

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 14, 2010

R. Despres
RD-IPtech
July 13, 2009

Scalable Multihoming across IPv6 Local-Address Routing Zones
Global-Prefix/Local-Address Stateless Address Mapping (SAM)
draft-despres-sam-03

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 January 14, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

The continuous growth of routing tables in the core of Internet is a challenge. It would become overwhelming if each multihomed customer

site would need a provider independent prefix to take full advantage of its multihoming. IPv6 has the potential to solve this problem, but a complete specification is still missing. This draft proposes an approach for a solution.

The Stateless Address Mapping (SAM) model, introduced for this, is applicable to a hierarchy of routing zones with multihoming permitted at each level, and with each zone using local addresses for its internal routing plan. End-to-end transparency of the Internet is maintained across these local-address zones, thanks to a systematic encapsulation of global-address packets into local-address packets. Local addresses are statelessly derived from prefixes found in global addresses, and from static parameters of traversed zones. Global prefixes delegated by a zone to its child interfaces can be obtained by autoconfiguration, thanks to a bidirectional correspondence between SAM local addresses and SAM global prefixes.

Deployment can be incremental.

Table of Contents

1.	Introduction	4
2.	Problem statement	5
2.1.	Multihoming in Hierarchized Routing Zones	5
2.2.	Routing Zones in which Addressing is Local	7
2.3.	Anti-Spoofing Compatibility	7
2.4.	Autoconfiguration of Global Prefixes	9
2.5.	IPv6 hierarchical addressing beyond 64 bits	10
3.	SAM proposed specification	10
3.1.	SAM Parameters advertised to SAM Child Interfaces	10
3.2.	SAM Formats for Local Addresses and Global infixes	11
3.2.1.	SAM Subnet Prefixes	11
3.2.2.	SAM IIDs	12
3.2.3.	SAM Global Infixes	13
3.3.	Autoconfiguration Procedure for SAM Interfaces	14
4.	Application Example	15
5.	Security considerations	18
6.	IANA Considerations	18
7.	Acknowledgements	19
8.	References	19
8.1.	Normative References	19
8.2.	Informative References	19
	Author's Address	20

1. Introduction

The continuous growth of routing tables in the core of Internet is one of the important challenges to be faced. This growth would become an overwhelming problem if each multihomed customer site would need a provider independent prefix to take full advantage of its multihoming. As analyzed in particular in [[RFC3582](#)], [[RFC3582](#)], and [[RFC4219](#)] , IPv6 should help to solve this problem, but a complete solution has yet to be proposed. Such a solution is needed in not too far a future because of the increasing variety of access technologies, both terrestrial and by radio, and because of the increasing number of usages that require service continuity, both of these tendencies leading to more and more multihoming.

This draft describes an attempt at filling this gap, using for this a Stateless Address Mapping model (SAM), between local addresses and global prefixes.

The SAM model applies to hierarchies of independently administered routing zones where multihoming is possible anywhere in the hierarchy (each zone may have several parent zones). Each zone in the hierarchy has its internal routing based on local addresses.

End-to-end Internet transparency, as defined in [[RFC2775](#)] is important in particular for IPsec security and for various address referrals. Despite traversals of zones where addressing is local, SAM maintains end-to-end transparency, using for this a systematic encapsulation of global-address packets in local-address packets.

To have a unique routing plan for both local addresses and delegated global prefixes, and to permit autoconfiguration of delegated prefixes, a bidirectional correspondence is established between local addresses and global prefixes. This correspondence depends statelessly on only a few zone parameters.

With its encapsulations and address mappings, SAM can be viewed as a generalization, to IPv6 in IPv6 and to routing-zone hierarchies with multihoming, of techniques used, for IPv4 in IPv6, in the ISATAP of [[RFC5214](#)], the 6to4 of [[RFC3056](#)], and the 6rd of [[RFC 5569](#)].

In SAM's multihoming support, precaution is taken to guarantee that routes toward the global Internet can also be taken in the reverse direction. This is for compatibility with ingress filtering, the basic anti-spoofing mechanism of [[RFC3704](#)].

In an incremental deployment of SAM, SAM-capable hosts that are in SAM-capable sites take advantage of SAM-specific benefits independently of when other hosts and sites become SAM capable.

These benefits include end-to-end Internet transparency, continued Internet connectivity as long as at least one interface to the Internet of the site is operational, and possible load sharing or fast recovery if several such interfaces are operational. In this respect, SAM is a complement to STTP [[RFC4960](#)] and to SHIM6 [[RFC5533](#)]: these protocols exploit several source addresses but have no control by themselves on which gateways to the global Internet packets go through; without such a control, packets sent by Shim6 or SCTP can be routed through a gateway that is incompatible with ingress filtering, and can be systematically lost (a black hole situation).

Previous versions of this draft had a wider scope, which included some port-range extension of IPv4 addresses, some privacy protection in IPv6, and a discussion of IPv6 NATs. This version is purposely much more focused. The scope is now limited to IPv6 multihoming in a hierarchy of local-address routing zones. The idea is that this is an application case needing a rather short term solution, while other subjects, of more debatable interest, would delay acceptability of the limited-scope solution.

