IPv6 Rapid Deployment on IPv4 infrastructures (6rd)
draft-despres-v6ops-6rd-ipv6-rapid-deployment-00

Status of this Memo

Copyright Notice

Abstract

IPv6 rapid deployment (6rd) builds upon mechanisms of 6to4 (RFC3056) to enable a service provider to rapidly deploy IPv6 unicast service to its existing IPv4 sites. Like 6to4, it utilizes stateless IPv6 in IPv4 encapsulation in order to transit IPv4-only network infrastructure. Unlike 6to4, 6rd requires a service provider to use one of its own IP prefixes rather than the fixed 6to4 prefix. A service provider has used this mechanism for its own "rapid deployment" of IPv6 (five weeks from first exposure to "opt-in" deployment for 1,500,000 residential sites).

Table of Contents

1. Introduction and 6rd purpose

After having had a succinct presentation of the 6rd idea, a major ISP in France, Free Telecom, did all of the following in only five weeks (November 7th to December 11th 2007): (1) obtain its first IPv6 prefix from its RIR; (2) make the needed 6rd software modifications to all CPE that it provides to its customers; (3) provision a 6to4 gateway, duly modified to support 6rd; (4) test IPv6 operation with a few applications; (6) finish deployment; (7) announce IPv6 availability, at no extra charge, for all customers accepting to have it. More than 1,500,000 residential customers thus became able to use IPv6 in their sites, with a look and feel as that of other native IPv6 solutions, at the only condition to consciously activate the function in their CPEs.

This story is not reported to suggest that other ISPs should be able to do the same. It illustrates however that, under some circumstances, the following vicious circle has been broken:

- o ISPs that have large IPv4 installed bases tend to wait for more customer demand before bearing the cost of generalized IPv6 support.

- o Customers tend to wait for more IPv6 depending applications before requiring IPv6 support by ISPs.

- o Application developers tend to wait for more IPv6 support by ISPs before investing substantially on IPv6 dependent applications.

6rd has been designed to drastically simplify first IPv6 deployments on IPv4 installed infrastructures.

Ideal conditions for 6rd deployment, which were satisfied at Free, are that:

- o The ISP controls CPEs of all its sites.

- o It can easily modify CPE software and have it downloaded.

- o It can quickly install gateways of its own, with high bandwidth in both IPv4 and IPv6.

- o It can adapt the software of these gateways to 6rd (preferably starting with an existing 6to4 relay software to minimize the effort).

A less ideal but workable situation is one where CPEs are provisioned by customers themselves (hosts , or site entrance routers). For these sites to benefit from 6rd being deployed by their ISP, they need 6rd support also in their CPEs. The necessary software upgrade then depends on their CPE vendor to implement 6rd, and to whoever manages theses CPEs to download the new releases. The 6rd specification is proposed to become a standard for this to be possible.

The purpose of this draft is that, with it:

- The community of IETF experts can critically evaluate its soundness.

- ISPs that offer only IPv4 services can determine whether, based on their own constraints, they wish to use 6rd to accelerate their own IPv6 deployment.

- CPE manufacturers can determine whether they wish to support 6rd in their products. If they do, their clients whose ISPs support 6rd will have native IPv6 operational in their sites.

- Manufacturers of routers used in ISP infrastructures can determine whether they wish to include 6rd support in their products. (Note that in which products to do it may depend on whether address parsing hardware is used, and on whether it is suitable for parse 6rd prefix parsing.) If they do, 6rd relays at the IPv4-IPv6 border will no longer have to be devices external to these routers, and their number will be easily increased. The added value will then move from external devices to these manufacturer's routers.

Readers are supposed to be familiar with IPv6 address formats and notation.

It is understood that the wording of this draft 00 can be much improved, in particular with respect to the English language. But it is also felt that presenting the 6rd to the IETF shouldn't be delayed further, so that debates on it can start as soon as possible.

