Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: January 6, 2011

R. Despres
RD-IPtech

July 5, 2010

Rapid Deployment of Native IPv6 Behind IPv4 NATs (6rd+)
draft-despres-softwire-6rdplus-00

Abstract

The [6rd] mechanism permit IPv6 "rapid deployment" across IPv4 infrastructures of Internet Service Providers, but only in cases where customer-premise equipments can be upgraded to support it. 6rd+ extends this IPv6 rapid deployment capability to hosts behind customer-premise equipments that cannot be upgraded (provided these hosts can be upgraded themselves to support 6rd+). Like [6rd], 6rd+ provides native IPv6 addresses, so that quality of service can be guaranteed by Internet Service Providers; it operates statelessly, which ensures simplicity and scalability; and it transmits encapsulated IPv6 packets along optimized paths, which contributes to efficiency and acceptability.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 6, 2011.

Table of Contents

1. Introduction 3
2. Terminology 4
3. The 6rd+ Protocol 5
3.1. Participating Entities and IPv6 Paths 5
3.2. 6rd+C Locator Formats 7
3.3. 6rd+ Datagram Formats 9
3.4. Mapping Rules of 6rd+Cs and 6rd+Ps 11
3.5. Acquisition of 6rd+ parameters by 6rd+Cs 14
3.6. Anti-Spoofing and Anti-Routing-Loop Protections 16
4. Security considerations 17
5. IANA Considerations 17
6. Acknowledgments 17
7. References 17
7.1. Normative References 17
7.2. Informative References 17
Author's Address 19

1. Introduction

Although most operating systems now support IPv6, the large installed base of IPv4-only customer-premise equipments (CPEs) acts as a deterrent to IPv6 rapid deployment. Mechanisms such as [Tunnel Broker] and [Teredo] are designed to provide IPv6 connectivity behind such CPEs, but with too severe limitations to be generalized: [Tunnel Broker] necessitates that hosts be individually registered in tunnel-broker servers, and tends to produce unnecessarily long IPv6 paths; [Teredo] does not work with all types of customer-site NATs, and, because it uses IPv6 addresses that start with a Teredo-specific prefix, prevents Internet Service Providers (ISPs) to control the IPv6 quality of service (QoS) experienced by their customers.

6rd+ is designed to avoid these limitations:

- All types of NATs are supported.
- IPv6 addresses obtained behind IPv4-only CPEs are "native" (unicast starting with ISP-allocated prefixes), so that ISPs can control the IPv6 QoS experienced by their customers.
- IPv6 paths are optimized.
- Hosts don't need to be individually registered in any specific server.

The name 6rd+ is derived from that of [6rd], to express that it extends its IPv6 "rapid deployment" to cases not covered so far. Like [6rd], 6rd+ uses IPv6-packet encapsulations and stateless address mappings, but where [6rd] necessitates CPEs to be upgraded, 6rd+ does not. The counterpart is that hosts behind non-upgraded CPEs have to support 6rd+ to obtain their native IPv6 connectivity. Their upgrade is expected to be feasible with a downloadable module, at least with Linux for which similar stateless functions have been shown to be downloadable.

Other approaches to extend [6rd] to traverse IPv4-only CPEs have been proposed in [6rd UDP] and [SAMPLE], but without optimization of IPv6 paths. As this optimization is expected to be important for a wide use of native IPV6, 6rd+ is proposed as a more complete solution.

Formats 6rd+ based IPv6 addresses presuppose that the rule of

[RFC4291] that IPv6 addresses must contain 64-bit-IIDs is partially relaxed. This relaxation is needed to keep formats of 6rd+ based addresses simple, and does not conflict with the technical rationale for this rule. The 64-bits IID is indeed justified for addresses that, on some link, are subject to the neighbor discovery protocol [ND], but only for these. As 6rd+ addresses that don't have 64-bits IIDs are only assigned to host interfaces on virtual links, where the [ND] protocol is not used, the conflict doesn't exist. (Note that the viability of virtual links having no [ND] has been abundantly validated with [6to4], and that an independent proposal to partially relax the 64-bits-IID rule is available in [IPv6 /127 prefix].)

At this stage, the 6rd+ design is so new that some fixes are likely to be needed, in particular after running code has been developed. This is the reason why it is proposed as experimental. But the intent is to progress it to the standard track as soon as possible: the urgency of IPv6 actual use continues to grow.

2. Terminology

IPv4+: The extended-address family whose addresses have 48 bits, 32 of unicast IPv4 addresses plus 16 of a port numbers.

Locator: Any routable address or prefix, in a specified address space among IPv4, IPv6, and IPv4+.

IPv4+ datagram : An IPv4 datagram that contains a 6rd+ message, or an IPv6 packet in which at least one address has been obtained with 6rd+. The protocol above IPv4 is either UDP or the SAM protocol whose ID has to be assigned by IANA (see Section 3.3. If it is UDP, at least one of the two UDP ports is one of the two 6rd+ well-known ports, to be also assigned by IANA.

NAT: In this document, only IPv4-to-IPv4 network address translations are considered (NAT44s). They typically translate address plus ports couples, as described in [RFC2663].

