Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: September 25, 2009

R. Despres
March 24, 2009

Stateless Address Mapping (SAM)
Avoiding NATs and restoring the end-to-end model in IPv6
draft-despres-sam-02

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on September 25, 2009.

Copyright Notice

Abstract

Stateless Address Mapping (SAM) is a generic mechanism to support global addressing across network zones where routing is based on a

different address space. With it, the end-to-end model, lost in IPv4 with the deployment of NATs, can be restored without losing services that NAT44s offer beyond address-space extension (private addressing, basic firewall, site multihoming, privacy protection, host-rooted subnets). Global-address packets are encapsulated in local-address packets to traverse SAM zones, and global prefixes are statelessly mapped into local addresses. For the IPv6-IPv4 coexistence period, port-restricted IPv4 addresses are used to extend the global IPv4 address space.

Table of Contents

1. Introduction

In IPv4, Network Address Translations have been extensively deployed (NAT44s). They are key to mitigate the IPv4 address shortage. But they also offer various auxiliary services, described in Section 2 : private addressing, basic firewall, site multihoming, privacy protection, host-rooted subnets.

In counterpart to these auxiliary services, these NAT44s have introduced two drawbacks:

- Non compliance with the end-to-end model of the Internet where addresses and ports are unchanged end to end (e2e).

  Negative consequences include incompatibility with the IPsec security mechanism, and difficulties for hosts to know their own global addresses, which they need for connection redirections, for host referrals, and, in sites having several site entrance routers, for multihoming support mechanisms like the SCTP of [RFC4960] and [Shim6].

- Stateful operation.

  Most NAT44s are in fact stateful NAPTs as defined in [RFC2663]: to support more local addresses than they have external addresses, they maintain per-transport-connection states. Negative consequences include limited scalability, and the risk of denial of service attacks that go with it, as well as single points of failures.

Since no global address shortage is in view in IPv6, the following questions have to be asked:

- Which NAT44 services can, in IPv6, be offered statelessly and without breaking the e2e model?

- How?

This draft proposes to answer these questions, more completely and with more technical details than in [RFC4864], the most advance document on the subject so far.

For this, a Stateless Address Mapping generic mechanism is introduced (SAM).

The conclusion is that, provided SAM is supported in nodes at borders of independently administered routing zones, the e2e model can be restored in IPv6, for all identified useful functions of NAT44s.

(This conclusion needs however to be confirmed after further work on SAM details, after criticisms by other experts, after some possible bug corrections, and after validations with running code.)

Thus, traversal of NATs in ISP infrastructures can be avoided. (These NATs do provide useful connectivity to some non-SAM-capable nodes, but have the drawback of breaking the e2e model, with the mentioned consequences on security, referrals, multihoming, scalability, and reliability.)

2. NAT44 services that remain desirable in IPv6

2.1. Private addressing (easy renumbering)

With NAT44s, when a prefix assigned by an ISPs to a customer site is modified, local IP addresses in the site can remain unchanged.

2.2. Basic firewall (by default, no incoming connections)

Most NAT44s, being NAPTs, and therefore maintaining states for all TCP and UDP connections, have as a byproduct a protection against incoming connections (unless some "holes" are "punched" in this protection, under explicit customer control). This level of security protection is largely relied upon.

2.3. Site multihoming (automatic fallback)

In a site is multi-homed, and if it has a NAT device supporting all its ISP interfaces, its hosts can take advantage of multihoming without having to support any multihoming-specific function. This level of multihoming support is better than none.

(For this, a NAT44 needs only to make sure that, for each transport connection, all outgoing packets go through the same ISP. Thus, if an ISP access fails, current TCP and UDP connections that go via this ISP are broken, but they can immediately be replaced by new ones.)

2.4. Privacy protection

From outside a site where a NAT44 operates in NAPT mode, it is difficult to determine which hosts establish which connections. This level of privacy protection, in particular for some web requests, is an added value.

2.5. Host-rooted subnets

Behind a host that is assigned a single IPv4 address, it is possible, with a NAT44 in the host, to deploy a private subnet. As modern operating systems include a router function with a NAT44, a computer can serve as a root for a LAN.

Thus, the distinction between hosts and a routers is no longer a distinction between types of devices. It has become only a distinction between functions within nodes.

3. SAM specification

3.1. Local zones - Root SAMs - Branch SAMs

As presented in Figure 1, the SAM mechanism applies to a SAM "local zone" Z. Routing within this zone is independently administered, and is based on a "local address space".

Each SAM zone has one or several "root interfaces" (Ri), that give access to the global Internet. Each one has, in the global Internet, one or several "global prefixes" (gZij) exclusively assigned to zone Z.

SAM global prefixes can be global IPv6 and/or global IPv4. SAM local address spaces can be IPv6 or IPv4, global or private. If both IPv4 and IPv6 are routed in the zone, one of the two is chosen for SAM. (SAM is in this respect an extension of the 6to4 of [RFC3056], of the ISATAP of [RFC5214], and of [6rd], where all global prefixes are IPv6 and all local address spaces are IPv4).

As explained in Section Section 3.7, global IPv4 addresses can be extended beyond 32 bits to deal with the IPv4 address shortage during the IPv4-IPv6 coexistence period.

