Behave and Softwires WGs	D. Wing	
Internet-Draft	D. Ward	
Intended status: Informational	Cisco	
Expires: April 3, 2009	A. Durand	
	Comcast	
	September 30, 2008	

TOC

A Comparison of Proposals to Replace NAT-PT
draft-wing-nat-pt-replacement-comparison-02

Status of this Memo

Abstract

As we approach IPv4 address depletion, the IETF must provide for IPv4 and IPv6 coexistence: a way for ISPs and enterprises to reduce public IPv4 address consumption and a way for hosts to migrate to IPv6 connectivity -- while providing reasonable access for those IPv6 hosts to access the IPv4 Internet.
This draft compares eight proposals for IPv6 and IPv4 coexistence.

Table of Contents

---

1. Terminology

TOC

The following terms are used throughout this document.

Address Family Translation (AFT): The function of translating from one IP address family (IPv4 or IPv6) to another (IPv6 or IPv4).

Carrier Grade NAT (CGN): A NAT device used by many subscribers (homes or end sites), where 'many' would be on the order of dozens to hundreds of thousands of subscribers. This might NAT between any combination of IPv4 and IPv6. Typically, the end user does not have the ability to adjust the behavior of the CGN (i.e., no ability to create static port mapping).

CPE router: Customer Premise Equipment router. A device that performs routing functions, located at the customer's premise. This device does not perform NAT functions. (Some referenced specifications use the term 'Home GateWay' (HGW) to mean the same thing. We use CPE router because the subscriber might not be a business and not a 'home'.)

DNS rewriting: The generalized function of synthesizing a DNS AAAA response from a DNS A response. The term "DNS rewriting" instead of "DNS-ALG" because in NAT-PT (Tsirtsis, G. and P. Srisuresh, “Network Address Translation - Protocol Translation (NAT-PT),” February 2000.) [RFC2766] DNS-ALG meant a DNS rewriting function with an interface to the NAT function.

NAT: Network Address Translation. This translates IP addresses, one-for-one, between two networks, without changing transport protocol ports. The two networks might be IPv4 (NAT44), IPv6 and IPv4 (NAT64), or IPv4 and IPv6 (NAT46). This document follows

(the unfortunate) common usage that "NAT" can also mean "NAPT" (Network Address and Port Translator).

Softwire A tunnel for carrying IPv4 and IPv6 traffic over IPv6 and IPv4 networks [RFC4925] (Li, X., Dawkins, S., Ward, D., and A. Durand, “Softwire Problem Statement,” July 2007.).

---

2. Introduction

TOC

As the Internet approaches IPv4 address depletion, it will be necessary for Internet Service Providers to continue to simultaneously provide their users with access to the IPv4 Internet, reduce the number of IPv4 public addresses consumed by each subscriber, and provide a way for subscribers to migrate to IPv6.
The proposals have several high-level attributes in common:

Provide access to the IPv4 Internet: There are two approaches to provide access to the v4 Internet. One approach is to have a dual-stack host with some modifications from the classic design [RFC4213] (Nordmark, E. and R. Gilligan, “Basic Transition Mechanisms for IPv6 Hosts and Routers,” October 2005.). The other approach is to have an IPv6-only host and operate an address family translation (AFT) device between the IPv6-only host and the IPv4 Internet.

Reduce consumption of global IPv4 addresses: Network address and port translator (NAPT) technology, and NAPT itself, allows more than one host to simultaneously use a single IPv4 address. NAPT technology is used in all proposals to conserve IPv4 public address space.

IPv6 migration: it is important that a migration path for users and content providers to move to IPv6 is enabled and encouraged. This is necessary because operating a NAT device in order to reduce per-subscriber IPv4 address consumption is not a viable long-term solution: we will still exhaust the IPv4 public address space, and operating NATs is expensive and reduces the reliability of the Internet.

Port limitations: All proposals use a NAPT to provide access to the IPv4 Internet, which reduces the number of ports each subscriber can use. This has negative impacts on some applications (e.g., Apple iTunes, Google Maps). This problem is resolved by the content provider and the subscriber both using IPv6.

This draft is discussed on the [v4v6interm-interest] (IETF, "v4v6interm-interest mailing list," .) mailing list. Individual proposals are discussed on the mailing list indicated in this document.

---

3. Overview of Proposals

TOC

This document classifies the proposals into two categories. The first category provides IPv4 and IPv6 access to the subscriber, and the second category provides only IPv6 access to the subscriber. In both categories, IPv4 addresses are conserved by using a NAT device. This NAT device is placed in the carrier's network ("Carrier Grade NAT") or (in the case of A+P and SAM) in the CPE router. In all proposals (except NAT444) a host can obtain native IPv6 connectivity with native IPv6 hosts without regard to the co-existence proposal.
The descriptions below provide a very brief overview of each proposal, in alphabetical order.

3.1. IPv4 hosts in Customer Premise

TOC

For Internet access, the following proposals allow for IPv4 hosts in the customer premise.

3.1.1. Address Plus Port (A+P)

TOC

Address Plus Port (Maennel, O., Bush, R., Cittadini, L., and S. Bellovin, "A Better Approach than Carrier-Grade-NAT," .) [A+P] uses a NAT in (or close to) the customer premise for access to the IPv4 Internet. The Service Provider conveys both an IPv4 address (as done today) and a range of TCP/UDP ports to the NAT. Outgoing IPv4 traffic is NATted to that range of TCP/UDP ports, and the Service Provider routes packets to the appropriate customer using both the destination IPv4 address (as done today) and destination TCP/UDP port.
One of A+P's architectures is depicted in Figure 1 (Address Plus Port, v4-capable ISP (A+P-v4)), where the ISP's network supports both IPv4 and IPv6. In this architecture, the NAT's job is straight forward: it NATs to a limited port range and sends the packet upstream to the ISP. Because multiple customers share a single IPv4 address, the aggregation router needs to route return packets to the appropriate customer's NAT using destination IPv4 and destination port. This architecture is denoted as "A+P-v4".

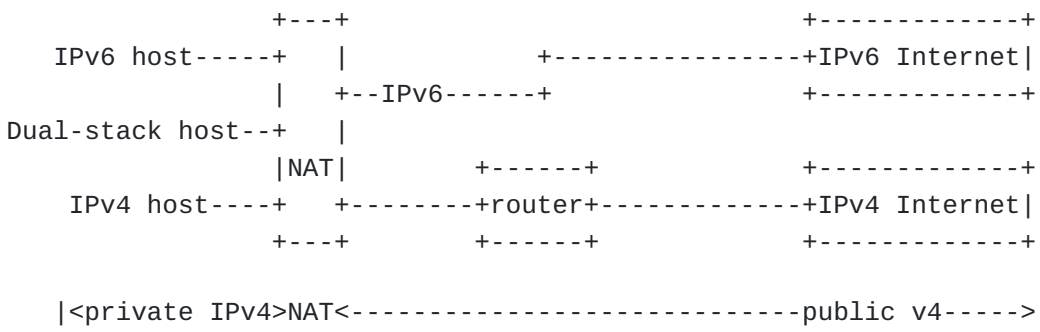

Figure 1: Address Plus Port, v4-capable ISP (A+P-v4)

Another of A+P's architectures is depicted in Figure 2 (Address Plus Port, v6-only ISP (A+P-v6)), where the ISP's network runs only IPv6. In this architecture, the NAT encapsulates the IPv4 packet into an IPv6 packet, where the IPv6 packet has a source address that corresponds to that NAT, and a destination address that corresponds to a well-known prefix which routes to a IPv6 tunnel concentrator. When the packet arrives at the IPv6 tunnel concentrator the IPv6 source address is used to construct the IPv4 source address and TCP/UDP port, and the IPv6 destination address is used to construct the IPv4 destination address; the IPv6 destination TCP/UDP port is taken from the IPv6 packet's destination TCP/UDP header. This architecture is denoted as "A+P-v6".