EIM, EDM: An EIM NAT is one that does endpoint-independent mappings as specified in [RFC4787] (also known as a "full-cone" NAT). An EDM NAT is one that does endpoint-dependent mappings (the opposite of an EIM NAT).

6rd+C, 6rd+P: A 6rd+C is a 6rd+ "customer" function, and a 6rdP is a 6rd+ "provider" function. Each one has one lower-layer interface, and two upper-layer interfaces. The lower-layer interface is that of an IPv4 link layer. One of its upper-layer interfaces is a replication of the lower-layer interface? the other is an IPv6

pseudo interface (as in [6to4]. Between two 6rd+Cs, or between a 6rd+C and a 6rd+P, IPv6 packets are encapsulated in IPv4+ datagrams.

6rd+ domain: A network domain delimited by 6rd+Ps and 6rd+Cs of a given ISP.

Local, Interior: The "local" address space of a 6rd+C is that available at its lower-layer interface. Its interior address space is that available after traversal of all NATs that may exist between this 6rdC and 6rd+Ps. It is the same for all 6rd+Cs of the domain. (A 6rd+C that has no NAT between itself and 6rd+Ps has identical local an interior address spaces.)

Shortcut: An path, between two 6rd+Cs of the same 6rd+ domain, that doesn't go via any domain 6rd+P.

3. The 6rd+ Protocol

3.1. Participating Entities and IPv6 Paths

In a 6rd+ domain, participating entities are the following [Figure 1]:

- 6rd+Ps (6rd+ Provider functions)
- 6rd+Cs (6rd+ Customer functions)
- CPE NATs of two types
  - EIM CPE NAT (doing endpoint-independent mapping)
  - EDM CPE NAT (doing endpoint-dependent mapping)
- ISP NATs

6rd+Ps are stateless. They can be duplicated in any number of nodes. Their interior unicast address is then routed like any other anycast address.

6rd+C functions are also stateless. However, if there are NATs between them and 6rd+Ps, their operation depends on states that are maintained in these NATs. In this case, they cannot be duplicated in several nodes.

