TOC

Internet Engineering Task Force	M. Ford, Ed.
Internet-Draft	Internet Society
Intended status: Informational	M. Boucadair
Expires: April 29, 2010	France Telecom
	A. Durand
	Comcast
	P. Levis
	France Telecom
	P. Roberts
	Internet Society
	October 26, 2009

Issues with IP Address Sharing
draft-ford-shared-addressing-issues-01

Status of this Memo

Copyright Notice

Abstract

The completion of IPv4 address allocations from IANA and the RIRs is causing service providers around the world to question how they will continue providing IPv4 connectivity service to their subscribers when there are no longer sufficient IPv4 addresses to allocate them one per subscriber. Several possible solutions to this problem are now emerging based around the idea of shared IPv4 addressing. These solutions give rise to a number of issues and this memo attempts to identify those common to all such address sharing approaches. Solution specific discussions are out of scope.

Table of Contents

1. Introduction
2. Shared Addressing Solutions
3. Address Space Multiplicative Factor
4. Port Allocation
    4.1. Outgoing Ports
    4.2. Incoming Ports
        4.2.1. Port Negotiation
        4.2.2. Connection to a Well-Known Port Number
5. Impact on Applications
6. ICMP
7. Fragmentation
8. Support of Multicast
9. Mobile-IP
10. Introduction of Single Points of Failure
11. Security
    11.1. Port Randomisation
    11.2. Abuse Logging and Penalty Boxes
    11.3. Spam
    11.4. IPsec
    11.5. Policing Forwarding Behaviour
12. Geo-location and Geo-proximity
13. Authentication
14. Traceability
15. IPv6 Transition Issues
16. IANA Considerations
17. Security Considerations
18. Contributors
19. Acknowledgements
20. Informative References
§ Authors' Addresses

1. Introduction

TOC

Allocations of IPv4 addresses from the Internet Assigned Numbers Authority (IANA) are currently forecast to be complete during 2011 [IPv4_Report] (Huston, G., “IPv4 Address Report,” 2009.). Allocations from some 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 causing service providers around the world to start to question how they will continue providing IPv4 connectivity service to their subscribers when there are no longer sufficient IPv4 addresses to allocate them one per subscriber. Several possible solutions to this problem are now emerging based around the idea of shared IPv4 addressing. These solutions give rise to a number of issues and this memo attempts to identify those common to all such address sharing approaches. Over the long term, deploying IPv6 is the only way to ease pressure on the public IPv4 address pool and thereby mitigate the need for address sharing mechanisms that give rise to the issues identified herein. In the short term, maintaining growth of IPv4 services in the presence of IPv4 address depletion will require address sharing. Address sharing will cause issues for end-users, service providers and third parties such as law enforcement agencies and content providers. This memo is intended to highlight these issues.

In the presence of continued network growth, and in the absence of very widespread dual-stack deployment, increased IP address sharing is inevitable. A restricted type of IPv4 connectivity service is going to operate in parallel with the existing IPv4 Internet of today. This restricted Internet service isn't going to be the same as existing services - some applications aren't going to work and third-parties will also be impacted.

Increased IPv6 deployment should reduce the burden being placed on an address-sharing solution, and should reduce the costs of operating that solution. Increasing IPv6 deployment should cause a reduction in the number of concurrent IPv4 sessions per subscriber. If the percentage of end-to-end IPv6 traffic significantly increases, so that the volume of IPv4 traffic begins decreasing, then the number of IPv4 sessions will decrease. The smaller the number of concurrent IPv4 sessions per subscriber, the higher the number of subscribers able to share the same IPv4 public address, and consequently, the lower the number of IPv4 public addresses required. However, this effect will only occur for subscribers who have both an IPv6 access and a shared IPv4 access. This motivates a strategy to systematically bind a shared IPv4 access to an IPv6 access. It is difficult to foresee to what extent growing IPv6 traffic will reduce the number of concurrent IPv4 sessions, but in any event, IPv6 deployment and use should reduce the pressure on the available public IPv4 address pool.