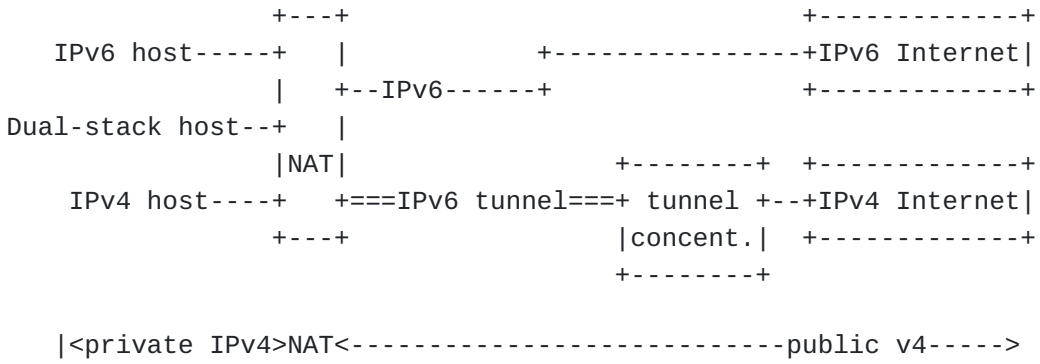

Figure 2: Address Plus Port, v6-only ISP (A+P-v6)

TOC

3.1.2. Stateless Address Mapping (SAM) (previously APB-Revised)

Stateless Address Mapping (SAM) [I-D.despres-sam] (Despres, R., "Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM)," July 2009.) shares each IPv4 address amongst several subscribers through a tunnel aggregation device. The static mapping avoids the need for the service provider equipment to NAT.
SAM can be implemented with subscriber site tunnel endpoints either in a router (CPE router or other router) or in a SAM host. In both implementations, each subscriber site is assigned a shared IPv4 address (shared with other subscribers) and a port range. Figure 3 (SAM-CPE, tunnel between CPE and tunnel concentrator) shows the SAM architecture in the case where the tunnel is established between the CPE router (upgraded to support SAM) and the SAM-capable tunnel concentrator. Any IPv4 traffic from hosts behind the CPE router is NAT'd (using classic NAT44) and forwarded through the tunnel to the SAM tunnel concentrator. The customer premise NATs using the external port range it is 'borrowing' from the SAM concentrator. This is abbreviated SAM-CPE in this document.
This proposal is discussed in [Softwires] (IETF, "Softwires working group mailing list," .).

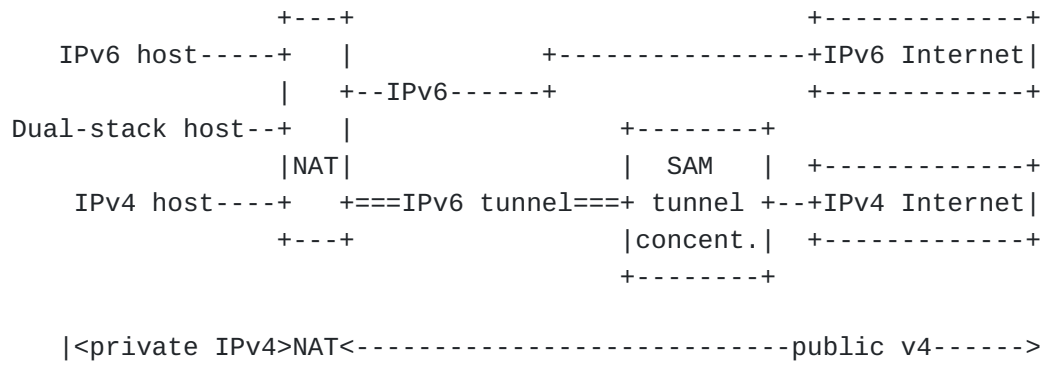

Figure 3: SAM-CPE, tunnel between CPE and tunnel concentrator

In the figure above, the IPv6 tunnel is an IPv4-over-IPv6 tunnel. Figure 4 (SAM-host - tunnel between host and tunnel concentrator) shows another SAM architecture where the tunnel is established directly between the host (upgraded to support SAM) and the SAM tunnel endpoint. Any IPv4 traffic from the SAM host is routed through the tunnel to the SAM-capable tunnel concentrator. Tunnelling is sufficient; no NAT device is needed between the host and the public IPv4 network. This is abbreviated SAM-host in this document. In Figure 4 (SAM-host - tunnel between host and tunnel concentrator), the customer premise NAT does

not NAT traffic to/from the SAM host; however, it does NAT traffic to/ from the IPv4-only host to support a non-SAM-capable IPv4 host.

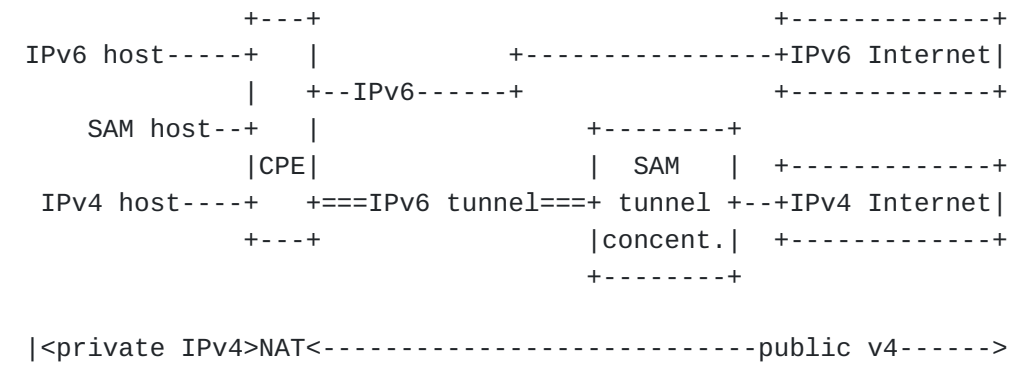

Figure 4: SAM-host - tunnel between host and tunnel concentrator

Figure 5 (SAM-HC - host and CPE tunnels) shows the SAM architecture with two tunnels. One tunnel is established between the CPE router and the SAM endpoint, and a second tunnel between the subscriber host and the CPE router. In this architecture, the CPE router is upgraded to establish a tunnel to the SAM-capable tunnel concentrator (external side) and to accept a tunnel from the host (internal side); the SAM-capable host IP stack is upgraded to establish a tunnel to the CPE router. Any traffic from the SAM-capable host is routed by the host's SAM stack and forwarded through the tunnel to the CPE router. Tunnelling is sufficient; no NAT device is needed between the host and the core IPv4 network. This is abbreviated SAM-HC (Host and CPE router) in this document.