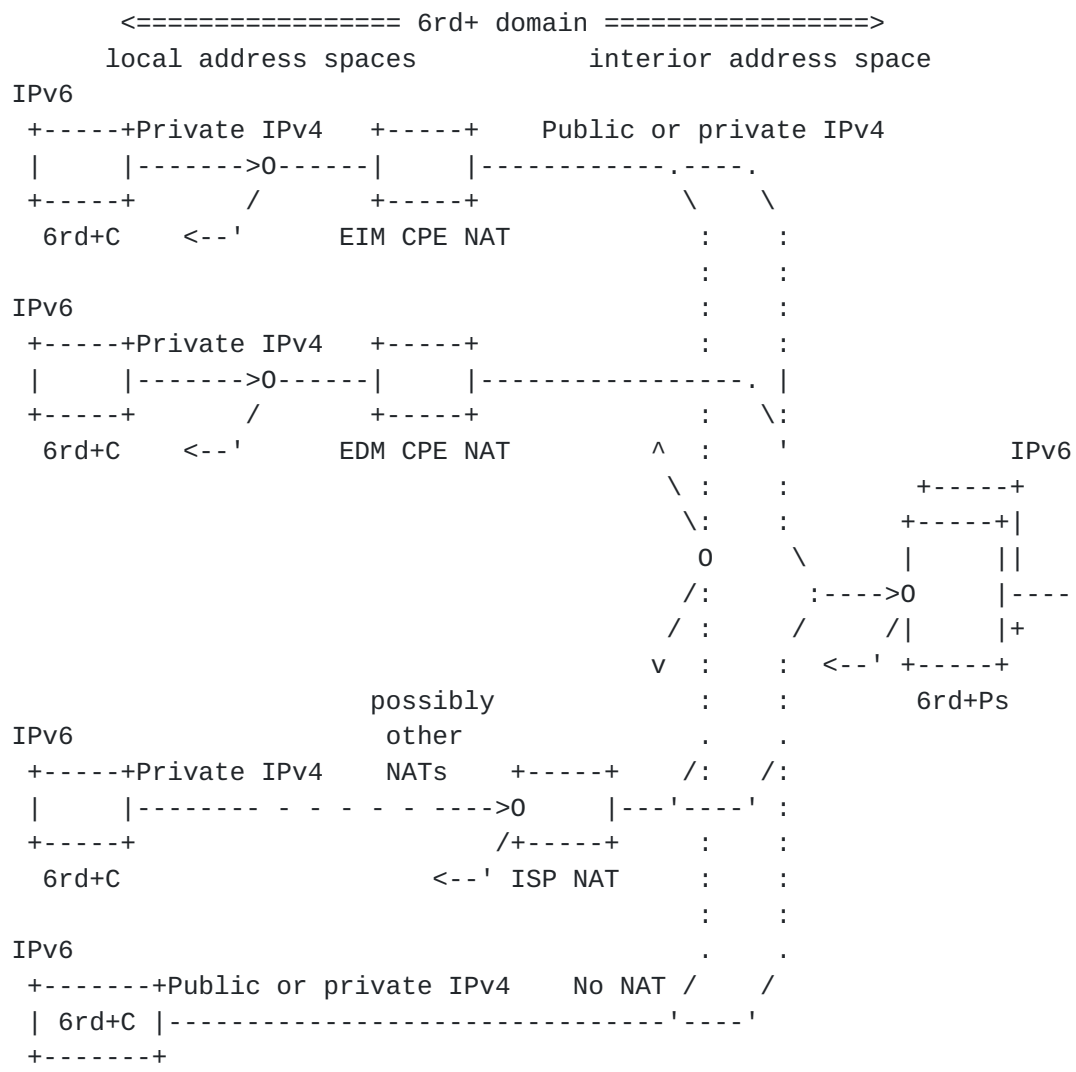

6rd+ PARTICIPATING ENTITIES AND IPv6 PATHS

Figure 1

The distinction between EIM and EDM NATs is key, like in [STUN] and in [Teredo], to find some shortcuts that are possible only with EIM NATs. CPE NATs, being in legacy devices of unknown types, may be either EIM or EDM. ISP NATs are assumed to be EIM and to support hairpinning, as specified in [RFC4787].

Paths that can be followed by IPv6 packets are illustrated in Figure 1. Possible shortcuts are the following:

Intra-site shortcut: This shortcut concerns two 6rd+Cs that are on the same LAN behind a common CPE NAT (EIM CPE NAT or EDM CPE NAT). IPv6 packets are exchanged via the LAN, without going to the CPE.

Intra-ISP shortcut: This shortcut concerns 6rd+Cs that are behind EIM CPE NATs, or ISP NATs, or no NAT. IPv6 packets between them, unless intra-site shortcuts which are shorter are possible, traverse the IPv4 routing infrastructure of the domain without going to any of its 6rd+Ps.

ISP-NAT hairpin shortcut: If two 6rd+Cs are behind the same ISP NAT, IPv6 packets they exchange follow an hairpin path whose turning point is in this ISP NAT.

3.2. 6rd+C Locator Formats

Formats of IPv6 locators that 6rd+Cs obtain depend on the types of NATs that exist between themselves and 6rd+Ps of the domain. Each such IPv6 locator is chosen to contain all addressing information that is needed to reach it via with or without shortcut, and nothing more. The different formats that result from this principle are detailed in Figure 2 (where the dot is used as concatenation operator).

The three NAT-type codes, C, C' and C", because they must be distinguishable in a prefix-match search, have to be be disjoint prefixes (no one starting with another). Besides that, they may have different lengths:

- NAT-type code C", which is that of 6rd+Cs that have no NAT to traverse, should typically be short so that IPv6 locators E=D.C".X remain within a 64-bits limit. These locators can thus be used as link prefix on ordinary IPv6 links (subject to [ND]).

- NAT-type codes C and C', which are those of 6rd+Cs behind NATs, appear in IPv6 locators E=D.C.X.Z, E=D.C.X.Z.Y, and E=D.C'.X.Z.Y, which contain too much information to have a chance to fit in 64-bits. In order to limit the amount of ISP IPv6 address space devoted to 6rd+, C and C' may therefore be chosen significantly longer than C". These locators, which are too long to be used as link prefixes on ordinary IPv6 links, can be used to obtain native-IPv6 host addresses (duly completed with trailing 0s if needed to reach 128 bits). Other possible uses of these locators are beyond the scope of this draft.

