

Internet Engineering Task Force (IETF)
Request for Comments: 7157
Category: Informational
ISSN: 2070-1721

O. Troan, Ed.
Cisco
D. Miles
Google Fiber
S. Matsushima
Softbank Telecom
T. Okimoto
NTT West
D. Wing
Cisco
March 2014

IPv6 Multihoming without Network Address Translation

Abstract

Network Address and Port Translation (NAPT) works well for conserving global addresses and addressing multihoming requirements because an IPv4 NAPT router implements three functions: source address selection, next-hop resolution, and (optionally) DNS resolution. For IPv6 hosts, one approach could be the use of IPv6-to-IPv6 Network Prefix Translation (NPTv6). However, NAT and NPTv6 should be avoided, if at all possible, to permit transparent end-to-end connectivity. In this document, we analyze the use cases of multihoming. We also describe functional requirements and possible solutions for multihoming without the use of NAT in IPv6 for hosts and small IPv6 networks that would otherwise be unable to meet minimum IPv6-allocation criteria. We conclude that DHCPv6-based solutions are suitable to solve the multihoming issues described in this document, but NPTv6 may be required as an intermediate solution.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7157>.

Copyright Notice

Copyright (c) 2014 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|--|----|
| 1. Introduction | 3 |
| 2. Terminology | 5 |
| 3. IPv6 Multihomed Network Scenarios | 6 |
| 3.1. Classification of Network Scenarios for Multihomed Host | 6 |
| 3.2. Multihomed Network Environment | 8 |
| 3.3. Problem Statement | 9 |
| 4. Requirements | 11 |
| 4.1. End-to-End Transparency | 11 |
| 4.2. Scalability | 11 |
| 5. Problem Analysis | 11 |
| 5.1. Source Address Selection | 11 |
| 5.2. Next Hop Selection | 12 |
| 5.3. DNS Recursive Name Server Selection | 13 |
| 6. Implementation Approach | 13 |
| 6.1. Source Address Selection | 14 |
| 6.2. Next Hop Selection | 14 |
| 6.3. DNS Recursive Name Server Selection | 15 |
| 6.4. Other Algorithms Available in RFCs | 16 |
| 7. Considerations for MHMP Deployment | 16 |
| 7.1. Non-MHMP Host Consideration | 16 |
| 7.2. Coexistence Considerations | 17 |
| 7.3. Policy Collision Consideration | 17 |
| 8. Security Considerations | 18 |
| 9. Contributors | 19 |
| 10. References | 20 |
| 10.1. Normative References | 20 |
| 10.2. Informative References | 20 |

1. Introduction

In this document, we analyze the use cases of multihoming, describe functional requirements, and describe the problems with IPv6 multihoming. There are two ways to avoid the problems of IPv6 multihoming:

1. using IPv6-to-IPv6 network prefix translation (NPTv6) [RFC6296], or;
2. refining IPv6 specifications to resolve the problems with IPv6 multihoming.

This document concerns itself with the latter and explores the solution space. We hope this will encourage the development of solutions to the problem so that, in the long run, NPTv6 can be avoided.

IPv6 provides enough globally unique addresses to permit every conceivable host on the Internet to be uniquely addressed without the requirement for Network Address Port Translation (NAPT) [RFC3022], offering a renaissance in end-to-end transparent connectivity.

Unfortunately, this may not be possible in every case, due to the possible necessity of NAT even in IPv6, because of multihoming. Though there are mechanisms to implement multihoming, such as BGP multihoming [RFC4116] at the network level and multihoming based on the Stream Control Transmission Protocol (SCTP) [RFC4960] in the transport layer, there is no mechanism in IPv6 that serves as a replacement for NAT-based multihoming in IPv4. In IPv4, for a host or a small network, NAT-based multihoming is easily deployable and is an already-deployed technique.

Whenever a host or small network (that does not meet minimum IPv6 allocation criteria) is connected to multiple upstream networks, an IPv6 address is assigned by each respective service provider resulting in hosts with multiple global scope IPv6 addresses with different prefixes. As each service provider is allocated a different address space from its Internet Registry, it, in turn, assigns a different address space to the end-user network or host. For example, a remote access user's host or router may use a VPN to simultaneously connect to a remote network and retain a default route to the Internet for other purposes.