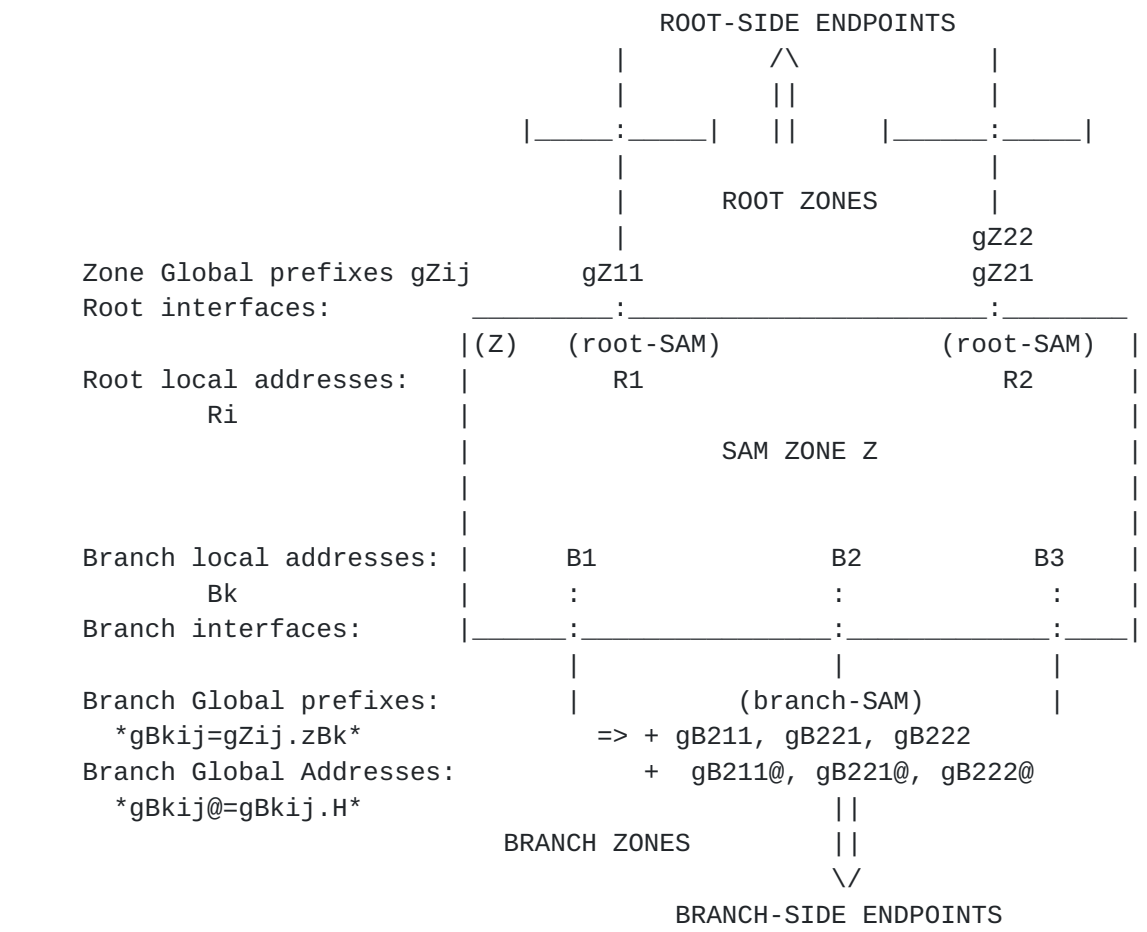

ROOT AND BRANCH INTERFACES AND SAMs

Figure 1

Each root interface that supports a root-SAM function has a local address (Rk), and each "branch interface" has a local address (Bk).

If a "branch SAM" function is supported at a branch interface Bk, this interface gets, in addition to its local address, global prefixes (gBkij). Each of these prefixes is made of a global prefix of the zone (gZij) followed by an identifier (zBk) of the branch in its zone.

For each each of its global prefixes gBkij, a branch interface has also a host global address (gBkij@), derived from the prefix by appending a standard host suffix (H) to complete the address length.

Thus, if a zone D is accessible from the global Internet via a zone hierarchy A, B, C, it has at least gA.aB.bC.cD as a global prefix gD, and gA.aB.bC.cD.H as a global address gD@. SAM is thus an application of the locator-identifier separation principle. (It

differs however from [LISP], in that no new protocol is needed for SAM (only new options in existing protocols such as DHCP [RFC2131], DHCPv6 [RFC3315], or ND [RFC4861], to advertise SAM parameters to branch interfaces.)

3.2. Encapsulation of global packets in local packets

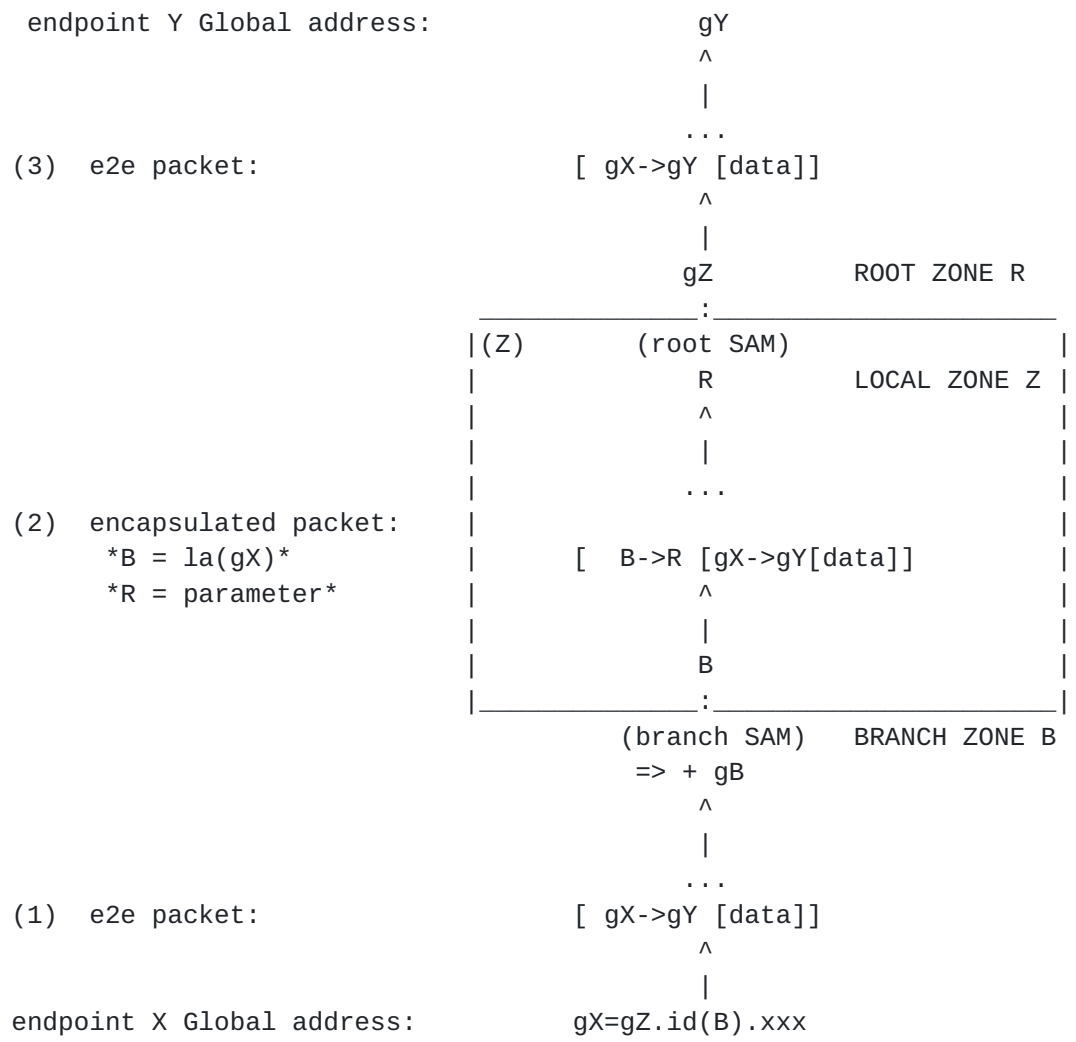

PACKET ENCAPSULATION AND ADDRESS MAPPING - BRANCH SIDE TO ROOT SIDE

Figure 2

To traverse a SAM local zone, global-address packets are encapsulated into local address packets, as illustrated in Figure 2 and Figure 3.