2. Abbreviations and terminology

6rd: a mechanism for IPv6 rapid deployment by independent ISPs on their existing IPv4 infrastructures

6rd CPE: CPE which supports 6rd (host or site entrance router)

6rd ISP: ISP which supports 6rd (operates 6rd relays)

6rd site: IPv4 customer site of a 6rd ISP

6rd ISP anycast address: IPv4 anycast address chosen by a 6rd ISP

6rd ISP prefix: IPv6 prefix chosen a 6rd ISP

6rd ISP relay: a stateless encpasulation-decapsulation function between IPv4 and IPv6 routing infrastructures of an ISP

6to4: Connection of IPv6 Domains via IPv4 clouds [RFC3056][RFC3068]

CPE: Customer Premise Equipment (host or router at site entrance)

DHCPv4: IPv4 Dynamic Host Configuration Protocol [RFC2131]

IANA: Internet Assigned Numbers Authority

ICMPv4: IPv4 Internet Control Message Protocol [RFC0792]

ICMPv6: IPv6 Internet Control Message Protocol [RFC0792]

IPv4: Layer 3 Internet Protocol of 1981 [RFC0791]

IPv6: Layer 3 Internet Protocol version 6 [RFC2460]

ISP: Internet Service Provider

MTU: Maximum Transfer Unit [RFC1191]

NAT: Network Address Translator [RFC2663]

Pure IPv6: IPv6 with addresses that contain no IPv4 address

RIR: Regional Internet Registry

Site: Customer site (has a CPE at its entrance)

3. 6rd specification

3.1. General principles

The 6rd specification is based on the following principles :

a. For RAPIDITY of IPv6 deployment by ISPs that are still IPv4-only, this deployment should be feasible on their UNMODIFIED IPv4 INFRASTRUCTURES.

b. For COMPLETENESS of the IPv6 unicast service offered to their customers, ISPs should make NO ASSUMPTION on how hosts of other ISPs obtain their IPv6 service.

c. For SCALABILITY of IPv6 deployment on IPv4 infrastructures, encapsulation-decapsulation functions of IPv6 packets in IPv4 ones should be STATELESS (load sharing between a number of distributed processors should be feasible).

d. For EFFICIENCY, IPv6 packet between two IPv4 sites of a same ISP should follow the SAME ROUTES as those of IPv4 packets between the same sites.

The rapidity principle implies that 6rd functions should be introduced only at the periphery of IPv4 infrastructures. Routing on these infrastructures should be that of IPv4, with no need for an independent address assignment and routing policy for IPv6.

The completeness principle implies that IPv6 prefixes of 6rd sites MUST start with prefixes that belong to the IPv6 address spaces. (In particular, the 6to4 prefix 2002::/16 would not satisfy this principle: packets destined to a 6to4 site reach their destination only if they come from sites where CPEs support 6to4, or from ISPs where the 6to4 Anycast Address is routable in IPv4 to at least one 6to4 relay. Neither of these two conditions can be guaranteed by the destination end ISP).

The scalability principle implies that encapsulation functions in 6rd relays can find IPv4 destination addresses without depending on some temporary states that would relate IPv4 destinations to IPv6 ones. For this, the most straightforward approach consists in having the IPv4 CPE address of a 6rd site explicitly contained in its IPv6 prefix. Since this approach happens to be compatible with satisfaction of other principles, it is adopted for 6rd.

Another implication of the scalability principle is that 6rd relays should be reachable, across the ISP IPv4 infrastructure, at an anycast address. Each 6rd ISP choses for this its "6rd ISP anycast address".

The efficiency principle implies that 6rd CPEs can recognize, in packets leaving their sites, which ones can be routed to their destination directly across the local ISP IPv4 infrastructure (i.e. without having to go through a 6rd relay). For this, a simple approach consists in each 6rd ISP to chose one and only one of its IPv6 unicast prefixes as the "6rd ISP prefix" which appears at the start of 6rd site addresses addresses, and to have this prefix known by 6rd CPEs.