In IPv4, a common solution to the multihoming problem is to employ NAPT on a border router and use private address space for individual host addressing. The use of NAPT allows hosts to have exactly one IP address visible on the public network, and the combination of NAPT with provider-specific outside addresses (one for each uplink) and destination-based routing insulates a host from the impacts of multiple upstream networks. The border router may also implement a DNS cache or DNS policy to resolve address queries from hosts.

It is our goal to avoid the IPv6 equivalent of NAT. So, the goals for IPv6 multihoming defined in [RFC3582] do not match the goals of this document. Also, regardless of what the NPTv6 specification is, we are trying to avoid any form of network address translation technique that may not be visible to either of the end hosts. To reach this goal, several mechanisms are needed for end-user hosts to have multiple address assignments and resolve issues such as which address to use for sourcing traffic to which destination:

- o If multiple routers exist on a single link, the host must select the appropriate next hop for each connected network. Each router is in turn connected to a different service provider network, which provides independent address assignment. Routing protocols that would normally be employed for router-to-router network advertisement seem inappropriate for use by individual hosts.
- o Source address selection becomes difficult whenever a host has more than one address of the same address scope. Current address selection criteria may result in hosts using an arbitrary or random address when sourcing upstream traffic. Unfortunately, for the host, the appropriate source address is a function of the upstream network for which the packet is bound. If an upstream service provider uses IP anti-spoofing or ingress filtering, it is conceivable that the packets that have an inappropriate source address for the upstream network would never reach their destination.
- o In a multihomed environment, different DNS scopes or partitions may exist in each independent upstream network. A DNS query sent to an arbitrary upstream DNS recursive name server may result in incorrect or poisoned responses.

In short, while IPv6 facilitates hosts having more than one address in the same address scope, the application of this causes significant issues for a host from routing, source address selection, and DNS resolution perspectives. A possible consequence of assigning a host multiple identically scoped addresses is severely impaired IP connectivity.

If a host connects to a network behind an IPv4 NAPT, the host has one private address in the local network. There is no confusion. The NAT becomes the gateway of the host and forwards the packet to an appropriate network when it is multihomed. It also operates a DNS cache server or DNS proxy, which receives all DNS inquiries, and gives a correct answer to the host.

2. Terminology

| | |
|-------|---|
| NPTv6 | IPv6-to-IPv6 Network Prefix Translation as described in [RFC6296]. |
| NAPT | Network Address Port Translation as described in [RFC3022]. In other contexts, NAPT is often pronounced "NAT" or written as "NAT". |
| MHMP | Multihomed with multi-prefix. A host implementation that supports the mechanisms described in this document; namely, source address selection policy, next hop selection, and DNS selection policy. |

3. IPv6 Multihomed Network Scenarios

In this section, we classify three scenarios of the multihoming environment.

3.1. Classification of Network Scenarios for Multihomed Host

Scenario 1:

In this scenario, two or more routers are present on a single link shared with the host(s). Each router is, in turn, connected to a different service provider network, which provides independent address assignment and DNS recursive name servers. A host in this environment would be offered multiple prefixes and DNS recursive name servers advertised from the two different routers.




Figure 1: Single Uplink, Multiple Next Hop, Multiple Prefix (Scenario 1)

Figure 1 illustrates the host connecting to rtr1 and rtr2 via a shared link. Networks 1 and 2 are reachable via rtr1 and rtr2, respectively. When the host sends packets to network 1, the next hop to network 1 is rtr1. Similarly, rtr2 is the next hop to network 2.

Example: multiple broadband service providers (Internet, VoIP, IPTV, etc.)

Scenario 2:

In this scenario, a single gateway router connects the host to two or more upstream service provider networks. This gateway router would receive prefix delegations and a different set of DNS recursive name servers from each independent service provider network. The gateway, in turn, advertises the provider prefixes to the host, and for DNS, may either act as a lightweight DNS cache server or advertise the complete set of service provider DNS recursive name servers to the hosts.




Figure 2: Single Uplink, Single Next Hop, Multiple Prefix (Scenario 2)

Figure 2 illustrates the host connected to GW rtr. GW rtr connects to networks 1 and 2 via port1 and 2, respectively. As the figure shows a logical topology of the scenario, port1 could be a pseudo-interface for tunneling, which connects to network 1 through network 2 and vice versa. When the host sends packets to either network 1 or 2, the next hop is GW rtr. When the packets are sent to network 1 (network 2), GW rtr forwards the packets to port1 (port2).