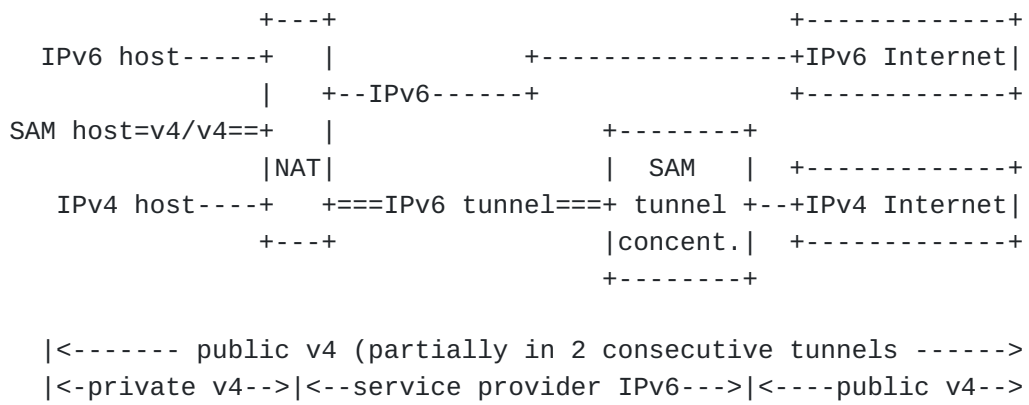

Figure 5: SAM-HC - host and CPE tunnels

3.1.3. Dual-Stack Lite (DS-Lite)

TOC

Dual-Stack Lite (DS-Lite) [I-D.durand-softwire-dual-stack-lite] (Durand, A., Droms, R., Haberman, B., and J. Woodyatt, "Dual-stack lite broadband deployments post IPv4 exhaustion," September 2008.) provides a global IPv4 address that is shared amongst several subscribers through a CGN. Each subscriber network is connected to the CGN through a tunnel, using IPv6 as the tunnel transport. All IPv4 traffic is sent inside of that tunnel. The tunnel endpoint implements Dual-Stack (Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers," October 2005.) [RFC4213]. This draft is discussed in [Softwires] (IETF, "Softwires working group mailing list," .).
DS-Lite can be implemented with the tunnel endpoints either in a router (CPE router or aggregation router) or in a host. In both cases, a single subscriber IPv4 address or IPv4 prefix may overlap, or even be identical for all subscribers. Addresses from overlapping address spaces are disambiguated by the tunnels between the subscriber networks and the CGN.
Figure 6 (Dual-Stack Lite, tunnel terminated on router (DS-Lite router)) shows the DS-Lite architecture in the case where the tunnel is terminated in a router, which could be the CPE router or an aggregation router. In the diagram, the router terminating the tunnel is a CPE router, but another router could be used as well (e.g., service provider's aggregation router). In this architecture, the router is upgraded to establish a tunnel to the CGN, and does not perform any NAT processing on subscriber traffic. The router provides DHCP service (addresses and other configuration information) to the subscriber hosts. This is abbreviated "DS-Lite router" in this document.

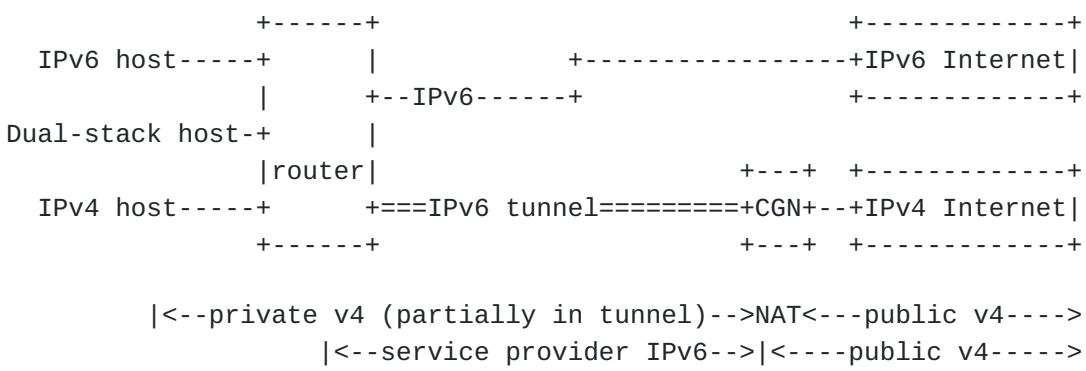

Figure 6: Dual-Stack Lite, tunnel terminated on router (DS-Lite router)

The choice of encapsulation for the IPv6 tunnel is outside the scope of this document.
Figure 7 (Dual-Stack Lite, tunnel terminated on host (DS-Lite host)) shows the DS-Lite architecture when the tunnel is terminated in the subscriber host. In this architecture, the DS-Lite host IP stack is upgraded to establish a tunnel to the CGN, through an unmodified CPE router and across either IPv6 transport. IPv4 traffic from the DS-Lite host is routed through the tunnel to the CGN. This is abbreviated "DS-Lite host" in this document.

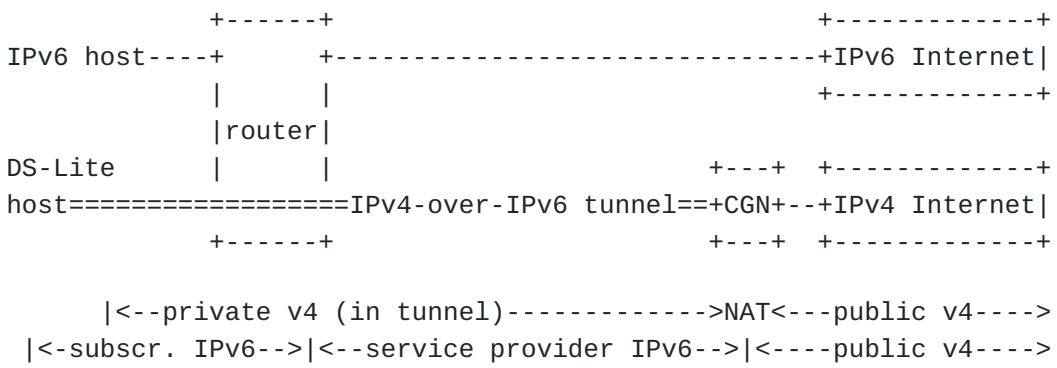

Figure 7: Dual-Stack Lite, tunnel terminated on host (DS-Lite host)

The choice of encapsulation for the IPv6 tunnel is outside the scope of this document.

3.1.4. NAT444

TOC

NAT444 (no written proposal) would NAT twice: first using a NAT device in the customer premise (as typically deployed today) and another NAT device in the ISP's network (a CGN). This proposal is discussed in [Behave] (IETF, “BEHAVE working group mailing list,” .).
The subscriber could access the IPv6 Internet using Teredo (Huitema, C., “Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs),” February 2006.) [RFC4380]. The Teredo service could be provided by the ISP (shown as "Teredo relay-1") or on the Internet (shown as "Teredo relay-2").

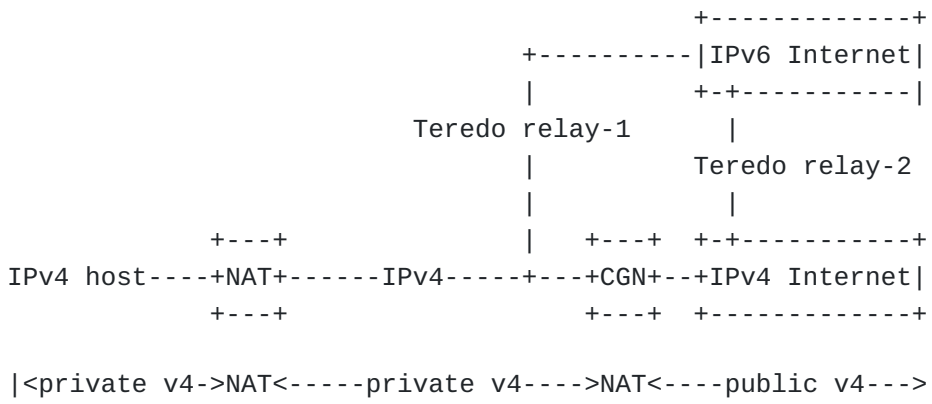

3.2. IPv6 hosts in Customer Premise

TOC

For access to the IPv4 Internet, the following proposals require IPv6 hosts in the customer premise, and do not support IPv4 hosts. These proposals provide access to the IPv4 Internet without requiring dual-stack on client equipment.