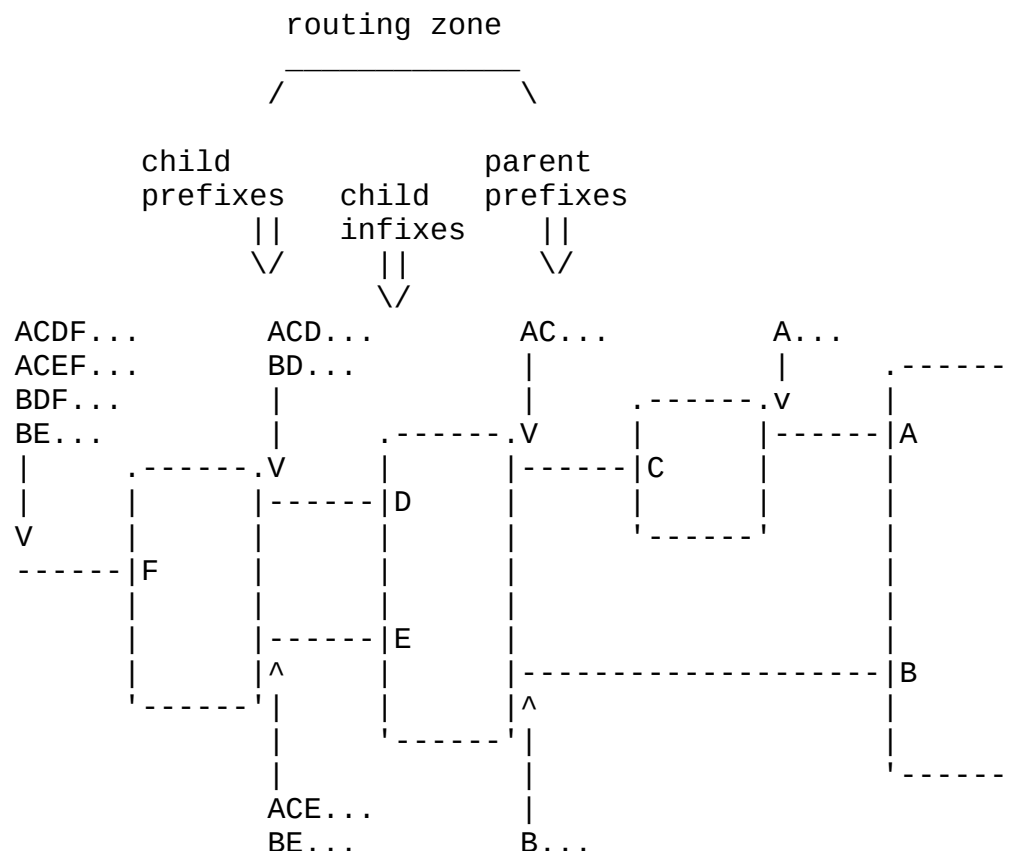
Future evolutions of the SAM proposal, and improvements of its presentation, are still expected to take place. This early-stage proposal is submitted to open it to collective work in a direction that is felt important.

[2.](#) Problem statement

[2.1.](#) Multihoming in Hierarchized Routing Zones

The principle of hierarchical addressing is that the hierarchy of a number of independently administered routing zones is directly reflected in global prefixes that are assigned to these zones. Formats of CIDR IPv4 addresses, and of IPv6 subnet indexes in the first 64 bits of IPv6 addresses, are particular applications of hierarchical addressing.

In the absence of multihoming, a routing-zone hierarchy is a tree. Each zone has only one global prefix. The global prefix delegated to a zone is that of its parent parent zone followed by the infix that, in its parent zone, identifies this particular child of his.



GLOBAL-PREFIX INHERITANCE EXAMPLE IN A HIERARCHY WITH MULTIHOMING

Figure 1

But with multihoming, zones may have several global prefixes. As illustrated in Figure 1, Prefixes delegated to a zone, at one of its parent interfaces, can be all those of its parent zone at this interface followed by the infix that, in this parent zone, identifies this particular child. This is illustrated in Figure 1 in which letters C to F are infixes, and ACDF..., for example, is a notation for a global prefix composed global prefix A followed by successive infixes C, D, and F.

The SAM model is devised to support multihoming in such routing-zone hierarchies.

2.2. Routing Zones in which Addressing is Local

Local address spaces in customer sites are largely used in IPv4 for a number of reasons, only one of which being the lack of enough global addresses for all hosts. Reasons to also use local addresses in IPv6 are documented in particular in [\[RFC4193\]](#). Particularly important is the stability of routing plans, in independently administered zones, when global prefixes that are assigned to these zones are modified, added, or deleted (renumbering simplicity).