2. Shared Addressing Solutions

TOC

In many networks today a subscriber is provided with a single public IPv4 address at their home or small business. For instance, in fixed broadband access, an IPv4 public address is assigned to each CPE (Customer Premises Equipment). CPEs embed a NAT function which is responsible for translating private IPv4 addresses ([RFC1918] (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, “Address Allocation for Private Internets,” February 1996.) addresses) assigned to hosts within the local network, to the public IPv4 address assigned by the service provider (and vice versa). Therefore, devices located with the LAN share the single public IPv4 address and they are all associated with a single small set of users, and a single subscriber account with a single network operator.

A number of proposals currently under consideration in the IETF rely upon the mechanism of multiplexing multiple subscribers' connections over a smaller number of shared IPv4 addresses. These proposals include Carrier Grade NAT [I-D.nishitani-cgn] (Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, “Common requirements for IP address sharing schemes,” March 2010.) , Dual-Stack-Lite [I-D.ietf-softwire-dual-stack-lite] (Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, “Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion,” March 2010.) , NAT64 [I-D.ietf-behave-v6v4-xlate-stateful] (Bagnulo, M., Matthews, P., and I. Beijnum, “Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers,” March 2010.) , IVI [I-D.ietf-behave-v6v4-xlate] (Li, X., Bao, C., and F. Baker, “IP/ICMP Translation Algorithm,” April 2010.) , Address+Port (A+P) proposals [I-D.ymbk-aplusp] (Bush, R., “The A+P Approach to the IPv4 Address Shortage,” October 2009.) , [I-D.boucadair-port-range] (Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, “IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture,” July 2009.) and SAM [I-D.despres-sam] (Despres, R., “Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM),” July 2009.) .

In these new proposals, a single public IPv4 address would be shared by multiple homes or small businesses (i.e. multiple subscribers) so the operational paradigm described above would no longer apply. All these proposals extend the address space by adding port information, they differ in the way they manage the port value.

IP address sharing solutions fall into two classes. Either a centralised, service-provider operated NAT function is introduced and subscribers are allocated addresses from [RFC1918] (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, “Address Allocation for Private Internets,” February 1996.) space, or public IPv4 addresses are shared across multiple subscribers by restricting the range of ports available to each subscriber. These classes of solution are described in a bit more detail below.

*CGN-based solutions: These solutions propose the introduction of a NAPT function in the service provider's network, denoted also as Carrier Grade NAT (CGN), or Large Scale NAT (LSN) [I-D.nishitani-cgn] (Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, “Common requirements for IP address sharing schemes,” March 2010.) , or Provider NAT. The CGN is responsible for translating private addresses to publicly routable addresses. Private addresses are assigned to subscribers, a pool of public addresses is assigned to the CGN, and the number of public addresses is smaller than the number of subscribers. A public IPv4 address in the CGN pool is shared by several subscribers at the same time. Solutions making use of a service provider-based NAT include [I-D.shirasaki-nat444] (Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444,” March 2010.) (two layers of NAT) and [I-D.ietf-softwire-dual-stack-lite] (Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, “Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion,” March 2010.) (a single layer of NAT).

*Port-range solutions: These solutions avoid the presence of a CGN function. A single public IPv4 address is assigned to several subscribers at the same time. A restricted port range is also assigned to each subscriber so that two subscribers with the same IPv4 address have two different port ranges that do not overlap. These solutions are called A+P (Address+Port) [I-D.ymbk-aplusp] (Bush, R., “The A+P Approach to the IPv4 Address Shortage,” October 2009.) , or Port Range [I-D.boucadair-port-range] (Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, “IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture,” July 2009.) , or SAM (Stateless Address Mapping) [I-D.despres-sam] (Despres, R., “Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM),” July 2009.) .

