

Internet Engineering Task Force (IETF)
Request for Comments: 8475
Category: Informational
ISSN: 2070-1721

J. Linkova
Google
M. Stucchi
RIPE NCC
October 2018

Using Conditional Router Advertisements for Enterprise Multihoming

Abstract

This document discusses the most common scenarios of connecting an enterprise network to multiple ISPs using an address space assigned by an ISP and how the approach proposed in "Enterprise Multihoming using Provider-Assigned Addresses without Network Prefix Translation: Requirements and Solution" could be applied in those scenarios. The problem of enterprise multihoming without address translation of any form has not been solved yet as it requires both the network to select the correct egress ISP based on the packet source address and hosts to select the correct source address based on the desired egress ISP for that traffic. The aforementioned document proposes a solution to this problem by introducing a new routing functionality (Source Address Dependent Routing) to solve the uplink selection issue. It also proposes using Router Advertisements to influence the host source address selection. It focuses on solving the general problem and covering various complex use cases, and this document adopts its proposed approach to provide a solution for a limited number of common use cases. In particular, the focus of this document is on scenarios in which an enterprise network has two Internet uplinks used either in primary/backup mode or simultaneously and hosts in that network might not yet properly support multihoming as described in RFC 8028.

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 candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8475>.

Copyright Notice

Copyright (c) 2018 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 (https://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
1.1. Requirements Language 4
2. Common Enterprise Multihoming Scenarios 4
2.1. Two ISP Uplinks, Primary and Backup 4
2.2. Two ISP Uplinks, Used for Load-Balancing 5
3. Conditional Router Advertisements 5
3.1. Solution Overview 5
3.1.1. Uplink Selection 5
3.1.2. Source Address Selection and Conditional RAs 5
3.2. Example Scenarios 8
3.2.1. Single Router, Primary/Backup Uplinks 8
3.2.2. Two Routers, Primary/Backup Uplinks 9
3.2.3. Single Router, Load-Balancing between Uplinks 12
3.2.4. Two Routers, Load-Balancing between Uplinks 12
3.2.5. Topologies with Dedicated Border Routers 13
3.2.6. Intrasite Communication during Simultaneous Uplinks Outage 15
3.2.7. Uplink Damping 15
3.2.8. Routing Packets When the Corresponding Uplink Is Unavailable 16
3.3. Solution Limitations 16
3.3.1. Connections Preservation 17
4. IANA Considerations 17
5. Security Considerations 18
5.1. Privacy Considerations 18
6. References 18
6.1. Normative References 18
6.2. Informative References 20
Acknowledgements 20
Authors' Addresses 21

1. Introduction

Multihoming is an obvious requirement for many enterprise networks to ensure the desired level of network reliability. However, using more than one ISP (and address space assigned by those ISPs) introduces the problem of assigning IP addresses to hosts. In IPv4, there is no choice but using address space [RFC1918] and NAT [RFC3022] at the network edge [RFC4116]. Using Provider Independent (PI) address space is not always an option, since it requires running BGP between the enterprise network and the ISPs. The administrative overhead of obtaining and managing PI address space can also be a concern. As IPv6 hosts can, by design, have multiple addresses of the global scope [RFC4291], multihoming using provider addresses looks even easier for IPv6: each ISP assigns an IPv6 block (usually /48), and hosts in the enterprise network have addresses assigned from each ISP block. However, using IPv6 provider-assigned (PA) blocks in a multihoming scenario introduces some challenges, including, but not limited to:

- o Selecting the correct uplink based on the packet source address;
- o Signaling to hosts that some source addresses should or should not be used (e.g., an uplink to the ISP went down or became available again).