Example: Internet + VPN / Application Service Provider (ASP)

Scenario 3:

In this scenario, a host has more than one active interface that connects to different routers and service provider networks. Each router provides the host with a different address prefix and set of DNS recursive name servers, resulting in a host with a unique address per link/interface.




Figure 3: Multiple Uplink, Multiple Next Hop, Multiple Prefix (Scenario 3)

Figure 3 illustrates the host connecting to rtr1 and rtr2 via a direct connection or a virtual link. When the host sends packets to network 1, the next hop to network 1 is rtr1. Similarly, rtr2 is the next hop to network 2.

Example: Mobile Wifi + 3G, ISP A + ISP B

3.2. Multihomed Network Environment

In an IPv6 multihomed network, a host is assigned two or more IPv6 addresses and DNS recursive name servers from independent service provider networks. When this multihomed host attempts to connect with other hosts, it may incorrectly resolve the next-hop router, use an inappropriate source address, or use a DNS response from an incorrect service provider that may result in impaired IP connectivity.

In many cases, multihomed networks in IPv4 have been implemented through the use of a gateway router with NAPT function (scenario 2 with NAPT). An analysis of the current IPv4 NAPT and DNS functions within the gateway router should provide a baseline set of

requirements for IPv6 multihomed environments. A destination prefix/route is often used on the gateway router to separate traffic between the networks.

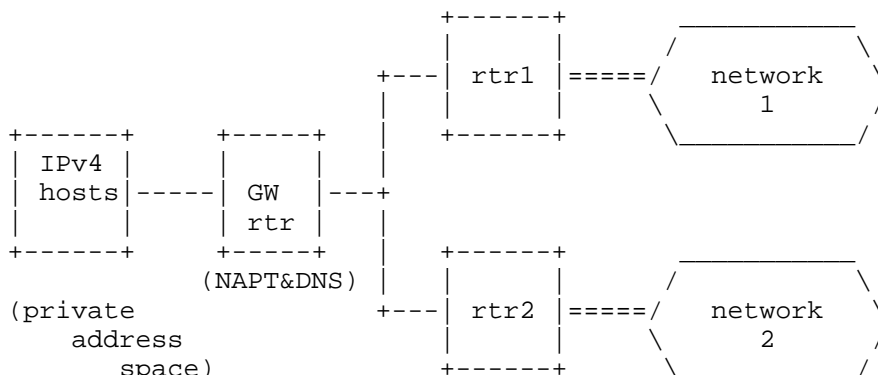


Figure 4: IPv4 Multihomed Environment with Gateway Router Performing NAT

3.3. Problem Statement

A multihomed IPv6 host has one or more assigned IPv6 addresses and DNS recursive name servers from each upstream service provider, resulting in the host having multiple valid IPv6 addresses and DNS recursive name servers. The host must be able to resolve the appropriate next hop, the correct source address, and the correct DNS recursive name server to use based on the destination prefix. To prevent IP spoofing, operators will often implement ingress filtering to discard traffic with an inappropriate source address, making it essential for the host to correctly resolve these three items before sourcing the first packet.

IPv6 has mechanisms for the provision of multiple routers on a single link and multiple address assignments to a single host. However, when these mechanisms are applied to the three scenarios described in Section 3.1, a number of connectivity issues are identified:

Scenario 1:

The host has been assigned an address from each router and recognizes both rtr1 and rtr2 as valid default routers (in the default routers list).

- o The source address selection policy on the host does not deterministically resolve a source address. Ingress filtering or filter policies will discard traffic with source addresses that the operator did not assign.
- o The host will select one of the two routers as the active default router. No traffic is sent to the other router.

Scenario 2:

The host has been assigned two different addresses from the single gateway router. The gateway router is the only default router on the link.

- o The source address selection policy on the host does not deterministically resolve a source address. Ingress filtering or filter policies will discard traffic with source addresses that the operator did not assign.
- o The gateway router does not have an autonomous mechanism for determining which traffic should be sent to which network. If the gateway router is implementing host functions (i.e., processing Router Advertisement (RA)), then two valid default routers may be recognized.

Scenario 3:

A host has two separate interfaces, and each interface has a different address assigned. Each link has its own router.

- o The host does not have enough information to determine which traffic should be sent to which upstream routers. The host will select one of the two routers as the active default router, and no traffic is sent to the other router. The default address selection rules select the address assigned to the outgoing interface as the source address. So, if a host has an appropriate routing table, an appropriate source address will be selected.

