Internet Engineering Task Force
Internet-Draft
Intended status: Best Current Practice
Expires: 5 September 2024

T. Lemon
Apple Inc.
J. Hui
Google LLC

4 March 2024

Automatically Connecting Stub Networks to Unmanaged Infrastructure
draft-ietf-snac-simple-04

Abstract

This document describes a set of practices for connecting stub networks to adjacent infrastructure networks. This is applicable in cases such as constrained (Internet of Things) networks where there is a need to provide functional parity of service discovery and reachability between devices on the stub network and devices on an adjacent infrastructure link (for example, a home network).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 September 2024.

Copyright Notice

Table of Contents

1. Introduction

This document describes a set of practices for connecting stub networks to adjacent infrastructure networks. There are several use cases for stub networks. Motivating factors include:

* Incompatible speed: for example, an 802.15.4 network could not be easily bridged to a WiFi network because the data rates are so dissimilar. So either it must be bridged in a very complicated and careful way to avoid overwhelming the 802.15.4 network with irrelevant traffic, or the 802.15.4 network needs to be a separate subnet.

* Incompatible media: for example, a constrained 802.15.4 network connected as a stub network to a WiFi or ethernet infrastructure network. In the case of an 802.15.4 network, it is quite possible that the devices used to link the infrastructure network to the stub network will not be conceived of by the end user as routers. Consequently, we cannot assume that these devices will be on all the time. A solution for this use case will require some sort of commissioning process for stub routers, and can’t assume that any particular stub router will always be available; rather, any stub router that is available must be able to adapt to current conditions to provide reachability.

* Incompatible mechanisms: the medium of the stub network may not, for example, use neighbor discovery to populate a neighbor table. If the infrastructure network (as is typical) does use neighbor discovery, then bridging the two networks together would require some way of translating between neighbor discovery and whatever mechanism is used on the stub network, and hence complicates rather than simplifying the problem of connecting the two networks.

* Incompatible framing: if the stub network is a 6lowpan [RFC4944] network, packets on the stub network are expected to use 6lowpan header compression [RFC6282]. Making this work through a bridge would be very difficult.

* Convenience: end users often connect devices to each other in order to extend networks

* Transitory connectivity: a mobile device acting as a router for a set of co-located devices could connect to a network and gain access to services for itself and for the co-located devices. Such a stub network is unlikely to have more than one stub router.

What makes stub networks a distinct type of network is simply that a stub network never provides transit between networks to which it is connected. The term "stub" refers to the way the network is seen by the link to which it is connected: there is reachability through a stub network router to devices on the stub network from the infrastructure link, but there is no reachability through the stub network to any link beyond that one.

Eliminating transit routing is not intended to be seen as a virtue in itself, but rather as a simplifying assumption that makes it possible to solve a subset of the general problem of automating multi-link networks. Stub networks may be globally reachable, or may be only locally reachable. A host on a locally reachable stub network can only interoperate with hosts on the network link(s) to which it is connected. A host on a globally reachable stub network should be able to interoperate with hosts on other network links in the same infrastructure as well as hosts on the global internet.

It may be noted that just as you can plug several Home Gateway devices together in series to form multi-layer NATs, there is nothing preventing the owner of a stub network router from attaching it to a stub network as if that network were its infrastructure network. In the case of an IoT wireless network, there may be no way to do this, nor would it be desirable, but a stub router that uses ethernet on both the infrastructure and stub network sides could be connected this way. Nothing in this document is intended to prevent this from being done, but neither do we attempt to solve the problems that this could create.

The goal of this document is to describe the minimal set of changes or behaviors required to use existing IETF specifications to support the stub network use case. The result is intended to be deployable on existing networks without requiring changes to those networks.

1.1. Interoperability Goals

The specific goal is for hosts on the stub network to be able to interoperate with hosts on the adjacent infrastructure link or links. What we mean by "interoperate" is that a host on a stub network:

* is discoverable by hosts attached to adjacent infrastructure links
* is able to discover hosts attached to adjacent infrastructure links
* is able to discover hosts on the Internet

* is able to acquire an IP address that can be used to communicate with hosts attached to adjacent infrastructure links

* has reachability to the hosts attached to adjacent infrastructure links

* is reachable by hosts on the adjacent infrastructure link

* is able to reach hosts on the Internet

Discoverability here means "discoverable using DNS, or DNS Service Discovery". DNS Service Discovery includes multicast DNS [RFC6762]. As an example, when one host connected to a specific WiFi network wishes to discover services on hosts connected to that same WiFi network, it can do so using multicast DNS. Similarly, when a host on some other network wishes to discover the same service, it must use DNS-based DNS Service Discovery [RFC6763]. In both cases, "discoverable using DNS" means that the host has one or more entries in the DNS that serve to make it discoverable.

We lump discoverability in with reachability and addressability, both of which are essentially Layer 3 issues. The reason for this is that it does us no good to automatically set up connectivity between stub network hosts and infrastructure hosts if the infrastructure hosts have no means to learn about the availability of services provided by stub network hosts. For stub network hosts that only consume cloud services this will not be an issue, but for stub networks that provide services, such as IoT devices on stub networks with incompatible media, discoverability is necessary in order for stub network connectivity to be useful.

Ability to acquire an IP address that can be used to communicate means that the IP address a host on the stub network acquires can be used to communicate with it by hosts not on the stub network.

Reachability to hosts on adjacent infrastructure links means that when a host (A) on the stub network has a datagram destined for the IP address of a host (B) on an adjacent infrastructure link, host (A) knows of a next-hop router to which it can send the datagram, so that it will ultimately reach host (B) on the infrastructure network.