3.2.1. IVI

TOC

IVI ([I-D.xli-behave-ivi] (Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, “The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition,” January 2010.), [I-D.baker-behave-ivi] (Li, X., Bao, C., Baker, F., and K. Yin, “IVI Update to SIIT and NAT-PT,” September 2008.)) uses an address and service architecture designed to facilitate transition from an IPv4 Internet to an IPv6 Internet. This service contains three parts: A DNS Application Layer Gateway, a stateful Network Address Translator that enables IPv6 clients to initiate connections to IPv4 servers and peers, and a stateless Network Address Translator that enables IPv4 and IPv6 systems to interoperate freely.
For an IPv6 host needing access to IPv4 hosts, IVI is similar to both SIIT (Nordmark, E., “Stateless IP/ICMP Translation Algorithm (SIIT),” February 2000.) [RFC2765] and NAT-PT (Tsirtsis, G. and P. Srisuresh, “Network Address Translation - Protocol Translation (NAT-PT),” February 2000.) [RFC2766] but with a different address format. IVI's DNS rewriting function (A to AAAA) returns an IPv6 address that routes

to a specific translation gateway that advertises that IPv6 prefix in the service provider's network. The DNS server may be in the IVI gateway or in a separate system related to it.
IVI also allows IPv4 hosts to access a IPv6 host, using a stateless NAT. This is accomplished by providing the IPv6 host an IVI address, which is simply an IPv6 address from a pool of IPv6 addresses. This pool of IPv6 addresses has a fixed IPv4-to-IPv6 mapping algorithm applied to translate between the two address families and the translation is implemented by an IVI gateway. The IPv6 address would be advertised in DNS with an A record, pointing to the IVI gateway. This allows IPv6-only hosts to have a presence on the IPv4 Internet. In this scheme, subsets of the IPv4 addresses are embedded in prefix-specific IPv6 addresses and these IPv6 addresses can therefore communicate with the global IPv6 networks directly and can communicate with the global IPv4 networks via stateless (or almost stateless) gateways. DNS rewriting is not used, or necessary, for this fixed mapping of IPv4 addresses to IPv6 address.
This proposal is discussed in [Behave] (IETF, “BEHAVE working group mailing list,” .).

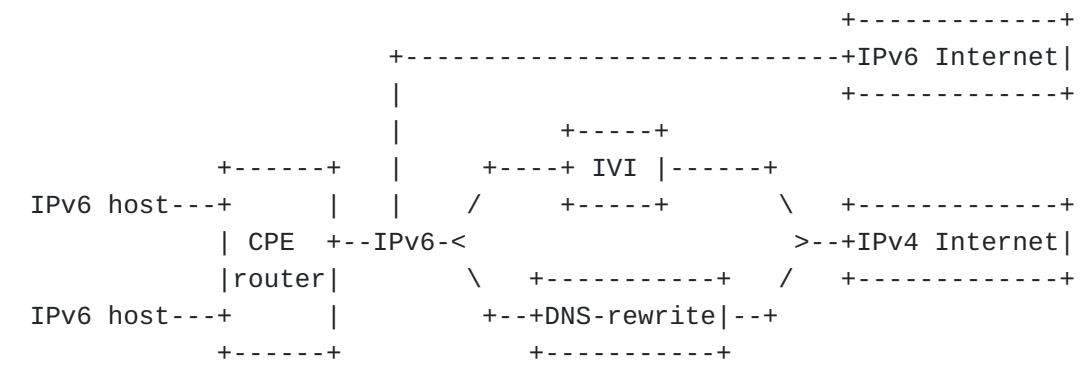

Figure 9: IVI

3.2.2. NAT6

TOC

NAT6 (Jennings, C., “NAT for IPv6-Only Hosts,” November 2008.)
[I-D.jennings-behave-nat6] encourages IPv6 host itself to provide necessary DNS rewriting functions (appending a configured IPv6 prefix to an IPv4 address to create an IPv6 address) and have the NAT function (from IPv6 to IPv4) performed in the network. By having the host provide the DNS rewriting function, a DNS rewriting function in the

network is avoided. This proposal is discussed in [Behave] (IETF, "BEHAVE working group mailing list," .).

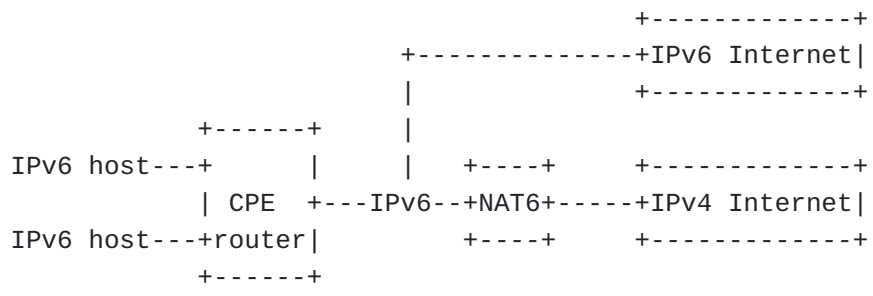

Figure 10: NAT6

3.2.3. NAT64

TOC

For an IPv6 host needing access to IPv4 hosts, NAT64 (Bagnulo, M., Matthews, P., and I. Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers," March 2009.) [I-D.bagnulo-behave-nat64] uses the logic of SIIT (Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)," February 2000.) [RFC2765] and NAT-PT (Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)," February 2000.) [RFC2766] but with a different address format. This removes the relationship between the NAT function and DNS function. Rather than using the DNS-ALG described in [RFC2766] (Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)," February 2000.), the DNS service simply advertises DNS A and AAAA records specifying the IPv6 address of the NAT64 device (which is a CGN device). The DNS server may be in the CGN or in a separate system related to it. This proposal is discussed in [Behave] (IETF, "BEHAVE working group mailing list," .).

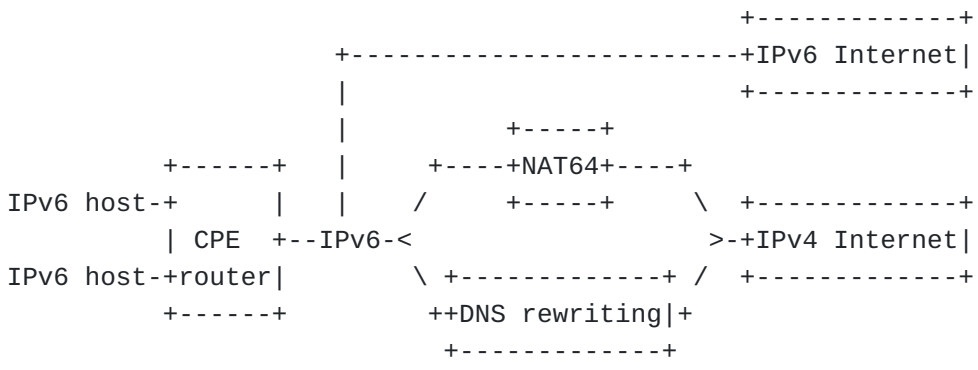

Figure 11: NAT64

It is also possible to utilize NAT64 to access private IPv4 address (Figure 12 (NAT64 to Private IPv4 Addresses)). To perform this function, NAT64 allows using a locally-assigned IPv6 prefix out of the address block of the site running the NAT64 device, and allows using a well-known prefix assigned to this purpose.