The architecture which results from these principles is that of Figure 1.

```
                                   6rd ISP                    Native
                      ______________/\________________       IPv6
               6rd   /                                \      routing
               CPEs          unchanged            Stateless     |
                |       IPv4 infrastructure     6rd ISP relays  |
                |                |                    |          |
                V                V                    V          V
                ___    ______________________         ___
        IPv6   |   |  |                      |       |   |
               |   |--|-.     .--------------|-------|   |--------
               |___|  |  \   /               |       |___|
             Customer |   \ /    6rd ISP  => |
               Sites  |    O  anycast address|            <= 6rd ISP
                ___   |   / \                |         ___    prefix
               |   |  |  /   \               |        |   |
        IPv6   |   |--|-'     '--------------|--------|   |--------
               |___|  |                      |        |___|
                      |______________________|
```

6rd ARCHITECTURE

Figure 1

3.2. 6rd ISP prefix and 6rd addresses

A 6rd address is a particular case of IPv6 address, with the following format:

```
<-- Link prefix -- 64 bits -->.
|                             |
|  16, 24             16, 8   |
|    or                 or    |<----------- 64 bits --------->.
|  32 bits   32 bits   0 bits |                               |
|      |        |          |  |                               |
|      V        V          V  |                               |
+---//----+------------+--//--+-------------------------------+
|"6rd ISP |    IPv4    |Subnet|        Interface ID           |
| prefix" |Site address|  ID  |                               |
+---//----+------------+--//--+-------------------------------+
<---- Site prefix ---->
```

6rd ADDRESS FORMAT

Figure 2

According to the completeness principle above, the 6rd ISP prefix, being in the first bits of the address, MUST belong to the ISP public unicast address space.

As far as the protocol is concerned, 6rd ISP prefixes could have any number of bits up to 32. But implementations of 6rd, necessary in both 6rd relays and 6rd CPEs, is simpler, and more amenable hardware implementation where appropriate, if 6rd ISP prefixes must be multiple of 8 bits. It is therefore proposed to be a MUST in 6rd.

The choice between /16, /24 and /32 6rd ISP prefixes depends on whether IPv6 prefixes ISPs obtain from their RIRs are short enough:

- With /32 RIR provided prefixes, the minimum generally recommended today for initial ISP assignments, 6rd ISPs can only assign /64 site prefixes to their 6rd customer sites. These sites are then limited to only one IPv6 link. This is expected to be satisfactory initially for the vast majority of residential sites, where the number of hosts is small. But this is a real restriction which would advantageously be avoided with RIRs cooperation (see below).

- With /16 RIR provided prefixes, 6rd ISPs can assign /48 site prefixes to their 6rd customer sites, i.e. prefixes the length of which is that RIRs recommend today for all customer sites. But RIRs may be expected to be reluctant to distribute /16s, especially since generalized /48 prefixes are more generous than really needed, at least initially.

- With /24 RIR provided prefixes, 6rd ISPs can assign /56 site prefixes to their 6rd customer sites. These sites can then support up to 256 subnets. This is expected to be more than sufficient for the most demanding residential sites, and largely sufficient for almost all professional sites pending deployment of native IPv6 infrastructures.

In view of the potential of 6rd to facilitate IPv6 availability, RIRs could be advised to consider assigning /24 prefixes to ISPs that deploy 6rd, and to endorse that these ISPs make their first IPv6 deployments with /56s distributed to their customers.

3.3. Encapsulation and decapsulation in 6rd CPEs and 6rd relays

3.3.1. Packet format

To traverse ISP IPv4 infrastructures, IPv6 packets are encapsulated in IPv4 ones in conformance with [RFC2893]. The IPv4 protocol field is set to 41, as specified by IANA for this encapsulation [Protocol Numbers].

In order to avoid that encapsulated packets would ever be subject to IPv4 fragmentation, 6rd ISPs MUST support path MTUs of at least 1300 octets on their complete IPv4 infrastructures. (1300 = 1280, the minimum IPv6 MTU, plus 20, the header length of the IPv4 encapsulating packet.) Fragmentation has to be avoided because of its incompatibility with stateless functions that operate on complete packets and that may be implemented in several distributed instances.

3.3.2. 6rd prefix in non 6rd addresses - the 6rd pure IPv6 tag

If an ISP has only one IPv6 prefix assigned by its RIR, as typical for a first assignment, and first uses it to deploy IPv6 in 6rd, it should, to avoid wasting its address space, be able to also use this same prefix to build "pure IPv6" site prefixes. (IPv6 prefixes are "pure" if they don't contain IPv4 addresses).

