Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 4, 2009

A. Durand
Comcast
M. Ford
P. Roberts
Internet Society

March 3, 2009

Issues with ISP Responses to IPv4 Address Exhaustion
draft-ford-shared-addressing-issues-00

Status of this Memo

Copyright Notice

Abstract

The looming completion of IPv4 address allocations from IANA and the

RIRs is already causing ISPs around the world to start to question how they will continue providing IPv4 service to IPv4-speaking customers when there are no longer sufficient IPv4 addresses to allocate them one per customer. Several possible solutions to this problem are now emerging and this memo identifies important criteria to be borne in mind when evaluating these solutions. We also seek to identify serious issues that remain even when mechanisms meeting our criteria are adopted. We wish to stress that these solutions have a number of common, and potentially serious, issues.

Table of Contents

1. Introduction

Allocations of IPv4 addresses from the Internet Assigned Numbers Authority (IANA) are currently forecast to be complete during the first half of 2011 [IPv4 Report]. Allocations from the Regional Internet Registries (RIRs) are anticipated to be complete around a year later, although the exact date will vary from registry to registry. This is already causing ISPs around the world to start to question how they will continue providing IPv4 service to IPv4-speaking customers when there are no longer sufficient IPv4 addresses to allocate them one per customer. Several possible solutions to this problem are now emerging and the following discussion identifies important criteria to be borne in mind when evaluating the merits and demerits of these solutions. These criteria are derived from the development and acceptance of a modern understanding and consistent implementation of the end-to-end principle of the Internet.

This memo is divided into two principal sections. The first deals with what we consider to be guiding principles that it is important to bear in mind when evaluating address-sharing proposals. The second section is concerned with identifying and discussing the operational implications of adopting any address-sharing mechanism. We wish to stress that these solutions have a number of common, and potentially serious, issues. Address sharing amongst multiple subscribers will inevitably result in a degraded experience of the network for many users, and increased operating costs for ISPs. Content providers are encouraged to consider carefully the potential impact of shared-addressing on their business and operational practices.

2. Guiding principles

"The end-to-end principle is the core architectural guideline of the Internet." Section 2 of RFC 3724 [RFC3724] provides a concise history of the end-to-end principle. While the original articulation

was concerned with where best to place functionality in a communication system, the growth and development of the Internet has resulted in an expansion of the scope of the end-to-end principle so that it now encompasses the question of where best to locate the state associated with Internet applications. This expanded principle is well-articulated in RFC 1958 [RFC1958],

> "An end-to-end protocol design should not rely on the maintenance of state (i.e., information about the state of the end-to-end communication) inside the network. Such state should be maintained only in the endpoints, in such a way that the state can only be destroyed when the endpoint

> itself breaks (known as fate-sharing)."

The end-to-end principle is arguably the fundamental principle of the Internet architecture. In a sense the Internet is the embodiment of the principle. By allowing either tacit or explicit erosion of the principle as we apply our understanding to the construction and operation of the global network, we allow the dismantling of the utility itself. Unfortunately the approaches being proposed to address the looming IPv4 address shortage threaten just such erosion.

2.1. IPv6 is the goal

While we are discussing solutions to allow continued operation of the IPv4 Internet and the continued provision of services to IPv4-speaking customers, it is absolutely not the intention of this discussion to in any way advocate the prolongation of the life of IPv4 or to (further) delay the widespread adoption of IPv6. Rather, the discussion herein is intended to guide reactions to proposals that are being made as a pragmatic response to some very real problems looming for operators who need to be able to continue providing service to customers who do not have any IPv6 capable equipment and/or who want to access services that are only available via IPv4.

2.2. Criteria for judging ISP responses

On that basis, we believe that solutions to the problem of continuing to provide IPv4 service post-IPv4-address-completion should be judged on two primary criteria:

1) The proximity of the end user to control over the impact of the solution on the end-to-end communication, and;

2) The extent to which the solution affords a natural progression to widespread IPv6 deployment.

2.3. Potential responses

Assuming ISPs wish to continue growing their businesses post-IPv4-address-completion there are a number of possible responses they can take:

- Obtain previously-allocated IPv4 addresses through some unspecified means

- Deploy CGN and allocate customers with private addresses

- Deploy a shared-address solution - customers get public addresses with fewer available ports

Let's deal with each of these in turn.

2.3.1. Obtaining previously-allocated addresses

Acquisition of previously-allocated IPv4 addresses by whatever means is a strategy with currently unknown (but definitely limited) viability. It is also impossible to estimate in advance the cost of such an approach, so it does nothing to minimise business risk. Acquiring previously-allocated addresses may provide a short-term tactical solution where a relatively small number of addresses are required urgently to address a specific need. It cannot be a solution for long-term network business growth. It is likely that previously-allocated blocks acquired by whatever means will be small and obtaining lots of contiguous small blocks may be impossible. This would inevitably lead to operational complexity and associated cost for the network taking this approach. It is operationally unsustainable in anything but the short term.