All scenarios:

- o In network deployments utilizing local namespaces, the host may choose to communicate with a "wrong" DNS recursive server unable to serve a local namespace.

4. Requirements

This section describes requirements that any solution multi-address and multi-uplink architectures need to meet.

4.1. End-to-End Transparency

One of the major design goals for IPv6 is to restore the end-to-end transparency of the Internet. If NAT is applied to IP communication between hosts, NAT traversal mechanisms are required to establish bidirectional IP communication. In an environment with end-to-end transparency, a NAT traversal mechanism does not need to be implemented in an application (e.g., ICE [RFC5245]). Therefore, the IPv6 multihoming solution should strive to avoid NPTv6 to achieve end-to-end transparency.

4.2. Scalability

The solution will have to be able to manage a large number of sites/nodes. In services for residential users, provider edge devices have to manage thousands of sites. In such environments, sending packets periodically to each site may affect edge system performance.

5. Problem Analysis

The problems described in Section 3 can be classified into these three types:

- o Wrong source address selection
- o Wrong next hop selection
- o Wrong DNS server selection

This section reviews the problem statements presented above and the proposed functional requirements to resolve the issues.

5.1. Source Address Selection

A multihomed IPv6 host will typically have different addresses assigned from each service provider on either the same link (scenarios 1 and 2) or different links (scenario 3). When the host wishes to send a packet to any given destination, the current source address selection rules [RFC6724] may not deterministically select the correct source address. [RFC7078] describes the use of the policy table (as discussed in [RFC6724]) to resolve this problem, using a DHCPv6 mechanism for host policy table management.

Again, by employing DHCPv6, the server could restrict address assignment (of additional prefixes) only to hosts that support policy table management.

Scenario 1: Host needs to support the solution for this problem.

Scenario 2: Host needs to support the solution for this problem.

Scenario 3: If Host supports the next hop selection solution, there is no need to support the address selection functionality on the host.

It is noted that the network's DHCP server and DHCP-forwarding routers must also support the Address Selection option [RFC7078].

5.2. Next Hop Selection

A multihomed IPv6 host or gateway may have multiple uplinks to different service providers. Here, each router would use Router Advertisements [RFC4861] to distribute default route/next-hop information to the host or gateway router.

In this case, the host or gateway router may select any valid default router from the default routers list, resulting in traffic being sent to the wrong router and discarded by the upstream service provider. Using the above scenarios as an example, whenever the host wishes to reach a destination in network 2 and there is no connectivity between networks 1 and 2 (as is the case for a walled-garden or closed service), the host or gateway router does not know whether to forward traffic to rtr1 or rtr2 to reach a destination in network 2. The host or gateway router may choose rtr1 as the default router, but traffic will fail to reach the destination server. The host or gateway router requires route information for each upstream service provider, but the use of a routing protocol between the gateway and the two routers causes both configuration and scaling issues. In IPv4, gateway routers are often pre-configured with static routes or use the Classless Static Route Options [RFC3442] for DHCPv4. An extension to Router Advertisements through Default Router Preference and More-Specific Routes [RFC4191] provides for link-specific preferences but does not address per-host configuration in a multi-access topology because of its reliance on Router Advertisements.

Scenario 1: Host needs to support the solution for this problem.

Scenario 2: GW rtr needs to support the solution for this problem.

Scenario 3: Host needs to support the solution for this problem.

5.3. DNS Recursive Name Server Selection

A multihomed IPv6 host or gateway router may be provided multiple DNS recursive name servers through DHCPv6 [RFC3646] or RA [RFC6106]. When the host or gateway router sends a DNS query, it would normally choose one of the available DNS recursive name servers for the query.

In the IPv6 gateway router scenario, the Broadband Forum (BBF) [TR-124] requires that the query be sent to all DNS recursive name servers and that the gateway wait for the first reply. In IPv6, given our use of specific destination-based policy for both routing and source address selection, it is desirable to extend a policy-based concept to DNS recursive name server selection. Doing so can minimize DNS recursive name server load and avoid issues where DNS recursive name servers in different networks have connectivity issues, or the DNS recursive name servers are not publicly accessible. In the worst case, a DNS query for a name from a local namespace may not be resolved correctly if sent towards a DNS server not aware of said local namespace, resulting in a lack of connectivity.

It is not an issue of the Domain Name System model itself, but an IPv6 multihomed host or gateway router should have the ability to select appropriate DNS recursive name servers for each service based on the domain space for the destination, and each service should provide rules specific to that network. [RFC6731] proposes a solution for distributing DNS server selection policy using a DHCPv6 option.