[PROVIDER-ASSIGNED] discusses these and other related challenges in detail in relation to the general multihoming scenario for enterprise networks. It proposes a solution that relies heavily on Rule 5.5 of the default address selection algorithm [RFC6724]. Rule 5.5 makes hosts prefer source addresses in a prefix advertised by the next hop and, therefore, is very useful in multihomed scenarios when different routers may advertise different prefixes. While [RFC6724] defines Rule 5.5 as optional, the recent [RFC8028] recommends that multihomed hosts SHOULD support it. Unfortunately, that rule has not been widely implemented at the time of writing. Therefore, network administrators in enterprise networks can't yet assume that all devices in their network support Rule 5.5, especially in the quite common BYOD ("Bring Your Own Device") scenario. However, while it does not seem feasible to solve all the possible multihoming scenarios without relying on Rule 5.5, it is possible to provide IPv6 multihoming using PA address space for the most common use cases. This document discusses how the general approach described in [PROVIDER-ASSIGNED] can be applied to solve multihoming scenarios when:

- o An enterprise network has two or more ISP uplinks;

- o Those uplinks are used for Internet access in active/backup or load-sharing mode without any sophisticated traffic engineering requirements;
- o Each ISP assigns the network a subnet from its own PA address space; and
- o Hosts in the enterprise network are not expected to support Rule 5.5 of the default address selection algorithm [RFC6724].

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Common Enterprise Multihoming Scenarios

2.1. Two ISP Uplinks, Primary and Backup

This scenario has the following key characteristics:

- o The enterprise network uses uplinks to two (or more) ISPs for Internet access;
- o Each ISP assigns IPv6 PA address space for the network;
- o Uplink(s) to one ISP is a primary (preferred) one. All other uplinks are backup and are not expected to be used while the primary one is operational;
- o If the primary uplink is operational, all Internet traffic should flow via that uplink;
- o When the primary uplink fails, the Internet traffic needs to flow via the backup uplinks;
- o Recovery of the primary uplink needs to trigger the traffic switchover from the backup uplinks back to the primary one;
- o Hosts in the enterprise network are not expected to support Rule 5.5 of the default address selection algorithm [RFC6724].

2.2. Two ISP Uplinks, Used for Load-Balancing

This scenario has the following key characteristics:

- o The enterprise network is using uplinks to two (or more) ISPs for Internet access;
- o Each ISP assigns an IPv6 PA address space;
- o All the uplinks may be used simultaneously, with the traffic flows being randomly (not necessarily equally) distributed between them;
- o Hosts in the enterprise network are not expected to support Rule 5.5 of the default address selection algorithm [RFC6724].

3. Conditional Router Advertisements

3.1. Solution Overview

3.1.1. Uplink Selection

As discussed in [PROVIDER-ASSIGNED], one of the two main problems to be solved in the enterprise multihoming scenario is the problem of the next-hop (uplink) selection based on the packet source address. For example, if the enterprise network has two uplinks, to ISP_A and ISP_B, and hosts have addresses from subnet_A and subnet_B (belonging to ISP_A and ISP_B, respectively), then packets sourced from subnet_A must be sent to the ISP_A uplink while packets sourced from subnet_B must be sent to the ISP_B uplink. Sending packets with source addresses belonging to one ISP address space to another ISP might cause those packets to be filtered out if those ISPs or their uplinks implement antispoofing ingress filtering [RFC2827][RFC3704].

While some work is being done in the Source Address Dependent Routing (SADR) (such as [DESTINATION]), the simplest way to implement the desired functionality currently is to apply a policy that selects a next hop or an egress interface based on the packet source address. Currently, most SMB/Enterprise-grade routers have such functionality available.

3.1.2. Source Address Selection and Conditional RAs

Another problem to be solved in the multihoming scenario is the source address selection on hosts. In the normal situation (all uplinks are up/operational), hosts have multiple global unique addresses and can rely on the default address selection algorithm [RFC6724] to pick up a source address, while the network is responsible for choosing the correct uplink based on the source

address selected by a host, as described in Section 3.1.1. However, some network topology changes (i.e., changing uplink status) might affect the global reachability for packets sourced from particular prefixes; therefore, such changes have to be signaled back to the hosts. For example:

- o An uplink to ISP_A went down. Hosts should not use addresses from an ISP_A prefix;
- o A primary uplink to ISP_A that was not operational has come back up. Hosts should start using the source addresses from an ISP_A prefix.