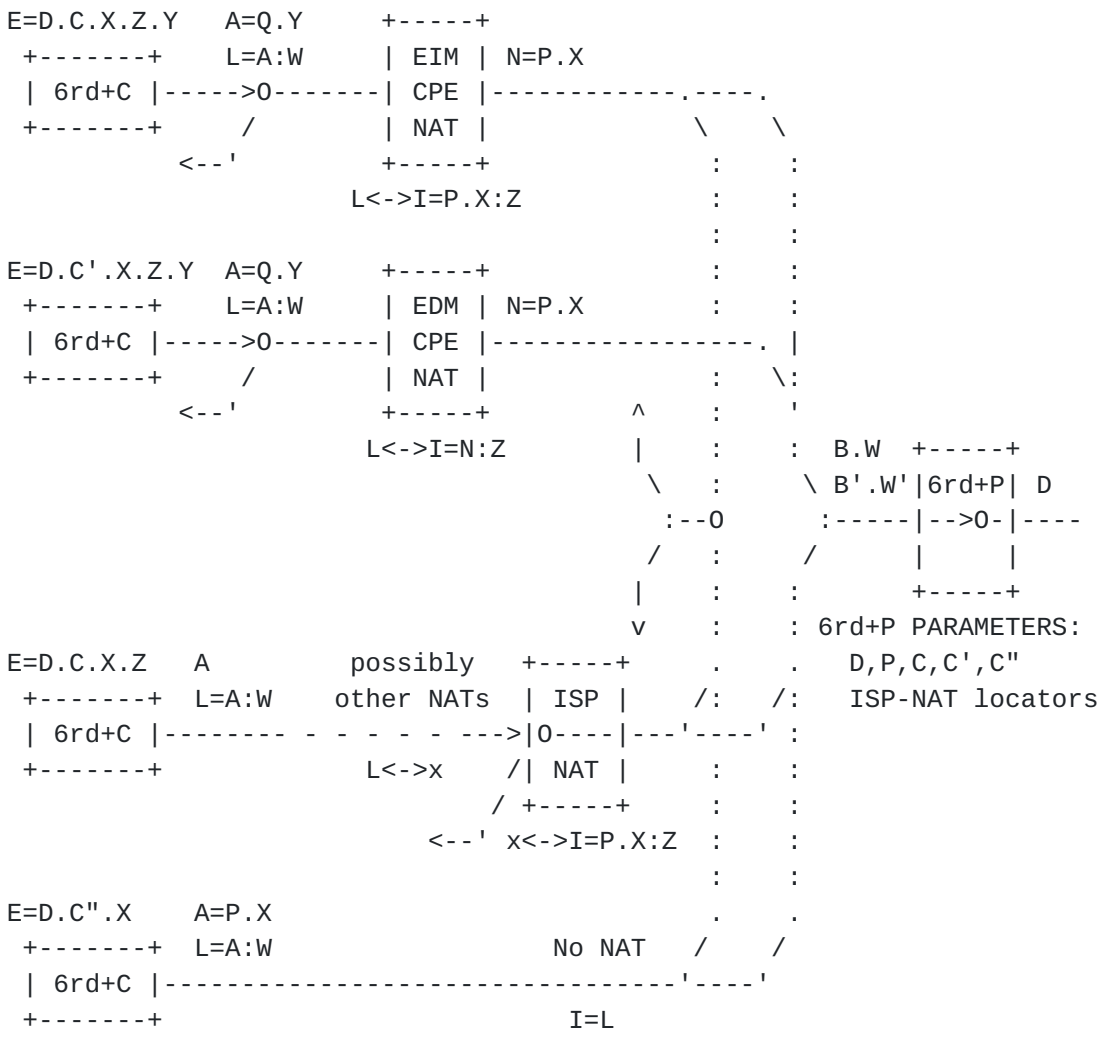


```
E       : 6rd+C IPv6 locator
D       : ISP prefix assigned to 6rd+
C,C',C" : NAT-type code (resp. EIN or ISP, EDM, NoNAT)
X       : complement of P in N
Z       : complement to N in I (port number)
Y       : complement to Q in L
A       : 6rd+C local IPv4 address (= Q.Y)
W,W'    : 6rd+ well-known UDP ports (to be assigned by IANA)
Q       : prefix of RFC 1918 in A (10/8, 172.16/12, or 192.168/16)
L       : 6rd+C local locator (= A.W)
I       : 6rd+C interior locator (IPv4+)
N       : IPv4 part of the interior locator I
P       : common prefix of the interior address space (possibly /0)
B,B'    : 6rd+ well-known IPv4 public addresses (to be assigned by IANA)
```

6RD+ LOCATOR FORMATS

Figure 2

3.3. 6rd+ Datagram Formats

6rd+Cs and 6rd+Ps encapsulate IPv6 packets they have to forward across their 6rd+ domain. For this they have to derive the interior destination address of each encapsulating packet from addresses contained in the IPv6 packet it encapsulates. If this interior destination has 32bits, encapsulation is in an IPv4 datagram, with the protocol field set to 41 (IP in IP). If it has 48 bits, encapsulation is in a UDP datagram. Its destination port is copied from the last 16 bits of the interior destination, and its source port is the 6rd+ well-known prefix W. Formats of IPv4+ datagrams are shown on Figure 3 and Figure 4.

