated".

The CIR model proposes IANA allocates to the ITU, with contiguous country-based blocks reserved for each CIR. However, the NAv6 paper states it would be possible to determine the IPv6 needs of a country for the foreseeable future, including its potential for growth, using a formula based on population, IPv6 utilization growth rate, and business and organizational growth. Current practices to determine suitable allocation sizes use the number of prospective connected devices, rather than business growth, as the parameter for measuring need. Strong business growth may, or may not, be directly correlated with the growth of the number of devices needing Internet connections.

[4] http://tools.ietf.org/html/rfc4632

The NAv6 paper states:

> *"The CIRs will take responsibility in allocating contiguous IP addresses to users when subsequent allocations are made from the distinct address space designated for a country. This may lead to better management of IPv6 addresses resulting in better aggregation and conservation if done properly. In conclusion to aggregation, the RIRs have allocated addresses fairly well till date. However, there is always room for improvement by both the RIRs and CIRs if newer and better algorithms are used"* (p. 20).

Under the model proposed by NAv6, if the originally reserved block size is inappropriate, the ITU will need to request multiple /12s over time to meet the growing needs of a country. This would have the effect of making it impossible for users to receive contiguous allocations from CIRs that have been allocated more than one block from the ITU. Such fragmentation of large IPv6 networks would be contrary to the NAv6 paper's stated primary goal of preserving aggregation (p. 23).

Additionally, we note that in case of multi-national ISPs, the CIR system would introduce a level of national-level fragmentation of current and future address space holdings, and a rapid escalation of peering and routing complexity as a result. Such providers, instead of being able to negotiate global network-to-network routing with their peers (as they can now), would be required to make all such arrangements on a national level, within each country of operation. Not only would this impose additional and artificial overheads in the process of peering negotiation and management, it would also result in the immediate expansion of Internet routing complexity, and a consequent performance penalty on the entire routing infrastructure of the Internet. Because an ISP's routing capacity is directly proportional to its infrastructure and financial capacities, the detrimental effect of any such routing expansion will be felt more acutely in developing parts of the Internet, thus raising barriers to entry and competitiveness for small ISPs and those in developing countries.

Sources used to form CIR model

On the cover page of the NAv6 paper is the following disclaimer:

> *"Disclaimer: Some of the thoughts and opinions invested in this study are elements of the research made by students towards their research studies at NAv6. The views here reflect the views of the authors at the time of writing this article. Even standards change over time, so may the views expressed here."*

Given the global impact of the Internet in today's economies, using a student research project to propose such fundamental modifications to an operational system of critical infrastructure would appear to be ill-advised.

The paper also makes predictions about the growth of the IPv6 routing table:

> *"In our observation, as a comparison, the size of the IPv6 BGP routing table growth rate in terms of fragmentation would be less when compared with IPv4. As the size of the IPv6 address (128 bits) is bigger than IPv4 (32 bits), the minimum size of the prefix allocated to the ISP would be large. Eventually these route prefixes added to the more specific prefixes would result in lengthier routing tables*" (p. 57).

This claim, based on the "observation" of a very limited data set and at a very early stage in the practical deployment of IPv6, makes a critical prediction about future fragmentation trends, but is supported by no further explanation, or discussion of its mathematical models.

The paper refers to methods that could be used to change Internet addressing and routing architecture:

> *"Solutions to this end are defined in the direction of location identifier split methods that includes, Mike O Dells GSE proposal, 8+8 addressing architecture, LISP (Locator/Identifier Separation Protocol, ILNP (Identifier, locator network protocol), HIP (Host Identity Protocol), FARA (Forwarding directive, Association and Rendezvous Architecture, and ISLAY – A new routing & addressing architecture. All these models tend towards defining a new addressing architecture with changes to the existing Internet architecture itself"* (p. 51).

However, this list of methods includes a number of outdated and abandoned models that have not been accepted or implemented by the Internet technical community.

ISLAY is a 2002 Internet Draft that never gained community endorsement.[5] FARA was proposed in 2003 at ACM SIGCOMM, but has not been accepted by the Internet technical community.[6] 8+8 is a 1996 Internet Draft that, like ISLAY, has not gained community endorsement.[7] HIP is an RFC, but an experimental protocol.[8] LISP is actively under development, but still experimental.[9] ILNP has become an Internet Draft in 2011, but as yet has not reached a level of community acceptance needed to make it an RFC.[10]

The NAv6 paper appears to recognize real challenges in Internet routing, and to some extent recognize that its proposal stands to exacerbate some of these challenges. However, the alternative "solutions" that it presents by way of reassurance that such effects can be mitigated, are in fact outdated, abandoned or unproven solutions.