[PROVIDER-ASSIGNED] provides a detailed explanation of why Stateless Address Autoconfiguration (SLAAC) [RFC4862] and Router Advertisements (RAs) [RFC4861] are the most suitable mechanisms for signaling network topology changes to hosts, thereby influencing the source address selection. Sending an RA to change the preferred lifetime for a given prefix provides the following functionality:

- o Deprecating addresses by sending an RA with preferred_lifetime set to 0 in the corresponding Prefix Information option (PIO) [RFC4861]. This indicates to hosts that addresses from that prefix should not be used;
- o Making a previously unused (deprecated) prefix usable again by sending an RA containing a PIO with nonzero preferred lifetime. This indicates to hosts that addresses from that prefix can be used again.

It should be noted that only the preferred lifetime for the affected prefix needs to be changed. As the goal is to influence the source address selection algorithm on hosts rather than prevent them from forming addresses from a specific prefix, the valid lifetime should not be changed. Actually, changing the valid lifetime would not even be possible for unauthenticated RAs (which is the most common deployment scenario), because Section 5.5.3 of [RFC4862] prevents hosts from setting the valid lifetime for addresses to zero unless RAs are authenticated.

To provide the desired functionality, first-hop routers are required to:

- o Send RAs triggered by defined event policies in response to an uplink status change event; and

- o While sending periodic or solicited RAs, set the value in the given RA field (e.g., PIO preferred lifetime) based on the uplink status.

The exact definition of the "uplink status" depends on the network topology and may include conditions like:

- o Uplink interface status change;
- o Presence of a particular route in the routing table;
- o Presence of a particular route with a particular attribute (next hop, tag, etc.) in the routing table;
- o Protocol adjacency change.

In some scenarios, when two routers are providing first-hop redundancy via Virtual Router Redundancy Protocol (VRRP) [RFC5798], the master-backup status can be considered to be a condition for sending RAs and changing the preferred lifetime value. See Section 3.2.2 for more details.

If hosts are provided with the IPv6 addresses of ISP DNS servers via a Recursive DNS Server (RDNSS) (see "IPv6 Router Advertisement Options for DNS Configuration" [RFC8106]), it might be desirable for the conditional RAs to update the Lifetime field of the RDNSS option as well.

The trigger is not only forcing the router to send an unsolicited RA to propagate the topology changes to all hosts. Obviously, the values of the RA fields (like PIO Preferred Lifetime or DNS Server Lifetime) changed by the particular trigger need to stay the same until another event causes the value to be updated. For example, if an ISP_A uplink failure causes the prefix to be deprecated, all solicited and unsolicited RAs sent by the router need to have the preferred lifetime for that PIO set to 0 until the uplink comes back up.

It should be noted that the proposed solution is quite similar to the existing requirement L-13 for IPv6 Customer Edge Routers [RFC7084] and the documented behavior of homenet devices [RFC7788]. It is using the same mechanism of deprecating a prefix when the corresponding uplink is not operational, applying it to an enterprise-network scenario.

3.2. Example Scenarios

This section illustrates how the conditional RAs solution can be applied to the most common enterprise multihoming scenarios, described in Section 2.

3.2.1. Single Router, Primary/Backup Uplinks




Figure 1: Single Router, Primary/Backup Uplinks

Let’s look at a simple network topology where a single router acts as a border router to terminate two ISP uplinks and as a first-hop router for hosts. Each ISP assigns a /48 to the network, and the ISP_A uplink is a primary one, to be used for all Internet traffic, while the ISP_B uplink is a backup, to be used only when the primary uplink is not operational.