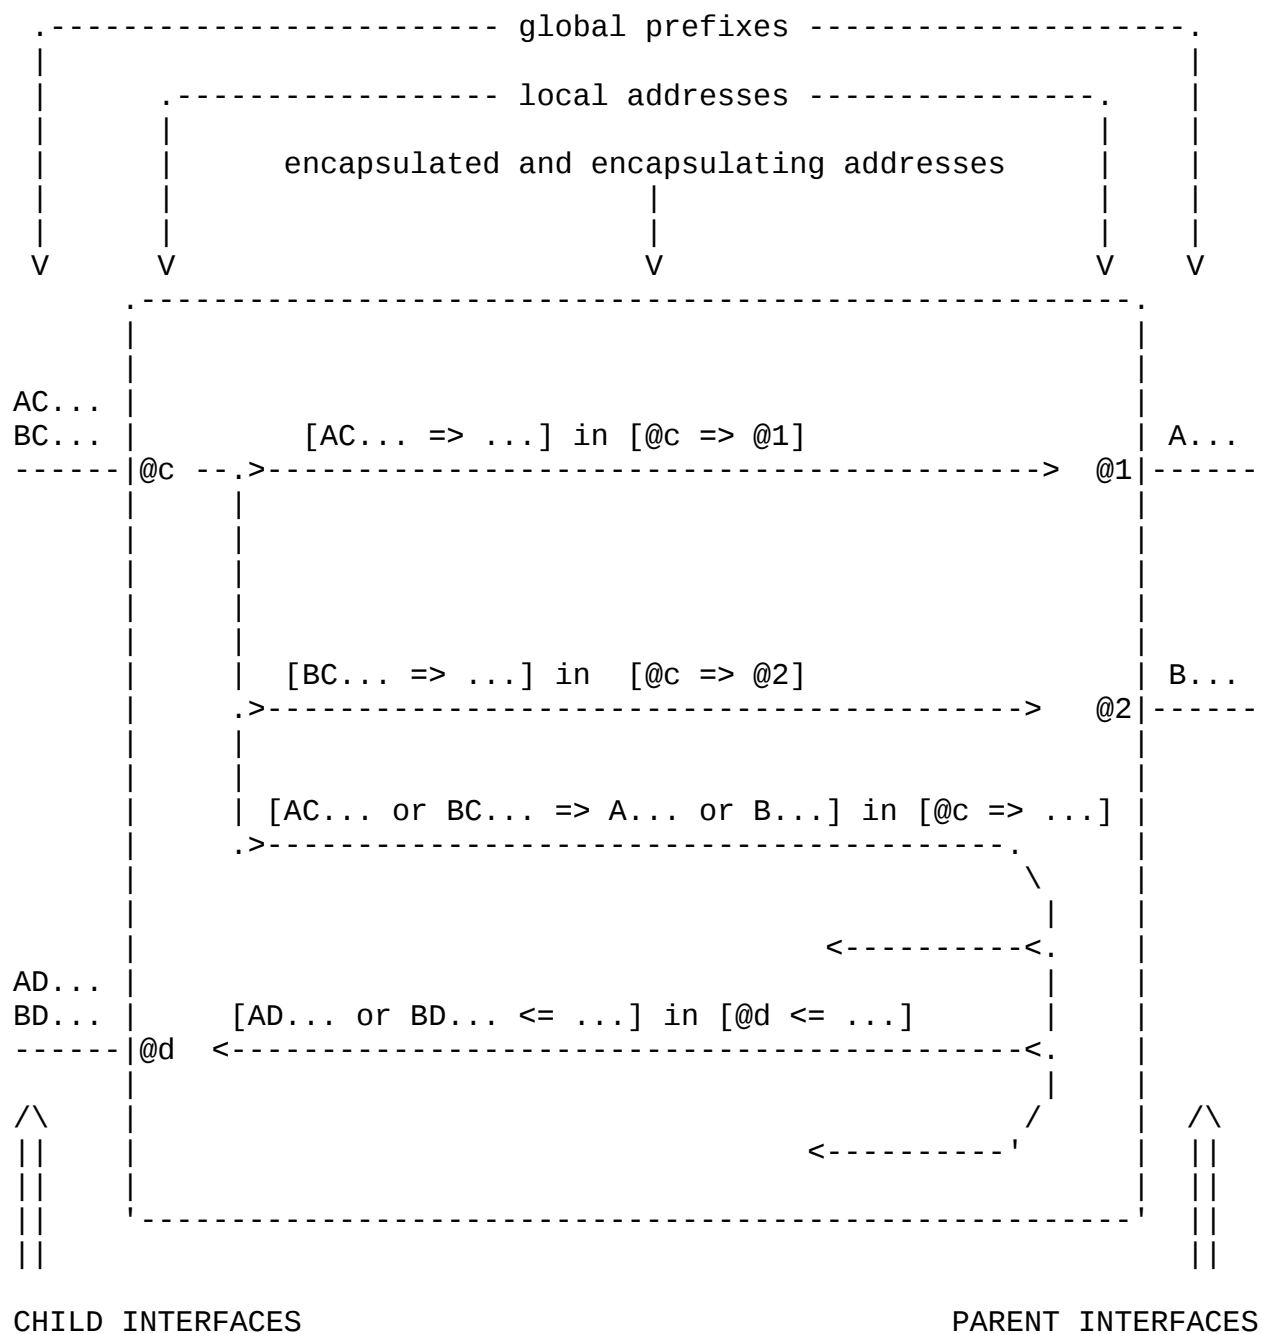
But local address spaces have a drawback: unless some precaution is taken, they conflict with Internet transparency as defined in [\[RFC2775\]](#): local addresses being prohibited in the core of Internet, a packet having a local-address source cannot traverse the global Internet, and some address translation is needed somewhere. To avoid this loss of transparency with SAM, local addresses are only used encapsulating packets in which global-address packets, to be transmitted end to end, are encapsulated.

2.3. Anti-Spoofing Compatibility

The anti-spoofing mechanism of [\[RFC3704\]](#), i.e. ingress filtering, requires that if a packet is routed across an interface in one direction, a packet with inverted source and destination has a valid route to traverse the same interface in the reverse direction.

In a multihomed routing zone, child nodes must therefore control via which parent interface they send packets toward the global Internet: if a packet transmitted by a child node has, at the beginning of its source address, a global prefix that has been assigned to the zone at its parent interface P, the packet must exit the zone via this interface P.

Figure 2 shows, for a generic multihomed site, which encapsulations are permitted to respect these constraints.