Reachability from hosts on adjacent infrastructure links means that when host (A) on an adjacent infrastructure link has a datagram destined for the IP address of a host (B) on the stub network, a next-hop router is known by host (A) such that, when the datagram is sent to that router, it will ultimately reach host (B) on the stub network.

To achieve the reachability goal described above, this document assumes hosts attempting to reach destinations on the stub network maintain a routing table - Type C hosts as defined in Section 3.1 of [RFC4191]). Type A and Type B hosts are out-of-scope for this document.

1.2. Usability Goals

In addition to the interoperability goals we've described above, the additional goal for stub networks is that they be able to be connected automatically, with no user intervention. The experience of connecting a stub network to an infrastructure should be as straightforward as connecting a new host to the same infrastructure network.

2. Glossary

Addressability: The ability to associate each node on a link with its own IPv6 address.

Reachability: Given an IPv6 destination address that is not on-link for any link to which a node is attached, the information required that allows the node to send packets to a router that can forward those packets towards a link where the destination address is on-link.

Adjacent Infrastructure Link (AIL): any link to which a stub network router is directly attached, that is part of an infrastructure network and is not the stub network.

Home Gateway: A device, such as a CE Router [RFC7084], that is intended to connect a single uplink network to a Local-Area Network. A CE router may be provided by an ISP and only capable of connecting directly to the ISP's means of service delivery, e.g. Cable or DSL, or it may have an ethernet port on the WAN side and one or more ethernet ports, plus WiFi, on the LAN side.

Infrastructure network: the network infrastructure to which a stub router connects. This network can be a single link, or a network of links. The network is typically formed by a Home Gateway, which may also provide some services, such as a DNS resolver, a DHCPv4 server, and a DHCPv6 prefix delegation server, for example.

Off-Stub-Network-Routable (OSNR) Prefix: a prefix advertised on the stub network that can be used for communication with hosts not on the stub network.

Stub Network: A network link that is connected by one or more Stub

Routers to an AIL an infrastructure network, but is not used for transit between that link and any other link. Section 2.1 of [RFC2328] describes the distinction between stub networks and transit networks from a topological perspective: a stub network is simply any network that does not provide transit within a routing fabric. There is reachability through a stub network router to hosts on the stub network, but there is no reachability through the stub network to any link beyond the stub network link.

Stub Router: A router that provides connectivity between a stub network and an infrastructure network. A stub router may also provide connectivity between other networks: the term "stub router" refers specifically to its role in providing connectivity to a stub network. For example, a Home Gateway may provide connectivity between a provider network (WAN) and a home network (LAN), while at the same time providing connectivity between the LAN and a stub network. What distinguishes the LAN from the stub network in this case is that the LAN is potentially a candidate to act as a transit network to reach other routers, whereas the stub network is not.

RA Beacon: A Router Advertisement (RA) that is multicast on a link so that hosts can see that the router is still present. This is in contrast to a unicast RA sent in response to the router solicit.

ULA Site Prefix: A Unique Local Address /48 prefix [RFC4193] randomly generated by each stub router for use in allocating ULA Link Prefixes to the stub network and the adjacent infrastructure link.

ULA Link Prefix: A Unique Local Address /64 prefix allocated from the ULA site prefix. Stub routers can use ULA Link prefixes to provide addressability on the stub network and/or adjacent infrastructure link as needed. If a stub router is doing NAT64, the NAT64 prefix is also a ULA Link Prefix. A total of 65,536 ULA link prefixes can be allocated from the ULA Site prefix.

3. Constants

This section describes the meaning of and gives default values for various constants used in this document.

STALE_RA_TIME (default: 10 minutes): The amount of time that can pass after the last time a router advertisement from a particular router has been received before we assume the router is no longer present. This is a stopgap in case the router is reachable but has silently stopped advertising a prefix; this situation is

unlikely, but if it does happen, new devices joining the infrastructure network will not be able to reach devices on the stub network until the stub router decides that the router that advertised the suitable prefix is stale.

STUB_PROVIDED_PREFIX_LIFETIME (default: 30 minutes): The valid and preferred lifetime the stub router will advertise. This should be long enough that a host is actually willing to use it, and obviously should also be long enough that a missed RA will not cause the host to stop using it. The values suggested here allow ten RAs to be missed before the host will stop using the prefix.

RA_BEACON_INTERVAL (default: 3 minutes): How often the stub router will transmit an RA beacon. This should be frequent enough that a missed Router Solicit (e.g. due to congestion on a WiFi link) will not result in an extremely long outage (assuming the congestion passes before the RA is sent, of course).

PREFIX_DELEGATION_INTERVAL (default: 30 minutes): The lifetime a stub router should request for a DHCPv6-delegated prefix. The longer this is, the more prefixes will be consumed on a network where stub routers are not stable. The lifetime here is chosen to be long enough that a reboot of the DHCP server will not prevent the prefix being renewed. It happens to coincide with the value of STUB_PROVIDED_PREFIX_LIFETIME, but the two should not be considered to be equivalent.

MAX_FLAGS_COPY_TIME (default: 150 minutes): The maximum time period, after receiving an RA, that a stub router can copy flag values from the header of this RA for use in its own transmitted RAs.

MAX_SUITABLE_REACHABLE_TIME (default: 60 seconds): The maximum ReachableTime value that a router can have in the Neighbor Table before any suitable prefixes it has advertised are no longer considered suitable.

4. Conventions and Terminology Used in This Document

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.