Security issues associated with NAT have long been documented (see [RFC2663] (Srisuresh, P. and M. Holdrege, “IP Network Address Translator (NAT) Terminology and Considerations,” August 1999.) and [RFC2993] (Hain, T., “Architectural Implications of NAT,” November 2000.)). However, sharing IPv4 addresses across multiple subscribers by any means, either moving the NAT functionality from the home gateway to the core of the service provider network, or restricting the port choice in the subscriber's NAT, creates additional issues for subscribers, content providers and network operators. All the proposals listed above share technical and operational issues and these are addressed in the sections that follow. These issues are common to any service-provider NAT, enterprise NAT, and also non-NAT solutions that share individual IPv4 addresses across multiple subscribers (e.g. A+P).

---

3. Address Space Multiplicative Factor

TOC

The purpose of sharing public IPv4 addresses is to increase the addressing space. A key parameter is the factor by which service providers want or need to multiply their IPv4 public address space; and the consequence is the number of subscribers sharing the same public IPv4 address. We refer to this parameter as the address space multiplicative factor, the inverse is called the compression ratio. The multiplicative factor can only be applied to the subset of subscribers that are eligible for a shared address. The reasons a subscriber cannot have a shared address can be:

- It would not be compatible with the service they are currently subscribed to (for example: business subscriber).

- Subscriber CPE is not compatible with the address sharing solution selected by the service provider (for example it does not handle port restriction for port-range solutions or it does not allow IPv4 in IPv6 encapsulation for the DS-lite solution), and its replacement is not easy.

Different service providers may have very different needs. A long-lived service provider, whose number of subscribers is rather stable, may have an existing address pool that will only need a small extension to cope with the next few years, assuming that this address pool can be re-purposed for an address-sharing solution (small multiplicative factor, less than 10). A new entrant or a new line of business will need a much bigger multiplicative factor (e.g. 1000). A mobile operator may see its addressing needs grow dramatically as the IP-enabled mobile handset market grows.
When the multiplicative factor is large, the average number of ports per subscriber is small. Given the large measured disparity between average and peak port consumption [CGN_Viability] (Alcock, S., “Research into the Viability of Service-Provider NAT,” 2008.) , this will create service problems in the event that ports are allocated statically. In this case, it is essential for port allocation to map to need as closely as possible, and to avoid allocating ports for longer than necessary. Therefore, the larger the multiplicative factor, the more dynamic the port assignment has to be.

TOC

4. Port Allocation

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. (While this is true today, in a port-range solution it is necessary for the source port to be within the allocated range). But for incoming connections, the specific port numbers allocated to subscribers matter because they are part of external referrals (used by third parties to contact services run by the subscribers).

The total number of subscribers able to share a single IPv4 address will depend upon assumptions about the average number of ports required per active subscriber, and the average number of simultaneously active subscribers.

Most of the time the source port selected by a client application will be translated (unless there is direct knowledge of a port-range restriction in the client's stack), either by a NAT in the subscriber's device, or by a CPE NAT, or by a CPE NAT and a CGN.

IANA has classified the whole port space into three categories (as defined in http://www.iana.org/assignments/port-numbers):

- The Well Known Ports are those from 0 through 1023.
- The Registered Ports are those from 1024 through 49151.
- The Dynamic and/or Private Ports are those from 49152 through 65535.

[RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.) notes that current NATs have different policies with regard to this classification; some NATs restrict their translations to the use of dynamic ports, some also include registered ports, some preserve the port if it is in the well-known range. [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.) makes it clear that the use of port space (1024-65535) is safe: "mapping a source port to a source port that is already registered is unlikely to have any bad effects". Therefore, for all address sharing solutions, there is no reason to only consider a subset of the port space (1024-65535) for outgoing source ports. In any case, limiting the number of ports available will limit the compression ratio.