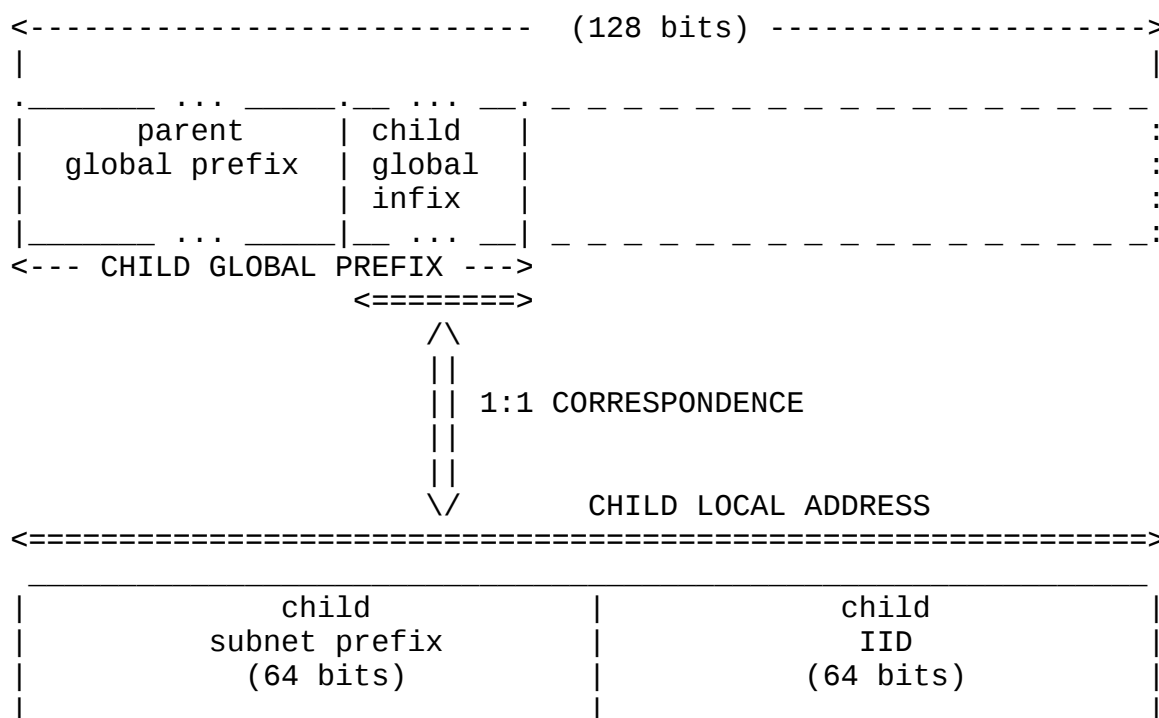
ENCAPSULATION CONSTRAINTS FOR MULTIHOMING WITH INGRESS FILTERING

Figure 2

2.4. Autoconfiguration of Global Prefixes

Stateless autoconfiguration has been specified for IPv6 to greatly simplify host address assignments [[RFC2462](#)] .

As mappings are necessary for SAM between local addresses and global prefixes, an extension of autoconfiguration to global prefixes can be envisaged. It's principle, illustrated in Figure 3, is to establish a 1:1 correspondence between the local address of a child interface and the global infix which, appended to a global prefixes of its parent zone, gives a global prefix of this interface.



CORRESPONDENCE BETWEEN LOCAL ADDRESSES AND GLOBAL PREFIXES

Figure 3

2.5. IPv6 hierarchical addressing beyond 64 bits

[RFC4291] currently imposes that ALL IPv6 addresses have a formatted IID in their 64 lower bits, although the role of formatted IIDs only appears the Stateless Autoconfiguration Procedure of [[RFC2462](#)].

Applying this constraint to hierarchical addresses of [Section 2.1](#) even if they are those of hosts in local-address zones of [Section 2.2](#) artificially prohibits to extend infix sequences beyond 64 bits. (Hosts that are in local-address zones never use their global addresses on any IPv6 link. The Stateless Autoconfiguration Procedure only applies to local addresses that are derived from these global addresses.)

Relaxing the 64-bit constraint for hierarchical addresses of [Section 2.1](#) if lower zones of the hierarchy have local addressing is possible without interfering with the role of IIDs on IPv6 links.

This remark is a contribution to 6man, the Working Group in charge maintaining the IPv6 specification. It is a request that compliance with IID formats in the 64 lower bits of IPv6 addresses be not imposed to addresses that cannot be directly used on IPv6 links.

3. SAM proposed specification

3.1. SAM Parameters advertised to SAM Child Interfaces

For a child of a zone to know the parent interface of the zone via which it has to send a packet toward the global Internet, it needs to know local addresses of all parent interfaces of the zone with, for each of them, the list of global prefixes that are delegated to the zone at this interface.

In the example of Figure 2, child interfaces that are respectively at local addresses @c and @d need to know local addresses @1 and @2 of the two parent interfaces, and the lists of zone prefixes inherited at these interfaces, i.e. {A...} and {B...} respectively. (Each of the two lists has only one element in this example). No other SAM parameter of the zone is needed.

Various ways to communicate these parameters to child interfaces can be envisaged. Router advertisements of [[RFC2462](#)] (RAs) have the advantage to be receivable by child interfaces before they have finalized their local address autoconfiguration. They can consequently recognize whether they are in a SAM-capable site early enough to adapt their autoconfiguration procedures.

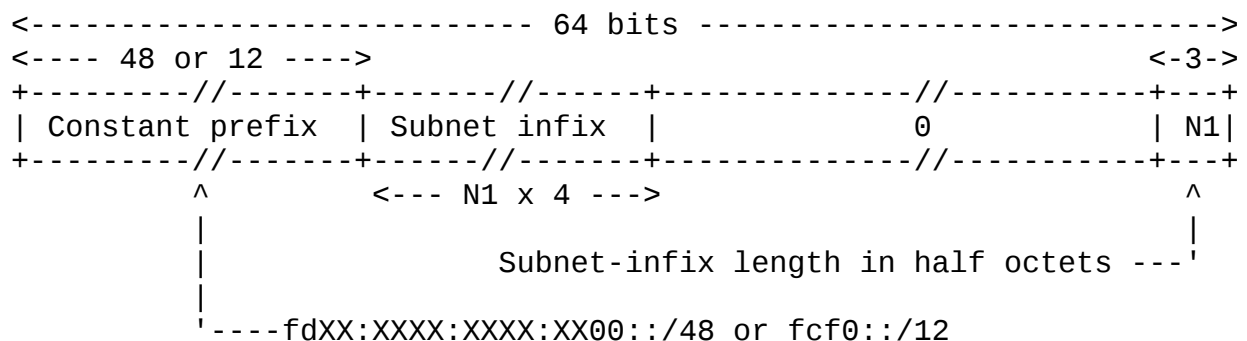
In SAM zones that have multiple links, the procedure for routers that have to send these RAs to obtain SAM parameters is left for a further version of this draft.