5. Support for adjacent infrastructure links

We assume that the AIL supports Neighbor Discovery [RFC4861], and specifically that routers and on-link prefixes can be advertised using router advertisements and discovered using neighbor solicits. The stub network link may also support this, or may use some different mechanism. This section specifies how advertisement of the on-link prefix for such links is managed. In this section we will use the term "Advertising Interface" as described in Section 6.2.2 of [RFC4861].

Support for AILS on networks where Neighbor Discovery is not supported is out of scope for this document. Stub routers do not provide routing between AILs when connected to more than one such link.

5.1. Managing addressability on an adjacent infrastructure link

In order to provide IPv6 routing to the stub network, IPv6 addressing must be available on each AIL. Ideally such addressing is already present on these links, and need not be provided. However, if it is not present, the stub router must provide it.

5.1.1. Suitable On-Link Prefixes

Stub routers must evaluate prefixes that are advertised on-link as to their suitability for use in communicating with devices on the stub network. If no suitable prefix is found, a stub router MUST advertise one.

An on-link prefix is considered suitable if it is advertised on the link in a Prefix Information option ([RFC4861], Section 4.6.2) with the following Prefix Information option header values:

* Prefix Length value is 64,
* ’L’ bit is set,
* ’A’ bit is set, and
* Preferred Lifetime of 30 minutes or more.

A prefix is not considered a suitable on-link prefix if the ’L’ bit is set, but the ’A’ bit is not set. This indicates that node addressability is being managed using DHCPv6. Nodes are not required to use DHCPv6 to acquire addresses, so a prefix that requires the use of DHCPv6 can’t be considered "suitable"not all hosts can actually use it.

A prefix is considered to be advertised on the link if, when a Router Solicit message ([RFC4861], Section 4.1) is sent, a Router Advertisement message is received in response which contains a prefix information option ([RFC4861], Section 4.6.2) for that prefix.

After an RA message containing a suitable prefix has been received, it can be assumed for some period of time thereafter that that prefix is still valid on the link. However, prefix lifetimes and router lifetimes are often quite long. In addition to knowing that a prefix has been advertised on the link in the past, and is still valid, we must therefore ensure that at least one router that has advertised this prefix is still alive to respond to router advertisements.

5.1.2. State Machine for maintaining a suitable on-link prefix on an infrastructure link

The possible states of an interface connected to an AIL are described here, along with actions required to be taken to monitor the state. The purpose of the state machine described here is to ensure that at all times, when a new host arrives on the AIL, it is able to acquire an IPv6 address on that link.

5.1.2.1. Status of IP addressability on adjacent infrastructure link unknown (STATE-UNKNOWN)

When the stub router interface first connects to the AIL, it MUST begin router discovery.

If, after router discovery has completed, no suitable on-link prefix has been found, the router moves this interface to STATE-BEGIN-ADVERTISING (Section 5.1.2.3).

If, during router discovery, a suitable on-link prefix is found, the router moves the interface to STATE-SUITABLE (Section 5.1.2.2).

In this state, the stub router MUST NOT treat this interface as an advertising interface as described in Section 6.2.2 of [RFC4861].

5.1.2.2. IP addressability already present on adjacent infrastructure link (STATE-SUITABLE)

When entering this state, if the router MUST discontinue treating the interface as an Advertising Interface as described in Section 6.2.2 of [RFC4861], if it has been doing so.

When a new host appears on the AIL and sends an initial router solicit, if it does not receive a suitable on-link prefix, it will not be able to communicate. Consequently, the stub router MUST

monitor router solicits and advertisements on the interface in order to determine whether a prefix that has been advertised on the link is still being advertised. To accomplish this we have two complementary methods: router staleness detection and neighbor unreachability detection.

5.1.2.2.1. Router staleness detection

The stub router MUST listen for router advertisements on the AIL to which the interface is attached, and record the time at which each router advertisement was received. The router MUST NOT consider any router advertisement that is older than STALE_RA_TIME to be suitable. When the last non-stale router advertisement containing a suitable prefixes on the link is marked stale, the stub router MUST move the interface to STATE-BEGIN-ADVERTISING.

5.1.2.2.2. Router Unreachability Detection

For each suitable route, the stub router MUST monitor the state of reachability to the router(s) that advertised it as described in ([RFC4861], Section 7.3.1) using a ReachableTime value of no more than MAX_SUITABLE_REACHABLE_TIME. The reason for this is that if no router providing the on-link prefix on the AIL is reachable, then when a new host joins the network, it will have no suitable on-link prefix to use for autoconfiguration, and thus will be unable to communicate with hosts on the stub network.

Whenever the ReachableTime for a router advertising a suitable prefix exceeds MAX_SUITABLE_REACHABLE_TIME, the stub router MUST send unicast neighbor solicits as described in Section 7.2.2 of [RFC4861] until either a response is received, which resets ReachableTime to zero, or the maximum number of retransmissions has been sent.

The stub router MUST listen for router solicits on the AIL. When a router solicit is received, if none of the on-link routers on the AIL are marked reachable, the stub router MUST move this interface to the STATE-BEGIN-ADVERTISING state (Section 5.1.2.3).

If a RA beacon interval arrives, and there are no routers advertising suitable prefixes that have a ReachableTime that is less than MAX_SUITABLE_REACHABLE_TIME, then the router MUST move this interface to the STATE-BEGIN-ADVERTISING state.

5.1.2.3. IP addressability not present on adjacent infrastructure link (STATE-BEGIN-ADVERTISING)

In this state, the stub router generates its own on-link prefix for the interface. This prefix has a valid and preferred lifetime of STUB_PROVIDED_PREFIX_LIFETIME seconds. The stub router sends a router advertisement (RA) containing this prefix in a Prefix Information Option (PIO). In the PIO, the A (autonomous configuration) flag Section 4.6.2 of [RFC4861] MUST be set and the L (on-link) flag SHOULD be set. The exception cases where the L flag can be cleared is where the specific link-layer technology and/or configuration requires clearing the L flag.