Thus, compatibility is ensured, within the local zone, with the ingress filtering for multihomed networks of [RFC3704], the basic anti-spoofing mechanism.

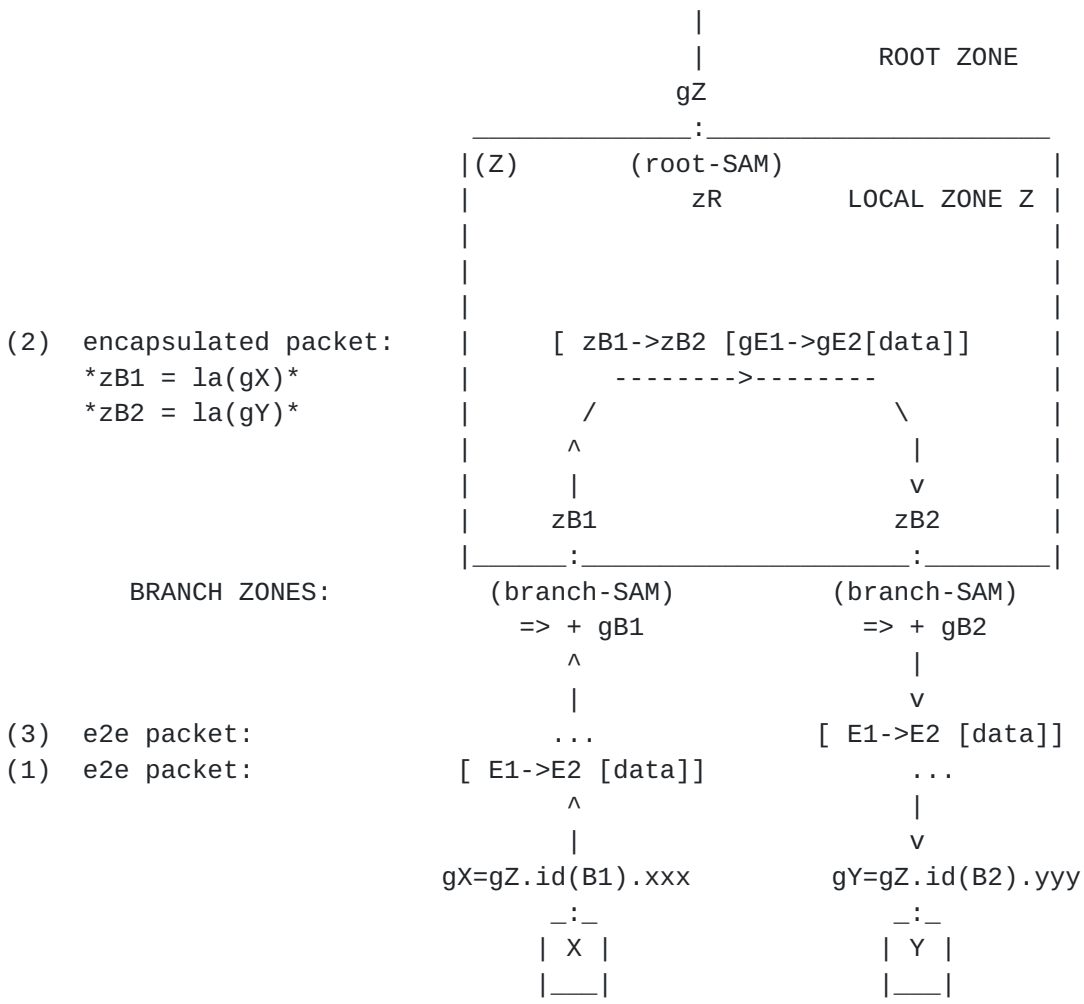

ADDRESS MAPPING AND PACKET ENCAPSULATION - BRANCH SIDE TO BRANCH SIDE

Figure 3

For the IP-in-IP encapsulation, the IPv6 next header or the IPv4 protocol id which indicates the type of IP payload is set to 41 (the same value as for 6to4, ISATAP, and 6rd).

Local addresses are determined as follows (illustrated in Figure 2 and Figure 3):

1. If an endpoint global address gE, indifferently source or destination, is that of a branch-side endpoint, this is recognized by the fact that it starts with one of the global prefixes of the zone. Then, the local address B is obtained by a function B=la(gX), completely determined by SAM parameters of the zone (details in Section 3.4).

2. If an endpoint global address gE, indifferently source or destination, is that of a root-side endpoint, this is recognized by the fact that it doesn't start with any of the global prefixes of the zone. In this case, the other address gX of the packet, destination or source respectively, is necessarily that of a branch-side endpoint (otherwise the packet would not traverse the local zone). Then, local address Ri is that of the root interface that has, in its assigned global prefixes, the global prefix present at the beginning of the branch-side address gX.

In multihomed sites, the second of these rules ensures compatibility with the ingress filtering of [RFC3704] in root zones (if it does apply, as necessary for anti-spoofing protection).

In Figure 2 and Figure 3, packets in the reverse direction, not shown, would have the same addresses but with sources and destinations inverted, and with encapsulations and decapsulations made at inverted interfaces.

Decapsulation functions MUST verify, for anti-spoofing protection, that local addresses present in headers of encapsulating packets are consistent with global addresses present in headers of encapsulated packets.

3.3. Global prefixes - global addresses - local addresses

Internal structures of SAM global prefixes, global addresses, and local addresses are detailed in Figure 4.

A branch-interface global prefix necessarily starts with a global prefix of the zone Z. Its remaining bits are a "branch identifier" in the zone (gBkij = gZij.zB).