To ensure that packets with source addresses from ISP_A and ISP_B are only routed to ISP_A and ISP_B uplinks, respectively, the network administrator needs to configure a policy on R1:

```

IF (packet_source_address is in 2001:db8:1::/48)
  and
  (packet_destination_address is not in
  (2001:db8:1::/48 or 2001:db8:2::/48))
  THEN
    default next hop is ISP_A_uplink
  
```



```
IF (packet_source_address is in 2001:db8:2::/48)
  and
  (packet_destination_address is not in
  (2001:db8:1::/48 or 2001:db8:2::/48))
  THEN
    default next hop is ISP_B_uplink
```

Under normal circumstances, it is desirable that all traffic be sent via the ISP_A uplink; therefore, hosts (the host H1 in the example topology figure) should be using source addresses from 2001:db8:1:1::/64. When or if the ISP_A uplink fails, hosts should stop using the 2001:db8:1:1::/64 prefix and start using 2001:db8:2:1::/64 until the ISP_A uplink comes back up. To achieve this, the RA configuration on the R1 device for the interface facing H1 needs to have the following policy:

```
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}

prefix 2001:db8:2:1::/64 {
  IF (ISP_A_Uplink is up)
    THEN
      preferred_lifetime = 0
    ELSE
      preferred_lifetime = 604800
}
```

A similar policy needs to be applied to the RDNSS lifetime if ISP_A and ISP_B DNS servers are used.

3.2.2. Two Routers, Primary/Backup Uplinks

Let's look at a more complex scenario where two border routers are terminating two ISP uplinks (one each), acting as redundant first-hop routers for hosts. The topology is shown in Figure 2.




Figure 2: Two Routers, Primary/Backup Uplinks

In this scenario, R1 sends RAs with PIO for 2001:db8:1:1::/64 (ISP_A address space), and R2 sends RAs with PIO for 2001:db8:2:1::/64 (ISP_B address space). Each router needs to have a forwarding policy configured for packets received on its hosts-facing interface:

```

IF (packet_source_address is in 2001:db8:1::/48)
  and
  (packet_destination_address is not in
  (2001:db8:1::/48 or 2001:db8:2::/48))
  THEN
    default next hop is ISP_A_uplink
  
```

```

IF (packet_source_address is in 2001:db8:2::/48)
  and
  (packet_destination_address is not in
  (2001:db8:1::/48 or 2001:db8:2::/48))
  THEN
    default next hop is ISP_B_uplink
  
```

In this case, there is more than one way to ensure that hosts are selecting the correct source address based on the uplink status. If VRRP is used to provide first-hop redundancy, and the master router is the one with the active uplink, then the simplest way is to use the VRRP mastership as a condition for RA. So, if ISP_A is the primary uplink, the routers R1 and R2 need to be configured in the following way:

R1 is the VRRP master by default (when the ISP_A uplink is up). If the ISP_A uplink is down, then R1 becomes a backup (the VRRP interface-status tracking is expected to be used to automatically

modify the VRRP priorities and trigger the mastership switchover). RAs on R1's interface facing H1 needs to have the following policy applied:

```
prefix 2001:db8:1:1::/64 {
  IF (vrrp_master)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

R2 is VRRP backup by default. RA on R2's interface facing H1 needs to have the following policy applied:

```
prefix 2001:db8:2:1::/64 {
  IF(vrrp_master)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

If VRRP is not used or interface status tracking is not used for mastership switchover, then each router needs to be able to detect the uplink failure/recovery on the neighboring router, so that RAs with updated preferred lifetime values are triggered. Depending on the network setup, various triggers can be used, such as a route to the uplink interface subnet or a default route received from the uplink. The obvious drawback of using the routing table to trigger the conditional RAs is that some additional configuration is required. For example, if a route to the prefix assigned to the ISP uplink is used as a trigger, then the conditional RA policy would have the following logic:

R1:

```
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

R2:

```
prefix 2001:db8:2:1::/64 {
  IF (ISP_A_uplink_route is present)
    THEN
      preferred_lifetime = 0
    ELSE
      preferred_lifetime = 604800
}
```

3.2.3. Single Router, Load-Balancing between Uplinks

Let's look at the example topology shown in Figure 1, but with both uplinks used simultaneously. In this case, R1 would send RAs containing PIOs for both prefixes, 2001:db8:1:1::/64 and 2001:db8:2:1::/64, changing the preferred lifetime based on particular uplink availability. If the interface status is used as an uplink availability indicator, then the policy logic would look like the following:

```
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
prefix 2001:db8:2:1::/64 {
  IF (ISP_B_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

R1 needs a forwarding policy to be applied to forward packets to the correct uplink based on the source address, similar to the policy described in Section 3.2.1.