The "alternative CIR model" is a duplication of the existing RIR model

The NAv6 paper states:

> *"The CIR model could facilitate in the formation of regional agencies or Internet Registries at a smaller regional level grouping together CIRs. For example, North African Countries can group together and form a North Africa Internet Registry (NAIR)"* (p. 23).

Within the CIR model, the ITU is proposed to act as an RIR, which reserves space for CIRs, which in turn group together to form additional RIRs. If such regional groupings of

---

[5] http://bgp.potaroo.net/ietf/idref/draft-irtf-routing-islay/
[6] http://portal.acm.org/citation.cfm?id=944770&CFID=9044077&CFTOKEN=54797370
[7] http://potaroo.net/ietf/all-ids/draft-odell-8+8-00.txt
[8] http://tools.ietf.org/html/rfc5201
[9] https://datatracker.ietf.org/wg/lisp/charter/
[10] http://tools.ietf.org/id/draft-rja-ilnp-intro-10.txt

CIRs into RIRs are needed, particularly in addition to the existing five RIRs, it suggests that the use of ITU as an RIR is not meeting the needs that this model proposes to solve.

In addition, the paper states:

> *"The users of CIR model would include ISPs and organizations. The Internet being a public infrastructure has become a source of revenue and also a system for transferring funds. Economic models developed by ITU would help in bringing fair revenues to all parties involved in utilizing the Internet infrastructure"* (p. 24).

The existing five RIRs also serve ISPs and organizations. The proposed CIR model would serve to duplicate the role of the existing RIR system, with added expenses associated with creating and maintaining close to 200 additional registry organizations. It also increases risks for fragmentation of the global and regional policy framework into a great many more policy systems, and the introduction of many additional policies that may not be accepted by the global Internet networks.

Underestimating the cost of Internet Registry operations

The NAv6 paper states:

> *"The CIR model would work with the existing network infrastructure provided by the ISPs and demands no additional hardware or software capabilities. Thus from an economic feasibility viewpoint, the introduction of the CIR model will not introduce any additional costs to the ISPs and those managing the Internet"* (p. 16).

> *"Cost Structures: The cost structures of the way the CIRs charge their LIRs/ISPs for the IPv6 addresses must be competitive. It must, in general, be equivalent or less than the current price structure that is provided by the RIRs. Only then will users find this to be a viable alternative, as most users are price sensitive"* (p. 28).

The NAv6 paper's proposition that the CIR model could provide cheaper services is based on uninformed assumptions about the costs of registry operations and the potential for cost savings within a competitive system. As noted in ICP-2, there are many services associated with address allocation that Internet registries must perform. The NAv6 does not take such services into account.

The cost to RIR organizations of supporting their policy development processes is relatively high. This is an operational cost that does not seem to have been considered when developing the CIR model and it is one that would be magnified very significantly by the need to coordinate policy development among a large number of countries.

The costs of running many small CIRs, including duplication of roles and software and hardware in around 200 locations, will be, in aggregate, higher than the cost of running five RIRs and NIRs (National Internet Registries) created through community consensus.

Furthermore, since the NAv6 paper itself states that it is a technical study, it appears contradictory for it to make any assertions about the economic viability of the CIR model.

The paper also states:

> *"Similar to the existing RIR model a "Whois Database" would be made available"* (p. 21).

Registration of allocations and assignments, whois database management, and associated functions such as reverse DNS are complex operations. The CIR model proposed in the NAv6 paper does not explain which organizations would maintain whois registries. If it were to be the ITU, please note that there will be significant complexities involved in creating a compatible whois system to allow almost 200 CIRs to register ISP allocations and for ISPs under all those CIRs to be able to add their customer assignment details to the database. It must also be noted that reverse DNS is a critical 24-hour function associated with IP address management and needs to be staffed accordingly and technically provisioned to high levels.

The paper also suggests that:

> *"Most small CIRs can be a 3-6 person operation, once a fully automated web system is in place for managing the address resources"* (p. 28).

This in conflict with the paper's suggestion that "Proper and detailed evaluation process, so that the address space is not wasted or depleted." It also underestimates the need to provide 24 hour assistance for services like reverse DNS, liaison and coordination on policy and other matters with RIRs, with the ITU, and with other CIRs.

Working together in the spirit of Resolution 180 (Guadalajara, 2010)

Thank you for allowing ITU sector members, AfriNIC, APNIC, ARIN and the RIPE NCC, to contribute these perspectives on behalf of the Number Resource Organization, representing all five of the Regional Internet Registries. We hope that they will be useful during the proceedings of the meeting.