[3.2.](#) SAM Formats for Local Addresses and Global infixes

[3.2.1.](#) SAM Subnet Prefixes

SAM subnet prefixes must be suitable for the 1:1 correspondence of [Section 2.4](#) between SAML local addresses and SAM global infixes.

Figure 4 presents a proposed format for this.



SAM SUBNET PREFIX FORMAT

Figure 4

For IPv6 local addresses to be distinguishable from global-scope addresses, [\[RFC4193\]](#) reserves the 7-bit prefix `fc00::/7`. It also specifies 48-bit prefixes for ULAs (unique local IPv6 unicast addresses), which start with `fd00::/8`.

For convenience, we propose that the constant-prefix part of SAM-zone local address may not only be a 48-bit ULA, differing from a zone to another, but may also be shorter and the same for different SAM zones. As shown in Figure 4, we propose for this to reserve for prefix `fcf0::/12`. The `f` it contains after to `fc00::/8` is to permit other uses of the `fc00::/8` prefix in the future (forward compatibility precaution).

Since routing is in general based on longest-prefix matching, subnet infix bits, which identify a particular subnet in a SAM zone, are preferably be placed immediately after the constant prefix.

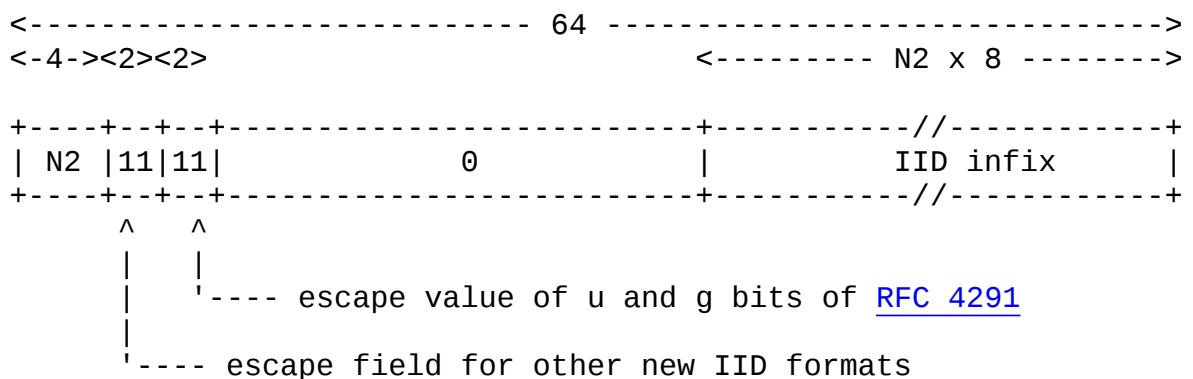
For the length of subnet infixes to be recognizable in SAM subnet prefixes themselves, we propose to place a length indicator in the last half octet of the subnet prefix, and to express it in number of half octets [Figure 4]. Restricting subnet-infix lengths to multiple of 4 bits should not be inconvenient in practice since as it facilitates interpretation of IPv6 addresses in their standard notation (thus facilitating network maintenance). Using only 3 bits permits commonality with the length indicator of SAM global infixes of [Section 3.2.3](#), in which field lengths should preferably be short.

As an example, fcf9:0:0:1 is a SAM subnet prefix that comprises, from left to right, the short constant prefix 0xFCF, the subnet infix 0x9 (1 half-octet long in this case), an unused field of 5.5 octets, and the length indicator 0x1 of the subnet infix (for the half octet 0x9).

[3.2.2](#). SAM IIDs

SAM IIDs have also to be suitable for the 1:1 correspondence of [Section 2.4](#) between SAM local addresses and SAM global infixes.

Figure 5 presents a proposed format for this.



SAM IID FORMAT

Figure 5

For SAM IIDs to be distinguishable from already specified IIDs of [[RFC4291](#)], whose u and g bits are 00, 01 or 10, we propose SAM IIDs to have 11 in u and g bits and, to permit other uses of this 11 pattern in the future, 11 in the two preceding bits (forward compatibility precaution).

Since the Duplicate Address Detection procedure of[RFC2462] uses the 24 last bits of IIDs to discriminate multicast groups that it uses, bits of IID infix bits, which identify a particular SAM interface on its link, are preferably placed at the end of IIDs.

For the length of subnet infixes to be recognizable in SAM IIDs, we propose to place a length indicator in the first half octet of the IID, and to express it in number of octets. Using only 4 bits permits commonality with the length indicator of global infixes of [Section 3.2.3](#), in which all field lengths should preferably be short.

As an example, 2f00:0:0:6666 is a SAM IID that comprises, from left to right, the length indicator 0x2 of the IID infix (indicating 2 octets), the SAM-IID tag 0xF, an unused field of 5 octets, and the IID infix 0x6666 (2 octets long as specified in the length indicator).

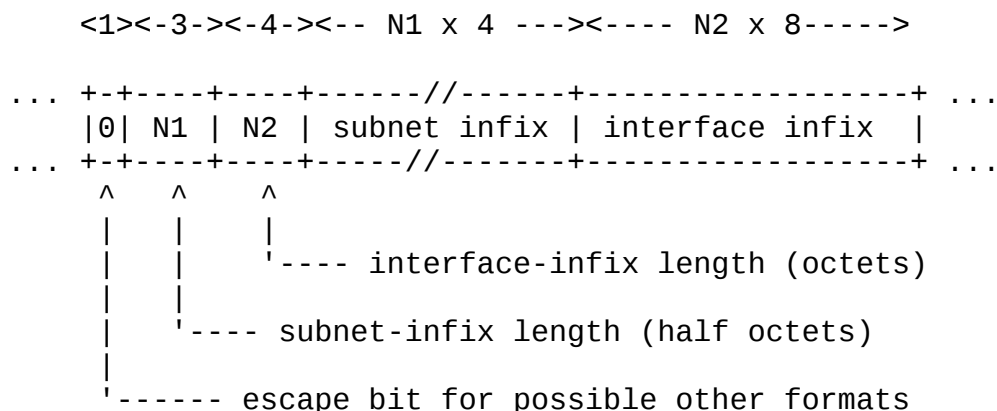
NOTE: [RFC4291], where current IID formats are specified and leave unused the 11 value of u and g bits, has a sentence that may be interpreted as limiting uses of this 11 value to IIDs that have universal scope: "The use of the universal/local bit in the Modified EUI-64 format identifier is to allow development of future technology that can take advantage of interface identifiers with universal scope." This interpretation would prohibit using the proposed SAM IID format because SAM IIDs have local scopes. This note is therefore a contribution to the 6man Working Group. It is a request to clarify that the 11 combination of u and g bits may be used for all new IIDs, including those that have a local scope.