TOC

4.1. Outgoing Ports

According to measurements the average number of outgoing ports consumed per active subscriber is much, much smaller than the maximum number of ports a subscriber 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] (Alcock, S., “Research into the Viability of Service-Provider NAT,” 2008.) . This means that an algorithm that dynamically allocates outgoing port numbers from a central pool will typically allow more subscribers to share a single IPv4 address 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 subscribers wishing to use a relatively high number of ports.

It is important to note here that the desire to dynamically allocate outgoing port numbers will make a service provider's job of maintaining records of subscriber port number allocations considerably more onerous (see Section 14 (Traceability)). The number of records per subscriber will increase from 1 in a scheme where ports are statically allocated, to a much larger number equivalent to the total number of outgoing ports consumed by that subscriber during the time period for which detailed logs must be kept.

A potential problem with dynamic allocation 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. It is therefore necessary to impose limits on the total number of ports available to an individual subscriber to ensure that the shared resource (the IPv4 address) remains available in some capacity to all the subscribers using it.

4.2. Incoming Ports

TOC

It is desirable to ensure that incoming ports remain stable over time. This is challenging as the network doesn't know anything in particular about the applications that it is supporting and therefore has no real notion of how long an application/service session is still ongoing and therefore requiring port stability.

Early measurements [CGN_Viability] (Alcock, S., “Research into the Viability of Service-Provider NAT,” 2008.) also seem to indicate that, on average, only very few ports are used by subscribers 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 incoming ports to each subscriber. 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. The bulk of incoming ports can be reserved as a centralized resource shared by all subscribers using a given public IPv4 address.

4.2.1. Port Negotiation

TOC

In current deployments, one important and widely used feature of many CPE devices is the ability to open incoming ports (port forwarding) either manually, or with a protocol such as UPnP IGD. If a CGN is present, the port must also be open in the CGN. The situation may be alleviated somewhat if the CGN architecture is composed of only one NAT level (no NAT in the CPE) as for DS-lite, although a service provider operating this solution will still be required to offer some means for configuring of incoming ports by their subscribers. This may be either via a UPnP or NAT-PMP relay over a tunnelled direct connection between CPE and CGN or a web interface to configure the incoming port on the CGN. Note, that such an interface effectively makes public what was previously a private service interface and this may raise security concerns.

For port-range solutions, port forwarding capabilities may still be present at the CPE, with the limitation that the open incoming port must be within the allocated port-range (for instance it is not possible to open port 5002 for incoming connections if port 5002 is not within the allocated port-range).

4.2.1.1. Universal Plug and Play (UPnP)

TOC

Using the UPnP semantic, an application asks "I want to use port number X, is that ok?" and the answer is yes or no. If the answer is no, the application will typically try the next port in sequence, until it either finds one that works or gives up after a limited number of attempts. UPnP has, currently, no way to redirect the application to use another port number. UPnP IGD 2.0, currently being defined, should improve this and allow for allocation of any available port.

4.2.1.2. NAT Port Mapping Protocol (NAT-PMP)

TOC

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

4.2.2. Connection to a Well-Known Port Number

TOC

Once an IPv4 address sharing mechanism is in place, connections to well-known port numbers will not work in the general case. Any application that is not port-agile cannot be expected to work. Some workaround (e.g. redirects to a port-specific URI) 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.
For example, the use of the DNS SRV records [RFC2782] (Gulbrandsen, A., Vixie, P., and L. Esibov, “A DNS RR for specifying the location of services (DNS SRV),” February 2000.) provides a potential solution for subscribers wishing to host services in the presence of a shared-addressing scheme. SRV records make it possible to specify a port value related to a service, thereby making services accessible on ports other than the Well-Known ports (e.g. a web server accessible on a port other than port 80).