If a 6rd ISP prefix has such a mixed use, CPEs must distinguish between destinations that are 6rd addresses of the same ISP, and other destinations. Packets having the former have to be routed directly to their destinations across local ISP IPv4 infrastructures; packets having the latter have to be routed toward their destinations via 6rd ISP relays of the local ISP). To facilitate this determination, the 6rd specification requires that all pure IPv6 addresses that start with the 6rd ISP prefix contain a special value in lieu of the beginning of an IPv4 field, the "6rd pure IPv6 tag".

The value of this tag has to be chosen so that it never appears at the beginning of IPv4 unicast addresses assigned by the ISP. Its proposed binary value is 1110 (OxE in hexadecimal, 240.0.0.0/4 in IPv4 notation), which has the desired property: IANA has reserved the set of all 224/8 to 239/8 consecutive prefixes, i.e. the 224/4 prefix, for IPv4 multicast addresses [IPv4 addresses].

```
<-- Link prefix -- 64 bits -->.
|                              |
|           "6rd pure          |
|  16, 24    IPv6 tag"         |
| or 32 bits  4 bits           .<----------- 64 bits --------->.
|     |       |                |                               |
|     V       V                |                               |
+----//-----+----+-----//-----+-------------------------------+
|  6rd ISP  |1110|             |         Interface ID          |
|   prefix  |    |             |                               |
+----//-----+----+-----//-----+-------------------------------+
```

6rd PURE IPv6 ADDRESS FORMAT

Figure 3

NOTE: If an ISP wants to use hardware which works only, or better, on octet boundaries, it MAY decide to assign pure IPv6 addresses that, after the 6rd ISP prefix all start with OxE0 (binary 11100000). With this restriction on assigned addresses, presence tests of the 6rd prefix can be performed indifferently on 4 or 8 bits.

3.3.3. Relationship between IPv6 and IPv4 addresses in encapsulation packets

```
     IPv6 address (source or destination)         IP4 address
          in an encapsulated packet        in the encapsulating packet


                             6rd site addresses

                 4 bits              \
      +----//-----+-|--+-----         |       +---------------+
      |  6rd ISP  |not :               \       |     IPv4      |
      |   prefix  |1110:               /       | site address  |
      +----//-----+----+-----         |        +---------------+
                   <-- IPv4 ---      /
                    site address


                            Non 6rd site addresses

      +----//-----+----------          \
      |  not 6rd  |                     |
      |ISP prefix |                     |
      +----//-----+----------           |        +---------------+
                   6rd                   \       |   6rd ISP     |
              pure IPv6 tag              /       |anycast address|
      +----//-----+-|--+-----           |        +---------------+
      |  6rd ISP  |1110|                |
      |   prefix  |    |                |
      +----//-----+----+-----          /
                  4 bits

            RELATIONSHIP BETWEEN IPv6 AND IPv4 ADDRESSES
                IN ENCAPSULATION PACKETS OF AN ISP
```

Figure 4

In an encapsulating packet, IPv4 addresses, source or destination, have the following relationship with IPv6 addresses of the encapsulated packets, respectively source or destination:

- o If the IPv6 address starts with the 6rd ISP prefix, and if this prefix is not followed by the 6rd pure IPv6 tag, this site prefix is that of a 6rd site of the local ISP. The IPv4 address MUST then be equal to the 32 bits which follow the 6rd ISP prefix in the IPv6 site prefix.

- o Contrarily, if the ISP prefix is followed by the 6rd pure IPv6 tag, or if the IPv6 address does not start with the ISP prefix, then this address does not belong to a 6rd site of the local ISP.

These relationships MUST be implemented in encapsulation functions of 6rd relays and 6rd CPEs.

In addition, to protect against source address spoofing, 6rd CPEs and 6rd ISP relays SHOULD silently discard packets they receive that have unrealistic source and destination addresses or unrealistic IPv4-IPv6 address relationships. The pseudo-code of Section 3.6 presents details of these verifications.

3.4. ICMP consideration

In 6rd decapsulation functions of 6rd CPEs and 6rd relays, ICMPv4 packets that are received to signal unreachable destinations (ICMPv4 type field = 3) SHOULD be converted into ICMPv6 packets (ICMPv6 type field = 1 and code field = 0).