3.2.4. Two Routers, Load-Balancing between Uplinks

In this scenario, the example topology is similar to the one shown in Figure 2, but both uplinks can be used at the same time. This means that both R1 and R2 need to have the corresponding forwarding policy to forward packets based on their source addresses.

Each router would send RAs with PIO for the corresponding prefix, setting preferred_lifetime to a nonzero value when the ISP uplink is up and deprecating the prefix by setting preferred_lifetime to 0 in the case of uplink failure. The uplink recovery would trigger another RA with a nonzero preferred lifetime to make the addresses from the prefix preferred again. The example RA policy on R1 and R2 would look like:

```
R1:
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

```
R2:
prefix 2001:db8:2:1::/64 {
  IF (ISP_B_uplink is up)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

3.2.5. Topologies with Dedicated Border Routers

For simplicity, all topologies above show the ISP uplinks terminated on the first-hop routers. Obviously, the proposed approach can be used in more complex topologies when dedicated devices are used for terminating ISP uplinks. In that case, VRRP mastership or interface status cannot be used as a trigger for conditional RAs. Route presence as described in Section 3.2.2 should be used instead.

Let's look at the example topology shown in Figure 3:

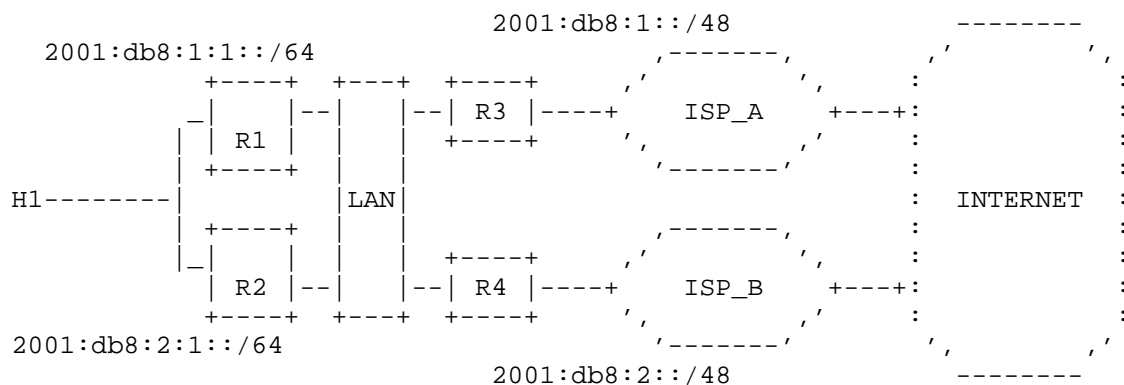


Figure 3: Dedicated Border Routers

For example, if ISP_A is a primary uplink and ISP_B is a backup, then the following policy might be used to achieve the desired behavior (H1 is using ISP_A address space, 2001:db8:1:1::/64, while the ISP_A uplink is up and only using the ISP_B 2001:db8:2:1::/64 prefix if the uplink is non-operational):

R1 and R2 policy:

```
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink_route is present)
  THEN
    preferred_lifetime = 604800
  ELSE
    preferred_lifetime = 0
}

prefix 2001:db8:2:1::/64 {
  IF (ISP_A_uplink_route is present)
  THEN
    preferred_lifetime = 0
  ELSE
    preferred_lifetime = 604800
}
```