5. Impact on Applications

TOC

Address sharing solutions will have an impact on the following types of applications:

- Applications that establish inbound communications - these applications will have to ensure that ports selected for inbound communications are either within the allocated range (for port-range solutions) or are forwarded appropriately by the CGN (for CGN-based solutions). See Section 4.2 (Incoming Ports) for more discussion of this;

- Applications that carry address and/or port information in their payload - where translation of port and/or address information is performed at the IP and transport layers by the address-sharing solution, an ALG will also be required to ensure application layer data is appropriately modified;

- Applications that use fixed ports (e.g. well-known ports) - see Section 4.2.2 (Connection to a Well-Known Port Number) for more discussion of this;

- Applications that do not use any port (e.g. ICMP) - where address sharing solutions map subscribers to (private) IP addresses on a one-to-one basis this will not be an issue, otherwise such applications will require special handling - see Section 6 (ICMP) for more discusion of this;

- Applications that assume the uniqueness of source addresses (e.g. IP address as identifier) - such applications will fail to operate correctly in the presence of multiple, discrete, simultaneous connections from the same source IP address;

- Applications that explicitly prohibit concurrent connections from the same address - such applications will fail when multiple subscribers sharing an IP address attempt to use them simultaneously.

Applications already frequently implement mechanisms in order to circumvent the presence of NATs (typically CPE NATs):

- Application Layer Gateways (ALGs): Many CPE devices today embed ALGs that allow applications to behave correctly despite the presence of NAT on the CPE. When the NAT belongs to the subscriber, the subscriber has flexibility to tailor the device to his or her needs. For CGNs, subscribers will be dependent on the set of ALGs that their service provider makes available. A service provider-based NAT may, or may not, support [RFC3947] (Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, “Negotiation of NAT-Traversal in the IKE,” January 2005.) for example. For port-range solutions, ALGs will require modification to deal with the port-range restriction, but will otherwise have the same capabilities as today.

- NAT Traversal Techniques: ICE, STUN, TURN, etc.

6. ICMP

TOC

ICMP does not carry any port information and is consequently problematic for address-sharing mechanisms. Sourcing ICMP from hosts behind an address-sharing solution does not pose problems. For inbound ICMP there are two cases. The first case is that of ICMP sourced from outside the network of the address-sharing solution provider. Several applications make use of this, e.g. P2P applications, and measurements derived by such applications in the presence of an address-sharing solution will be erroneous. Responses to outgoing ICMP should make use of the ICMP identifier value to route the response appropriately. The second case is that of ICMP sourced from within the network of the address-sharing solution provider (e.g. for network management and diagnostic purposes). In this case ICMP can be routed normally for CGN-based solutions owing to the presence of discrete private IP addresses for each CPE device. For port-range solutions, ICMP will will not be routable without special handling, e.g. placing a port number in the ICMP identifier field, and having port-range routers make routing decisions based upon that field. Alternatively another protocol could be used for diagnostic purposes, e.g UDP ping.

7. Fragmentation

TOC

When a packet is fragmented, transport-layer port information (either UDP or TCP) is only present in the first fragment. Subsequent fragments will not carry the port information and so will require special handling.

8. Support of Multicast

TOC

[RFC5135] (Wing, D. and T. Eckert, “IP Multicast Requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT),” February 2008.) specifies requirements for a NAT that supports Any Source IP Multicast or Source-Specific IP Multicast. Port-range routers that form part of port-range solutions will need to support similar requirements if multicast support is required.
[Placeholder for more details of impact of address-sharing on multicast deployments.]

9. Mobile-IP

TOC

IP address sharing within the context of Mobile-IP deployments (in the home network and/or in the visited network), will require Home Agents and/or Foreign Agents to be updated so as to take into account the relevant port information. There may also be issues raised when an additional layer of encapsulation is required thereby causing, or increasing the need for, fragmentation and reassembly.
Issues for Mobile-IP in the presence of NAT are discussed in [I-D.haddad-mext-nat64-mobility-harmful] (Haddad, W. and C. Perkins, “A Note on NAT64 Interaction with Mobile IPv6,” April 2010.)
[Placeholder for more details of impact of address-sharing on mobility deployments.]

