

behave	X. Li, Ed.
Internet-Draft	C. Bao, Ed.
Intended status: Standards Track	CERNET Center/Tsinghua University
Expires: October 24, 2009	F. Baker, Ed.
	Cisco Systems
	April 22, 2009

IPv4/IPv6 Translation Prefix Recommendation
draft-xli-behave-v4v6-prefix-00

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 October 24, 2009.

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

This document is part of a series of the IPv4/IPv6 translation

documents. In this document, the address format and the corresponding prefix are recommended for representing IPv4 addresses in IPv6 and/or for representing IPv6 addresses in IPv4.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Embedded Address Format	4
4.	Uses of the Embedded Address Format	5
4.1.	Representing the IPv4 addresses in IPv6	5
4.2.	Representing the relationship between IPv4 and IPv6 addresses	6
4.2.1.	Stateless Translation	6
4.2.2.	Stateful Translation	7
5.	Parameter Analysis of the Embedded Address Format	8
5.1.	PREIX Analysis	8
5.1.1.	IPv6 Routing System Scalability	8
5.1.2.	Other Issues	10
5.2.	Prefix Length Analysis	12
5.2.1.	Routing Policy	12
5.2.2.	IPv6 Address Consumption	12
5.2.3.	Forwarding Efficiency	13
5.2.4.	SUFFIX	13
5.2.5.	EUI-64 format	13
5.3.	SUFFIX Analysis	13
6.	Conclusions	14
6.1.	PREFIX Recommendation	14
6.2.	Prefix Length Recommendation	14
6.3.	SUFFIX Recommendation	15
7.	IANA Considerations	15
8.	Security Considerations	15
9.	Acknowledgements	15
10.	References	15
10.1.	Normative References	15
10.2.	Informative References	16
	Authors' Addresses	18

1. Introduction

This document is part of a series of the IPv4/IPv6 translation documents.

In order to perform the translation function between the IPv4 and IPv6, the translator needs to represent the IPv4 addresses in the IPv6 Internet and the IPv6 addresses in the IPv4 Internet.

In this document, the embedded address format and the corresponding prefix are recommended.

The address format and the corresponding prefix selection are related to the translation scenarios and the operation mode, discussed in the framework document [[I-D.baker-behave-v4v6-framework](#)]. For the stateless translator, this address format is used to represent the IPv4 addresses in IPv6 and to represent the IPv6 addresses in IPv4. For the stateful translator, this address format is used to represent the IPv4 addresses in IPv6 only and the session initiated states are used to represent the IPv6 addresses in IPv4.

The technical specification of the translation is in the translation document [[I-D.baker-behave-v4v6-translation](#)], which uses the recommendations in this document.

The DNS ALG document [[I-D.bagnulo-behave-dns64](#)] uses the recommendation in this document.

2. Terminology

The following terminology is used in this document and other documents related to it.

Embedded Address: The IPv6 address with IPv4 address embedded.

PREFIX: The most significant bits before the embedded IPv4 address.

SUFFIX: The least significant bits after the embedded IPv4 address.

IPG4: The global IPv4 addresses.

ISP4: The ISP's IPv4 prefix.

IPv4 pool: The ISP4 configured in the stateful translator.

ISP6: The ISP's IPv6 prefix.

IPG4prefix: The IPv6 address representation of IPG4.

ISP4prefix: The IPv6 address representation of ISP4.

ISP6prefix: Same as ISP6. Note that IPG4prefix is a subset of ISP6prefix and ISP4prefix is a subset of IPG4prefix.

Stateless Translation: A translation algorithm that is not "stateful" is "stateless". It may require configuration of a static translation table, or may derive its needed information algorithmically from the messages it is translating.

Stateful Translation: A translation algorithm may be said to "require state in a network element" or be "stateful" if the transmission or reception of a packet creates or modifies a data structure in the relevant network element.

LIR (LIR Prefix): The IPv6 prefix assigned by the network operator for embedding IPv4 addresses into IPv6 addresses. In this case, each network running a translator will create a representation of the whole IPv4 address space in the IPv6 address space.