For the load-balancing case, the policy would look slightly different: each prefix has a nonzero preferred_lifetime only if the corresponding ISP uplink route is present:

```
prefix 2001:db8:1:1::/64 {
  IF (ISP_A_uplink_route is present)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}

prefix 2001:db8:2:1::/64 {
  IF (ISP_B_uplink_route is present)
    THEN
      preferred_lifetime = 604800
    ELSE
      preferred_lifetime = 0
}
```

3.2.6. Intrasite Communication during Simultaneous Uplinks Outage

Prefix deprecation as a result of an uplink status change might lead to a situation in which all global prefixes are deprecated (all ISP uplinks are not operational for some reason). Even when there is no Internet connectivity, it might be still desirable to have intrasite IPv6 connectivity (especially when the network in question is an IPv6-only one). However, while an address is in a deprecated state, its use is discouraged, but not strictly forbidden [RFC4862]. In such a scenario, all IPv6 source addresses in the candidate set [RFC6724] are deprecated, which means that they still can be used (as there are no preferred addresses available), and the source address selection algorithm can pick up one of them, allowing intrasite communication. However, some operating systems might just fall back to IPv4 if the network interface has no preferred IPv6 global addresses. Therefore, if intrasite connectivity is vital during simultaneous outages of multiple uplinks, administrators might consider using Unique Local Addresses (ULAs) [RFC4193] or provisioning additional backup uplinks to protect the network from double-failure cases.

3.2.7. Uplink Damping

If an actively used uplink (a primary one or one used in a load-balancing scenario) starts flapping, it might lead to the undesirable situation of flapping addresses on hosts: every time the uplink goes up, hosts receive an RA with a nonzero preferred PIO lifetime, and every time the uplink goes down, all addresses in the affected prefix

become deprecated. This would, undoubtedly, negatively impact the user experience, not to mention the impact of spikes of duplicate address detection traffic every time an uplink comes back up. Therefore, it's recommended that router vendors implement some form of damping policy for conditional RAs and either postpone sending an RA with a nonzero lifetime for a PIO when the uplink comes up for a number of seconds or (even) introduce accumulated penalties/exponential backoff algorithm for such delays. (In the case of multiple simultaneous uplink failure, when all but one of the uplinks are down and the last remaining one is flapping, it might result in all addresses being deprecated for a while after the flapping uplink recovers.)

3.2.8. Routing Packets When the Corresponding Uplink Is Unavailable

Deprecating IPv6 addresses by setting the preferred lifetime to 0 discourages but does not strictly forbid its usage in new communications. A deprecated address may still be used for existing connections [RFC4862]. Therefore, when an ISP uplink goes down, the corresponding border router might still receive packets with source addresses belonging to that ISP address space while there is no available uplink to send those packets to.

The expected router behavior would depend on the uplink selection mechanism. For example, if some form of SADR is used, then such packets will be dropped as there is no route to the destination. If policy-based routing is used to set a next hop, then the behavior would be implementation dependent and may vary from dropping the packets to forwarding them based on the routing table entries. It should be noted that there is no return path to the packet source (as the ISP uplink is not operational). Therefore, even if the outgoing packets are sent to another ISP, the return traffic might not be delivered.

3.3. Solution Limitations

It should be noted that the proposed approach is not a "silver bullet" for all possible multihoming scenarios. It would work very well for networks with relatively simple topologies and straightforward routing policies. The more complex the network topology and the corresponding routing policies, the more configuration would be required to implement the solution.

Another limitation is related to the load-balancing between the uplinks. In the scenario in which both uplinks are active, hosts would select the source prefix using the Default Address Selection algorithm [RFC6724]; therefore, the load between two uplinks most likely would not be evenly distributed. (However, the proposed

mechanism does allow a creative way of controlling uplinks load in software-defined networks where controllers might selectively deprecate prefixes on some hosts but not others to move egress traffic between uplinks). Also, the prefix selection does not take into account any other properties of uplinks (such as latency), so egress traffic might not be sent to the nearest uplink if the corresponding prefix is selected as a source. In general, if not all uplinks are equal, and some uplinks are expected to be preferred over others, then the network administrator should ensure that prefixes from non-preferred ISP(s) are kept deprecated (so primary/backup setup is used).