The Stub Router flag ([I-D.hui-stub-router-ra-flag]) MUST be set in the RA flags field. The values of the M and O flags MUST be copied from the respective M/O flag values seen in the most recent (unicast or multicast) RA received from a non-stub-router. For the selection of the most recent RA, the following RAs MUST be excluded:

* An RA received from a router longer ago than the Router Lifetime period indicated in the RA header. This only applies for a non-zero Router Lifetime value.

* An RA received more than MAX_FLAGS_COPY_TIME ago.

If there is no recent RA from a non-stub-router, both M and O flags MUST be cleared, unless the stub router rebooted recently. After a reboot, if no recent RA is received from a non-stub router, but a recent RA has been received from a stub router, the values for the M and O flags provided by that stub router MUST be copied. After MAX_FLAGS_COPY_TIME after reboot, the stub router MUST go back to the regular behavior defined above. This avoids a situation where a stub router that has rebooted starts to advertise different M/O flag values than other stub routers present on the same link.

The sent router advertisement MUST also include a Route Information option (Section 2.3 of [RFC4191]) for each routable prefix advertised on the stub network. If the stub router is also a normal router (e.g. a home WiFi router), it SHOULD include all other routes that it is advertising in the RA, if there is space.

After having sent the initial router advertisement, the stub router moves the interface into the STATE-ADVERTISING-SUITABLE state (Section 5.1.2.4).

5.1.2.4. IP addressability not present on adjacent infrastructure link (STATE-ADVERTISING-SUITABLE)

When entering this state, if the router MUST begin treating the interface as an Advertising Interface as described in Section 6.2.2 of [RFC4861] if it is not already doing so.

The stub router sends a router advertisement message, as described in Section 5.1.2.3, every RA_BEACON_INTERVAL seconds.

The stub router may receive a router advertisement containing one or more suitable on-link prefixes on the AIL. If any of these prefixes are different than the prefix the stub router is advertising as the on-link suitable prefix, and the Stub Router flag is not set in in the Router Advertisement flags field, the stub router moves the interface to STATE-DEPRECATING (Section 5.1.2.5).

If the stub router bit is set in the RA header flags field, then one of the following must be true in order for that prefix to be considered suitable:

* The prefixes are equal. In this case, the interface remains in STATE-ADVERTISING-SUITABLE.

* The prefix the stub router is advertising is a ULA prefix [RFC4193], and the received prefix is a non-ULA prefix. In this case, the interface moves into the STATE-DEPRECATING (Section 5.1.2.5) state.

* Both prefixes are ULA prefixes, and the received prefix, considered as a 128-bit big-endian unsigned integer, is numerically lower, then the interface moves to STATE-DEPRECATING (Section 5.1.2.5.

* Otherwise the interface remains in STATE-ADVERTISING-SUITABLE.

5.1.2.5. Stub router deprecating the on-link prefix it is advertising (STATE-DEPRECATING)

On entry to this state, the stub router has been treating the interface as an Advertising Interface as described in Section 6.2.2 of [RFC4861], and MUST continue to do so.

When the stub router has detected the availability of suitable on-link prefix on the AIL to which the interface is attached, and that prefix is preferable to the one it is advertising, it continues to advertise its own prefix, but deprecates it:

* the preferred lifetime for its prefix should be set to zero in subsequent router advertisement messages.

* the valid lifetime for its prefix should be reduced with each subsequent router advertisement messages.

* the usability of the infrastructure-provided on-link prefix should be monitored as in the STATE-SUITABLE state; if during the deprecation period, the stub router detects that there are no longer any suitable prefixes on the link, as described in Section 5.1.2.2.1 or in Section 5.1.2.2.2, it MUST return the interface to the STATE-BEGIN-ADVERTISING (Section 5.1.2.4) state and resume advertising its prefix with the valid and preferred lifetimes described there.

In this state, the valid lifetime (VALID) is computed based on three values: the current time when a router advertisement is being generated (NOW), the time at which the new suitable on-link prefix advertisement was received (DEPRECATE_TIME), and STUB_PROVIDED_PREFIX_LIFETIME. All of these values are in seconds. VALID is computed as follows:

VALID = STUB_PROVIDED_PREFIX_LIFETIME - (NOW - DEPRECATE_TIME)

If VALID is less than RA_BEACON_INTERVAL, the stub router does not include the deprecated prefix in the router advertisement. Note that VALID could be less than zero. Otherwise, the prefix is provided in the advertisement, but with a valid lifetime of VALID.

5.2. Managing addressability on the stub network

How addressability is managed on stub networks depends on the nature of the stub network. For some stub networks, the stub router can be sure that it is the only router. For example, a stub router that is providing a Wi-Fi network for tethering will advertise its own SSID and use its own joining credentials; in this case, it can assume that it is the only router for that network, and advertise a default route and on-link prefix just like any other router.

However, some stub networks are more cooperative in nature, for example IP mesh networks. On such networks, multiple stub routers may be present and be providing addressability and reachability.

In either case, some stub router connected to the stub network MUST provide a suitable on-link prefix (the OSNR prefix) for the stub network. If the stub network is a multicast-capable medium where Router Advertisements are used for router discovery, the same mechanism described in Section 5.1.2 is used.

Stub networks that do not support the use of Router Advertisements for router discovery must use some similar mechanism that is compatible with that type of network. Describing the process of establishing a common OSNR prefix on such networks is out of scope for this document.