2.3.2. Deploy CGN, allocate private addresses

In light of the two criteria for judging solutions to the address allocation completion problem that we identified above, so-called 'carrier-grade' NAT (CGN) proposals [I-D.nishitani-cgn] raise several issues. Centralisation of NAT functionality in the network core will reduce the ability of end-users to deploy applications as they wish without reference to the network operator. This means that unadorned CGN solutions will struggle to meet the first criterion. Providing mechanisms for end-users to control their treatment by the CGN may go some way to mitigate this concern, however those mechanisms would need to be very carefully engineered to avoid raising additional scalability and resilience concerns of their own. CGNs may create a single point of failure for all their clients and decrease the resilience of the network from an end-user's perspective. CGN implementations may also struggle when considering our second criterion as there is no requirement to make use of IPv6 technology as part of the solution. For these reasons there is a real risk that CGNs will do nothing to progress the state of IPv6 deployment and will only serve to degrade the utility of the current network.

While the subject of CGN deployment has arisen recently in the context of IPv4 address depletion, some operators, particularly mobile network operators, have a long history of allocating private addresses to their subscribers. Recent discussions have indicated

that the increasing sophistication of both mobile handsets and the applications that run on them is driving operators of mobile networks towards public addressing solutions, including IPv6 deployment, to improve scalability and minimise operating expenses. This suggests that those operators with real-world experience of CGN technology are already choosing to migrate away from it as a solution to their addressing needs.

2.3.3. Shared address solutions

How could we do better? There are proposals currently in the IETF that address one or both of the criteria we identify as critical. These alternative proposals are simplified by using IPv6 as a transport substrate for the legacy traffic [I-D.durand-softwire-dual-stack-lite], thereby motivating IPv6

deployment, and may also ensure that control over the fate of end-user applications is kept as close to the end-user as possible by distributing the NAT functionality towards the CPE [I-D.ymbk-aplusp].

However, some reduction of utility for IPv4-speaking Internet users is unavoidable in the future. It is inevitable that a reduced number of ports will be available for individual end-user applications. Running servers on well-known ports will most probably be an activity that is restricted to users willing to pay a premium for a higher tier of service contract. These may turn out to be good incentives for end-users to migrate to IPv6.

The remainder of this memo deals with issues that arise when shared address solutions are deployed by ISPs.

3. Issues with shared address solutions

A number of proposals currently under consideration for standardization or contribution to some future standard rely upon the concept of address sharing across multiple subscribers to achieve their goals. These proposals include Carrier Grade NAT [I-D.nishitani-cgn], Dual-Stack-Lite [I-D.durand-softwire-dual-stack-lite], NAT64 [I-D.bagnulo-behave-nat64], IVI [I-D.baker-behave-ivi], Address+Port proposals [I-D.ymbk-aplusp], and SAM [I-D.despres-sam].

In many operator networks today a subscriber receives a single public IPv4 address at their home or small business. Within that home or small business there is a NAT function that issues private addresses (RFC1918 addresses) to devices within the home. All of those devices share the single public IPv4 address and they are all associated with a single small set of users, and a single operator subscriber

account.

With these new proposals a single public IPv4 address would be shared by a number of homes or small businesses (i.e. multiple subscribers) so the operational paradigm described above would no longer apply.

All the proposals listed above share a number of technical or operational issues and these are addressed in the subsections that

follow.

3.1. Port Distribution, Port Reservation, Port Negotiation

When we talk about port numbers we need to make a distinction between outgoing connections and incoming connections. For outgoing connections, the actual source port number used is usually irrelevant. But for incoming connections, the specific port numbers allocated to customers matter because they are part of external referrals (used by third parties to contact services run by the customers). It is desirable to make sure those incoming ports remain stable over time. This is challenging as the network doesn't know anything in particular about the applications which it is supporting and therefore has no real notion of how long an application/service session is still ongoing and therefore requiring port stability.

According to actual measurements the average number of outgoing ports per customer is much, much smaller than the maximum number of ports a customer can use at any given time. However, the distribution is heavy-tailed, so there are typically a small number of subscribers who use a very high number of ports [CGN_Viability]. This means that an algorithm that dynamically allocates outgoing port numbers from a central pool is much more efficient than algorithms that statically divide the resource by pre-allocating a fixed number of ports to each subscriber. Similarly, such an algorithm should be more able to accommodate users wishing to use a relatively high number of ports.

Early measurements also seem to indicate that, on average, only very few ports are used by customers for incoming connections. However, a majority of subscribers accept at least one inbound connection.

This means that it is not necessary to pre-allocate a large number of ports to each subscriber, but that it is possible to either pre-allocate a small number of ports for incoming connections or do port allocation on demand when the application wishing to receive a connection is initiated and reserve the bulk of ports as a centralized resource shared by all subscribers using a given public IPv4 address.

A potential problem with this approach occurs when one of the

subscriber devices behind such a port-shared IPv4 address becomes

infected with a worm, which then quickly sets about opening many outbound connections in order to propagate itself. Such an infection could rapidly exhaust the shared resource of the single IPv4 address for all connected subscribers. Poor network hygiene of one subscriber now threatens the connectivity for all immediate network neighbors.

