

Internet Engineering Task Force
Internet-Draft
Intended status: Best Current Practice
Expires: 27 October 2022

T. Lemon
Apple Inc.
25 April 2022

Connecting Stub Networks to Existing Infrastructure
draft-lemon-stub-networks-03

Abstract

This document describes a set of practices for connecting stub networks to adjacent infrastructure networks, as well as to larger network fabrics. 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).

The stub networks use case is intended to fully address the need to attach a single network link to an infrastructure network, where the attached link provides no through routing and in cases where integration to the infrastructure routing fabric (if any) is not available.

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 27 October 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	3
2.	Glossary	3
3.	Support for adjacent infrastructure links	4
3.1.	Managing addressability on the adjacent infrastructure link	4
3.1.1.	IP addressability already present on adjacent infrastructure link	4
3.1.2.	IP addressability not present on adjacent infrastructure link	6
3.1.3.	Resolving contention over which prefix to deprecate	6
3.1.4.	Handling the presence of multiple stub routers	7
3.2.	Managing addressability on the stub network	7
3.2.1.	Maintenance across stub router restarts	8
3.2.2.	Generating a ULA prefix to provide addressability	9
3.3.	Managing reachability on the adjacent infrastructure link	9
3.4.	Managing reachability on the stub network	9
3.5.	Providing discoverability of stub network hosts on the adjacent infrastructure link	10
3.6.	Providing discoverability of adjacent infrastructure hosts on the stub network	12
4.	Providing reachability to IPv4 services to the stub network	12
4.1.	NAT64 provided by infrastructure	12
4.2.	NAT64 provided by stub router(s)	13
5.	Handling partitioning events on a stub network	14
6.	Support for non-adjacent links	14
6.1.	Acquiring an off-stub-network-routable prefix for the stub network	15
6.2.	Arranging for routing to a stub network's off-stub-network routable prefix	16

6.3.	Making service advertisements available on non-adjacent infrastructure	16
6.4.	Making service advertisements available on the internet	16

6.5.	Distinction between non-adjacent infrastructure and global internet connectivity	17
7.	Normative References	17
	Author's Address	18

[1.](#) Introduction

This document describes a set of practices for connecting stub networks to adjacent infrastructure networks, as well as to larger network fabrics. 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).

The stub networks use case is intended to fully address the need to attach a single network link to an infrastructure network, where the attached link provides no through routing and in cases where integration to the infrastructure routing fabric (if any) is not available.

[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.

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 may also provide some services, such as a DNS resolver, a DHCPv4 server, and a DHCPv6 prefix delegation

server, for example.

Infrastructure link any link in a network infrastructure that is managed by a single entity.

Adjacent infrastructure link (AIL) an infrastructure link to which a stub router is directly connected.

Non-adjacent infrastructure link (NAIL) an infrastructure link to which a stub router is not directly connected.

Non-adjacent link (NAL) any link to which the stub router is not

Lemon

Expires 27 October 2022

[Page 3]

Internet-Draft

Connecting Stub Networks

April 2022

directly connected, whether within an infrastructure or elsewhere on the Internet.

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.

[3.](#) Support for adjacent infrastructure links

We assume that adjacent infrastructure link supports Router and Prefix Discovery using router advertisements. Adjacent infrastructure links on networks where this is not supported are out of scope for this document.

[3.1.](#) Managing addressability on the adjacent infrastructure link

In order to provide IPv6 routing to the stub network, IPv6 addressing must be available on the adjacent infrastructure link. In the ideal case, such addressing is already present on the link, and need not be provided. In this case, the stub router **SHOULD NOT** provide addressability on the adjacent infrastructure link.

[3.1.1.](#) IP addressability already present on adjacent infrastructure link

IPv6 addressing is considered to be present on the link if a usable on-link prefix is advertised on the adjacent infrastructure link. A usable on-link prefix could be a prefix advertised on the link that is on-link and allows autonomous configuration. A prefix is also a

usable on-link prefix if it is advertised on the link as on-link, and if the 'm' bit is set in the Router Advertisement message header ([\[RFC4861\], Section 4.2](#)) that contains the Prefix option. This indicates that node addressability is being managed using DHCPv6.

A prefix is 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 such an RA message has been received, it can be assumed for some period of time thereafter that the prefix is still valid on the link. However, prefix lifetimes and router lifetimes are often quite long. The mere fact that a prefix that has been advertised is still within its valid lifetime does not mean that that prefix is still being advertised on the link.