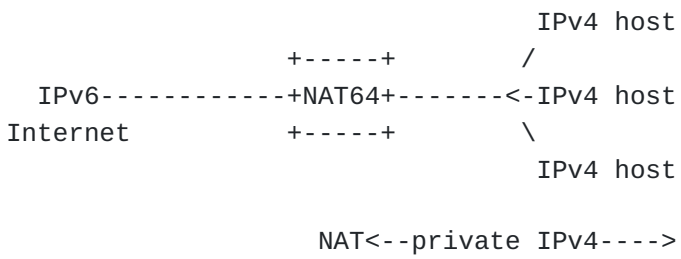

Figure 12: NAT64 to Private IPv4 Addresses

Note that in this scenario, DNS rewriting is not necessary as all of the IPv4 addresses could be given AAAA records.

3.2.4. NAT-PT

TOC

This section is provided for reference only, because NAT-PT has been deprecated by [RFC4966] (Aoun, C. and E. Davies, “Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status,” July 2007.).

NAT-PT (Tsirtsis, G. and P. Srisuresh, “Network Address Translation - Protocol Translation (NAT-PT),” February 2000.) [RFC2766] provides a combined DNS-ALG and NAT function, which share state. This is typically implemented in a single device.

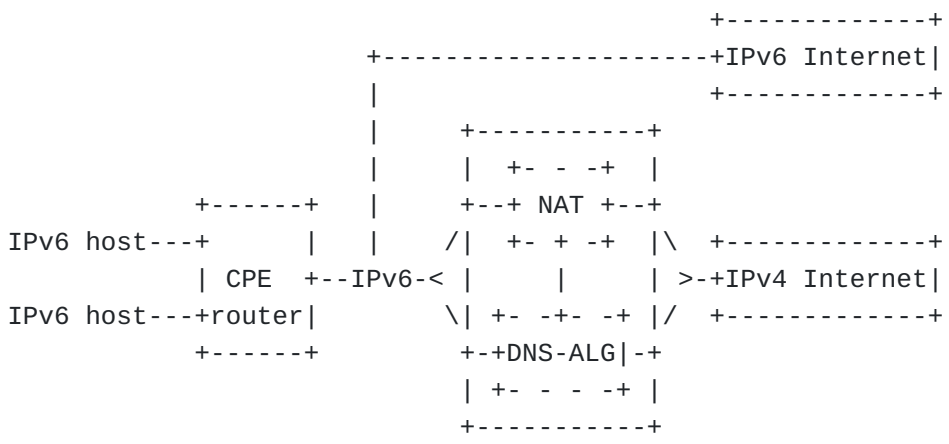

Figure 13: NAT-PT

NAT-PT (Tsirtsis, G. and P. Srisuresh, “Network Address Translation - Protocol Translation (NAT-PT),” February 2000.) [RFC2766] and [RFC4966] (Aoun, C. and E. Davies, “Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status,” July 2007.) can be discussed in [Behave] (IETF, “BEHAVE working group mailing list,” .).

3.2.5. sNAT-PT

TOC

For an IPv6 host needing access to IPv4 hosts, sNAT-PT (Miyata, H. and M. Endo, “sNAT-PT: Simplified Network Address Translation - Protocol Translation,” September 2008.) [I-D.miyata-v6ops-snatpt] provides DNS rewriting and NAT functionality. The DNS rewriting component is described in [I-D.endo-v6ops-dnsproxy] (Endo, M. and H. Miyata, “Translator Friendly DNS Proxy,” October 2008.).
This proposal is discussed in [Behave] (IETF, “BEHAVE working group mailing list,” .).

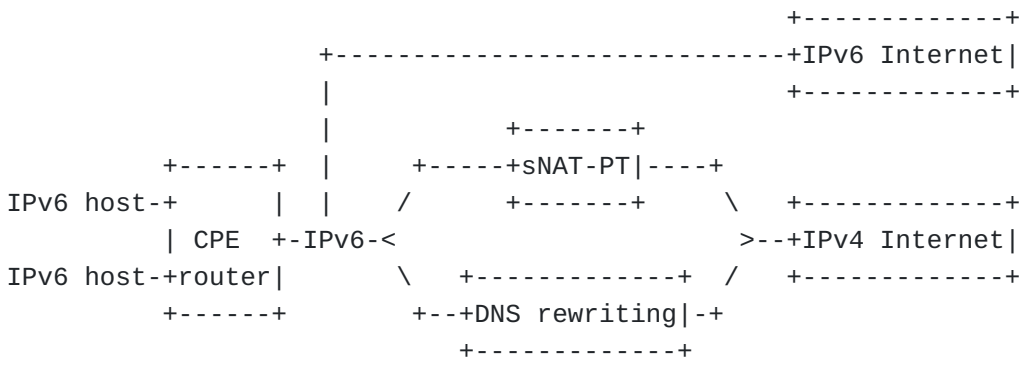

Figure 14: sNAT-PT

---

sNAT-PT also provides access from the IPv4 Internet to IPv6 hosts. This can be done with a 1-for-1 mapping or with a 1-for-N mapping using IPv4 ports. These do not require a DNS rewriting function.

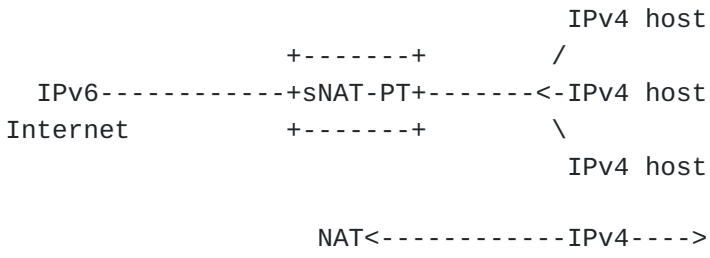

Figure 15: sNAT-PT

---

4. Changes Required in Network Elements

TOC

This section describes changes to network elements for various scenarios. In all cases, the content provider's DNS and content provider's network does not need to change (except due to the problem of port limitations as described in Section 2 (Introduction)).

TOC

4.1. IPv4 and IPv6 Hosts Accessing the IPv4 Internet

For the case of an IPv4 host, IPv6 host, or dual-stack host that need to connect to IPv4 hosts on the Internet, the following table summarizes the changes required to subscriber's hosts (when CPE routers are present and when CPE routers are not present) and to some network elements:

Proposal	Subscriber Hosts w/CPE router	Subscriber Hosts w/o CPE router	CPE router	ISP Access Edge Network
A+P-v4	no change	no change (A+P NAT would be performed by SP)	A+P support	route using destination port
A+P-v6	no change	no change (A+P NAT would be performed by SP)	A+P support	tunnel concentrator
SAM-CPE	no change	(not applicable)	SAM CPE	tunnel concentrator
SAM-host	SAM-host	SAM-host	no change	tunnel concentrator
SAM-HC	SAM support	(not applicable)	SAM CPE internal & external	tunnel concentrator
NAT444	no change	no change	no change	NAT v4v4
DS-Lite router	no change	(not supported; use DS-Lite host)	DS-Lite CPE	NAT v4v4 w/ tunnel
DS-Lite host	(not supported; use DS-Lite router)	DS-Lite v6	no change	NAT v4v4 w/ tunnel
IVI	move to v6	move to v6	move to v6	IVI + DNS rewriting
NAT6	move to v6	move to v6	move to v6	NAT6
NAT64	move to v6	move to v6	move to v6	NAT64 + DNS rewriting
NAT-PT	move to v6	move to v6	move to v6	NAT-PT + DNS-ALG
sNAT-PT	move to v6	move to v6	move to v6	sNAT-PT + DNS rewriting

Table 1: Changes Required to Network Elements

For IPv6 hosts that access the IPv4 Internet, the following table describes the high-level technologies used by each proposal.