WKP (Well-Known Prefix): The IPv6 prefix assigned by IANA for embedding IPv4 addresses into IPv6 addresses. In this case, there would be a single representation of a public IPv4 address in the IPv6 address space.

3. Embedded Address Format

Embedding IPv4 address in IPv6 address (defined as IPv4-embedded IPv6 address) will be formed by concatenating a prefix to the IPv4 address and optionally a suffix. The prefix is called the PREFIX and the suffix is called SUFFIX. The resulting IPv6 representation is depicted in the figure below.




Figure 1: Embedded Address Format

As shown in Figure 1, the embedded address format has three components:

bits 0..n-1 (PREFIX): An LIR-specified prefix, either 32..63 bits long or 96 bits long,

bits n..n+31 An embedded IPv4 address. Except in the case of a 96 bit prefix, this address intentionally straddles the boundary between [\[RFC4291\]](#)'s 64 bit "subnet" locator and its 64 bit host identifier. The intention is that the /64 be used in routing and the bits in the host part be used for host identification as described in the address architecture.

bits n+32..127 (SUFFIX): Entirely zero; note that if n=96, this is null.

The selection of the PREFIX, the prefix length and SUFFIX is discussed in the following sections.

4. Uses of the Embedded Address Format

The embedded address format is used both for the stateless translator and the stateful translator. For the stateless translator, this address format is used to represent the IPv4 addresses in IPv6 and to represent the IPv6 addresses in IPv4. For the stateful translator, this address format is used to represent the IPv4 addresses in IPv6 only and the session initiated states are used to represent the IPv6 addresses in IPv4.

4.1. Representing the IPv4 addresses in IPv6

To represent the IPv4 addresses in IPv6, a PREFIX is selected and the global IPv4 addresses is embedded in this PREFIX, as shown in the following figure. This special IPv6 prefix (as IPG4prefix) is the image of the global IPv4 addresses and the IPv6 hosts will communicate with the global IPv4 addresses via this IPv6 prefix. The SUFFIX are all zeros.

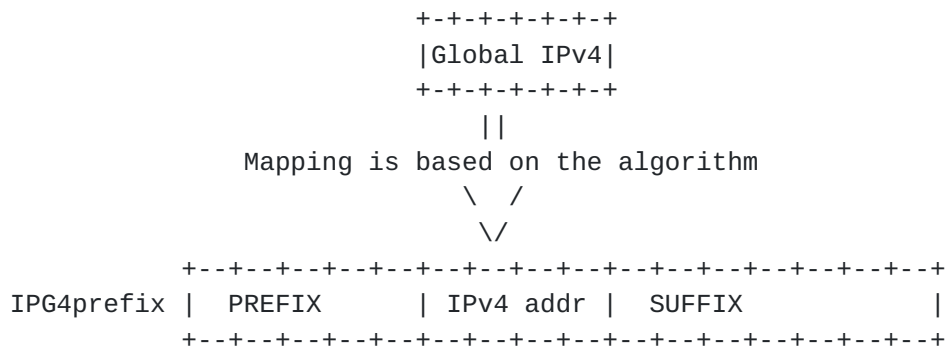


Figure 2: Represent the IPv4 addresses in IPv6

4.2. Representing the relationship between IPv4 and IPv6 addresses

In the IPv6 domain, the addresses of IPv4 systems are represented using IPv4 embedded addresses, which can be translated without the maintenance of dynamic state. However, addresses in the IPv6 domain are different depending on the function of the system using them; IPv4-accessible servers in the IPv6 domain use the IPv4 embodied address format (4.2.1), while client systems that are inaccessible from the IPv4 domain can use stateful translation (4.2.2) to access IPv4 services.

4.2.1. Stateless Translation

To represent the IPv6 addresses in IPv4, each provider can borrow a portion of its IPv4 addresses (ISP4) and maps them into IPv6 (as ISP4prefix) based on the above embedded address format, as shown in the following figure. These special IPv6 addresses will be physically used by IPv6 hosts. The original IPv4 form of the borrowed addresses is the image of these special IPv6 addresses and the global IPv4 addresses will communicate with ISP4 via this more specific prefix. Note that ISP4prefix is "more specifics" of IPG4prefix in the IPv6 Internet. The SUFFIX can either be all zeros or for the future extensions.