[3.2.3](#). SAM Global Infixes

SAM global-infixes must also be suitable for their 1:1 correspondence of [Section 2.4](#) with SAM local addresses.

To derive a child local address from a child global infix, the subnet infix part and the IID infix part of the global infix in must be extractible. Their lengths must therefore be recognizable when parsing bits that follow parent-zone global prefixes.

For this, the format proposed in Figure 6 has, before subnet-infix and IID-infix fields themselves, their length indicators. They are chosen to be in the short format described for local addresses. Before these fields, a 1-bit field set to 0 is included so that different formats can be defined in the future (forward compatibility precaution).



PROPOSED FORMAT FOR SAM GLOBAL INFIXES

Figure 6

If $N_1 = N_2 = 0$, the infix contains neither subnet infix nor an IID infix. While it therefore cannot designate any child interface at a lower level, it naturally provides a global address for the child interface itself, at its level. Bits that follow, which are unused, are set to 0.

3.3. Autoconfiguration Procedure for SAM Interfaces

As we have seen in [Section 2.4](#), the correspondence between SAM local addresses and global prefixes must be such that the autoconfiguration of local addresses can also be that of global prefixes.

For this, a first point to be noted is that, since SAM IIDs are distinct from IIDs having a classic format, SAM-capable hosts can coexist on a link with classic IIDs. As necessary for incremental deployment, there is no risk that hosts that don't support SAM would miss IID duplicates caused by SAM-capable neighbors.

The second point is that SAM-capable nodes can have a SAM-specific Duplicate Address Detection procedure. The SAM duplicate address procedure is the same as usual except that two SAM local addresses are considered duplicate if one of the two IID infixes is a prefix of the other.

Thus, if a SAM-capable node recognizes that it is in a SAM zone (finding SAM parameters in received RAs), it can proceed as follows to get an n -bit infix (from which it derives its global prefixes): pick at random a SAM prefix having an n -bit long IID infix; test it to see whether it has a duplicate on the link; if there is one, try again with a different randomly selected IID; if there is none,

derive the delegated global prefix from the obtained local address (and from SAM parameters of the zone).

4. Application Example

To detail how SAM applies to a simple practical example, we now consider the customer site of Figure 7.

The site has gateways to the Internet from service providers ISP-A and ISP-B, with global prefixes assigned to the site being A... and B... respectively. The ISP-A gateway is physically connected to a WiFi LAN and to the site Ethernet LAN, with a bridge between them. The ISP-B gateway is physically connected to the Ethernet LAN. IIDs of these gateways are m and n respectively.

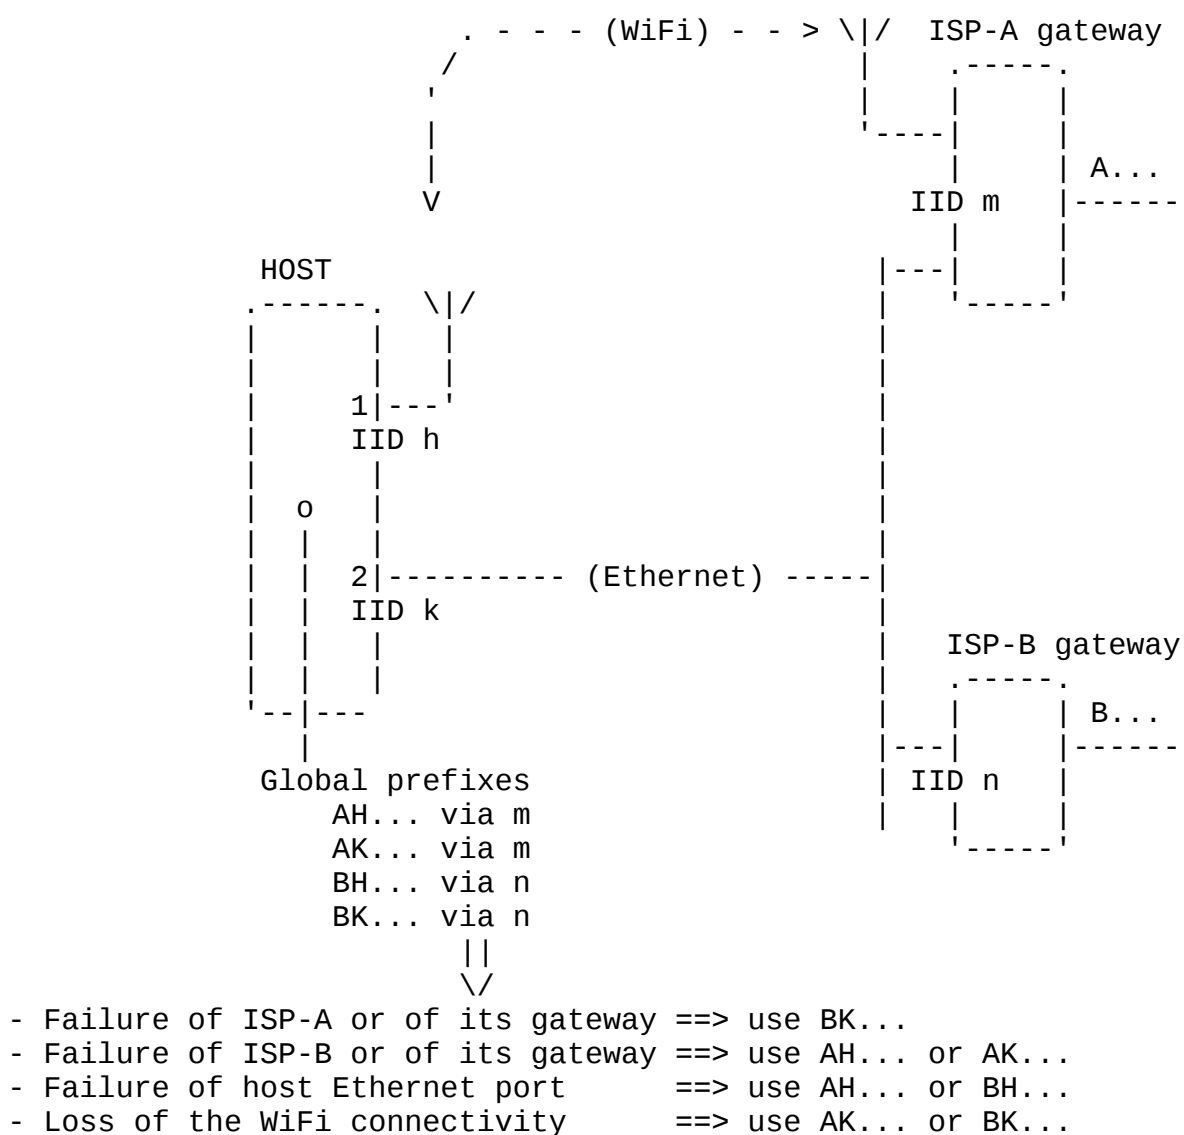
To keep the example simple, it is assumed that there is only one subnet in the site (but it could be easily extended to include several administratively configured subnets).