3.3.1. Connections Preservation

The proposed solution is not designed to preserve connection state after an uplink failure. If all uplinks to an ISP go down, all sessions to/from addresses from that ISP address space are interrupted as there is no egress path for those packets and there is no return path from the Internet to the corresponding prefix. In this regard, it is similar to IPv4 multihoming using NAT, where an uplink failure and failover to another uplink means that a public IPv4 address changes and all existing connections are interrupted.

However, an uplink recovery does not necessarily lead to connections interruption. In the load-sharing/balancing scenario, an uplink recovery does not affect any existing connections at all. In the active/backup topology, when the primary uplink recovers from the failure and the backup prefix is deprecated, the existing sessions (established to/from the backup ISP addresses) can be preserved if the routers are configured as described in Section 3.2.1 and send packets with the backup ISP source addresses to the backup uplink, even when the primary one is operational. As a result, the primary uplink recovery makes the usage of the backup ISP addresses discouraged but still possible.

It should be noted that in IPv4 multihoming with NAT, when the egress interface is chosen without taking packet source address into account (as internal hosts usually have addresses from [RFC1918] space), sessions might not be preserved after an uplink recovery unless packet forwarding is integrated with existing NAT sessions tracking.

4. IANA Considerations

This document has no IANA actions.

5. Security Considerations

This memo introduces no new security considerations. It relies on RAs [RFC4861] and the SLAAC [RFC4862] mechanism and inherits their security properties. If an attacker is able to send a rogue RA, they could deprecate IPv6 addresses on hosts or influence source-address-selection processes on hosts.

The potential attack vectors include, but are not limited to:

- o An attacker sends a rogue RA deprecating IPv6 addresses on hosts;
- o An attacker sends a rogue RA making addresses preferred while the corresponding ISP uplink is not operational;
- o An attacker sends a rogue RA making addresses preferred for a backup ISP, steering traffic to an undesirable (e.g., more expensive) uplink.

Therefore, the network administrators SHOULD secure RAs, e.g., by deploying an RA guard [RFC6105].

5.1. Privacy Considerations

This memo introduces no new privacy considerations.

6. References

6.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.

- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <<https://www.rfc-editor.org/info/rfc3022>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, DOI 10.17487/RFC4116, July 2005, <<https://www.rfc-editor.org/info/rfc4116>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <<https://www.rfc-editor.org/info/rfc6105>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <<https://www.rfc-editor.org/info/rfc8028>>.
- [RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <<https://www.rfc-editor.org/info/rfc8106>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

[DESTINATION]

Lamparter, D. and A. Smirnov, "Destination/Source Routing", Work in Progress, draft-ietf-rtgwg-dst-src-routing-06, October 2017.

[PROVIDER-ASSIGNED]

Baker, F., Bowers, C., and J. Linkova, "Enterprise Multihoming using Provider-Assigned Addresses without Network Prefix Translation: Requirements and Solution", Work in Progress, draft-ietf-rtgwg-enterprise-pa-multihoming-07, June 2018.

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.

- [RFC5798] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, DOI 10.17487/RFC5798, March 2010, <<https://www.rfc-editor.org/info/rfc5798>>.

- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.

- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.

Acknowledgements

Thanks to the following people (in alphabetical order) for their review and feedback: Mikael Abrahamsson, Lorenzo Colitti, Marcus Keane, Erik Kline, David Lamparter, Dusan Mudric, Erik Nordmark, and Dave Thaler.

Authors' Addresses

Jen Linkova
Google
Mountain View, California 94043
United States of America

Email: furry@google.com

Massimiliano Stucchi
RIPE NCC
Stationsplein, 11
Amsterdam 1012 AB
The Netherlands

Email: mstucchi@ripe.net