For this to be possible, ISPs that support 6rd SHOULD ensure that all routers of their IPv4 infrastructures return ICMPv4 packets long enough to contain IPv6 source addresses of encapsulated packets. (This is automatically the case if these routers conform to [RFC1812] Section 4.3.2.3.)

ICMPv6 packets received by encapsulation functions of 6rd ISP relays and 6rd CPEs are encapsulated like any other IPv6 packets.

3.5. IPv4 routing precaution

For an ISP to guarantee proper routing of IPv6 packets going from its IPv4 sites to other ISPs, its IPv4 infrastructure SHOULD NOT accept from other ISPs IPv4 routes that include the 6rd ISP anycast address.

3.6. Pseudo code

Figure 5 details, in pseudo code and with intuitive notations, encapsulation and decapsulation functions, in both 6rd ISP relays and 6rd CPEs.

```
        CPE DECAPSULATION                ISP RELAY ENCAPSULATION
   IF v6Pkt = ICMPv4(Unreachable)     IF Packet > 1280 Octets
   THEN Forward ICMPv6(Pkt too big)   THEN Return ICMPv6(Pkt too big)
   ELSEIF                             ELSEIF (S6 <> 6rdPrf...
    (IF S4 = 6rdAnycast                      OR S6 = 6rdPrf.Pv6Tag...)
     THEN [S6 <> 6rdPrf...            THEN
           OR S6 = 6rdPrf.Pv6Tag...]     D4 <- bits pl to pl+32 of D6
     ELSE [S6 = 6rdPrf...                S4 <- 6rd Anycast
        AND S6 <> 6rdPrf.Pv6Tag...])     Encapsulate IPv6 packet
     AND D6 = 6rdPrf.v4SiteAdd...      ELSE Discard packet
   THEN Decapsulate IPv6 packet             |
   ELSE Discard packet                      |
           |      __________________        |        ______
   IPv6    V     |       IPv4       |       V       | IPv6
          ___    |MTU at least 1300 |      ___      |
         |   |   |                  |     |   |     |
    --<--|---|---|--<---      ---<--|-----|---|-----|--<---
         |   |   |<= CPE addresses  |     |   |     |
         |   |   |                  |     |   |     |       6rd
   0::/0 |   |   |   6rd ISP anycast|     |   |     | <= ISP prefix
      => |   |   |       address => |     |   |     |
    -->--|---|---|-->---      --->--|-----|---|-----|-->---
         |___|   |                  |     |___|     |
           A     |__________________|       A       |______
           |                                |
           |                                |
        CPE ENCAPSULATION                 ISP RELAY DECAPSULATION
   IF packet > 180 octets               IF v6Pkt = ICMPv4(unreachable)
   THEN Return ICMPv6(Pkt too big)      THEN Forward ICMPv6(Pkt too big)
   ELSEIF S6 = 6rdPrf.Pv6Tag...)        ELSEIF
      IF ( D6 = 6rdPrf...                  S4 <> 6rdAnycast
           AND D6 <> 6rdPrf.Pv6Tag...)     AND S6 = 6rdPrf...
      THEN D4 <- bits pl to pl+32 of D6    AND S6 <> 6rdPrf.Pv6tag...
      ELSE D4 <- 6rdAnycast                AND ( D6 = 6rdPrf.Pv6tag...
      S4 <- v4SiteAdd                            OR D6 <> 6rdPrf...
   ELSE Discard packet                  THEN Decapsulate IPv6 packet
                                        ELSE Discard packet

              PSEUDO CODE DESCRIPTION OF 6rd FUNCTIONS
            (lp being the bit length of the ISP 6rd prefix)
```

Figure 5

4. The 6rd DHCPv4 option

For full support of 6rd, ISPs that have sites with customer-supplied CPEs, and suppliers of these CPEs, SHOULD support the "6rd DHCPv4 option.

With the 6rd DHCPv4 option, 6rd CPEs obtains 6rd ISP prefixes and 6rd ISP anycast addresses, in a DHCPv4 message [RFC2131].