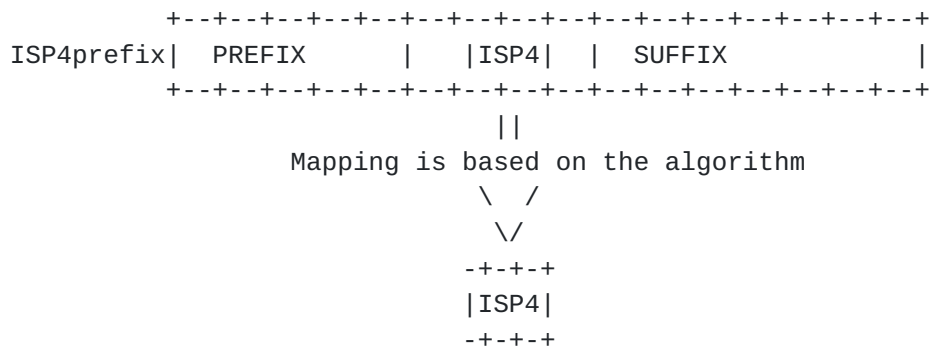


Figure 3: Represent the IPv6 addresses in IPv4 (stateless)

Note that in the stateless case,

- o This mode is suitable for the "an IPv6 Network connecting to the IPv4 Internet" scenarios, since only a subset of the IPv6 addresses can be represented by IPv4.
- o This subset of the IPv6 addresses supports both IPv6 initiated and IPv4 initiated communications. Therefore, it can be used for the IPv6 only servers.
- o When the IPv4 address sharing technique is used, this subset of the IPv6 addresses will be big enough to meet the IPv4 address requirements in the IPv4 to IPv6 transition stages. The details of the technical specification will be discussed in other documents.

4.2.2. Stateful Translation

The session initiated states are used to represent the IPv6 addresses (as ISP6prefix) in IPv4 for the stateful translator, as shown in the following figure.

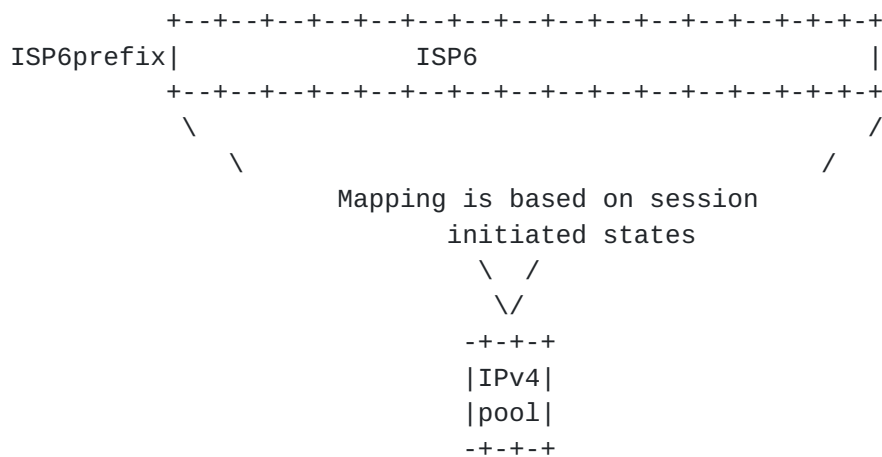


Figure 4: Represent the IPv6 addresses in IPv4 (stateful)

Note that in the stateful case,

- o This mode is suitable for both the "an IPv6 Network connecting to the IPv4 Internet" and "an IPv4 Network connecting to the IPv6 Internet" scenarios. But due to the stateful nature, it may have scalability problem and is suitable for small sized network.
- o It only supports IPv6 initiated communication. Therefore, it can be used for the IPv6 only clients.

5. Parameter Analysis of the Embedded Address Format

5.1. PREFIX Analysis

The PREFIX Recommendation Section discusses the selection of the PREFIX in the address format. The possible candidates are LIR (Local Internet Registry) and WKP (Well-Known Prefix). The major evaluation criterion is the IPv6 routing system scalability.