Proposal	ISP's Internal Network	DNS Impact	Carrier Grade NAT
A+P-v4	IPv4 destination port routing	no change	(no CGN, if subscriber's NAT support A+P NAT)
A+P-v6	IPv4/IPv6 tunnel	no change	(no CGN, if subscriber's NAT support A+P NAT)
SAM	IPv4/IPv6 tunnel	no change	(no CGN)
DS-Lite router	IPv4/IPv6 tunnel	no change	IPv4/IPv4
DS-Lite host	IPv4/IPv6 tunnel	no change	IPv4/IPv4
NAT444	(v6 not supported)	(v6 not supported)	(v6 not supported)
IVI	v4 NATted, native v6 address	DNS rewriting	IPv6/IPv4
NAT64	v4 NATted, native v6 address	DNS rewriting	IPv6/IPv4
NAT-PT	v4 NATted, native v6 address	DNS-ALG	IPv6/IPv4
sNAT-PT	v4 NATted, native v6 address	DNS rewriting	IPv6/IPv4

Table 2: IPv6 to IPv4 - technologies involved

4.2. IPv4 Hosts Accessing the IPv4 Internet

TOC

The following table compares the five mechanisms that support end hosts running IPv4 to access the IPv4 Internet: SAM, Dual-Stack Lite, NAT444.

Proposal	CPE router	ISP's Internal Network	Service Provider Equipment

A+P-v4	IPv6 support + A+P NAT44	IPv4 and IPv6	destination port routing
A+P-v6	IPv6 support + IPv4/ IPv6 tunnel + A+P NAT44	IPv6	IPv6 tunnel termination
SAM-CPE	IPv6 support + IPv4/ IPv6 tunnel + NAT44	IPv6	IPv6 tunnel termination
DS-Lite router	IPv6 support + IPv4/ IPv6 tunnel	IPv6	IPv6 tunnel termination, NAT44 (CGN)
DS-Lite host	IPv6 support (if using DS-Lite IPv6 tunneling)	IPv6 (if using DS-Lite IPv6 tunneling)	IPv6 tunnel termination, NAT44
NAT444	no change	multi-realm IPv4	NAT44 (CGN)

Table 3: IPv4 Hosts Accessing the IPv4 Internet

The proposals IVI, NAT6, NAT64, NAT-PT, and sNAT-PT are not shown in table Table 3 (IPv4 Hosts Accessing the IPv4 Internet) because those proposals provide no support for IPv4-only hosts to access the IPv4 Internet.

4.3. IPv4 Internet Accessing IPv6 hosts

TOC

IVI, NAT-PT, and sNAT-PT all provide mechanisms for IPv4 hosts on the Internet to access IPv6-only servers. Such mappings consume IPv4 address space.
IVI: IVI allows 1:1 mapping from an IPv4 client to an IPv6 server. IVI also allows 1:n mappings, by utilizing the TCP/UDP port number of the incoming IPv4 packet in the algorithm to determine the destination IPv6 host; this conserves IPv4 address space consumption for those hosts that need a few TCP/UDP ports available from the IPv4 Internet.
NAT-PT: NAT-PT allows 1:n mapping from an IPv4 client to an IPv6 server, which is accomplished by dynamically mapping an IPv4 address to an IPv6 address after a DNS "A" record query.
sNAT-PT: NAT-PT allows 1:1 mapping from an IPv4 client to an IPv6 server.

TOC

5. Port Forwarding

Some applications require accepting incoming UDP or TCP traffic. When the remote host is on IPv4, the incoming traffic will be directed towards an IPv4 address. The applications are separated into two broad categories: those requiring static incoming ports and those requiring dynamic incoming ports.
Due to IPv4 NATs and IPv4 firewalls, some applications use [UPnP-IGD] (UPnP Forum, “Universal Plug and Play Internet Gateway Device,” 2000.) (e.g., XBox) or ICE (Rosenberg, J., “Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols,” October 2007.) [I-D.ietf-mmusic-ice] (e.g., SIP, Yahoo!/Google/Microsoft chat networks), other applications have all but completely abandoned incoming connections (e.g., most FTP transfers use passive mode). But some applications rely on ALGs, UPnP IGD, or manual port configuration. Further discussion in the IETF community is necessary to decide how to proceed on this issue.

> Note: Placing application awareness (i.e., ALG) in the CGN will cause bug fixes and new features to be delayed by development, testing, and deployment. To prevent such delays, application awareness should be placed elsewhere (e.g., in the CPE router or in the end host).

> Note: Extending NAT-PMP (Cheshire, S., “NAT Port Mapping Protocol (NAT-PMP),” April 2008.) [I-D.cheshire-nat-pmp] to support IPv6 could provide provide static port forwarding and dynamic port forwarding for IPv4 and IPv6 hosts needing access from IPv4 hosts.

5.1. Static Incoming Ports

TOC

Static incoming ports are used by applications for multiple sessions.

> Note: Some applications (e.g., BitTorrent) can use UPnP IGD to control IPv4 NATs and open a static incoming port. However, technical limitations of UPnP IGD appear to prevent UPnP IGD from being directly implemented in a CGN. The most significant technical limitation is that UPnP IGD expects the control point (the host) to be able to specify the public port; with hundreds of subscribers utilizing the same public IP address, this is untenable. Other UPnP IGD technical limitations may be surmountable (e.g., UPnP IGD's ability to create and destroy mappings for other IP addresses).

Examples of applications that require static incoming ports include:

*HTTP

*SMTP (must be on TCP/25)

*ssh

*BitTorrent

*games (of particular note is that XBox uses UPnP IGD)

The solutions proposed for static ports are:

> A+P: The subscriber's customer premise NAT can forward ports within the allocated port range. This port could be advertised by the subscriber using DNS SRV resource records or other means.
>
> SAM-host and SAM-HC: assign a port in the available port range; advertise it with the IPv4 address using a DNS SRV resource record.
>
> Dual-Stack Lite: none
>
> NAT444: none
>
> IVI: assign IPv6 IVI address to IPv6 hosts that require incoming IPv4 connections
>
> NAT6: none
>
> NAT64: none
>
> sNAT-PT: assign IPv4 address to IPv6 hosts that require incoming IPv4 sessions.

5.2. Dynamic Incoming Ports

TOC

Dynamic incoming ports are, generally, used by applications for a single session. Examples of applications that require dynamic incoming ports include:

*applications that use real-time transport protocol (RTP)

  -SIP, RTSP, H.323, MGCP, H.248/Megaco

*non-passive FTP client

*games (of particular note is that XBox uses UPnP IGD)

The solutions proposed for dynamic ports are:

> A+P: An ALG can be incorporated into the subscriber's A+P-aware NAT, as done today with subscriber's NAT44 devices.
>
> SAM-host and SAM-HC: assign a port in the available port range.
>
> Dual-Stack Lite: none
>
> NAT444: none (although it is reasonable to expect that ALGs, as they exist in today's IPv4 NATs, might be utilized)
>
> IVI: assign IPv6 IVI address to IPv6 hosts that require incoming IPv4 connections
>
> NAT6: none
>
> NAT64: applications could be modified to support STUN (for TCP and UDP) to learn their public IPv4 address and TCP/UDP port.
>
> sNAT-PT: assign IPv4 address to IPv6 hosts that require incoming IPv4 connections.

---

6. Transport Protocol Support

TOC

[[Placeholder: discuss how DCCP and SCTP (and other transport protocols) are supported by each proposal. Although existing IPv4 NATs do not support DCCP or SCTP, it is reasonable to expect that new NATs could support those transport protocols if we want those protocols to work between address families.

---

7. Analysis with V6OPS's NAT64 Problem Statement

TOC

This section analyzes how each proposal maps to the requirements in [I-D.ietf-v6ops-nat64-pb-statement-req] (Bagnulo, M., Baker, F., and I. Beijnum, “IPv4/IPv6 Coexistence and Transition: Requirements for solutions,” May 2008.).
[[Placeholder until [I-D.ietf-v6ops-nat64-pb-statement-req] (Bagnulo, M., Baker, F., and I. Beijnum, “IPv4/IPv6 Coexistence and Transition: Requirements for solutions,” May 2008.) becomes more stable.]]