The host we consider is physically connected via its interface 1 to the WiFi LAN of the ISP-A, on which it has h as IID, and via its interface 2 to the Ethernet LAN on which it has k as IID. It may have to act as a router, with one or several subnets behind it, so that it may need to know its global prefixes.

ISP gateways and the host are assumed to be SAM capable.

The host can derive its 4 permitted global prefixes from the local addresses it has at its two interfaces, and from SAM parameters of the zone. They are indicated in Figure 7 with a notation where, for example, prefix AH..., stands for prefix A... followed by the SAM infix derived from the local address whose IID is h.

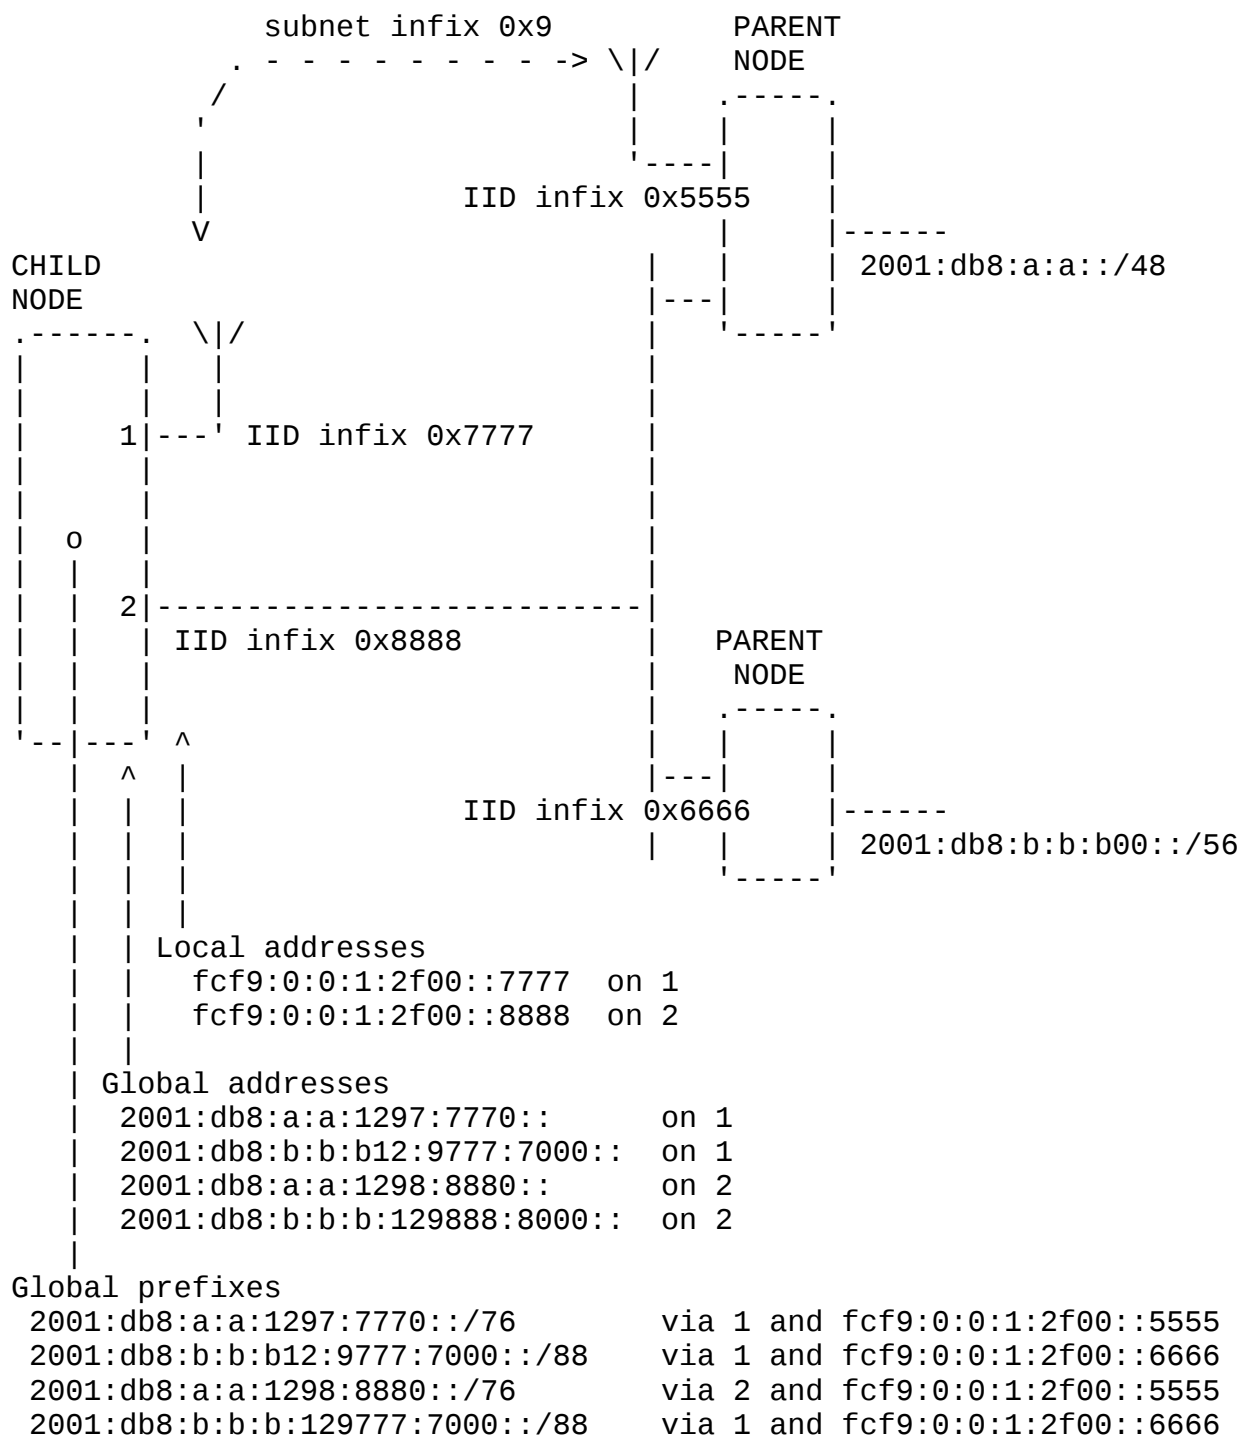
Depending on which of these prefixes appears in the source address of a global-address packet, the local destination to reach the global Internet has to be gateway A or gateway B. Host parameters that govern this choice are shown on the Figure.