5.1.1. IPv6 Routing System Scalability

5.1.1.1. An IPv4 Network connecting to the IPv6 Internet

For "an IPv4 Network connecting to the IPv6 Internet" scenario, only stateful translation can be used, as shown in the following figure.




Figure 5: An IPv4 Network connecting to the IPv6 Internet

With private IPv4 addresses, the WKP doesn't work, since there is no distinction among IPv4 hosts using the private IPv4 blocks.

With public IPv4 addresses, each IPv4 site will inject a route in the IPv6 routing table in border A, i.e. importing IPv4 routing table entropy into IPv6 routing table.

Therefore, the LIR should be selected in "An IPv4 Network connecting to the IPv6 Internet" scenario.

[5.1.1.2.](#) An IPv6 Network connecting to the IPv4 Internet

For "an IPv6 Network connecting to the IPv4 Internet" scenario, as shown in the following Figure, the stateless and stateful mode should be discussed separately.




Figure 6: An IPv6 Network connecting to the IPv4 Internet

With private IPv4 addresses, the WKP doesn't work, since there is no distinction among IPv4 hosts using the private IPv4 blocks.

[5.1.1.2.1.](#) Stateless Mode

If Xlate is stateless, "an IPv6 network" will consist of hosts using IPv6 addresses of ISP4prefix (more specifics).

- o In border A, "an IPv6 network" will advertise ISP4prefix to the Xlate and the Xlate will advertise IPG4prefix (corresponding to 0.0.0.0) to "an IPv6 network".

- o In border B, "an IPv6 network" will advertise ISP4prefix addresses with proper aggregation to the IPv6 Internet and the IPv6 Internet will advertise the global IPv6 routing table or 2000::/3 to "an IPv6 network". In other words, "an IPv6 network" should advertise the aggregated prefix of "an IPv6 network" and should neither advertise the IPG4prefix corresponding to 0.0.0.0/0 nor ISP4prefix. This can be achieved easily if LIR is used, since ISP4prefix is "a more specific" in IPG4prefix and IPG4prefix is "a more specific" in ISP6prefix. However, it cannot be achieved if WKP is used, since WKP is independent of the ISP6prefix of "an IPv6 network". In the case when WKP is used, the ISP6prefix must be advertised to the IPv6 Internet in order to communicate with other IPv6 hosts in the IPv6 Internet and this advertisement will eventually result in the injecting of the global IPv4 routing table into the global IPv6 routing system.

5.1.1.2.2. Stateful Mode

If Xlate is stateful, "an IPv6 network" will consist of ordinary IPv6 addresses and Xlate maintains session-initiated states to map between ordinary IPv6 addresses and the IPv4 addresses in the IPv4 address pool.

- o In border A, "an IPv6 network" will advertise the aggregated prefix (ISP6prefix) of "an IPv6 network" to the Xlate and the Xlate will advertise the IPv4 embedded IPv6 addresses corresponding to 0.0.0.0/0 (IPG4prefix) to "an IPv6 network".
- o In border B, "an IPv6 network" will advertise the aggregated prefix of "an IPv6 network" (ISP6prefix) to the IPv6 Internet and the IPv6 Internet will advertise the global IPv6 routing table or 2000::/3 to "an IPv6 network". There is no need to advertise the IPv4 embedded IPv6 addresses corresponding to 0.0.0.0/0 (IPG4prefix) to the IPv6 Internet, if translation service of "an IPv6 Internet" is not supported to other IPv6 networks. Therefore, the selection of WKP versus LIR is mainly the operational consideration.

5.1.2. Other Issues

When the public IPv4 addresses are used in the "an IPv6 network connecting to the IPv4 Internet" scenario and it is in stateful operation mode, there are several issues concerning the selection of LIR or WKP. These issues are: the referral support, the native connectivity preference in communications involving dual stack nodes, the DNS ALG configuration and the support for multiple translators.

5.1.2.1. Referral support

A referral operation is when a host A passes the IP address of a Host B to a third Host C as application data. The Host C will then initiate a communication towards the Host B using the IP address received.