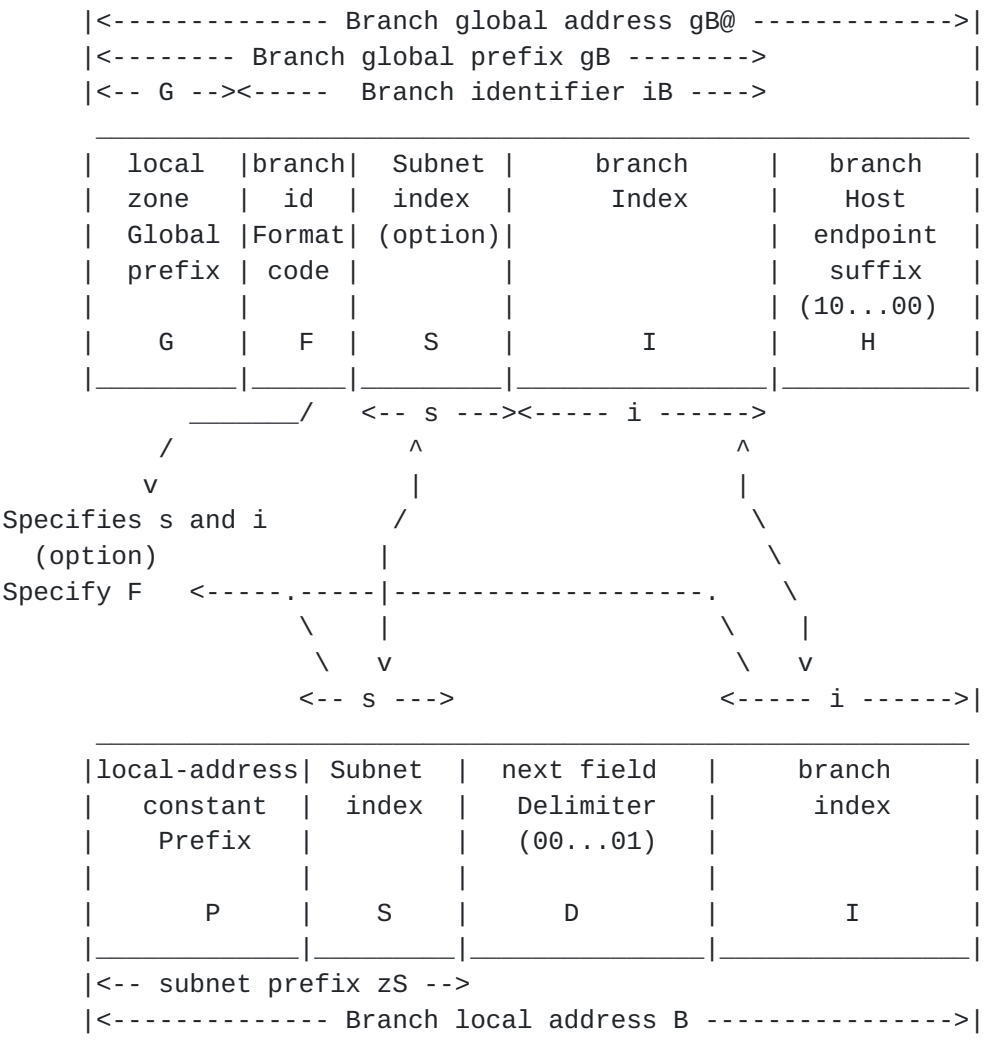

SAM GLOBAL PREFIXES - GLOBAL ADDRESSES - LOCAL ADDRESSES

Figure 4

Principles that influence the internal structure of branch identifiers proposed for SAM are the following:

1. To permit a flexible hierarchy of local zones, branch identifiers should be kept rather short. They should, at least to some extent, be proportionate to the maximum number of branches supported in their zone.

2. Several subnets must be possible in the zone. For this, a branch identifier contain an optional "subnet index" (S), followed the "branch index" (I) which identifies the branch in its subnet. (The word "index" is chosen to express that these fields have no further internal structure.)

3. For the efficiency of routing tables, intra-zone subnet indexes have to be in the upper part of local addresses, just behind the "constant prefix" (P) that is common to all local addresses. (In IPv6, this constant prefix is typically an ULA prefix of [RFC4193]; in IPv4, it is typically a private-address prefix of [RFC1918].)

4. For efficiency of the neighbor discovery protocol of [RFC2461], branch indexes B have on the contrary to be in the lowest part of branch local addresses B.

5. Consequently, it must be possible to extract separately, from a intra-zone branch identifier iB, the subnet index S and the interface index I, and for this to know their lengths (s and i).

6. In order to permit to configure several subnet-index lengths, and/or several interface index lengths, in SAM zones, an optional branch-identifier "format code" (F) is placed at the beginning of a branch identifier B (just before the optional subnet index S and the branch index I). Each format codes specifies a subnet-index length s and an interface-index length i. To be recognized, format codes that have different lengths must be non overlapping prefixes.

Since the local address B of a branch interface starts with a constant prefix P followed by the interface subnet index S , and is terminated by the interface-index of the interface, space is left between them. It is filled with a next-field delimiter (D). Its format, a series of 0s followed by a 1, i.e. 00...01 with a minimal length of 1 bit, is chosen so that knowing the constant prefix P and the subnet prefix of a branch interface, lengths s and i of the its subnet index S and of its interface index I can be determined. Then, the identifier format F to be placed in global prefixes of B can be derived from these lengths s and i.

3.4. Endpoint global address to branch local address mapping

Detailed steps by which a branch local address B is derived from the global address of a branch-side endpoint are presented in Figure 5.