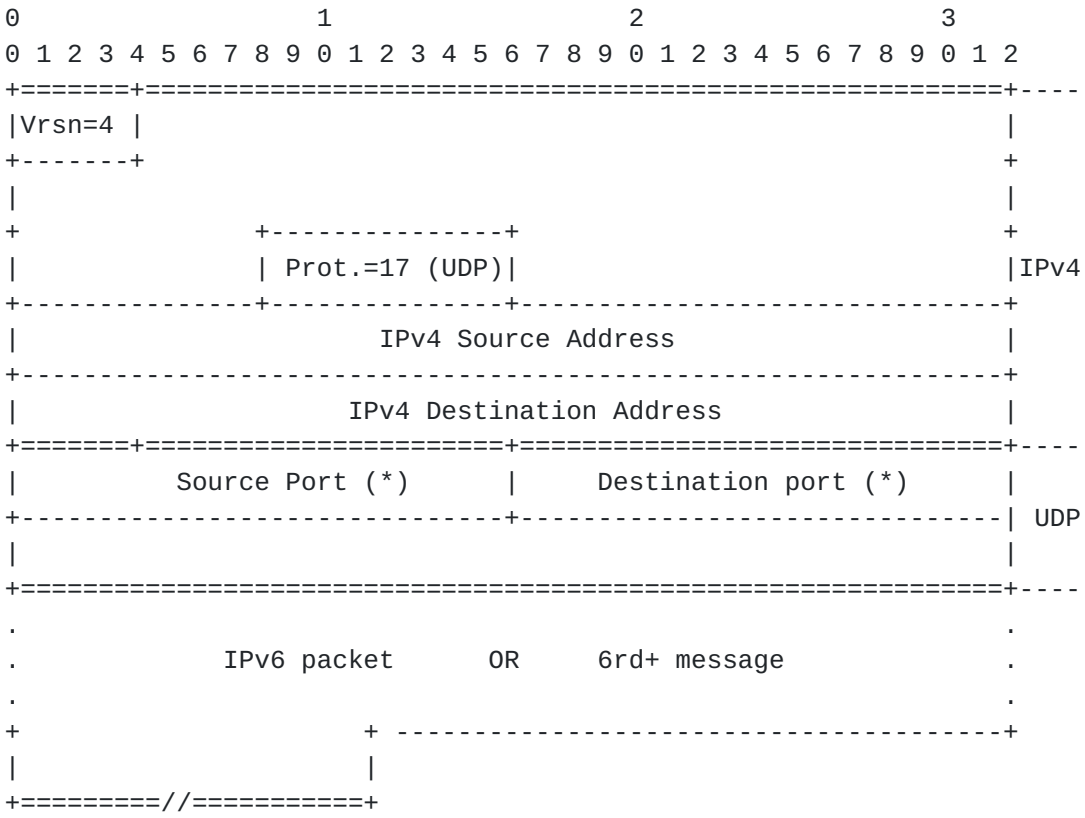

(*) At least one of the two ports is one of the two 6rd+ well-known prefixes W and W'.