5.2.1. Maintenance across stub router restarts

Stub routers may restart from time to time; when a restart occurs, the stub router may have been advertising state to the network which, following the restart, is no longer required.

For example, suppose there are two stub routers connected to the same infrastructure link. When the first stub router is restarted, the second takes over providing an on-link prefix. Now the first router rejoins the link. It sees that the second stub router's prefix is advertised on the infrastructure link, and therefore does not advertise its own.

This behavior can cause problems because the first stub router no longer sees the on-link prefix it had been advertising on infrastructure as on-link. Consequently, if it receives a packet to forward to such an address, it will forward that packet directly to a default router, if one is present; otherwise, it will have no route to the destination, and will drop the packet.

To address this problem, stub routers SHOULD remember the last time a prefix was advertised across restarts. On restart, the router configures the prefix on its interface but does not advertise it in Router Advertisements. Devices that are still using that prefix will be seen as on-link to the router, and so packets will be delivered using ND on-link rather than forwarded to the default router.

When a stub router has only flash memory with limited write lifetime, it may be inappropriate to do a write to flash every time an RA beacon containing a prefix is sent. In this case, the router SHOULD record the set of prefixes that have been advertised on infrastructure and the maximum valid lifetime that was advertised. On restart, the router should assume that hosts on the infrastructure link have received advertisements for any such prefixes.

When possible, it is best if all stub routers serving a particular stub network use the same 64-bit prefix on the AIL. For example, Thread stub routers use bits from the Thread Extended PAN ID to generate the ULA prefix's Global ID and Subnet ID. The Global ID generation conforms to [RFC4193] because the Extended PAN ID is generated randomly using the same mechanism that is specified in RFC 4193 for the ULA prefix bits.

5.2.2. Generating a per-stub-router ULA Site Prefix

In order to be able to provide addressability either on the stub network or on an adjacent infrastructure network, a stub router MUST allocate its own ULA Site Pefix. ULA prefixes, described in Unique Local IPv6 Unicast Addresses ([RFC4193]) are randomly allocated prefixes. A stub router MUST allocate a single ULA Site Prefix for use in providing on-link prefixes to the stub network and the adjacent infrastructure link, as needed.

Any ULA Link Prefixes allocated by a stub router SHOULD be maintained across reboots, and SHOULD remain stable over time. (TBD: mention the SHOULD exception cases) However, for privacy reasons, a stub router that roams from network to network SHOULD allocate a different ULA Link Prefix each time it connects to a different infrastructure network, unless configured to behave otherwise.

5.2.3. Using DHCPv6 Prefix Delegation to acquire a prefix to provide addressability

If DHCPv6 prefix delegation and IPv6 service are both available on the infrastructure link, then the stub router MUST attempt to acquire a prefix using DHCPv6 prefix delegation. Using a prefix provided by the infrastructure DHCPv6 prefix delegation service means (assuming the infrastructure is configured correctly) that routing is possible between the stub network links and all links on the infrastructure network, and possibly to the general internet.

By contrast, if the prefix generated by the stub router is used, reachability is only possible between the stub network and the AIL. The OSNR prefix in this case is not known to the infrastructure network routing fabric, so even though packets might be able to be forwarded to the intended destination, there would be no return path. So when the only prefix that is available is the one provided by the stub router, cloud services will not be reachable via IPv6, and infrastructure-provided NAT64 will not work. Therefore, when the stub router is able to successfully acquire a prefix using DHCPv6 PD, it MUST use DHCPv6 PD rather than the ULA Link prefix it allocated for the stub network out of its ULA Site Prefix.

A stub router SHOULD request stub network prefixes with length 64. If the stub router obtains a prefix with length less than 64, it SHOULD generate a /64 from the obtained prefix by padding with zeros. If the stub router obtains a prefix with length greater than 64, the stub router MUST treat the prefix as unsuitable and allocate a ULA Link Prefix out of its ULA Site Prefix instead.

5.3. Managing reachability on the adjacent infrastructure link

Stub routers MUST advertise reachability to stub network OSNR prefixes on any AIL to which they are connected. If the stub router is advertising a suitable prefix on any interface, any such prefixes MUST be advertised on that interface in the same router advertisement that is advertising the suitable prefix, to avoid unnecessary multicast traffic.

Each stub network will have some set of prefixes that are advertised as on-link for that network. A stub router connected to that stub network SHOULD advertise reachability to all such prefixes on any AIL to which it is attached using router advertisements.

A stub router SHOULD NOT advertise itself as a default router on an AIL by setting a non-zero Router Lifetime value in the header of its Router Advertisements. The exception to this rule is the case where the stub router itself is the default router for a particular AIL: for example, it may be the home router providing connectivity to an ISP.

5.4. Managing reachability on the stub network

The stub router MAY advertise itself as a default router on the stub network, if it itself has a default route on the AIL. In some cases it may not be desirable to advertise reachability to the Internet as a whole; in this case the stub router is not required to advertise itself as a default router.

If the stub router is not advertising itself as a default router on the stub network, it MUST advertise reachability to any prefixes that are being advertised as on-link on AILs to which it is attached. This is true for prefixes it is advertising, and for other prefixes being advertised on that link.

Note that in some stub network configurations, it is possible for more than one stub router to be connected to the stub network, and each stub router may be connected to a different AIL. In this case, a stub router advertising a default route may receive a packet destined for a link that is not an AIL for that router, but is an AIL for a different router. In such a case, if the infrastructure is not capable of routing between these two AILs, a packet which could have been delivered by another stub router will be lost by the stub router that received it.