This is important because when a new host appears on the adjacent infrastructure link and sends an initial router solicit, if it does not receive a usable on-link prefix, it will not be able to communicate. Consequently, the stub router MUST monitor router solicits and advertisements on the link in order to determine whether a prefix that has been advertised on the link is still being advertised.

There are several methods that can be used to accomplish this:

The stub router MUST listen for router advertisements on the adjacent infrastructure link, and record the time at which each router advertisement was received. A router advertisement that is more than STALE_RA_TIME seconds old MUST be assumed to no longer be advertised on the link. When the last non-stale router advertisement containing a usable prefixes on the link is marked stale, the stub router should begin Router Discovery ([\[RFC4861\], Section 6.3](#)).

The stub router MUST listen for router solicits on the adjacent infrastructure link. When a router solicit is received, the router SHOULD set a timer for VICARIOUS_SOLICIT_TIME seconds. If, after that amount of time, no router advertisements are received that

contain a usable on-link prefix, the stub router MUST begin router discovery. This is necessary in case the response to the router solicit was unicast, since in this case the stub router would not see that response. When the stub router first connects to the adjacent infrastructure link, it MUST begin router discovery.

When router discovery completes, the stub router evaluates whether or not a usable on-link prefix has been seen in a non-stale router advertisement during router discovery. If no usable on-link prefix has been seen, then the stub router MUST begin to provide a usable on-link prefix.

As an alternative to the vicarious router discovery process described here, the stub router could monitor the presence of the router advertising the on-link prefix in the neighbor cache. If the neighbor cache entry becomes stale, this could be an indication that the prefix is also stale. If the neighbor cache entry goes stale, the router would need to try to refresh it, and if that fails, then the stub router must begin advertising its own on-link prefix on the stub network.

[3.1.2.](#) IP addressability not present on adjacent infrastructure link

When there is no usable on-link prefix on the adjacent infrastructure network, the stub router provides its own on-link prefix. This prefix has a valid and preferred lifetime of STUB_PROVIDED_PREFIX_LIFETIME seconds. This prefix MUST allow for autonomous configuration (SLAAC).

The stub router must advertise this prefix every BEACON_INTERVAL seconds. When the stub router is advertising reachability to the stub network, the on-link prefix advertisement and the route information advertisement must be contained in the same router advertisement.

When the stub router is advertising an on-link prefix on the AIL, it

may receive a router advertisement containing a usable on-link prefix for the AIL with a non-zero preferred lifetime. In this case, the stub router should begin to deprecate the on-link prefix it is advertising on the AIL. The preferred lifetime for this prefix should be set to zero in subsequent advertisements.

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 usable 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:

$$\text{VALID} = \text{STUB_PROVIDED_PREFIX_LIFETIME} - (\text{NOW} - \text{DEPRECATE_TIME})$$

If VALID is less than 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.

[3.1.3.](#) Resolving contention over which prefix to deprecate

It is also possible that all routers on the link that are capable of advertising prefixes might be following the same protocol of deprecating their own prefix when a valid prefix shows up. To prevent a situation where all routers deprecate their prefix and wait until there are no valid prefixes being advertised before advertising a prefix, each stub router must detect the situation where, having deprecated its own prefix, all of the other prefixes being advertised on the link have also been deprecated.

When this situation occurs, each router on the link MUST compare the valid lifetimes of all the prefixes that have been seen. If the router's own prefix expires last, then that router should immediately resume publishing its prefix as a preferred prefix.

If a router observes this situation and its prefix is not the one that expires last, it MUST set a timer for UNDEPRECATE_WAIT seconds, while continuing to observe prefix advertisements on the link. If,

when the timer expires, the prefix that expires last has not been re-published as a preferred prefix, then that prefix is marked as 'really deprecated', and no longer considered a candidate for deprecation.

Using the remaining list of prefixes, the router should then apply the same algorithm. It should continue to apply this algorithm until either its prefix becomes the one to re-publish as preferred, or some other router has re-published its prefix as preferred.

[3.1.4.](#) Handling the presence of multiple stub routers

When multiple stub routers are connected to the same AIL, and no usable on-link prefix is being provided on that link by the infrastructure, there will be a competition between routers to provide a usable on-link prefix. In order to avoid duplication, stub routers **MUST** include a random offset in the time interval across which router discovery is performed. This ensures that after a power failure, not all stub routers will exit router discovery at the exact same time, and so one stub router should advertise a usable on-link prefix before the others. This should prevent the other stub routers from advertising additional on-link prefixes.