FORMATS OF 6RD+ DATAGRAMS FOR IPV4+ INTERIOR DESTINATIONS

Figure 3

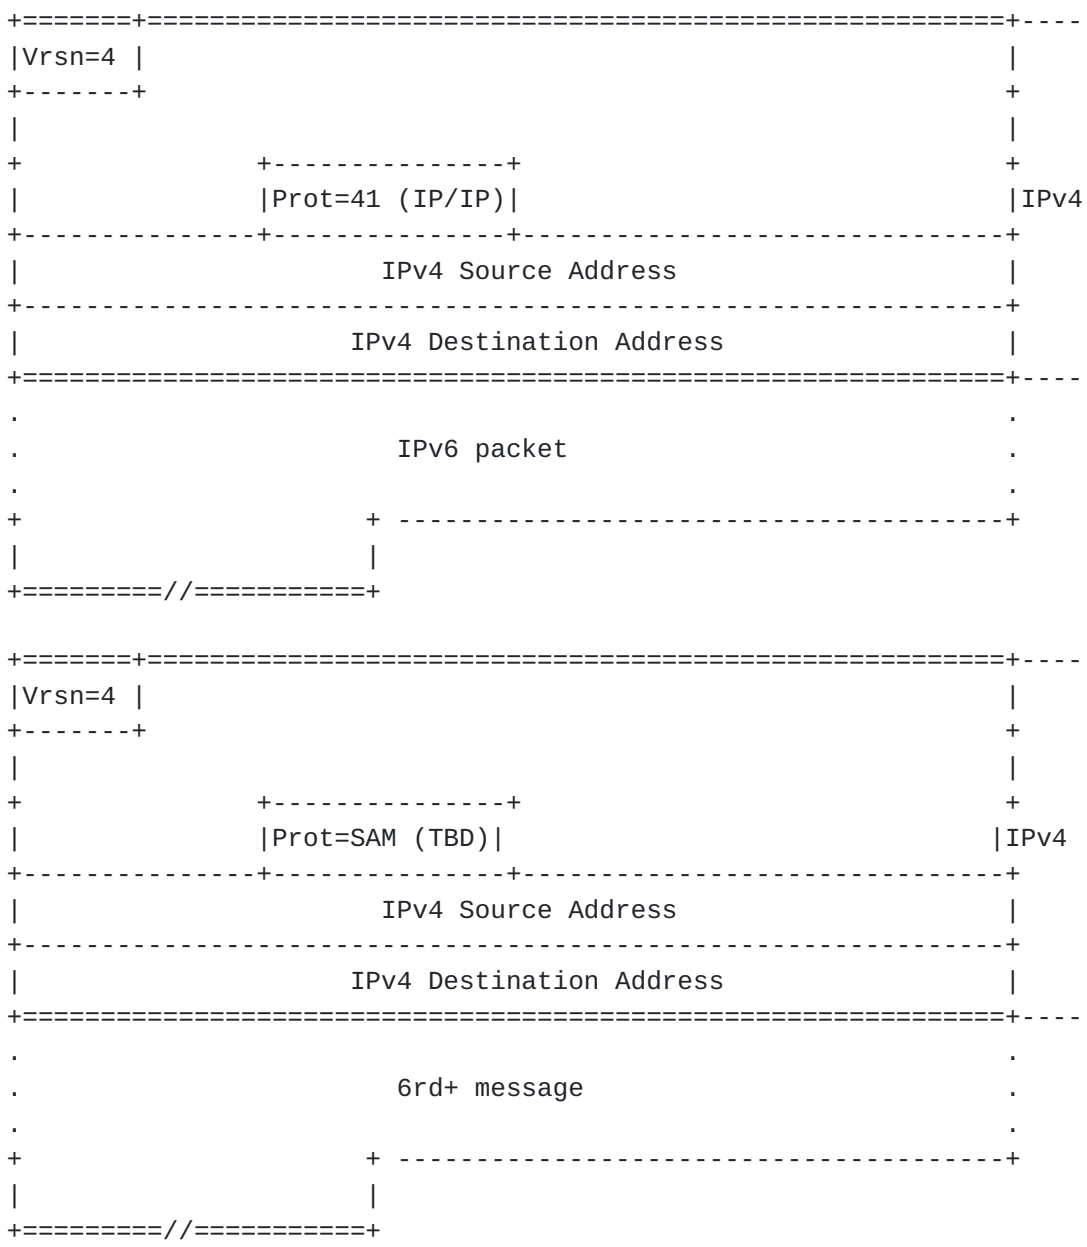

FORMATS OF 6RD+ DATAGRAMS FOR IPV4 INTERIOR DESTINATIONS

Figure 4

Recommendations of [RFC4213] that concern these encapsulations have to be followed.

Since IPv6 packets are encapsulated in IPv4 datagrams which may be fragmented on their way, a large IPv6 packet can in general be encapsulated in a single IPv4 datagram, independently of path MTUs of the traversed domain. However, reassembly of multi-fragment datagrams tends to consume processing power and can bring some

buffer-size problems. It is therefore advisable to limit the size of accepted IPv6 packets (and to return ICMPv6 packet-too-big error messages when received IPv6 packets are too big to be forwarded). Forwarding only IPv6 packets that don't exceed the 1280 octets, the size that [RFC2460] requests every link to support, is the choice that minimizes the risk of such reassembly problems. Yet, accepting somewhat larger sizes may be preferred to obtain a different trade-off between reassembly avoidance and accepted IPv6 packet sizes.

The protocol ID to be used for 6rd+ messages exchanged directly in IPv4 datagrams is that of the more general [SAM] protocol, of which the 6rd+ protocol is a subset.

3.4. Mapping Rules of 6rd+Cs and 6rd+Ps

Mapping rules that, in 6rd+Cs, determine interior destinations iDSTs are detailed in Table 1. For each NAT type, the listed rules have to be tested in the indicated order, until one is found to apply. If none applies, the IPv6 packet is invalid, and silently discarded.

The logic of these mapping rules directly results from shortcuts that have to be supported, and from IPv6 locator formats of Section 3.2.

In Table 1:

- Lower-case letters represent lengths of fields whose symbols are upper-case letters (e.g. d is the number of bits of D).
- Column 1 indicates the NAT type of the 6rd+C (EIM CPE NAT, EDM CPE NAT, ISP NAT, No NAT).
- In column 2, a prefix indicates that the rule applies only if the IPv6 destination address eDST starts with this prefix.
- In column 3, a number of bits n indicates that the rule applies only if both the IPv6 destination address and the 6rd+C IPv6 locator start with the same n bits.
- In column 4, a prefix indicates that the rule applies if the IPv6 source address eSRC starts with this prefix. A rule has contents in one and only one of the first three columns.
- In column 5, a prefix indicates that, if the rule applies, this prefix has to be included the interior destination iDST.
- In column 6, a number of bits n indicates that, if the rule applies, n bits in the IPv6 destination address eDST have to be skipped, after those that may have already been matched if any.

-1-	-2-	-3-	- 4 -	-5-	-6-	-7-	-8-
NAT Tpe	eDST prfx	length of common prefix of eDST & IPv6-lctr	eSRC prfx	iDST prfx	Skip in eDST	Add to iDST from eDST	iDST sffx
EIM CPE NAT		d+c+32-p		Q	16	32-q	
"	D.C			P		48-p	
"	D.C"			P		32-p	W
"			D	B:W			
EDM CPE NAT		d+c'+32-p		Q	16	32-q	
"			D	B:W			
ISP NAT	D.C			P		48-p	
"	D.C"			P		32-p	W
"			D	B:W			
No NAT	D.C			P		48-p	
"	D.C"			P		32-p	W
"			D	B			

6rd+C MAPPING RULES

Table 1

- In column 7, a number of bits n indicates that, if the rule applies, n bits of the IPv6 destination eDST, after those that may have already been matched and/or skipped, have to be included in the interior destination iDST, after any bits it may have already been included.

- A column-8 square may contain a suffix. If the rule applies, it is included in the interior destination iDST, after any bits it may already contain.

6rd+P mapping rules are straightforward, as no possible shortcuts have to be considered. Using the same notation as in Table 1, they are detailed in Table 2 (without any implication on how they can be implemented).