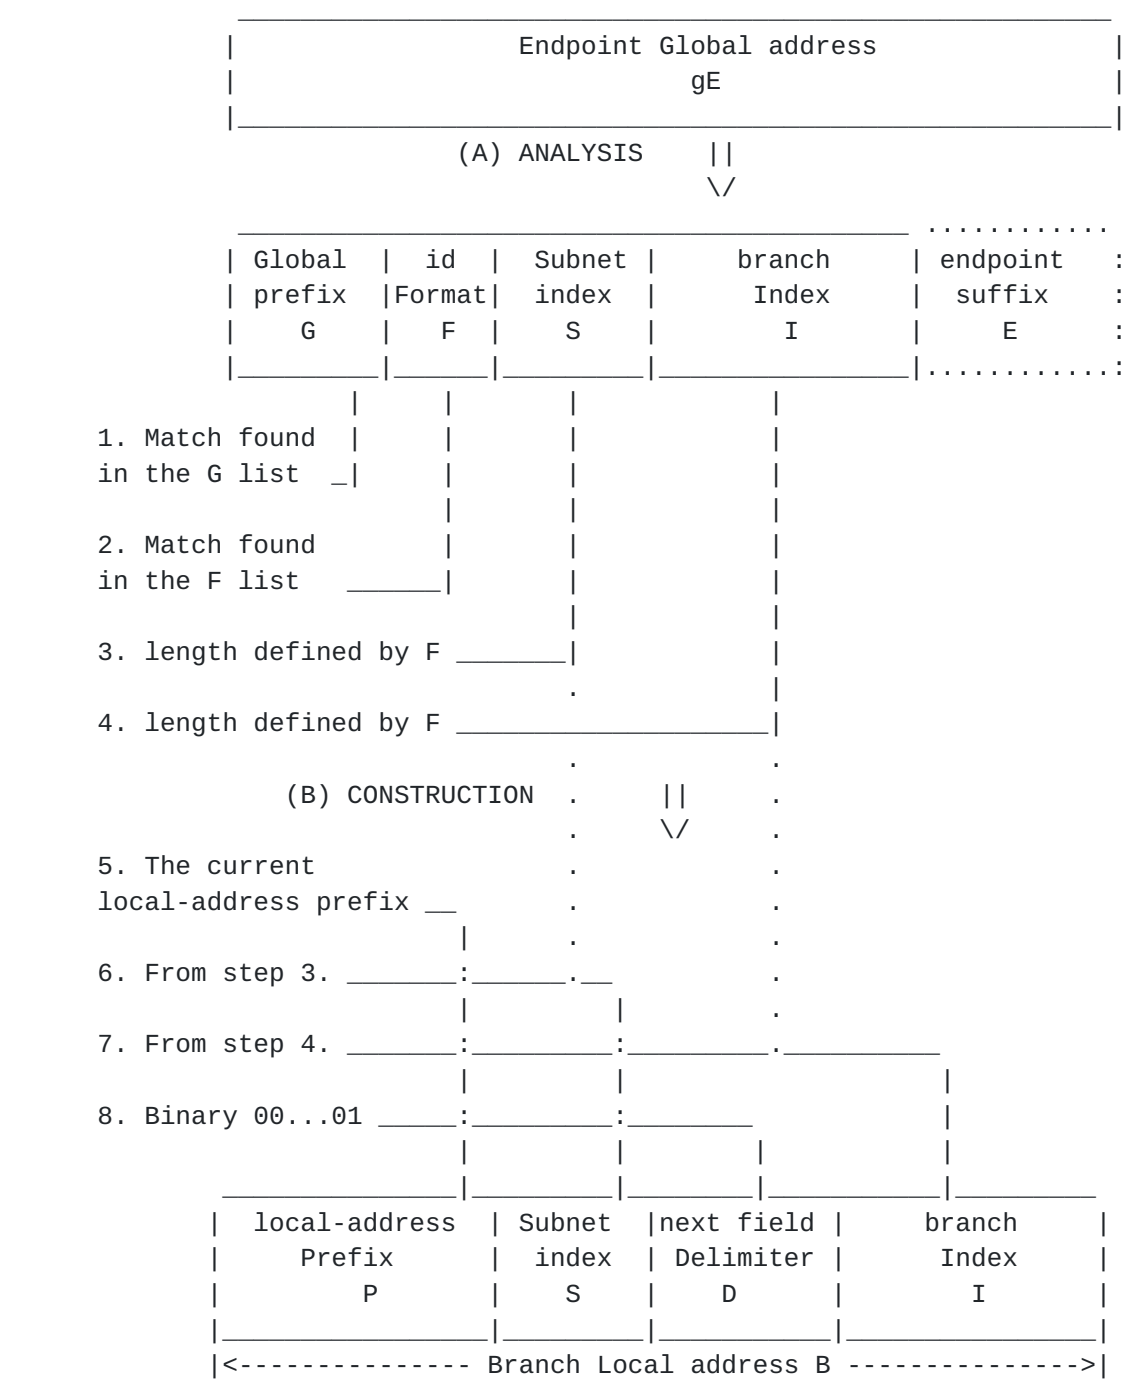

DERIVING A BRANCH LOCAL ADDRESS FROM AN ENDPOINT GLOBAL ADDRESS

Figure 5

3.5. Privacy protection

In a zone where privacy protection is desired, the privacy option can be turned on. Principles of this option are the following:

1. Fields that identify branch-side IP endpoints in privacy protected zones, or transport endpoints if endpoints are at this layer, are obfuscated in e2e packets that traverse the global Internet.

2. This obfuscation is stateless and reversible.

3. Branch SAMs of a privacy-protected zone are informed of parameters of this obfuscation. They can thus know which "hidden" addresses (or addresses plus ports), appear on the global Internet in place of their "clear" addresses (or address plus ports). These clear addresses are those from which local addresses are derived in the privacy-protected zone and in zones that are lower in the hierarchy.

4. In these lower zones, all branch SAMs are informed that a root SAM in the global-Internet direction has activated a privacy option, and are informed of parameters of this option. They can thus derive a clear address (or address plus port) from an obfuscated address (or address plus port), and conversely. They can also avoid to activate the privacy option so that obfuscation is never done more than once.

Parameters of a privacy option are a privacy global prefix (PPm) and a scrambling multiplier (PMm). The prefix is that which, at the beginning of global addresses, is not obfuscated in the global Internet. The multiplier is an odd constant.

Obfuscation consists in a modulo 2^n multiplication by the scrambling multiplier, where n is the number of bits to be obfuscated. De-obfuscation is the modulo 2^n multiplication by the inverse of the scrambling multiplier (for odd numbers, such an inverse modulo 2^n always exists).

In hosts in which the branch SAM is informed of an active privacy option, applications that ask for their addresses and their ports at their socket interface, get them in hidden form, that which appears in the global Internet. The e2e model is thus preserved despite the fact that the topology of the privacy-protected zone and that of lower zones in the hierarchy are all hidden, and despite the fact that successive transport connections from a same host cannot, in the global Internet, be related to a single host.

Ports that are concerned with the privacy option are only the IANA dynamic and/or private ports (ports 49152 to 65535, those starting with binary 11). Well known ports and registered ports, which have an e2e meaning not to be lost, must not be obfuscated.

Since some applications, e.g. active mode FTP of [RFC0959], work on port pairs rather than on individual ports, port bits to be obfuscated must exclude the last one. Port bits that are part of obfuscated endpoint identifiers are then bits 2 to 14.