3.2. Connection to a Well-Known Port Number

Once a port address-mapping scheme is in place, connections to well-known port numbers will not work in the general case. Some workaround (e.g. redirects to a port-specific URL) could always be deployed given sufficient incentives. There exist several proposals for 'application service location' protocols which would provide a means of addressing this problem, but historically these proposals have not gained much deployment traction.

3.3. Universal Plug and Play

Using the UPnP semantic, a client is asking "I want to use port number X, is that ok?" and the answer is yes or no. If the answer is no, the client will typically try the next port, until it either finds one that works or gives up after a limited number of attempts. UPnP has, currently, no way to redirect the client to use another port number.

NAT-PMP has a better semantic here, enabling the NAT to redirect the client to an available port number.

3.4. Security and Subscriber Identification with IPv4

When an abuse is reported today, it is usually done in the form: IPv4 address X has done something bad at time T0. This is not enough information to uniquely identify the subscriber responsible for the abuse when that IPv4 address is shared by more than one subscriber. This particular issue can be fixed by logging port numbers.

A number of application servers on the network today log IPv4 addresses in connection attempts to protect themselves from certain attacks. For example, if a server sees too many login attempts from the same IPv4 address, it may decide to put that address in a penalty box for a certain time. If an IPv4 address is shared by multiple subscribers, this would have unintended consequences in a couple of ways. First it may become the natural behavior to see many login attempts from the same address because it is now shared across a potentially large number of users. Second and more likely is that one user who fails a number of login attempts may block out other

users who have not made any previous attempts but who will now fail on their first attempt.

Moreover, the assumption that a single IPv4 address maps to a single user may be used for other purposes like geolocation, counting the number of individual users of a service, etc. All those things may become more complicated when an IPv4 address is shared by several subscribers at the same time.

To some extent these problems of shared addressing are already with us due to the prevalence of dyamically assigned addresses to domestic broadband subscribers and the use of CPE NAT, but the point we wish to make here is that the widespread adoption of port-shared addresses by service providers will make these complications considerably more widespread and severe.

4. Concluding remarks

Of the various options that are now available to service providers as we approach the completion of IPv4 address allocations from the IANA, there are some shared-address solutions that seem to offer an approach consistent with a long-term goal of IPv6 deployment and maximal preservation of the end-to-end principle. Nevertheless, it must be stressed that these solutions have a number of common, and potentially serious, issues. Address sharing amongst multiple subscribers will inevitably result in a degraded experience of the network for many users, and increased operating costs for ISPs. Content providers are encouraged to consider carefully the potential impact of shared-addressing on their business and operational practices.

5. Acknowledgements

This memo was largely inspired by conversations that took place as part of an Internet Society hosted roundtable event for operators deploying IPv6. Participants in that discussion included John Brzozowski, Leslie Daigle, Tom Klieber, Yiu Lee, Kurtis Lindqvist, Wes George, and Christian Jacquenet.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

Section 3.4 discusses some of the security and identity-related implications of address sharing.

8. Informative References

[CGN_Viability]
Alcock, S., "Research into the Viability of Service-Provider NAT", 2008.

[I-D.bagnulo-behave-nat64]
Bagnulo, M., Matthews, P., and I. Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", draft-bagnulo-behave-nat64-02 (work in progress), November 2008.

[I-D.baker-behave-ivi]
Li, X., Bao, C., Baker, F., and K. Yin, "IVI Update to SIIT and NAT-PT", draft-baker-behave-ivi-01 (work in progress), September 2008.

[I-D.despres-sam]
Despres, R., "Stateless Address Mappings (SAMs) IPv6 & extended IPv4 via local routing domains - possibly multihomed", draft-despres-sam-01 (work in progress), November 2008.

[I-D.durand-softwire-dual-stack-lite]
Durand, A., Droms, R., Haberman, B., and J. Woodyatt, "Dual-stack lite broadband deployments post IPv4 exhaustion", draft-durand-softwire-dual-stack-lite-01 (work in progress), November 2008.

[I-D.nishitani-cgn]
Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Functions of Large Scale NAT (LSN)", draft-nishitani-cgn-01 (work in progress), November 2008.

[I-D.ymbk-aplusp]
Maennel, O., Bush, R., Cittadini, L., and S. Bellovin, "The A+P Approach to the IPv4 Address Shortage", draft-ymbk-aplusp-02 (work in progress), January 2009.

[IPv4_Report]
Huston, G., "IPv4 Address Report", 2009, <http://www.potaroo.net/tools/ipv4/index.html>.

[RFC1958] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, June 1996.

[RFC3724] Kempf, J., Austein, R., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, March 2004.

Authors' Addresses

Alain Durand
Comcast

Email: Alain_Durand@cable.comcast.com

Mat Ford
Internet Society
Geneva
Switzerland

Email: ford@isoc.org

Phil Roberts
Internet Society
Reston, VA
USA

Email: roberts@isoc.org