Scenario 1: Host needs to support the solution for this problem.

Scenario 2: GW rtr needs to support the solution for this problem.

Scenario 3: Host needs to support the solution for this problem.

It is noted that the network's DHCP server and DHCP-forwarding routers must also support the Address Selection option [RFC6731].

6. Implementation Approach

As mentioned in Section 5, in the multi-prefix environment, we have three problems: source address selection, next hop selection, and DNS recursive name server selection. In this section, possible solutions for each problem are introduced and evaluated against the requirements in Section 4.

6.1. Source Address Selection

The problems of address selection in multi-prefix environments are summarized in [RFC5220]. When solutions are examined against the requirements in Section 4, the proactive approaches, such as the policy table distribution mechanism and the routing hints mechanism, are more appropriate in that they can propagate the network administrator's policy directly. The policy distribution mechanism has an advantage with regard to the host's protocol stack impact and the static nature of the assumed target network environment.

6.2. Next Hop Selection

As for the source address selection problem, both a policy-based approach and a non-policy-based approach are possible with regard to the next hop selection problem. Because of the same requirements, the policy propagation-based solution mechanism, whatever the policy, should be more appropriate.

Routing information is a typical example of policy related to next hop selection. If we assume source-address-based routing at hosts or intermediate routers, pairs of source prefixes and next hops can be another example of next hop selection policy.

The routing-information-based approach has a clear advantage in implementation and is already commonly used.

The existing proposed or standardized routing information distribution mechanisms are routing protocols (such as Routing Information Protocol Next Generation (RIPng) and OSPFv3), the RA extension option defined in [RFC4191], and the CPE WAN Management Protocol (CWMP) [TR069] standardized at BBF.

The RA-based mechanism doesn't handle distribution of per-host routing information easily. Dynamic routing protocols are not typically used between residential users and ISPs, because of their scalability and security implications. The DHCPv6 mechanism does not have these problems and has the advantage of relay functionality. It is commonly used and is thus easy to deploy.

[TR069], mentioned above, defines a possible solution mechanism for routing information distribution to customer premises equipment (CPE). It assumes, however, that IP reachability to the Auto Configuration Server (ACS) has been established. Therefore, if the CPE requires routing information to reach the ACS, CWMP [TR069] cannot be used to distribute this information.

6.3. DNS Recursive Name Server Selection

Note: Split-horizon DNS is discussed in this section. Split-horizon DNS is known to cause problems with applications to allow information leakage. The discussion of split-horizon DNS is not condoning its use, but rather acknowledging that split-horizon DNS is used and that its use is another justification for network address translation. The goal of this document is to encourage building solutions that do not need network address translation. Two solutions appear possible: improve the function of split-horizon DNS (which is discussed below) or meet network administrators' requirements without split-horizon DNS (which is out of scope of this document).

As in the above two problems, a policy-based approach and a non-policy-based approach are possible. In a non-policy-based approach, a host or a home gateway router is assumed to send DNS queries to several DNS recursive name servers at once or to select one of the available servers.

In the non-policy-based approach, by making a query to a DNS recursive name server in a different service provider to that which hosts the service, a user could be directed to an unexpected IP address or receive an invalid response, and thus it could not connect to the service provider's private and legitimate service. For example, some DNS recursive name servers reply with different answers depending on the source address of the DNS query, which is sometimes called "split-horizon". When the host mistakenly makes a query to a different provider's DNS recursive name server to resolve a Fully Qualified Domain Name (FQDN) of another provider's private service, and the DNS recursive name server adopts the split-horizon configuration, the queried server returns an IP address of the non-private side of the service. Another problem with this approach is that it causes unnecessary DNS traffic to the DNS recursive name servers that are visible to the users.

The alternative to a policy-based approach is documented in [RFC6731], where several pairs of DNS recursive name server addresses and DNS domain suffixes are defined as part of a policy and conveyed to hosts in a new DHCP option. In an environment where there is a home gateway router, that router can act as a DNS recursive name server, interpret this option, and distribute DNS queries to the appropriate DNS servers according to the policy.

6.4. Other Algorithms Available in RFCs

The authors of this document are aware of a variety of other algorithms and architectures, such as Shim6 [RFC5533] and HIP [RFC5206], that may be useful in this environment. At the time of this writing, there is not enough operational experience on which to base a recommendation. Should such operational experience become available, this document may be updated in the future.