There is no particular harm caused by advertising multiple on-link prefixes, but it is preferable to minimize this, because each on-link prefix consumes space in every on-link host's routing table, and consumes time when making source address selection and routing decisions.

[3.2.](#) 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 usable 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 [Support for adjacent infrastructure links] 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.

[3.2.1.](#) 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 can immediately begin deprecating the prefix, and can stop after the prefix valid lifetime goes to zero, based on the recorded time that the last advertisement was sent.

When a stub router has only flash memory with limited write lifetime, it may be inappropriate to do a write to flash every time a prefix beacon happens. 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, and should immediately deprecate them, and continue to do so until the maximum valid lifetime has elapsed after restart.

[3.2.2.](#) Generating a ULA prefix to provide addressability

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 prefix. ULA prefixes, described in Unique Local IPv6 Unicast Addresses ([[RFC4193](#)]) are randomly allocated prefixes. A stub router **MUST** allocate a single ULA prefix for use in providing on-link prefixes to the stub network and the infrastructure network, as needed.

The ULA prefix allocated by a stub router **SHOULD** be maintained across reboots, and **SHOULD** remain stable over time. For privacy reasons, a stub router that roams from network to network may wish to allocate a different ULA prefix each time it connects to a different infrastructure network.

If IPv6 prefix delegation is available, which implies that IPv6 service is also available on the infrastructure link, then the stub router **MAY** use IPv6 prefix delegation to acquire a prefix to advertise on the stub network, rather than allocating one out of its ULA prefix.

[3.3.](#) 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.

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

[3.4.](#) 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 need not advertise itself as a default router.

If the stub router is not advertising itself as a default 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. Stub routers **SHOULD** be configurable to advertise NAIL prefixes on the stub network; such configuration would include a list of NAIL prefixes to advertise. This list may be configured in a management interface or as a result of these routes being delivered in a routing protocol or through router discovery. The mechanisms by which such configuration can be accomplished are out of scope for this document.

[3.5.](#) Providing discoverability of stub network hosts on the adjacent infrastructure link

In some cases it will be necessary for hosts on the adjacent infrastructure link 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 adjacent infrastructure link to be able to discover devices on a Wi-Fi tether, for example 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](#)].

Stub network routers cannot be assumed to be able to integrate into the DNS naming hierarchy of the infrastructure network. Therefore, stub networks must be able to rely on ad-hoc service advertisement protocols. Since mDNS is in wide use, this is a suitable protocol for this use case. This is not to say that mDNS is the only such protocol that could be used, but it is the one that we suggest implementing.

In order to provide mDNS discovery for devices on the stub network, one of two solutions is likely to be applicable, depending on the operational practicalities of the stub network. For a constrained stub network, on which battery operated devices may be attached, mass multicast traffic for service discovery is impractical, since every device needs to wake up for every service discovery, even if they don't offer that service, and since many such devices may be operating on battery power. For such a network, multicast DNS is not a good choice.

For such networks, a unicast service registration protocol such as

DNS-SD Service Registration Protocol (SRP) [[I-D.ietf-dnssd-srp](#)] is a good solution. The stub router can act as an SRP server on the stub network, accepting service advertisements from stub network devices. On the adjacent infrastructure network, it can advertise those services as multicast DNS Advertising Proxy [[I-D.sctl-advertising-proxy](#)].

For other stub networks, for example a Wi-Fi-based Personal Area Network provided as part of a tethering function on a mobile device, multicast DNS may be the only option. For Wi-Fi stub networks, there is such a large installed base of devices supporting mDNS that requiring some other service advertisement solution would be problematic simply because it would require new software for that entire installed base. For other networks, particularly constrained networks, where devices do not currently support mDNS, no such obstacle exists.

Because the primary use case for discovery of devices on a stub network is the use case where the stub network is joining a constrained network to an existing infrastructure link, we currently only describe a solution (DNS-SD SRP) for that use case. A solution for the use case where the stub router must provide discoverability for a stub network where mDNS advertising is preferred is out of scope for this document.

[3.6.](#) Providing discoverability of adjacent infrastructure hosts on the stub network

Hosts on the stub network may need to discover hosts on the adjacent infrastructure 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 adjacent infrastructure network. 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.

In the case of a Wi-Fi stub network, devices on the stub network will need to be able to access the Internet, and may also need to be able to access local services on the adjacent infrastructure link.