Consequently, stub routers SHOULD be configurable to not advertise themselves as default routers on the stub network. Stub routers SHOULD be configurable to explicitly advertise AIL prefixes on the

stub network even if they are advertising as a default router. The mechanisms by which such configuration can be accomplished are out of scope for this document.

It is also possible that stub routers for more than one stub network may be connected to the same AIL. In this case, the stub routers will be advertising Router Information options in their router advertisements for their OSNR prefixes. Stub routers MUST track the presence of such routes, and MUST advertise reachability to them on interfaces connected to stub networks.

5.5. Providing discoverability between stub network links and infrastructure network links

Since DNS-SD is in wide use, and provides for ad-hoc, self-configuring advertising using the mDNS transport, this is a suitable mandatory-to-implement protocol for stub networks, which must be able to attach to infrastructure networks without the help of new mechanisms provided by the infrastructure. Therefore, stub routers MUST provide DNS-SD service as described in this section.

5.5.1. Discoverability by hosts on adjacent infrastructure links

The adjacent infrastructure can be assumed to already enable some service discovery mechanism between hosts on the infrastructure network, and can be assumed to provide a local DNS resolver. Therefore, we do not need to define a stub-network-specific mechanism for providing these services on the infrastructure network.

In some cases it will be necessary for hosts on the AIL to be able to discover devices on the stub network. In other cases, this will be unnecessary or even undesirable. For example, it may be undesirable for devices on an AIL to be able to discover devices on a Wi-Fi tether provided by a mobile phone.

One example of a use case for stub networks where such discovery is desirable is the constrained network use case. In this case a low-power, low-cost stub network provides connectivity for devices that provide services to the infrastructure. For such networks, it is necessary that devices on the infrastructure be able to discover devices on the stub network.

The most basic use case for this is to provide feature parity with existing solutions like multicast DNS (mDNS). For example, a light bulb with built-in Wi-Fi connectivity might be discoverable on the infrastructure link to which it is connected, using mDNS, but likely is not discoverable on other links. To provide equivalent functionality for an equivalent device on a constrained network that

is a stub network, the stub network device must be discoverable on the infrastructure link (which is an AIL from the perspective of the stub network).

If services are to be advertised using DNS Service Discovery [RFC6763], there are in principle two ways to accomplish this. One is to present services on the stub network as a DNS zone which can then be configured as a browsing domain in the DNS ([RFC6763], Section 11). The second is to advertise stub network services on the AIL using multicast DNS (mDNS) [RFC6762].

Because this document defines behavior for stub routers connecting to infrastructure networks that do not provide any new mechanism for integrating stub networks, there is no way for a stub router to provide DNS-SD service on an infrastructure link in the form of a DNS zone in which to do discovery. Therefore, service on the infrastructure link MUST be provided using an Advertising Proxy, as defined in [I-D.ietf-dnssd-advertising-proxy].

One limitation of this solution is that it requires that hosts on the stub network use the DNS-SD Service Registration Protocol [I-D.ietf-dnssd-srp] to register their DNS-SD advertisements. This means that in the case of a stub network used for WiFi tethering, hosts on the stub network will not be discoverable by hosts on the infrastructure network. Any solution to this problem would require that the stub router provide a Discovery Proxy [RFC8766]. However, a discovery proxy is queried using DNS, not mDNS. This requires assistance from the infrastructure network, and is therefore out of scope for this document.

5.5.2. Providing discoverability of adjacent infrastructure hosts on the stub network

Hosts on the stub network may need to discover hosts on the AIL, or on the stub network. In the IoT network example we've been using, there might be a light switch on the stub network which needs to be able to actuate a light bulb connected to the AIL. In order to know where to send the actuation messages, the light switch will need to be able to discover the light bulb's address somehow.

Because the stub network is managed by stub routers, any DNS resolver that's available on the stub network will necessarily be provided by one or more stub routers. This means that the stub router can enable discovery of hosts on the infrastructure network by hosts on the stub network using a Discovery Proxy [RFC8766]. The Discovery Proxy can be advertised as available to hosts on the stub network through the DNS resolver provided on the stub network, as described in Section 11 of [RFC6763].

By implication, this means that stub routers MUST provide a DNS resolver. In addition, stub routers MUST provide DNS zones for each AIL, and MUST list these zones in the list of default browsing zones as defined in RFC6763. [[WG: we need to say how these zones are named. Or refer to the Advertising Proxy doc and have that doc say how they are named.]]

The stub router MUST also maintain an SRP registrar and use registrations made through that registrar to populate a DNS zone which is advertised as a default browsing domain, as above. This SRP registrar MUST be advertised on the stub network either using the dnssd-srp and/or dnssd-srp-tls service names or some stub-network-specific mechanism, the details of which are out of scope for this document.

6. Providing reachability to IPv4 services to the stub network

Stub Network routers must be capable of providing NAT64 themselves, and must be capable of discovering the availability of NAT64 service on the infrastructure network and providing it when it is available and suitable.

Some network media may provide their own mechanisms for advertising NAT64 service to the stub network. If such a mechanism is available, stub routers MUST use the mechanism provided by the network medium used on the stub network to advertise NAT64 service. Otherwise, NAT64 service MUST be advertised using the PREF64 Router Advertisement option [RFC8781].

There are four possible combinations of circumstances in which to consider how to provide NAT64 service:

1. Infrastructure provides DHCPv6 PD support, and the infrastructure network provides NAT64

2. Infrastructure provides no DHCPv6 PD support, Infrastructure is providing NAT64, and there is no IPv4 on infrastructure

3. Infrastructure provides no DHCPv6 PD support, Infrastructure is providing NAT64, and there is IPv4 on infrastructure

4. Infrastructure provides no DHCPv6 PD support, infrastructure is not providing NAT64 (and may also not be providing IPv6), and there is IPv4 on infrastructure

In the first case, infrastructure-provided NAT64 is preferred, and the stub router MUST advertise this service to the stub network.