---

8. Comparison of Proposals with NAT-PT Problems

TOC

The following sections analyze how proposals fare against the problems caused by NAT-PT (Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)," February 2000.) [RFC2766] as documented in [RFC4966] (Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status," July 2007.):

8.1. Issues Unrelated to an DNS-ALG

TOC

8.1.1. Issues with Protocols Embedding IP Addresses

TOC

NAT6 requires applications to handle NAT6 traversal themselves.
The other proposals are silent on this issue, but in general using an application layer gateway (ALG), in some device in the network, appears to be the only solution to this problem. See also Section 5 (Port Forwarding).

8.1.2. NAPT-PT Redirection Issues

TOC

All proposals are silent on this issue.

8.1.3. NAT-PT Binding State Decay

TOC

NAT6 and NAT64 discuss binding lifetimes.
The other proposals are silent on this issue.

8.1.4. Loss of Information through Incompatible Semantics

TOC

All proposals are silent on this issue.

TOC

8.1.5. NAT-PT and Fragmentation

[[NAT64, NAT6, DS-Lite, and IVI all mention fragmentation. Need to analyze how they differ.]]

8.1.6. NAT-PT Interaction with SCTP and Multihoming

TOC

IVI supports multi-homing if there is a 1:1 mapping between IPv4 and IPv6 addresses. However, 1:1 mapping is not sustainable as we approach IPv4 exhaustion.
SAM (both SAM-host and SAM-HC) support SCTP.
sNAT-PT explicitly indicates SCTP is out-of-scope.
The other proposals are silent on this issue. All proposals seem to be considering only TCP, UDP, and ICMP.

8.1.7. NAT-PT as a Proxy Correspondent Node for MIPv6

TOC

All proposals are silent on this issue.

8.1.8. NAT-PT and Multicast

TOC

IVI can support Source-Specific Multicast (Holbrook, H. and B. Cain, "Source-Specific Multicast for IP," August 2006.) [RFC4607] (see Section 7 of [I-D.xli-behave-ivi] (Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition," January 2010.)).
Dual-Stack Lite does not support multicast.
NAT6 does not specify how it can work with multicast.
In sNAT-PT, multicasting in either direction requires manual mapping.
The other proposals are silent on this issue.

> Note: it may be possible for IGMP messages to be propagated and proxied (Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying"), " August 2006.) [RFC4605] across their respective NAT device (Wing, D. and T. Eckert, "IP Multicast Requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT)," February 2008.) [RFC5135]. More study on this is needed.

8.2. Issues Exacerbated by the Use of DNS-ALG

TOC

8.2.1. Network Topology Constraints Implied by NAT-PT

TOC

The separation of NAT and DNS-rewriting reduces the impact of this issue. IVI, NAT64, and sNAT-PT separate the NAT and DNS-rewrite functions, and avoid this constraint.

8.2.2. Scalability and Single Point of Failure Concerns

TOC

The separation of NAT and DNS-rewriting reduces the impact of this issue. IVI, NAT64, and sNAT-PT all separate the NAT and DNS-rewrite functions.

8.2.3. Issues with Lack of Address Persistence

TOC

TBD.

8.2.4. DoS Attacks on Memory and Address/Port Pool

TOC

A CGN would only allow a certain subscriber to open a certain number of ports, thereby preventing a single subscriber from DoSing other subscribers ([I-D.nishitani-cgn] (Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, “Common requirements for IP address sharing schemes,” March 2010.), "a CGN SHOULD limit the number of the CGN external ports").

8.3. Issues Directly Related to Use of DNS-ALG

TOC

TOC

8.3.1. Address Selection Issues when Communicating with Dual-Stack End-Hosts

Unlike NAT-PT, all proposals that involve DNS-rewriting (IVI, NAT64, sNAT-PT) do not return synthetic AAAA records if a real AAAA record exists. This prevents the problem. A DNS timeout (of the AAAA query) will prevent the optimum DNS response from being returned (that is, the real AAAA record rather than the synthesized AAAA record pointing to the CGN device), but it is still possible to connect even when such a timeout occurs. In practice, such DNS timeouts are not a common occurance.
In [I-D.bagnulo-behave-nat64] (Bagnulo, M., Matthews, P., and I. Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers," March 2009.), EDNS0 (Vixie, P., "Extension Mechanisms for DNS (EDNS0)," August 1999.) [RFC2671] is proposed as the mechanism for a host to determine if the AAAA record is synthetic (that is, generated by the DNS rewriting function) or if the AAAA record is genuine (that is, was not synthesized).

8.3.2. Non-Global Validity of Translated RR Records

TOC

TBD.

8.3.3. Inappropriate Translation of Responses to A Queries

TOC

sNAT-PT avoids this problem with its stateful DNS proxy.

8.3.4. DNS-ALG and Multi-Addressed Nodes

TOC

The additional NAT binding state is not created if the DNS rewriting and NAT functions are separate. Thus, this problem is avoided by [I-D.xli-behave-ivi] (Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition," January 2010.), [I-D.bagnulo-behave-nat64] (Bagnulo, M., Matthews, P., and I. Beijnum, "NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers," March 2009.), and [I-D.miyata-v6ops-snatpt] (Miyata, H. and M. Endo, "sNAT-PT: Simplified Network Address Translation - Protocol Translation," September 2008.).

8.3.5. Limitations on Deployment of DNS Security Capabilities

TOC

DNSSEC is incompatible with synthesized DNS responses (DNS rewriting). NAT64 recommends that DNSSEC-capable IPv6-only hosts use the EDNS SAS option to ignore synthetic DNS responses. This would allow the IPv6 host to ignore synthetic DNS responses and allows DNSSEC to work for non- synthesized AAAA responses. This means, however, that DNSSEC only works for native IPv6 AAAA responses, and DNSSEC cannot be used for IPv4 A responses.
No other proposal discusses how it would work with DNSSEC.
Note: A proposal that does DNS rewriting only in a validating resolver (after validation), or construct records in the authoritative server, will work fine with DNSsec.

8.4. Impact on IPv6 Application Development

TOC

TBD.

---

9. Security Considerations

TOC

9.1. Address Sharing

TOC

When resources are shared it is important they are shared fairly. On today's Internet, the shared resource is bandwidth -- both the service provider's core bandwidth (sharing between subscribers) and subscriber access bandwidth (sharing between a subscriber's own hosts).
Subscribers are given an IP address(es) for their exclusive use. With all of the NAT44 and NAT64 mechanisms proposed, an IPv4 address is shared amongst several subscribers.
This address sharing raises some security considerations, including DoS potential (a subscriber might accidentally or purposefully use all available ports, denying ports to other subscribers [I-D.nishitani-cgn] (Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, “Common requirements for IP address sharing schemes,” March 2010.) and spoofing (a subscriber might send a packet with the correct IP address, but the port belongs to a different subscriber).
Address sharing causes false negatives and false positives for existing IP sddress spoofing mechanisms (DHCP snooping, ARP security, ingress filtering (Ferguson, P. and D. Senie, “Network Ingress Filtering:

Defeating Denial of Service Attacks which employ IP Source Address Spoofing," May 2000.) [RFC2827]).
For lack of a better identifier, many applications and systems use an IPv4 address as an end-host identifier and take action based on that identity. In the past, IP addresses sometimes provided additional privileges (e.g., the ability to login without a password using Berkeley "r services"). This persists today with some systems (e.g., DHCP snooping, ARP security, and email Sender Policy Framework (SPF)). Conversely, undesired behavior of a certain IP address can cause servers to refuse to provide service. For example, excessive connection attempts or excessive downloading can cause an HTTP server to delay (or refuse) providing service to that IP address. As another example, IP address blacklisting (e.g., DNSBL) might cause e-mail from that IP address to be considered more likely to be spam. Even with consumer NAT44, these systems work reasonably well because excessive connection attempts or spam originating from any host belonging to a subscriber is punished, without harming other subscribers of that ISP. (Of course, some such systems apply their rate limiting to entire subnets in order to purposefully punish other subscribers of that ISP.) However, when an ISP aggregates many subscribers behind the same public IPv4 address (such as used by all systems described in this paper), all of those subscribers will be appear as one identity to the rest of the Internet. This will cause problems with existing systems that equate an IPv4 address with an identity, and take action based on such identities.