The proposed format for the 6rd DHCPv4 option is as follows, where nn is a code to be assigned by IANA:

```
                              6rd prefix length
                                (16, 24 or 32)
  Code  Length                           |
 +-----+-----+-----+-----+-----+-----+--V--+-----+-----+-----+-----+
 |= nn | = 9 | 6rd ISP Prefix        | pl  |6rd ISP anycast address|
 +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

FORMAT OF THE 6rd DHCPv4 OPTION

Figure 6

5. 6rd Applicability to ISPs that use IPv4 private addresses

Some ISP support NAT functions [RFC2663] within their IPv4 infrastructures, so that they can assign IPv4 private addresses [RFC1918] to some of their sites. These ISPs can build IPv6 6rd site prefixes with such IPv4 addresses.

If an ISP supports several independent NAT functions (typically because of an insufficient scalability of NAT-supporting devices), it has to ensure, for 6rd, that address spaces behind these NATS are disjoint .

Figure 7 presents an example where IPv4 private addresses start with 10.0.0.0/8 [RFC1918], a typical choice since other private address prefixes leave room for less addresses.

```
                     ______________________________
                    | 10.x.x.x/8 addresses         |
                    |  <==                         |
              <-----|                              |----->
                    |      6rd ISP anycast address|
           6rd CPEs |                          ==> |  6rd relays
                    |                              |
              <-----|         0.0.0.0/0            |----->
                    |              :               |
                    |              V               |
                    |______________________________|
                                   |
                                 __|__
                      ISP       |     |
                    IPv4 NAT |_____|
                                   |
                                   |
                                   V
                      IPv4 public address infrastructure
```

EXAMPLE OF 6rd WITH ISP ASSIGNED IPV4 PRIVATE ADDRESSES

Figure 7

NOTE: This capability is another difference of scope with 6to4. It may increase the interest of 6rd, for rapid IPv6 deployment, where scarcity of IPv4 addresses has led ISPs to support NATs in their infrastructures.

6. Acknowledgements

The author would like to warmly acknowledge the major contribution of Rani Assaf to 6rd's credibility. He immediately appreciated 6rd's potential, and made the daring decision to implement it, and to rapidly deploy it on Free's operational network. He has also been first to point out that 6rd ISP prefixes can also be used for IPv6 addresses other than 6rd. Patrick Grossetete made useful suggestions on multi-subnet sites, and on 6rd anycast addresses. Mark Townsley advised on how to proceed in IETF.

Besides these direct contributions, acknowledgments are due to a few IPv6 confirmed experts who, when the author was still an IPv6 newcomer, taught him subtleties of IPv6. Without his past debates with Laurent Toutain, Francis Dupont and Alain Durand, this proposal would probably not have been possible.

7. Security Considerations

With the specification as is, and provided all recommended behaviors are supported, the author has identified one security risk beyond those that are common to all of IPv6 implementations. It results from possibilities of IPv4 address spoofing: If a 6rd site may receive packets with IPv4 spoofed source addresses, it may also receive, in IPv6 encapsulated packets, IPv6 spoofed source addresses.

Since this risk is generally lived with in IPv4, letting higher layers to ensure enough security when necessary, it is expected that it is acceptable in practice.

8. IANA Considerations

This memo implies a request to IANA for the code of the DHCPv4 6rd option presented of Section 4.

9. References

9.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.

[RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2463] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.

[RFC2893] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 2893, August 2000.

[RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

9.2. Informative References

[IPv4 addresses]
Internet Assigned Numbers Authority, "Internet Protocol v4 Address Space - http://www.nro.net/documents/pdf/nro46.pdf", February 2008.

[Protocol Numbers]

Internet Assigned Numbers Authority, "PROTOCOL NUMBERS - http://www.iana.org/assignments/protocol-numbers", January 2008.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.

[RIR Policy] Number resource Organization, "RIR Comparative Policy Overview - http://www.iana.org/assignments/ipv4-address-space", November 2007.

Author's Address

Remi Despres
RD-IPtech
3 rue du President Wilson
Levallois,
France

Phone: +33 6 72 74 94 88
Email: remi.despres@free.fr

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).