At this point in time, there are two widely-available protocols that operate on the IPv4 Internet which perform referrals: SIP and BitTorrent. The analysis in [[I-D.wing-behave-nat64-referrals](#)] of SIP (which does referrals between IPv4 and IPv6) shows that SIP needs to refer IPv4 addresses -- not IPv6 addresses. Thus, it doesn't matter if WKP or LIR is used, because an IPv6 address isn't referred anyway: the IPv4 address is referred.

5.1.2.2. Native connectivity preference in communications involving dual stack nodes

When dual stack nodes are involved in the communication, the potential issue is that they prefer translated connectivity over the native connectivity.

Communication initiated from an IPv6-only node towards a dual stack node: In this case, the IPv6 only node will query for the FQDN of the dual stack node. The DNS ALG function will try first to get the AAAA RR. Since there is one available, it will return it and no AAAA RR will be synthesized from the A RR of the dual stack node. The selection of the WKP or LIR will be discussed in the following section.

Communication initiated from a dual stack node toward an IPv4 only node. In this case, the dual stack node MUST use normal DNS (not the DNS ALG) and the native connectivity is ensured. Thus, it doesn't matter if WKP or LIR is used.

5.1.2.3. DNS ALG configuration

The DNS ALG function can be placed either in the DNS server or in the end host. In order to synthesize AAAA RR, the DNS ALG function needs to know the PREFIX. In the case that a WKP is used, the PREFIX information can be hardcoded in the DNS ALG code, requiring no additional tools for learning it. In the case that a LIR is used, the DNS ALG needs to discover the PREFIX information. In the case that the DNS ALG is located in the servers, it may be a viable option to manually configure the PREFIX in the DNS ALG for a few servers. However, in the case the DNS ALG is located in the hosts, the manual option seems inconvenient and alternative automatic means need to be provisioned. Moreover, since this information is used for DNSSEC

operations, the mechanism to configure the PREFIX need to be secure. The result is that the LIR option requires more tools than the WKP.

5.1.2.4. Support for multiple translators

This issue is somehow orthogonal on whether the prefix is WKP or LIR. In both cases, it is possible to use a single prefix for multiple translators or different prefixes for different translators. In any case, this would be achieved by inserting (or not) some subnet bits between the prefix and the embedded IPv4 address that would be used to identify the translator box. This issue does have implications on some of the different issues considered before. In particular, if a per translator prefix is used, then there is the need to configure the prefix in the DNS ALG, so the non configuration feature of the WKP is no longer achieved.

5.2. Prefix Length Analysis

There are three issues related to the prefix length selection, routing policy, IPv6 address consumption and the forwarding efficiency, etc.

5.2.1. Routing Policy

The major issue for selecting the prefix length is the routing policy. If the IPv4/IPv6 translation is implemented in a subnet, then a /96 should be fine. However, if the IPv4/IPv6 translation is implemented in an ISP's backbone, then the minimum prefix should be /64 and in some cases should be /48.

5.2.2. IPv6 Address Consumption

One issue that is worth considering is the one related to IPv6 address consumption. In particular, depending on the selected prefix length, IPv6 address consumption can become an issue.

For LIR case, the prefix must come out of the IPv6 allocation for the site running the translator. If the site running the translator is an ISP, then probably the allocation of the ISP is a /32 or shorter, so, it may be possible for the ISP to allocate a somehow short prefix for this, maybe a /40. However, if the translator is run by an end site, which normal allocation is a /48, then the LIR prefix for the translator should be much longer than that, possibly a /56. So, in the case the site needs to route based on the IPv4 prefix embedded in the IPv6 address (e.g. in order to access to different parts of the IPv4 space through different routes), then it is likely that it will need to route on the lower 64 bits of the IPv6 address.

For the WKP case, the prefix would be allocated by IANA for this particular purpose. As such, it seems reasonable that a short prefix can be obtained for this. Requesting for a /24 or even a few bits shorter seems feasible. The potential benefit of this is that IPv4 prefixes can be represented as IPv6 prefixes that are shorter than 64 bits. This would result in routing based on the upper 64 bits, which is compatible with current IPv6 practices. For instance, if we use a /24 for the WKP, an IPv4 /24 would result in an IPv6 /48, which seems somehow equivalent from the routing perspective.