7. Considerations for MHMP Deployment

This section describes considerations to mitigate possible problems in a network that implements MHMP (described in Section 6).

7.1. Non-MHMP Host Consideration

In a typical IPv4 multihomed network deployment, IPv4 NAPT is practically used and it can eventually avoid assigning multiple addresses to the hosts and solve the next hop selection problem. In a similar fashion, NPTv6 can be used as a last resort for IPv6 multihomed network deployments where one needs to assign a single IPv6 address to a non-MHMP host.

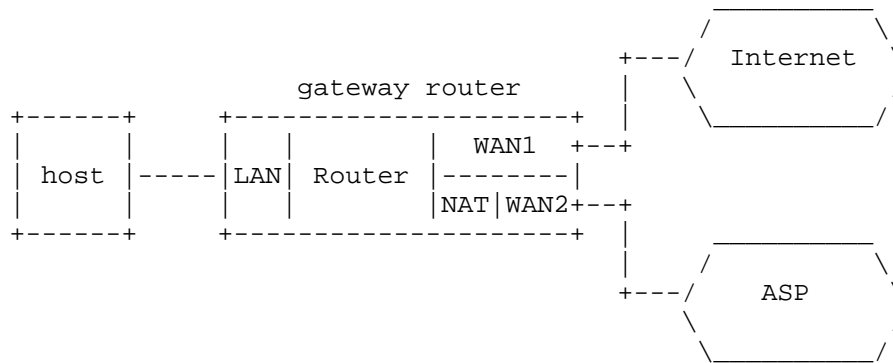


Figure 5: Legacy Host

The gateway router also has to support the two features, next hop selection and DNS server selection, shown in Section 6.

The implementation and issues of NPTv6 are out of the scope of this document, but are discussed in Section 5 of [RFC6296].

7.2. Coexistence Considerations

To allow the coexistence of non-MHMP hosts and MHMP hosts (i.e., hosts supporting multi-prefix with the enhancements for the source address selection), GW rtr may need to treat those hosts separately.

An idea for how to achieve this would be for GW rtr to identify the hosts, and then assign a single prefix to non-MHMP hosts and assign multiple prefixes to MHMP hosts. In this case, GW rtr can perform IPv6 NAT only for the traffic from non-MHMP hosts if its source address is not appropriate.

Another idea is that GW rtr could assign multiple prefixes to both hosts and perform IPv6 NAT for traffic from non-MHMP hosts if its source address is not appropriate.

In scenarios 1 and 3, the non-MHMP hosts can be placed behind the NAT box. In this case, the non-MHMP host can access the service through the NAT box.

The implementation of identifying non-MHMP hosts and NAT policy is outside the scope of this document.

7.3. Policy Collision Consideration

When multiple policy distributors exist, a policy receiver may not follow each of the received policies. In particular, when a policy conflicts with another policy, a policy receiver cannot implement both of the policies. To solve or mitigate this issue, a prioritization rule is required to align the policies with the preferences of a trusted interface. Another solution is to preclude the functionality of the acceptance of multiple policies at the receiver side. In this case, a policy distributor should cooperate with other policy distributors, and a single representative provider should distribute a merged policy.

This document does not presume specific recommendations for resolving policy collision. It is expected that the implementation will decide how to resolve the conflicts. If they are not resolved consistently by different implementations, that could affect interoperability and security trust boundaries. Future work is expected to address the need for consistent policy resolution to avoid interoperability and security trust boundary issues.

8. Security Considerations

In today's multihomed IPv4 networks, it is difficult to resolve or coordinate conflicts between the two upstream networks. This problem persists with IPv6, no matter if the hosts use IPv6 provider-dependent or provider-independent addresses.

This document requires that MHMP solutions have functions that provide policy controls. New security threats can be introduced depending on the kind and form of the policy. The threats can be categorized in two parts: the policy receiver side and the policy distributor side.

A policy receiver may receive an evil policy from a policy distributor. A policy distributor should expect that some hosts in its network will not follow the distributed policy. At the time of this writing, there are no known methods to resolve conflicts between the host's own policy (policy receiver) and the policies of upstream providers (policy provider). As this document is analyzing the problem space, rather than proposing a solution, we note the following problems:

Threats related to the policy distributor side:

The service provider should expect the existence of hosts that will not obey the received policy. A possible solution is to ingress-filter those packets that do not match the distributed policy and drop them. For route selection, packet forwarding or redirection can be another possible solution. For source address selection, IPv6 NAT can be another possible solution.