10. Introduction of Single Points of Failure

TOC

In common with all deployments of new network functionality, the introduction of new nodes or functions to handle the multiplexing of multiple subscribers across shared IPv4 addresses could create single points of failure in the network. Any IP address sharing solution should consider the opportunity to add redundancy features in order to alleviate the impact on the robustness of the offered IP connectivity service. The ability of the solution to allow hot swapping from one machine to another should be considered.

11. Security

TOC

11.1. Port Randomisation

TOC

A blind attack that can be performed against TCP relies on the attacker's ability to guess the 5-tuple (Protocol, Source Address, Destination Address, Source Port, Destination Port) that identifies the transport protocol instance to be attacked. [I-D.ietf-tsvwg-port-randomization] (Larsen, M. and F. Gont, "Transport Protocol Port Randomization Recommendations," April 2010.) describes a number of methods for the random selection of the source port number, such that the ability of an attacker to correctly guess the 5-tuple is reduced. With shared IPv4 addresses, the port selection space is reduced. Preserving port randomisation is important and may be more or less difficult depending on the address-sharing solution and the size of the port space that is being manipulated. Allocation of non-contiguous port ranges could help to mitigate this issue.

It should be noted that guessing the port information may not be sufficient to carry out a successful blind attack. The exact TCP Sequence Number (SN) should also be known. A TCP segment is processed only if all previous segments have been received, except for some Reset Segment implementations which immediately process the Reset as long as it is within the Window. If SN is randomly chosen it will be difficult to guess it (SN is 32 bits long); port randomisation is one protection among others against blind attacks.

TOC

11.2. Abuse Logging and Penalty Boxes

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. Law enforcement authorities may be particularly impacted because of this. This particular issue can be fixed by logging port numbers, although this will increase logging data storage requirements.
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 subscribers. 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.

11.3. Spam

TOC

Another case of identifying abusers has to do with spam blacklisting. When a spammer is behind a CGN or using a port-shared address, blacklisting of their IP address will result in all other subscribers sharing that address having their ability to source SMTP packets restricted to some extent.

11.4. IPsec

TOC

Even if IPSec is not deployed for mass market (e.g. residential), impacts of solutions based on shared IP addresses should be evaluated and assessed. [RFC3947] (Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, “Negotiation of NAT-Traversal in the IKE,” January 2005.) proposes a solution to solve issues documented in [RFC3715] (Aboba, B. and W. Dixon, “IPsec-Network Address Translation (NAT) Compatibility Requirements,” March 2004.) . The applicability of [RFC3947] (Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, “Negotiation of NAT-Traversal in the IKE,” January 2005.) in the context of shared IP address solutions should be evaluated.

11.5. Policing Forwarding Behaviour

TOC

[RFC2827] (Ferguson, P. and D. Senie, “Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing,” May 2000.) motivates and discusses a simple, effective, and straightforward method for using ingress traffic filtering to prohibit Denial-of-Service (DoS) attacks which use forged IP addresses. Following this recommendation, service providers operating shared-addressing mechanisms should ensure that source addresses, or source ports in the case of port-range schemes, are set correctly in outgoing packets from their subscribers or they should drop the packets.
If some form of IPv6 ingress filtering is deployed in the broadband network and DS-lite service is restricted to those subscribers, then tunnels terminating at the CGN and coming from registered subscriber IPv6 addresses cannot be spoofed. Thus a simple access control list on the tunnel transport source address is all that is required to accept traffic on the southbound interface of a CGN.