5.2.3. Forwarding Efficiency

According to current specifications, routers must handle routes containing prefixes of any valid length, from 0 to 128. However, some users have reported that routers exhibit worse performance when routing using long prefixes, in particular when using prefixes longer than 80 bits. This implies that using prefixes shorter than that would result in better performance in some cases.

5.2.4. SUFFIX

If the optional SUFFIX is required, the prefix should leave the space for the SUFFIX.

5.2.5. EUI-64 format

The selection of the prefix length may affect the EUI-64 format, since the subnet may not be in the 64 bit boundary. However, there is no contradiction to [[RFC4291](#)], which states that first, an interface identifier has to be unique on the LAN-or-whatever it is on. And second, when the interface identifier is a MAC address it should be in a format related to a MAC Address - except that it is different. So, especially given privacy addressing, we can't really assume that the router or neighboring host is going to extract a MAC address from the interface identifier and directly use it to direct a datagram to another host; rather, the relationship between a MAC address and an interface identifier has a level of indirection managed by Neighbor Discovery.

5.3. SUFFIX Analysis

In the current implementation of the stateless mode, the suffix is entirely zero. For the stateful mode when using Well-Known prefix, the suffix can be used to represent different NAT boxes.

6. Conclusions

The embedded address format is defined in this document, which can be used to represent IPv4 addresses in IPv6 for both stateless and stateful translations. The embedded address format is also used to represent IPv6 addresses in IPv4 for stateless translation. The PREFIX, prefix length and SUFFIX in the embedded address format are defined as well in this document.

6.1. PREFIX Recommendation

The PREFIX Recommendations are:

- o In the case when different sites are using same IPv4 addresses (for example, [[RFC1918](#)] space), the LIR MUST be used.
- o In the "an IPv4 network connecting to IPv6 Internet scenario, the LIR MUST be used.
- o In the stateless mode of large scale networks, the LIR MUST be used.
- o In other cases, the LIR is RECOMMENDED and all of the relevant protocols and software need to accommodate the ability to configure that LIR prefix.
- o If for some reason, you cannot use LIR (e.g. in an isolated network), you should use this WKP allocated by IANA for this purpose, rather than pulling one out of thin air, or using a prefix allocated for a different purpose.
- o When the WKP is used, it MUST be treated same as the 6to4 prefixes and the corresponding routing practice MUST be taken. (6to4 prefixes more specific than 2002::/16 must not be propagated in native IPv6 routing, to prevent pollution of the IPv6 routing table by elements of the IPv4 routing table. Therefore, a 6to4 site which also has a native IPv6 connection MUST NOT advertise its 2002::/48 routing prefix on that connection, and all native IPv6 network operators MUST filter out and discard any 2002:: routing prefix advertisements longer than /16.

6.2. Prefix Length Recommendation

For the prefix length selection, there are some obvious values that might be popular, including /40, /44, and /96, but there is no requirement than any of them be used; this is left to the operator's discretion.

6.3. SUFFIX Recommendation

For the SUFFIX selection, it is entirely zero at this time. However, it could be used for the future extension of the translation functions.

7. IANA Considerations

The future versions of this memo will require WKP assignment from IANA. It is an IPv6 block, but the prefix length of this block has not be determined. The possibilities are /16 (similar to 6to4 block), /32, /48 or /96.

Note to RFC Editor: This section will have served its purpose if it correctly tells IANA that no new assignments or registries are required, or if those assignments or registries are created during the RFC publication process. From the author's perspective, it may therefore be removed upon publication as an RFC at the RFC Editor's discretion.

8. Security Considerations

One "security" issue has been raised, with an address format that was considered and rejected for that reason. At this point, the editor knows of no other security issues raised by the address format that are not already applicable to the addressing architecture in general.