In the second case, there is no way to provide connectivity to the infrastructure: we don't have IPv6 routing other than to the adjacent infrastructure link, because we don't have a routable prefix, we don't have NAT64 for the same reason, and we don't have IPv4, so the stub router can't do NAT64 on its own. In this case, the stub router MUST NOT advertise NAT64 service.

In the third case, despite the infrastructure providing NAT64, we can't use it, so the stub router MUST provide its own NAT64 service.

In the fourth case, the stub router MUST provide its own NAT64 service.

An additional complication is that there may be more than one stub router connecting the stub network to infrastructure. In this case, it may be desirable to limit the number of stub routers providing NAT64 service, or it may be acceptable for all stub routers to provide it.

In the latter case, this should not be a problem: since each stub router is using its own ULA Site Prefix to provide NAT64, any 5-tuple that goes through a stub router's NAT64 translator will necessarily have as its destination an IPv6 address in a particular NAT64 prefix, and that address will select the correct stub router through which to send the packet for translation. This also works on the return path because each stub router has its own IPv4 address, and the return packet will be destined for that IPv4 packet, and hence will always return through the stub router that translated it on the way out.

A further complication is that in some cases, some stub routers connected to the stub network may not be able to advertise an infrastructure-provided NAT64 prefix, while others may. In this case, when the infrastructure-provided NAT64 service appears on the stub network, stub routers that are not able to advertise an infrastructure NAT64 service MUST NOT do so.

To differentiate between infrastructure-provided NAT64 service and stub router-provided NAT64 service, stub routers that advertise infrastructure-provided NAT64 service MUST use a preference of medium for this service. Stub routers advertising their own service MUST use a preference of low.

In some cases a stub router may be administratively configured with a NAT64 prefix. In this situation, the stub router MUST advertise the prefix with a preference of high.

Stub routers must monitor the advertisement of other NAT64 prefixes on the stub network. If a stub router is advertising a NAT64 prefix, and a NAT64 prefix is advertised on the stub network with a higher preference, the stub router SHOULD deprecate the prefix it is advertising.

6.1. NAT64 provided by infrastructure

Stub networks are defined to be IPv6-only because it would be difficult to implement a stub network using IPv4 technology. However, stub network devices may need to be able to communicate with IPv4-only services either on the infrastructure network, or on the global internet. Ideally, the infrastructure network fully supports IPv6, and all services on the infrastructure network are IPv6-capable. In this case, perhaps the infrastructure network provides NAT64 service to IPv4-only hosts on the internet. In this ideal setting, the stub router need do nothingthe infrastructure network is doing it all.

In this situation, if there are multiple stub routers, each connected to the same AIL, there is no need for special behavioreach stub router can advertise a default route, and any stub router may be used to route NAT64 traffic. If some stub routers are connected to different AILs than others, some of which support NAT64 and some of which do not, then the default route may not carry traffic to the correct link for NAT64 service. In this case, a more specific address to the infrastructure NAT64 prefix(es) MUST be advertised by those stub routers that are able to discover it.

In order for infrastructure-provided NAT64 to work, the stub network must have an OSNR prefix that is known to the infrastructure. Typically this means that the stub router must have acquired this prefix using DHCPv6 Prefix Delegation. Unless otherwise configured to do so, the stub router MUST NOT advertise infrastructure-provided NAT64 service on the stub network if it has not acquired the OSNR prefix through DHCPv6 Prefix Delegation.

6.2. NAT64 provided by stub router(s)

Most infrastructure networks at present do not provide NAT64 service. Many infrastructure networks do not provide DHCPv6 Prefix Delegation. In these cases it is necessary for stub routers to be able to provide NAT64 service if IPv4 hosts are to be reachable from the stub network. Therefore, stub routers MUST be capable of providing NAT64 service to the stub network. When infrastructure-provided NAT64 service is not present or is not usable, and when no other NAT64 service is already advertised on the stub network, stub routers MUST, by default, enable their own NAT64 service and advertise it on the

stub network.

To provide NAT64 service, a stub router must allocate a NAT64 prefix. For convenience, the stub network allocates a single prefix out of the ULA Site Prefix that it maintains. Out of the 2^16 possible subnets of the /48, the stub router SHOULD use the numerically highest /64 prefix.

If there are multiple stub routers providing connectivity between the stub network and infrastructure, each stub network uses its own NAT64 prefixthere is no common NAT64 prefix. The reason for this is that NAT64 translation is not stateless, and is tied to the stub router's IPv4 address. Therefore each NAT64 egress is not equivalent.

A stub network that services a Wi-Fi stub network SHOULD provide DNS64 translation: hosts on the stub network cannot be assumed to be able to do DNS64 synthesis in the stub resolver. In this case the DNS resolver on the stub router MUST honor the CD and DO bits if received in a request, since this indicates that the stub resolver on the requestor intends to do DNSSEC validation. In this case, the resolver on the stub router MUST NOT perform DNS64 synthesis.

On specific stub networks it may be desirable to require the stub network device to perform DNS64 synthesis. Stub network routers for such networks do not need to provide DNS64 synthesis. Instead, they MUST provide an ipv4only.arpa answer that advertises the NAT64 prefix for that stub router, and MUST provide an explicit route to that NAT64 prefix on the stub network using RA or whatever technology is specific to that stub network type.