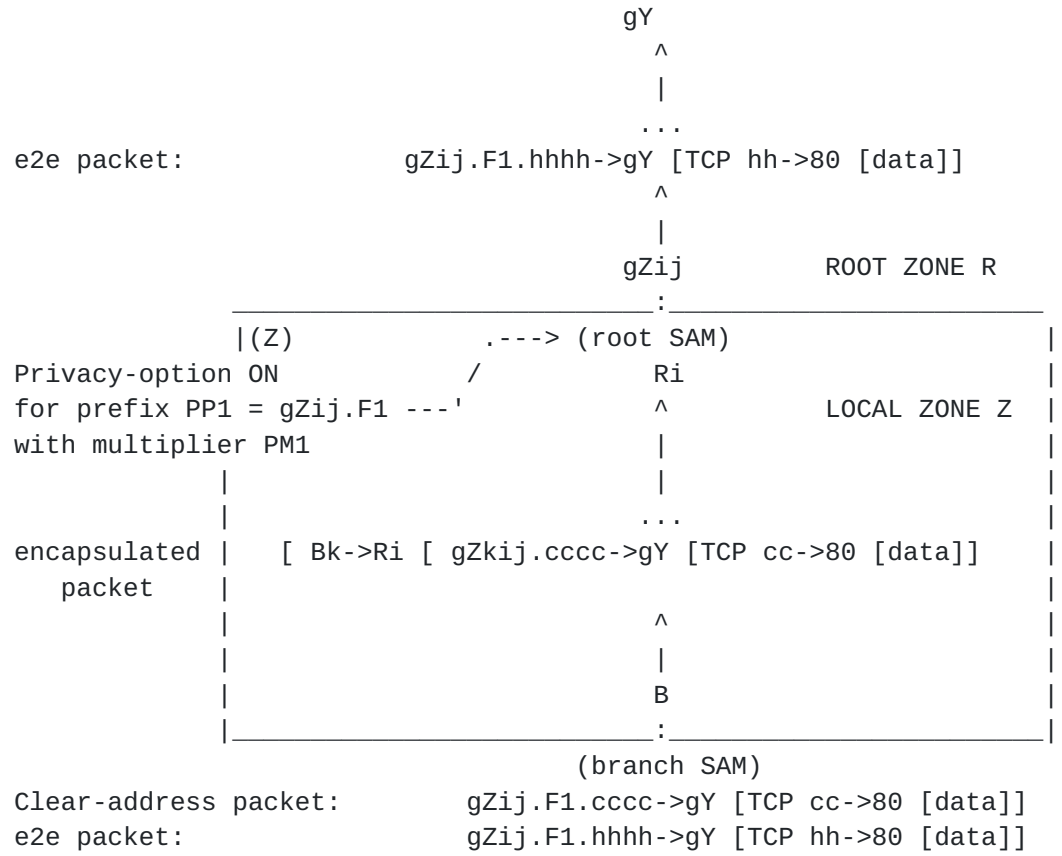

where . tmp = modulo 2^m (PM1 x (cccc . (bits 2 to 14 of cc))
where m = length of cccc + length of cc - 3
. hhhh = bits 0 to (length of hhhh - 1) of tmp
. hh = cc in which bits 2-15 are replaced by
bits(length of PP1 TO m - 1) of tmp

PRIVACY OPTION ILLUSTRATION

Figure 6

Figure 6 illustrates the effect of the privacy option. The option is supposed to be on in the root SAM of the zone, for its global prefix

gZij and its identifier format F1. The privacy-option prefix is therefore PP1 = gZj.F1. The scrambling multiplier is PM1.

3.6. SAM parameters

Table 1 to Table 4 present the complete set of SAM parameters described in previous sections.

Constant local prefix	TTL
...	...
Pm	PTm
...	...

LOCAL PREFIXES

Table 1

Identifier Format Code	TTL	Subnet-index Length	Interfacet-index Length
...	...	...	...
Fn	FTn	SLn	ILn
...	...	...	...

IDENTIFIER FORMATS

Table 2

Root local address	TTL	Global prefix 1	TTL1	...	Global prefix j	TTLj	...
...	...	...	...	...	...	...	...
Ri	RTi	gZi1	gZTi1	...	gZij	gZTij	...
...	...	...	...	...	...	...	...

ROOT PARAMETERS

Table 3

Privacy-option Prefix	TTL	Privacy-option Multiplier
...	...	...
PPp	PTp	PMp
...	...	...

PRIVACY OPTION

Table 4

3.7. Port range based extended IPv4 addressing

For a dual stack host not to break the e2e model when it establishes a connection with an remote endpoint that is still only reachable in IPv4, it must have a global IPv4 address. Because of the IPv4 address shortage, this address may however be shared with other hosts. For this, SAM accepts "port-extended" IPv4 prefixes, longer than 32 bits. Bits beyond the first 32 define a port range in the set of dynamic and/or private ports (those in which the two high order bits are binary 11). For example, a 3-bit prefix extension 010 imposes that branch-side hosts use only ports starting with binary 11010.

Note that, due to the systematic encapsulation of global packets in local packets of SAM, routing within SAM zones is not concerned with theses "port-extended" IPv4 addresses. Only root SAMs and branch SAMs have to know about port ranges.

The branch SAM in a host that is assigned a port-restricted IPv4 address has to inform its socket interface of the port range available to applications, and to inform its internal NAT if it has one. Consequences for applications, and for NATs, of restricted port ranges, are out of the scope of this SAM specification. Other documents are available on the subject, e.g. [Boucadair], which however requires further study.

4. SAM Application Examples

4.1. Private addressing in an IPv6 site

In the example of Figure 8, we consider a home or SOHO site in which an Ethernet and/or WiFi LAN is deployed. Its global IPv6 prefix gZ is 2001:0db8:9999::/48.

Local addressing is done in an IPv6 private space. To keep addresses

short in the figure, their constant prefix is fc00/8, the shortest prefix reserved for private IPv6 addressing in [RFC4193]. (Note that this prefix could be replaced by a full fdxx: xxxx:xxxx::/48 prefix, as recommended in [RFC4193] for ULAs, without changing the substance of the example.)

The site is configured to support 255 branch interfaces on the LAN (each branch being indifferently a host and/or a router). To facilitate future changes, a branch-identifier format code F1, set to 0/4, is used in branch global prefixes.

SAM parameters of the site are then following (ignoring TTLs):

Constant local prefix: P1 = fc00/8

Identifier format code: F1 = 0::/4

Subnet index length: SL1 = 0 (non applicable)

Interface index length: IL1 = 8

Root local address: R1 = fc00::0101

Zone Global prefix: gZ11 = 2001:0db8:9999::/48

Privacy option prefix: none in this example

We now consider a SAM-capable PC which serves as a router for a bluetooth link. On this link, a bluetooth mobile phone is active. (Configuring a root-SAM in the PC would permit the mobile phone, if acting as a SAM-capable router, to assign global prefixes and addresses to hosts behind it. But this would have been too much for the example).