MULTIHOMED-SITE EXAMPLE WITH HOST RESISTANCE TO SINGLE-POINT FAILURES

Figure 7

In Figure 8, the same site configuration is presented, but this time with detailed IPv6 addresses and prefixes in their standard format of [\[RFC4291\]](#).



EXAMPLE OF PREFIXES AND ADDRESSES FOR THE NETWORK EXAMPLE

Figure 8

In this example, the constant prefix of local addresses is chosen to be to 0xFCF. The subnet infix of the unique subnet is set to 0x9 (having only one subnet, the subnet infix could have been given a null length, but this would have imposed a reconfiguration if other subnets are added later on). Prefixes assigned by ISPs A and B are supposed to be 2001:db8:a:a::/48 and 2001:db8:b:b:b000::/56 respectively. IID infixes of m, n, h, and k interfaces, possibly obtained by autoconfiguration, are supposed to be 0x5555, 0x6666, 0x7777, and 0x8888 respectively.

SAM parameters of the zone that result are presented in Table 1. Global prefixes of the host, and via which interfaces and which gateways they can be used, are detailed in the Figure. Global addresses of the host itself, obtained with trailing 0s appended to global prefixes are also shown.

Parent local address	Zone global prefix
fcf9:0:0:1:2f00::5555	2001:db8:a:a::/48
fcf9:0:0:1:2f00::6666	2001:db8:b:b:b000::/56

SAM PARAMETERS TO BE ADVERTISED TO CHILD INTERFACES

Table 1

5. Security considerations

No security issue that appears to be specific of SAM has been identified so far.

In particular, provided consistency between local addresses and global prefixes are systematically checked at parent and child interfaces, no new address spoofing possibility seems to be introduced.

Also, SAM being stateless, its scalability is high. Resistance to denial of service attacks should therefore be possible even for very intense traffic, using if needed load balancers in front of parallel hardware devices.

6. IANA Considerations

As indicated in [Section 3.1](#), mechanisms to advertise SAM parameters of a SAM zone to its child interfaces will need some number

assignments by IANA. This is however beyond the scope of this version of the draft.

7. Acknowledgements

Although the substance of this draft, with its now restricted scope, is essentially the result of a personal work, the author expresses his gratitude to Mark Townsley. He took time to listen to intermediate stage presentations, and provided useful reactions. Thanks are also due to Dave Thaler who made a precious detailed review of the previous version. Beyond this, the open discussion environment of IETF in general has been a continuous encouragement.

8. References

8.1. Normative References

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

8.2. Informative References

- [RFC 5569] Despres, R., "IPv6 Rapid Deployment on IPv4 infrastructures (6rd) - Work in progress ([draft-despres-6rd-02](#)) *to soon be replaced by [RFC 5569*](#)", October 2008.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [RFC2775] Carpenter, B., "Internet Transparency", [RFC 2775](#), February 2000.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [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.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [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.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", [RFC 5533](#), June 2009.
- [[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
RD-IPtech
3 rue du President Wilson
Levallois,
France

Email: remi.despres@free.fr