12. Geo-location and Geo-proximity

TOC

IP addresses are frequently used to indicate, with some level of granularity and some level of confidence, where a host is physically located. Geo-location services are used by content providers to allow them to conform with regional content licensing restrictions, to target advertising at specific geographic areas, or to provide customised content. Geo-location services are also necessary for emergency services provision. In some deployment contexts (e.g. centralised CGN), shared addressing will reduce the level of confidence and level of location granularity that IP-based geolocation services can provide. Other forms of geo-location will still work as usual.
A slightly different use of an IP address is to calculate the proximity of a connecting host to a particular service delivery point. This use of IP address information impacts the efficient delivery of content to an end-user. If a CGN is introduced in communications and it is far from an end-user connected to it, application performance may be degraded insofar as IP-based geo-proximity is a factor.

13. Authentication

TOC

Simple address-based identification mechanisms that are used to populate access control lists will fail when an IP address is no longer sufficient to identify a particular subscriber. Including port numbers in access control list definitions may be possible at the cost of extra complexity, and may also require the service provider to make static port assignments, which conflicts with the requirement for dynamic assignments discussed in Section 4.1 (Outgoing Ports) .

14. Traceability

TOC

Legal obligations require a service provider to provide the identity of a subscriber upon request to the authorities. Where one public IPv4 address is shared between several subscribers, the knowledge of the IP address alone is not enough to identify the appropriate subscriber. The legal request should include the information: [IP address - Port - Protocol- Begin_Timestamp - End_Timestamp].

Address sharing solutions must record and store all mappings (typically during 6 months to one year, depending on the jurisdiction) that they create. If we consider one mapping per session, a service provider should record and retain traces of all sessions created by all subscribers during one year (if the legal storage duration is one year). This may be challenging due to the volume of data requiring storage, the volume of data to repeatedly transfer to the storage location, and the volume of data to search in response to a query.

Address sharing solutions may mitigate these issues to some extent by pre-allocating groups of ports. Then only the allocation of the group needs to be recorded, and not the creation of every session binding within that group. There are trade-offs to be made between the sizes of these groups, the ratio of public addresses to subscribers, whether or not these groups timeout, the impact on logging requirements and port randomisation security.

15. IPv6 Transition Issues

TOC

IPv4 address sharing solutions may interfere with existing IPv4 to IPv6 transition mechanisms, which were not designed with IPv4 shortage considerations in mind. With port-range solutions for instance, incoming 6to4 packets should be able to find their way from a 6to4 relay to the appropriate 6to4 CPE router, despite the lack of direct port range information (UDP/TCP initial source port did not pass through the CPE port range translation process). One solution would be for a 6to4 IPv6 address to embed not only an IPv4 address but also a port range value.

Subscribers allocated with private addresses will not be able to utilise 6to4 to access IPv6, but may be able to utilise Teredo.

16. IANA Considerations

TOC

This memo includes no request to IANA.

17. Security Considerations

TOC

This memo does not define any protocol and raises no security issues. Section 11 (Security) discusses some of the security and identity-related implications of address sharing.

18. Contributors

TOC

This document is based on sources co-authored by J.L. Grimault and A. Villefranque of France Telecom.

19. Acknowledgements

TOC

This memo was partly inspired by conversations that took place as part of Internet Society (ISOC) hosted roundtable events for operators and content providers deploying IPv6. Participants in those discussions included John Brzozowski, Leslie Daigle, Tom Klieber, Yiu Lee, Kurtis Lindqvist, Wes George, Lorenzo Colliti, Erik Kline, Igor Gashinsky, Jason Fesler, Rick Reed, Adam Bechtel, Larry Campbell, Tom Coffeen, David Temkin, Pete Gelbman, Mark Winter, Will Charnock, Martin Levy, Greg Wood and Christian Jacquenet. The authors are also grateful to Christian Jacquenet, Iain Calder, Joel Halpern, Brian Carpenter, Gregory Lebovitz, Bob Briscoe and Marcelo Bagnulo for their helpful comments and suggestions for improving this document.
This memo was created using the xml2rfc tool.

20. Informative References

TOC

[CGN_Viability]	Alcock, S., “Research into the Viability of Service-Provider NAT,” 2008.
[I-D.boucadair-port-range]	Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, “IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture,” draft-boucadair-port-range-02 (work in progress), July 2009 (TXT).