In a multihomed multiple-provider network, nodes in the network may be administered by different organizations. Administrators might need to control policies (and a node's behavior) independently of other administrators. Access control policies need to be in place to restrict the administrator's access to only the nodes it is authorized to control.

Threats related to the policy receiver side:

For the policy receiver side, who should be trusted to accept policies is a fundamental issue. How is the trust established? How can the network element be assured that it can establish that trust before the network is fully configured? If a policy receiver trusts an untrusted network, it will cause the distributing of the unwanted and unauthorized policy that is described below.

A policy receiver is exposed to the threats of unauthorized policy, which can lead to session hijack, falsification, DoS, wiretapping, and phishing. Unauthorized policy here means a policy distributed from an entity that does not have rights to do so. Usually, only a site administrator and a network service provider have rights to distribute these policies in addition to IP address assignment and DNS server address notification. Regarding source address selection, unauthorized policy can expose an IP address that will not usually be exposed to an external server, which can be a privacy problem. To solve or mitigate the problem of unauthorized policy, one approach is to limit the use of these policy distribution mechanisms, as described in the Section 4.4 of [RFC6731]. For example, a policy should be preferred or accepted if delivered over a secure, trusted channel such as a cellular data connection. The proposed solutions are based on DHCP, so the limitation of local site communication, which is often used in WiFi access services, should be another solution or mitigation for this problem. For the DNS server selection issue, DNS Security (DNSSEC) can be another solution. For source address selection, the ingress filter at the network service provider router can be a solution.

Another threat is the leakage of the policy and privacy issues resulting from that. Especially when clients receive different policies from the network service provider, that difference provides hints about the host itself and can be useful to uniquely identify the host. Encryption of the communication channel and separation of the communication channel per host can be solutions for this problem.

The security threats related to IPv6 multihoming are described in [RFC4218].

9. Contributors

The following people contributed to this document: Akiko Hattori, Arifumi Matsumoto, Frank Brockners, Fred Baker, Tomohiro Fujisaki, Jun-ya Kato, Shigeru Akiyama, Seiichi Morikawa, Mark Townsley, Wojciech Dec, Yasuo Kashimura, and Yuji Yamazaki. This document has greatly benefited from inputs by Randy Bush, Brian Carpenter, and Teemu Savolainen.

10. References

10.1. Normative References

- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.
- [RFC6731] Savolainen, T., Kato, J., and T. Lemon, "Improved Recursive DNS Server Selection for Multi-Interfaced Nodes", RFC 6731, December 2012.
- [RFC7078] Matsumoto, A., Fujisaki, T., and T. Chown, "Distributing Address Selection Policy Using DHCPv6", RFC 7078, January 2014.

10.2. Informative References

- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [RFC3442] Lemon, T., Cheshire, S., and B. Volz, "The Classless Static Route Option for Dynamic Host Configuration Protocol (DHCP) version 4", RFC 3442, December 2002.
- [RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, August 2003.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, July 2005.

- [RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions", RFC 4218, October 2005.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC5206] Nikander, P., Henderson, T., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol", RFC 5206, April 2008.
- [RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules", RFC 5220, July 2008.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.
- [TR-124] The Broadband Forum, "TR-124, Functional Requirements for Broadband Residential Gateway Devices", Issue: 2, May 2010, <http://www.broadband-forum.org/technical/download/TR-124_Issue-2.pdf>.
- [TR069] The Broadband Forum, "TR-069, CPE WAN Management Protocol v1.1", Version: Issue 1 Amendment 2, December 2007, <http://www.broadband-forum.org/technical/download/TR-069_Amendment-2.pdf>.

Authors' Addresses

Ole Troan (editor)
Cisco
Oslo
Norway

E-Mail: ot@cisco.com

David Miles
Google Fiber
Mountain View, CA
USA

E-Mail: davidmiles@google.com

Satoru Matsushima
Softbank Telecom
Tokyo
Japan

E-Mail: satoru.matsushima@g.softbank.co.jp

Tadahisa Okimoto
NTT West
Osaka
Japan

E-Mail: t.okimoto@west.ntt.co.jp

Dan Wing
Cisco
170 West Tasman Drive
San Jose
USA

E-Mail: dwing@cisco.com