9. Acknowledgements

This is under development by a large group of people. Those who have posted to the list during the discussion include Andrew Sullivan, Andrew Yourtchenko, Brian Carpenter, Congxiao Bao, Dan Wing, Ed Jankiewicz, Fred Baker, Hiroshi Miyata, Iljitsch van Beijnum, John Schnizlein, Keith Moore, Kevin Yin, Magnus Westerlund, Marcelo Bagnulo Braun, Margaret Wasserman, Masahito Endo, Phil Roberts, Philip Matthews, Remi Denis-Courmont, Remi Despres, William Waites and Xing Li.

10. References

10.1. Normative References

[I-D.bagnulo-behave-dns64]
Bagnulo, M., Sullivan, A., Matthews, P., Beijnum, I., and

M. Endo, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [draft-bagnulo-behave-dns64-02](#) (work in progress), March 2009.

[I-D.bagnulo-behave-nat64]

Bagnulo, M., Matthews, P., and I. Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [draft-bagnulo-behave-nat64-03](#) (work in progress), March 2009.

[I-D.baker-behave-v4v6-framework]

Baker, F., Li, X., and C. Bao, "Framework for IPv4/IPv6 Translation", [draft-baker-behave-v4v6-framework-02](#) (work in progress), February 2009.

[I-D.baker-behave-v4v6-translation]

Baker, F., "IP/ICMP Translation Algorithm", [draft-baker-behave-v4v6-translation-02](#) (work in progress), February 2009.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

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

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

[10.2. Informative References](#)

[I-D.baker-behave-ivi]

Li, X., Bao, C., Baker, F., and K. Yin, "IVI Update to SIIT and NAT-PT", [draft-baker-behave-ivi-01](#) (work in progress), September 2008.

[I-D.durand-softwire-dual-stack-lite]

Durand, A., Droms, R., Haberman, B., and J. Woodyatt, "Dual-stack lite broadband deployments post IPv4 exhaustion", [draft-durand-softwire-dual-stack-lite-01](#) (work in progress), November 2008.

[I-D.ietf-v6ops-addcon]

Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "IPv6 Unicast Address Assignment Considerations", [draft-ietf-v6ops-addcon-10](#) (work in progress), September 2008.

[I-D.miyata-v6ops-snatpt]

Miyata, H. and M. Endo, "sNAT-PT: Simplified Network Address Translation - Protocol Translation", [draft-miyata-v6ops-snatpt-02](#) (work in progress), September 2008.

[I-D.wing-behave-nat64-referrals]

Wing, D., "Referrals Across a NAT64", [draft-wing-behave-nat64-referrals-00](#) (work in progress), March 2009.

[I-D.xli-behave-ivi]

Li, X., Chen, M., Bao, C., Zhang, H., and J. Wu, "Prefix-specific and Stateless Address Mapping (IVI) for IPv4/IPv6 Coexistence and Transition", [draft-xli-behave-ivi-01](#) (work in progress), February 2009.

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

[RFC2428] Allman, M., Ostermann, S., and C. Metz, "FTP Extensions for IPv6 and NATs", [RFC 2428](#), September 1998.

[RFC2765] Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)", [RFC 2765](#), February 2000.

[RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", [RFC 2766](#), February 2000.

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

[RFC3142] Hagino, J. and K. Yamamoto, "An IPv6-to-IPv4 Transport Relay Translator", [RFC 3142](#), June 2001.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), February 2003.

[RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", [RFC 3879](#), September 2004.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for

Renumbering an IPv6 Network without a Flag Day", [RFC 4192](#), September 2005.

- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", [RFC 4213](#), October 2005.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", [RFC 4380](#), February 2006.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", [RFC 4864](#), May 2007.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), September 2007.
- [RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", [RFC 4966](#), July 2007.
- [RFC5211] Curran, J., "An Internet Transition Plan", [RFC 5211](#), July 2008.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", [RFC 5214](#), March 2008.

Authors' Addresses

Xing Li (editor)
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing, 100084
China

Phone: +86 62785983
Email: xing@cernet.edu.cn

Congxiao Bao (editor)
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing, 100084
China

Phone: +86 62785983
Email: cong Xiao@cernet.edu.cn

Fred Baker (editor)
Cisco Systems
Santa Barbara, California 93117
USA

Phone: +1-408-526-4257
Fax: +1-413-473-2403
Email: fred@cisco.com