As shown, packets exchanged between 6rd+Ps and 6rd+cs of the "No NAT" type are encapsulated in UDP/IP rather than directly in IPv4 as it would have been possible. This adds a little overhead to these packets, but has two advantages: hairpinning in 6rd+Ps for 6rd+Cs of the "No NAT" that exchange packets with 6rd+Cs of the EDM CPE NAT" type is facilitated, as received and transmitted datagram headers have the same length; routing loops with other tunneling mechanisms such as [6to4], [6rd], and [ISATAP], are made impossible (these three mechanisms encapsulate IPv6 in IPv4 while 6rd+Ps only encapsulate IPv6 in IPv4+ (see also Section 3.6.

eDST prefix	iDST prefix	Add to iDST from eDST
d+c	P	48-p
d+c'	P	48-p
d+c"	P	48-p

6rd+P MAPPING RULES

Table 2

3.5. Acquisition of 6rd+ parameters by 6rd+Cs

This 6rd+ protocol is a particular instance of the more general "stateless address mapping" protocol proposed in [SAM], where the exterior address family is IPv6 and the interior address family is IPv4+. Rules of the 6rd+ protocol, partially illustrated in Figure 5, are the following:

R1: If a 6rd+C that has an IPv4 address but no IPv6 address or prefix, it tries to obtain its parameters from a 6rd+P (its IPv6 locator an its mapping rules). For this, it sends two SAM messages containing a "Parameter Request" indication. The first one is sent to IPv4 address B. The second one is sent IPv4+ address B:W, with the local IPv4 address A as parameter. These two requests are also retransmitted from time to time to check whether newly received answers differ from previous ones, and for the second one to maintain mappings in all NATs that may exist between the 6rd+C and 6rd+Ps.

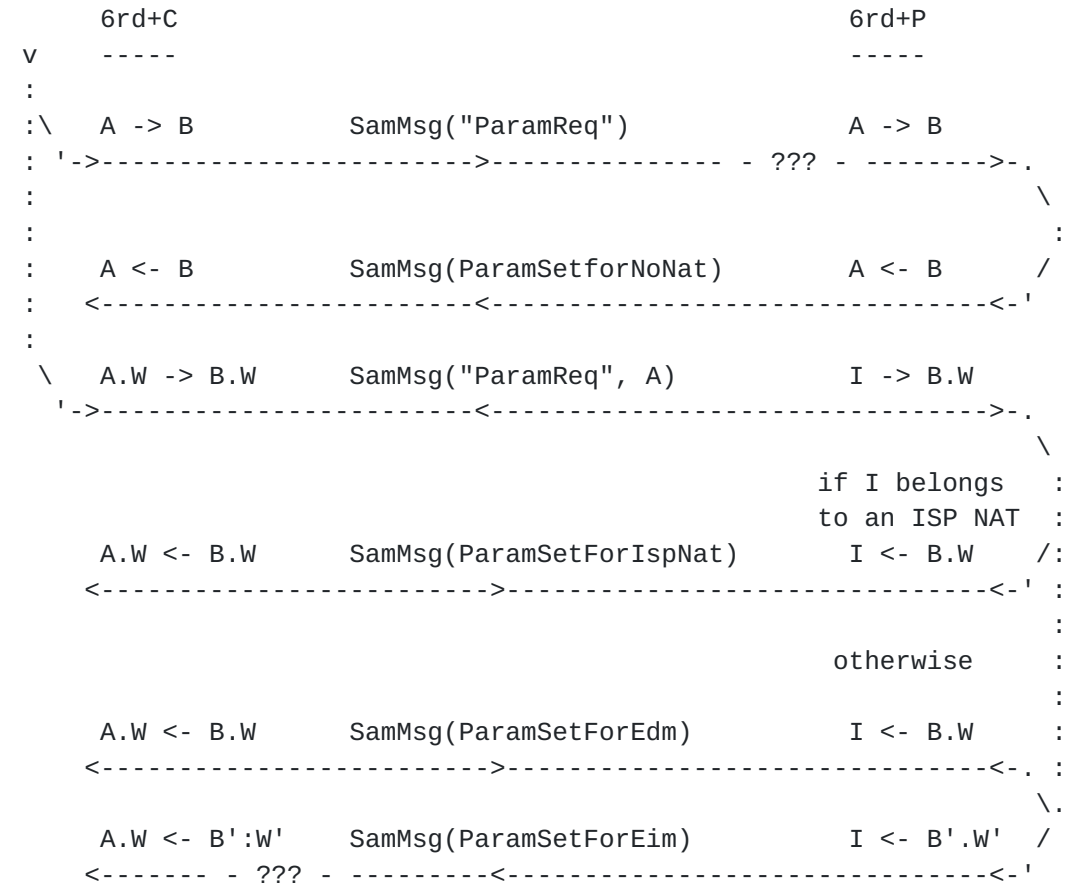

6rd+ PROTOCOL

Figure 5

R2: If a 6rd+P receives a SAM message at its IPv4 address B without a UDP header, it knows that this message crossed no NAT on its way. If the message contains a parameter request, it returns a SAM message containing parameters suitable for this 6rd+C, known to be of the "No NAT" type.

R3: If a 6rd+P receives a SAM message at its IPv4+ address B:W, it first tests whether its IPv4 source address starts with the IPv4 prefix of one ISP NATs of the domain. if yes, it returns an answer containing parameters suitable for this 6rd+C, known to be of the "ISP-NAT" type. Otherwise, it returns two answers. The first answer is sent from the B:W IPv4+ address so that it can reach its destination with on its way NATs of any types. It contains parameters suitable for this 6rd+C if it is of "EDM CPE NAT" type. The second answer is sent from the B':W' IPv4+ address so that it can reach its destination only if either there is no NAT on its way, or if the first traversed NAT is EIM. It contains parameters suitable for this 6rd+C if it is of the "EIM CPE NAT' type.