[I-D.despres-sam]	Despres, R., “Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM),” draft-despres-sam-03 (work in progress), July 2009 (TXT).
[I-D.haddad-mext-nat64-mobility-harmful]	Haddad, W. and C. Perkins, “A Note on NAT64 Interaction with Mobile IPv6,” draft-haddad-mext-nat64-mobility-harmful-01 (work in progress), April 2010 (TXT).
[I-D.ietf-behave-v6v4-xlate]	Li, X., Bao, C., and F. Baker, “IP/ICMP Translation Algorithm,” draft-ietf-behave-v6v4-xlate-19 (work in progress), April 2010 (TXT).
[I-D.ietf-behave-v6v4-xlate-stateful]	Bagnulo, M., Matthews, P., and I. Beijnum, “Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers,” draft-ietf-behave-v6v4-xlate-stateful-11 (work in progress), March 2010 (TXT).
[I-D.ietf-softwire-dual-stack-lite]	Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, “Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion,” draft-ietf-softwire-dual-stack-lite-04 (work in progress), March 2010 (TXT).
[I-D.ietf-tsvwg-port-randomization]	Larsen, M. and F. Gont, “Transport Protocol Port Randomization Recommendations,” draft-ietf-tsvwg-port-randomization-07 (work in progress), April 2010 (TXT).
[I-D.nishitani-cgn]	Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, “Common requirements for IP address sharing schemes,” draft-nishitani-cgn-04 (work in progress), March 2010 (TXT).
[I-D.shirasaki-nat444]	Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444,” draft-shirasaki-nat444-01 (work in progress), March 2010 (TXT).
[I-D.ymbk-aplusp]	Bush, R., “The A+P Approach to the IPv4 Address Shortage,” draft-ymbk-aplusp-05 (work in progress), October 2009 (TXT).
[IPv4_Report]	Huston, G., “IPv4 Address Report,” 2009.
[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, “Address Allocation for Private Internets,” BCP 5, RFC 1918, February 1996 (TXT).
[RFC2663]	Srisuresh, P. and M. Holdrege, “IP Network Address Translator (NAT) Terminology and Considerations,” RFC 2663, August 1999 (TXT).
[RFC2782]	

	Gulbrandsen, A., Vixie, P., and L. Esibov, “A DNS RR for specifying the location of services (DNS SRV),” RFC 2782, February 2000 (TXT).
[RFC2827]	Ferguson, P. and D. Senie, “Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing,” BCP 38, RFC 2827, May 2000 (TXT).
[RFC2993]	Hain, T., “Architectural Implications of NAT,” RFC 2993, November 2000 (TXT).
[RFC3715]	Aboba, B. and W. Dixon, “IPsec-Network Address Translation (NAT) Compatibility Requirements,” RFC 3715, March 2004 (TXT).
[RFC3947]	Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, “Negotiation of NAT-Traversal in the IKE,” RFC 3947, January 2005 (TXT).
[RFC4787]	Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” BCP 127, RFC 4787, January 2007 (TXT).
[RFC5135]	Wing, D. and T. Eckert, “IP Multicast Requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT),” BCP 135, RFC 5135, February 2008 (TXT).

Authors' Addresses

TOC

	Mat Ford (editor)
	Internet Society
	Geneva
	Switzerland
Email:	ford@isoc.org
	Mohamed Boucadair
	France Telecom
Email:	mohamed.boucadair@orange-ftgroup.com
	Alain Durand
	Comcast
Email:	Alain_Durand@cable.comcast.com
	Pierre Levis
	France Telecom
	42 rue des Coutures
	BP 6243
	Caen Cedex 4 14066
	France

Email: pierre.levis@orange-ftgroup.com

Phil Roberts
Internet Society
Reston, VA
USA
Email: roberts@isoc.org