In constrained networks it can be very useful if stub network resolvers provide the information required to do DNS64 translation in the answer to the AAAA query. If the answer to an AAAA query comes back with "no data" (not NXDOMAIN), this suggests that there may be an A record. In this case, the stub network's resolver SHOULD attempt to look up an A record on the same name. If such a record exists, the resolver SHOULD return no data in the Answer section of the DNS response, and SHOULD provide any CNAME records that were involved in returning the "no data" answer to the AAAA query, and SHOULD provide any A records that were ultimately returned, in the Additional section. The resolver should also include an ipv4only.arpa record in the Additional section.

7. Handling partitioning events on a stub network

Some technologies used for stub networks, for example Thread or 6LoWPAN mesh networks, can produce partitioned networks, where what is notionally the same stub network winds up looking like two or more discrete links. For mesh networks, such partitions can form and rejoin over time as a result of either changes in radio propagation or the addition of, or removal of, devices on the mesh.

On stub networks that can partition, one way of detecting that a partition has occurred is to notice that the stub router that has advertised the on-link prefix for the stub network is no longer reachable via the stub network. When this occurs, stub routers that notice this loss of reachability MUST advertise a ULA Link Prefix derived from their ULA Site Prefix on the stub network.

An implication of this is that when such a partition forms, the same ULA Link Prefix cant be advertised on both partitions, since this will result in ambiguous routing. We address this problem by requiring that each stub router generate its own ULA Site Prefix. This prefix is then available for two purposes: providing addressing on the AIL, if needed, and providing addressing on the stub network, if needed.

When partitions of this type occur, they may also heal. When a partition heals in a situation where two stub routers have both been advertising a prefix, it will now appear that there are two prefixes on the stub network.

When the time comes to deprecate one or more prefixes as a result of a network partition healing, only one prefix should remain. If there are any GUA prefixes, and if there is no specific configuration contradicting this, the GUA prefix that is numerically lowest should be kept, and all others deprecated. If there are no GUA prefixes, then the ULA Link Prefix that is numerically lowest should be kept, and the others deprecated. By using this approach, it is not necessary for the routers to coordinate in advance.

8. Services Provided by Stub Routers

In order to provide network access, stub routers must provide some network services to the stub network. We have previously discussed the following services:

DNS Resolver: The stub network MUST provide a DNS resolver. If RAs

are in use on the stub network, the DNS resolver is advertised in the Router Advertisement Recursive DNS Server option. If RAs are not in use on the stub network, then the mechanism whereby the DNS resolver is advertised by the stub router is specific to that type of stub network.

DHCPv6 Server: The use of DHCPv6 on the stub network is NOT RECOMMENDED. In some cases it may necessary, but should be disabled by default if the stub router provides this capability at all.

Discovery Proxy: In order to discover services on the AIL, stub routers MUST act as Discovery Proxies for any AILs to which they are attached.

SRP Registrar: Stub routers MUST provide SRP registrar service. This service MUST be advertised using DNS-SD in a legacy browsing domain that is discoverable through the stub router's resolver.

Legacy Browsing Domains: The stub resolver must advertise legacy browsing domains for each AIL, for the zone that is maintained by its SRP service, and in addition must list the legacy browsing domains provided by the infrastructure network, if any.

NAT64: As mentioned above, stub routers may need to provide NAT64 service so that devices on the stub network can communicate with IPv4 hosts on the infrastructure network and the global internet.

9. Normative References

[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>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <https://www.rfc-editor.org/info/rfc4191>.

[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>.

[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>.

[RFC5175] Haberman, B., Ed. and R. Hinden, "IPv6 Router Advertisement Flags Option", RFC 5175, DOI 10.17487/RFC5175, March 2008, <https://www.rfc-editor.org/info/rfc5175>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <https://www.rfc-editor.org/info/rfc6762>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <https://www.rfc-editor.org/info/rfc6763>.

[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>.

[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>.

[RFC8766] Cheshire, S., "Discovery Proxy for Multicast DNS-Based Service Discovery", RFC 8766, DOI 10.17487/RFC8766, June 2020, <https://www.rfc-editor.org/info/rfc8766>.

[RFC8781] Colitti, L. and J. Linkova, "Discovering PREF64 in Router Advertisements", RFC 8781, DOI 10.17487/RFC8781, April 2020, <https://www.rfc-editor.org/info/rfc8781>.

[I-D.ietf-dnssd-srp] Lemon, T. and S. Cheshire, "Service Registration Protocol for DNS-Based Service Discovery", Work in Progress, Internet-Draft, draft-ietf-dnssd-srp-25, 4 March 2024, <https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-dnssd-srp/>.

[I-D.ietf-dnssd-advertising-proxy] Cheshire, S. and T. Lemon, "Advertising Proxy for DNS-SD Service Registration Protocol", Work in Progress, Internet-Draft, draft-ietf-dnssd-advertising-proxy-03, 28 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-dnssd-advertising-proxy-03>.

[I-D.hui-stub-router-ra-flag]
Hui, J., "Stub Router Flag in ICMPv6 Router Advertisement Messages", Work in Progress, Internet-Draft, draft-hui-stub-router-ra-flag-02, 27 February 2024, <https://datatracker.ietf.org/doc/html/draft-hui-stub-router-ra-flag-02>.

10. Informative References

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <https://www.rfc-editor.org/info/rfc4944>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

Authors' Addresses

Ted Lemon
Apple Inc.
One Apple Park Way
Cupertino, California 95014
United States of America
Email: mellon@fugue.com

Jonathan Hui
Google LLC
1600 Amphitheatre Parkway
Mountain View, California 940432
United States of America
Email: jonhui@google.com