9.2. IPsec Compatibility

TOC

It is well known that IPSec AH (Kent, S., "IP Authentication Header," December 2005.) [RFC4302] does not work with NAT [RFC3715] (Aboba, B. and W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements," March 2004.). However, IPsec ESP (Kent, S., "IP Encapsulating Security Payload (ESP)," December 2005.) [RFC4303] can work with NATs because it does not include source or destination addresses in its keyed message integrity check. It is possible to carry IPsec ESP over UDP [RFC3948] (Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets," January 2005.), which survives well over NATs at the expense of a UDP header (8 bytes).
To avoid the UDP overhead and to allow for IPsec ESP endpoints that do not support IPsec over UDP, many deployed IPv4 NAT devices provide an "IPsec Passthru" feature, which uses the destination IP address and the IPsec ESP Security Parameters Index (SPI) field to perform its NAT function. However, "IPsec passthru" has some drawbacks (not described here).

---

10. Acknowledgements

TOC

Thanks to the authors of the contributions compared in this document, Cullen Jennings (NAT6); Marcelo Bagnulo, Philip Matthews, Iljitsch van Beijnum (NAT64); Xing Li, Maoke Chen, Congxiao Bao, Hong Zhang, Jianping Wu, Fred Baker (IVI); Alain Durand, Ralph Droms, Brian Haberman (DS-Lite); Tomohiro Nishitani, Shin Miyakawa (CGN); Remi Despres (SAM); Hiroshi Miyata, Masahito Endo (sNAT-PT); Olaf Maennel, Randy Bush, Luca Cittadini, Steven M. Bellovin (A+P).
Thanks to Fred Baker, Randy Bush, Wojciech Dec, Thomas Narten, Dave Thaler and Eric Vyncke for their review and suggested improvements to the document.

---

11. IANA Considerations

TOC

This document has no IANA actions.

---

12. References

TOC

12.1. Normative References

TOC

[A+P]	Maennel, O., Bush, R., Cittadini, L., and S. Bellovin, “A Better Approach than Carrier-Grade-NAT.”
[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 (TXT).
[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 (TXT).
[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-00 (work in progress), September 2008 (TXT).
[I-D.endo-v6ops-dnsproxy]	Endo, M. and H. Miyata, “Translator Friendly DNS Proxy,” draft-endo-v6ops-dnsproxy-01 (work in progress), October 2008 (TXT).

[I-D.ietf-v6ops-nat64-pb-statement-req]	Bagnulo, M., Baker, F., and I. Beijnum, "IPv4/IPv6 Coexistence and Transition: Requirements for solutions," draft-ietf-v6ops-nat64-pb-statement-req-00 (work in progress), May 2008 (TXT).
[I-D.jennings-behave-nat6]	Jennings, C., "NAT for IPv6-Only Hosts," draft-jennings-behave-nat6-01 (work in progress), November 2008 (TXT).
[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 (TXT).
[I-D.nishitani-cgn]	Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for IP address sharing schemes," draft-nishitani-cgn-04 (work in progress), March 2010 (TXT).
[I-D.xli-behave-ivi]	Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The CERNET IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition," draft-xli-behave-ivi-07 (work in progress), January 2010 (TXT).
[RFC4966]	Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status," RFC 4966, July 2007 (TXT).

12.2. Informative References

TOC

[Behave]	IETF, "BEHAVE working group mailing list."
[I-D.cheshire-nat-pmp]	Cheshire, S., "NAT Port Mapping Protocol (NAT-PMP)," draft-cheshire-nat-pmp-03 (work in progress), April 2008 (TXT).
[I-D.despres-sam]	Despres, R., "Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM)," draft-despres-sam-03 (work in progress), July 2009 (TXT).
[I-D.ietf-mmusic-ice]	Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," draft-ietf-mmusic-ice-19 (work in progress), October 2007 (TXT).
[RFC2671]	Vixie, P., "Extension Mechanisms for DNS (EDNS0)," RFC 2671, August 1999 (TXT).
[RFC2765]	Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)," RFC 2765, February 2000 (TXT).
[RFC2766]	

	Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)," RFC 2766, February 2000 (TXT).
[RFC2827]	Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," BCP 38, RFC 2827, May 2000 (TXT).
[RFC3715]	Aboba, B. and W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements," RFC 3715, March 2004 (TXT).
[RFC3948]	Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets," RFC 3948, January 2005 (TXT).
[RFC4213]	Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers," RFC 4213, October 2005 (TXT).
[RFC4302]	Kent, S., "IP Authentication Header," RFC 4302, December 2005 (TXT).
[RFC4303]	Kent, S., "IP Encapsulating Security Payload (ESP)," RFC 4303, December 2005 (TXT).
[RFC4380]	Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)," RFC 4380, February 2006 (TXT).
[RFC4605]	Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")," RFC 4605, August 2006 (TXT).
[RFC4607]	Holbrook, H. and B. Cain, "Source-Specific Multicast for IP," RFC 4607, August 2006 (TXT).
[RFC4925]	Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire Problem Statement," RFC 4925, July 2007 (TXT).
[RFC5135]	Wing, D. and T. Eckert, "IP Multicast Requirements for a Network Address Translator (NAT) and a Network Address Port Translator (NAPT)," BCP 135, RFC 5135, February 2008 (TXT).
[Softwires]	IETF, "Softwires working group mailing list."
[UPnP-IGD]	UPnP Forum, "Universal Plug and Play Internet Gateway Device," 2000.
[v4v6interm-interest]	IETF, "v4v6interm-interest mailing list."

Appendix A. Changes

TOC

A.1. Changes from 01 to 02

TOC

*Updated DS-Lite reference; no changes to text

*Updated from APB-Revised to SAM; changed text

*Updated sNAT-PT reference; added description of IPv4-to-IPv6 1:N port mapping

*Mentioned policing difficulties for shared addresses (DHCP snooping, ARP security, ingress filtering)

*Discuss IPsec compatibility

*Added explanation of how NAT444 can support IPv6 using Teredo

A.2. Changes from 00 to 01

TOC

*Added A+P

*Refined security considerations for sharing addresses

*"CPE" -> "CPE router"

*removed NAPT defintion (we use NAT to mean NAPT, as is done colloquially)

*fixed some text and figures for APB-Revised

*removed the DNS rewriting function from NAT64's figure showing IPv6 hosts accessing IPv4 servers.

Authors' Addresses

TOC

	Dan Wing
	Cisco Systems, Inc.
	170 West Tasman Drive
	San Jose, CA 95134
	USA
Email:	dwing@cisco.com
	David Ward

	Cisco Systems, Inc.
Email:	wardd@cisco.com
	Alain Durand
	Comcast
	1500 Market st
	Philadelphia, PA 19102
	USA
Email:	alain_durand@cable.comcast.com

Full Copyright Statement

TOC

Intellectual Property

Acknowledgment

This document was produced using xml2rfc v1.35 (of http://xml.resource.org/) from a source in RFC-2629 XML format.