R4: If a 6rd+C receives an answer, it first checks that its source address is valid (is B, B:W, or B':W'). If the received parameters are better than those of its current parameter set or if it has no current parameter set, it takes them as the new current parameter set. The decreasing order of best parameter sets is: (1) received from B (No NAT); (2) received from B':W' (EIM CPE NAT or ISP NAT); (3) received from B:W (EDM CPE NAT). If the received parameter set is less good as the current one, but has a lifetime longer than the the remaining lifetime of the current one, it should be memorized to possibly become current one if the current one becomes obsolete.

R5: If a 6rd+C receives at its IPv4 address Q.Y a 6rd+ datagram without UDP header, in which the IPv6 destination doesn't match its IPv6 locator D.X.Z.Y, it infers that an intra-site shortcut must have been attempted which cannot be taken. This may happen in the very specific case where: there is at least one NAT, internal to the site, between the source 6rd+C and the intended-destination 6rd+C; there is no NAT between the source 6rd+C and the receiving 6rd+C; the receiving 6rd+C has the same private IPv4 local address as the intended-destination 6rd+C. A SAM message is then returned to the IPv4 address of the source 6rd+C with a "SAM unreachability" indication, and the receiving 6rd+C disables its intra-site-shortcut mapping rule if it still has one (a rule leading to a 32-bits iDST).

R6: If a 6rd+C receives at its IPv4 address a SAM message containing a "SAM unreachability" indication, it disables its intra-site-shortcut mapping rule if it has one. With this rule ad the previous one, IPv6 connectivity between 6rd+C IPv6 locators is always ensured even in complex intra-site configuration with internal NATs, with IPv6 packets having in this case to go via the CPE.

Detailed formats of SAM messages used for 6rd+ are beyond the scope of this document. They should be specified in the wider context of [SAM].

3.6. Anti-Spoofing and Anti-Routing-Loop Protections

Anti-spoofing protection can be ensured by applying the general ingress-filtering principle: a packet received at an interface is valid only if the same packet with inverses source and destination may be transmitted at that same interface:

- An IPv6 packet received by a 6rd+C or by a 6rd+P in an IPv4+ datagram must be silently discarded if the source address of the datagram (IPv4 or IPv4+) is not that which mapping rules obtain as iDST if applied to a an IPv6 destination eDST equal to the IPv6 source address of the datagram.

- Routing-loop protection is necessary if an ISP operates two point-to-multipoint tunneling mechanisms on the same domain with the same exterior and interior address families. (Example with [ISATAP] and [6to4] are analyzed in [RoutingLoops].)

- As long as 6rd+ as specified here is the only point-to-multipoint tunneling mechanism used with IPv6 as exterior address space and IPv4+ as interior address space, no particular protection against routing loops is needed for IPv4+ tunnels.

4. Security considerations

With precautions taken in previous sections, no new security issue has been identified so far. More work on this specification is however desirable to improve confidence in this respect.

5. IANA Considerations

For plug-and-play operation of 6rd+, the following assignments have to be made by IANA:

- The two well-known IPv4 addresses B and B' and the two well-known UDP ports W and W' of Section 3.2;
- The SAM protocol ID of Section 3.3 (to be shared with [SAM]

6. Acknowledgments

A very early version of this proposal has been informally submitted to a number of delegates at IETF 77. Thanks are due to Dan Wing and Pascal Thubert for their encouragements to pursue in this direction, and also to Brian Carpenter for his strong support that the objective of 6rd+: native IPv6 across legacy CPEs.

7. References

7.1. Normative References

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

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

7.2. Informative References

[6rd] Townsley, M. and O. Troan, "IPv6 via IPv4 Service Provider Networks - draft-ietf-softwire-ipv6-6rd-10", February 2010.

[6rd UDP] Lee, Y. and P. Kapoor, "UDP Encapsulation of 6rd - draft-lee-softwire-6rd-udp-01", May 2010.

[6to4] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", February 2001.

[ARP] Plummer, D., "An Ethernet Address Resolution Protocol -- or -- Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", RFC 826, November 1992.

[IPv6 /127 prefix] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., and L. Colitti, "Using 127-bit IPv6 Prefixes on Inter-Router Links - draft-kohno-ipv6-prefixlen-p2p-01", March 2010.

[ISATAP] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

[ND] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

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

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

[RFC3053] Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.

[RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", RFC 5565, June 2009.

[RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4

infrastructures (6rd)", January 2010.

[RoutingLoops]
Nakibly, G. and F. Templin, "Routing Loop Attack using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations - draft-nakibly-v6ops-tunnel-loops-02", May 2010.

[SAM] Despres, R., "Stateless Address Mapping (SAM) for Softwire-Lite Solutions (NOTE: only the revised draft-despres-softwire-sam-01 will be consistent with this draft. It should be posted before the deadline for IETF 78)", July 2010.

[SAMPLE] Carpenter, B. and S. Jiang, "Legacy NAT Traversal for IPv6: Simple Address Mapping for Premises - Legacy Equipment (SAMPLE) - draft-carpenter-softwire-sample-00", June 2010.

[STUN] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.

[Teredo] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

[Tunnel Broker]
Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.

Author's Address

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

Email: remi.despres@free.fr