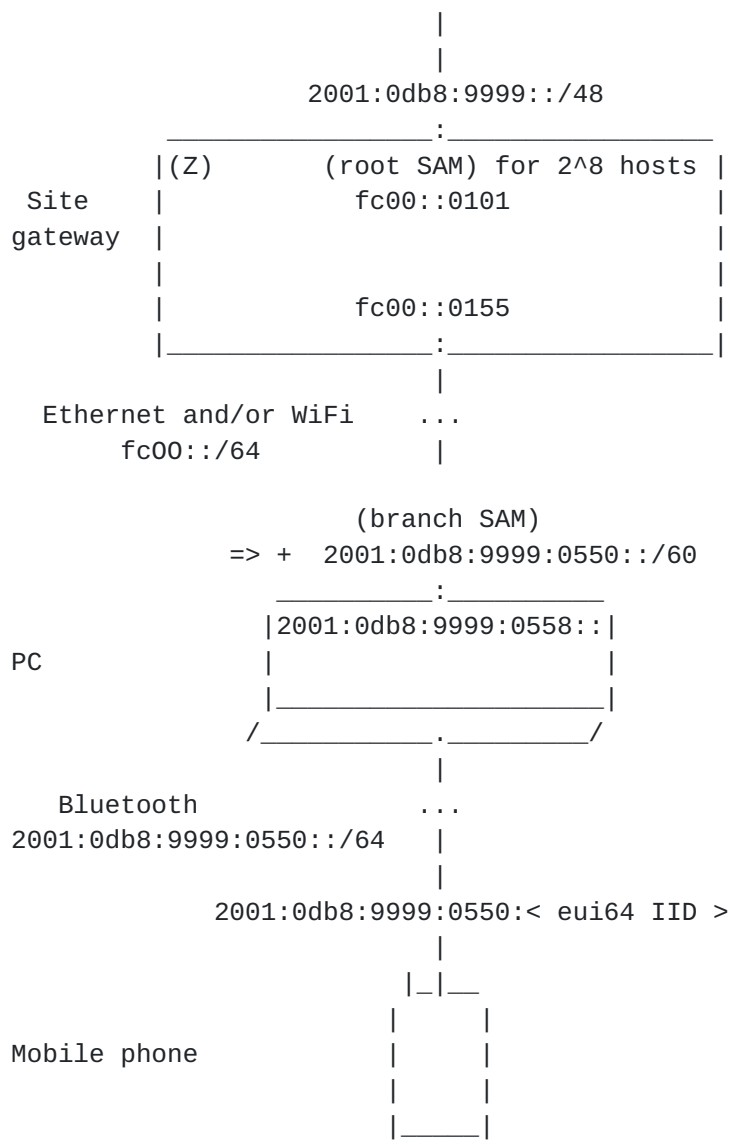

PRIVATE ADDRESSING IN AN IPV6 SITE

Figure 7

The PC local address B is fc00::0155, i.e. P.D.I where P is fc00::/8, where the 8 bits of I are supposed to be 55::/8, and where D is binary 00...01 with consequently (128 - 8 -8) = 112 bits.

The PC global prefix gB is therefore 2001:0db8:9999:0550::/60, i.e. G.F.I, where G is 2001:0db8:9999::/48, where F is 0::/4, and where I is 55::/8.

The PC global address is therefore 2001:0db8:9999:0558::, i.e. gB.E where E is binary 10...00 with (128 - 48 - 4 - 8) = 68 bits.

The bluetooth link is supposed to have 0::/4 as subnet ID in the PC.Its /64 subnet prefix is therefore 2001:0db8:9999:0550::/64.

This simple example illustrates how the SAM logic permits to establish a hierarchy of routing zones where each host can become a router, and where the e2e model is preserved.

4.2. Multihoming and Extended IPv4 addressing in a home site

In the example of Figure 8, we consider a home site S, multihomed with two ISPs A and B.

ISP A assigns to the site IPv6 prefix 2001:1111:1111:1110::/60, and IPv4 address 192.0.2.1.

ISP B can only assign port-restricted IPv4 addresses to its sites because it has to support up to 2^16 sites, and has only for this an IPv4 /18 prefix (namely 198.16.0.0/18, i.e. v4|c610:0000:/18), and since 18 + 16 = 34 which exceeds 32. Having 2001:0db8::/32 as its IPv6 prefix, it assigns /48s to its customer sites, in particular 2001:0db8:0202::/48 to site S.

Half of its IPv4 address space, namely v4|c610:2000/19 is allocated to a NAT, to support sites that are not SAM capable. The other half, i.e. v4|c610::/19, is allocated to a root SAM, the local address of which is supposed to be 2001:0db8::1.

SAM parameters of the zone of ISP B are then the following:

Constant local prefix: P1 = 2001:0db8::/32

Identifier format code: F1 = ::/0 (non applicable)

Subnet index length: SL1 = 0 (non applicable)

Interface index length: IL1 = 16

Root local address: R1 = 2001:0db8::1:1

Zone Global prefix: gZ11 = v4|c610::/19 (=198.16.0.0/19).

Privacy option prefix: none in this example (::/0)

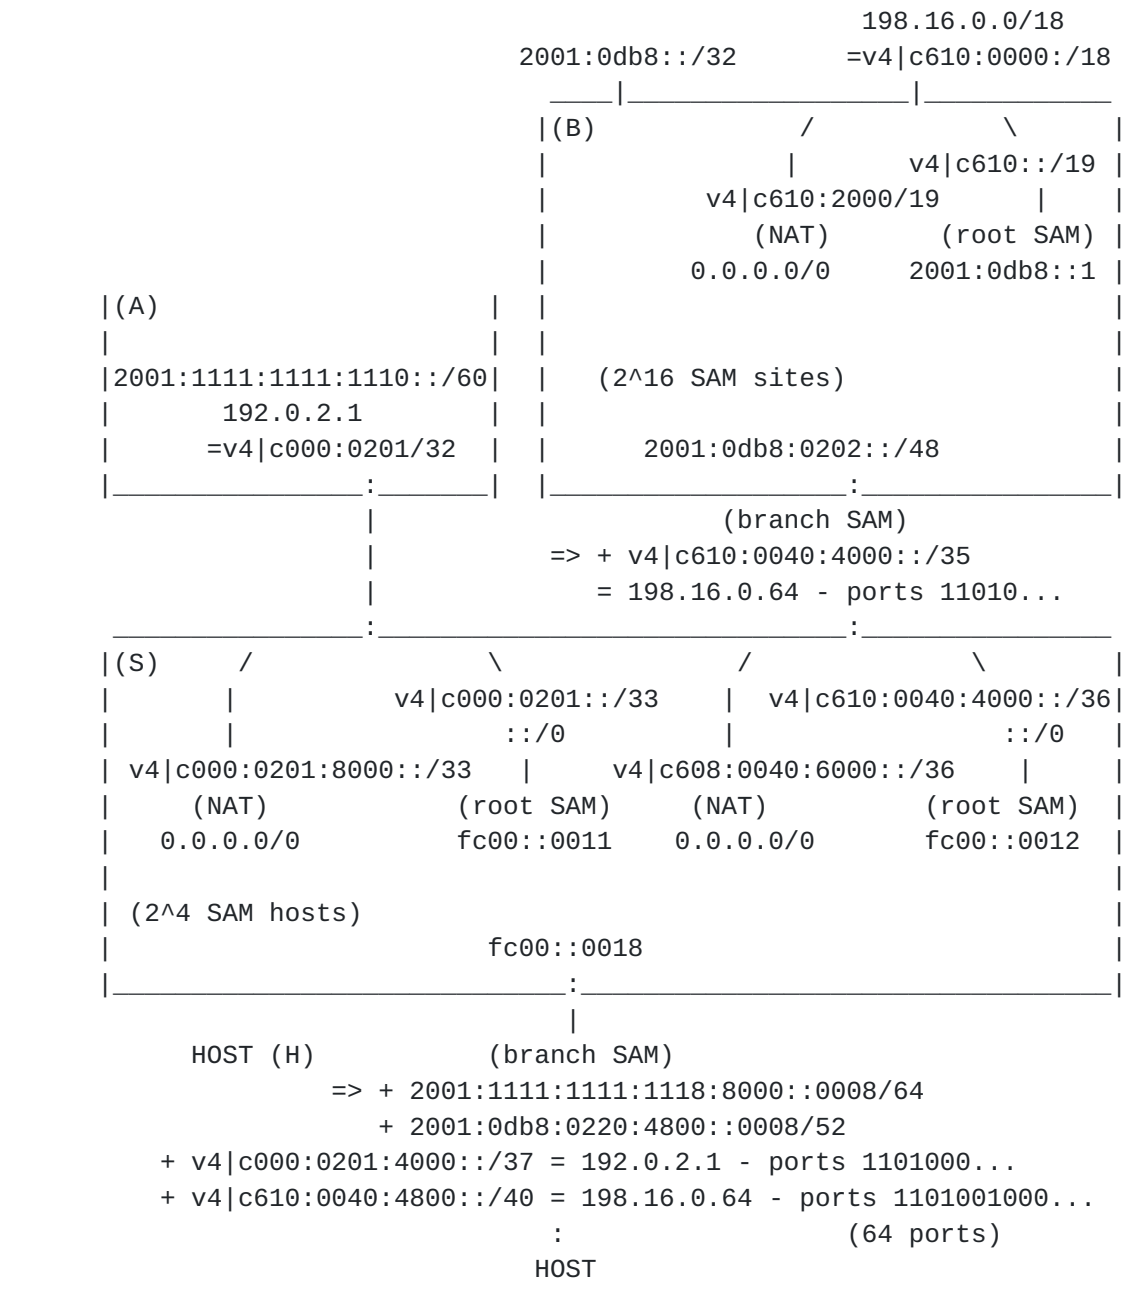

E2E IPV6 ADDRESSING AND SHARED IPV4 ADDRESSES IN A MULTIHOMED SITE

Figure 8

In site S, the branch SAM of its root interface with ISP B derives from its IPv6 prefix 2001:0db8:O2O2::/48, and from SAM parameters of ISP B, its IPv4 prefix v4|c610:2040:4000::/35, which is a port-restricted one.

The constant prefix of local addresses is fc00::/8. Two root SAMs and two NATs are configured, each one having half the available IPv4

address space.

Parameters of SAMs of site S are the following:

Constant local prefix: P1 = fc00::/8

Identifier format code: F1 = 0::/4

Subnet index length: SL1 = 0 (non applicable)

Interface index length: IL1 = 8

Root local addresses: R1 = fc00::0011; R2 = fc00::0012

Zone Global prefixes: gZ11 = 2001:1111:1111:1110::/60; gZ12 = v4| c000:0201/32; gz21 = 2001:0db8:0202::/48; gZ22 = v4| c610:0040: 4000::/35

Privacy option prefix: none in this example (::/0)

Among the 16 hosts of home site S, Host H is supposed to have local address fc00::0018. As shown on the figure, the branch SAM of host H then derives from this local address two IPv6 global prefixes, two IPv6 global host addresses starting with these prefixes, and two port-restricted IPv4 prefixes. With these prefixes, it can use, without breaking the e2e model, 512 ports for connections via ISP A, and 64 ports via ISP B.

5. Avoiding using NATs in IPv6 with SAM

With SAM as specified, all NAT44 services that have been listed in Section 2 can be offered in IPv6 without stateful processing and without breaking the e2e model:

1. In a private-addressing IPv6 site, hosts can know their global addresses to use them in e2e packets that are encapsulated in local packets to traverse the site. Renumbering is then automated simply by automating advertisement of SAM parameter changes (in DHCP and/or with router advertisements).

2. The fact that NAT44s are in general configured with by default rejection of all incoming calls can have a simple stateless equivalent in IPv6:

   * By default, reject all incoming packets that have a branch-side port in the well known or in the IANA defined registered port ranges.

   * By default, reject all TCP incoming packets that are attempts to open new incoming connections (SYN packets without ACK).

3. In a SAM-capable site, SAM-capable hosts can take advantage of site multihoming with full compatibility with ingress filtering of [RFC3704] in both the site itself and in ISP networks to which it is connected.

4. The privacy protection described in Section 3.5 maintains the e2e model. It is expected to be largely sufficient in practice. (Sophisticated hackers would probably find ways around it, and identify who does what in sites havin the privacy-protection option, but NAT44s are not perfect for privacy protection either).

5. As we have seen, SAM global addresses contain a flexible succession of branch identifiers, so that it becomes possible to set up a flexible hierarchy of private addressing zones. In particular, host-rooted subnets become possible without breaking the e2e model.

For information, no intellectual property right has been applied for by the author on any of SAM mechanisms. The intent is to facilitate IPv6 deployment with new mechanisms that enhance its potential.

6. Security considerations

Like any function where some parameters have to be configured, SAM introduces a risk of human errors.

Besides that, no security risk introduced by SAM has so far been identified. In particular, provided consistency between local addresses and global addresses are checked in root and branch SAMs, as they must be, no new address spoofing possibility is introduced with SAM.

SAM being stateless, its scalability is high. Prevention against denial of service attacks should therefore be possible even for very intense traffic (e.g. using load balancers in front of parallel devices).

7. IANA Considerations

Standardizing ways to advertise SAM parameters to branch SAMs will, in due time, imply some IANA number assignments.

8. Acknowledgements

As this specification has evolved during many months, precious encouragement and remarks were received from Mark Townsley. He has to be warmly thanked for it. Concerning what SAM can bring to port-restricted IPv4 addresses, stimulating discussions with Dan Wing, Teemu Savolainen, Gabor Bajko, Pierre Levis, Jean-Luc Grimault, and Alain Villefranque, have influenced progress of the work. Gratitude is due to them for this. Challenging remarks, and a few (deserved) criticisms from Alain Durand have also helped to better analyze how SAM will coexist with NATs. He deserves credit for it.

9. References

9.1. Normative References

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

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

9.2. Informative References

[6rd] Despres, R., "IPv6 Rapid Deployment on IPv4 infrastructures (6rd) - Work in progress (draft-despres-6rd-02)", October 2008.

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

[RFC3286] Ong, L. and J. Yoakum, "An Introduction to the Stream Control Transmission Protocol (SCTP)", RFC 3286, May 2002.

[RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, August 2003.

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

[RFC4219] Lear, E., "Things Multihoming in IPv6 (MULTI6) Developers Should Think About", RFC 4219, October 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

[draft-carpenter-renum-needs-work-01] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering still needs work - Work in progress", December 2008.

[shim6 fail detec] Arkko, J. and I. van Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming - Work in progress (draft-ietf-shim6-failure-detection-09)", July 2007.

[shim6 protocol] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6 - Work in progress (draft-ietf-shim6-failure-detection-09)", October 2007.

Author's Address

Remi Despres
3 rue du President Wilson
Levallois,
France

Email: remi.despres@free.fr