In order to address these use cases, the stub network router SHOULD provide a DNS-SD Discovery Proxy [[RFC8766](#)] and a DNS resolver. Since

these two functions are combined, if the stub router provides them, it MUST offer both services on the standard DNS UDP and TCP ports.

[4.](#) Providing reachability to IPv4 services to the stub network

[4.1.](#) 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 adjacent infrastructure, 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 nothing-the infrastructure network is doing it all.

In this situation, if there are multiple stub routers, each connected to the same adjacent infrastructure link, there is no need for special behavior-each stub router can advertise a default route, and any stub router will do to route NAT64 traffic. If some stub routers are connected to different adjacent infrastructure links 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.

[4.2.](#) NAT64 provided by stub router(s)

Most infrastructure networks at present do not provide NAT64 service. It is therefore necessary for stub routers to be able to provide NAT64 service if IPv4 hosts are to be reachable from 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 /48 ULA 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 prefix--there 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.

[5.](#) Handling partitioning events on a stub network

If a stub network is constructed using mesh technology, it may become partitioned. In such a situation, it may be one stub router is connected to one partition, and another stub router is connected to

the other partition. In this situation, in order for all nodes to be reachable, it is necessary that each partition of the stub network have its own prefix. When such a partition occurs, the stub routers must detect that it has occurred. If a stub router is currently providing a prefix on the stub network, it need take no action. If a stub router had not been providing a prefix on the stub network, and now discovers that there is no stub router providing a prefix on the network, it MUST begin to provide its own prefix on the stub network. It MUST also advertise reachability to that new prefix on its adjacent infrastructure link(s).

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. Since partition events may represent a recurring situation, stub routers SHOULD wait for at least PARTITION_HEAL_WAIT_TIME before deprecating one of these prefixes.

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 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.

[6.](#) Support for non-adjacent links

There are two ways that connectivity to non-adjacent links can be established. The first is that if the infrastructure network as a whole has a working IPv4 routing fabric, NAT64 can be used to enable hosts on the stub network to establish communications with hosts on non-adjacent links, including the Internet. In some cases, this is all that is needed.

However, if it will be necessary for nodes on non-adjacent networks to establish communications with nodes on the stub network, this will require a working IPv6 routing fabric connecting the stub network to any non-adjacent links from which communications will need to be

established.

In order for such routing to work, the stub network will also need to acquire a prefix that the infrastructure network is aware of and can route to. The ULA prefix that can work for communicating to adjacent infrastructure links will not work for communicating to non-adjacent links.

6.1. Acquiring an off-stub-network-routable prefix for the stub network

A prefix may be acquired by using DHCPv6 Prefix Delegation ([\[RFC8415\]](#), [Section 6.3](#)). The stub router then advertises this prefix as the on-link prefix for the stub network, as before. It also advertises reachability to this prefix using router advertisements, as before.

In the case where there is more than one stub router, it would be best if only one stub router requested a delegated prefix. This can be managed through the mechanism described earlier: the stub router only acquires a prefix to advertise when it has decided that it needs to advertise a prefix, and so in most cases only one stub router at a time will request a delegated prefix.

In order to avoid excessive consumption of delegated prefixes, stub routers connected to stub networks that support multiple stub routers SHOULD request short lifetimes for delegated prefixes and renew frequently. Stub routers SHOULD request a lifetime of PREFIX_DELEGATION_INTERVAL. Stub routers SHOULD record the time that a prefix was acquired in stable storage, and SHOULD release the prefix using a "DHCP Release" transaction when shutting down, or when it determines that a prefix is no longer needed (See "graceful shutdown" in Figure 9 of [\[RFC8415\]](#) for details). Stub routers SHOULD release any remembered still-valid prefix after reboot, if after rebooting it is discovered that another prefix is being advertised on the stub network.

[6.2.](#) Arranging for routing to a stub network's off-stub-network routable prefix

We can assume that a side effect of the prefix delegation process will be to establish routing to the stub router that requested the prefix. This should mean that any node that wishes to establish communication with a node on the stub network will be able to do so through the delegating router that provides the prefix or, if it is attached to an infrastructure link that is adjacent to the stub router, through the stub router itself by means of the router advertisement it is providing.

The case of multiple stub routers is more complicated however. Any routing that comes as a side-effect of DHCPv6 Prefix Delegation will only route through the stub router that acquired the prefix. Other stub routers can provide reachability on their respective adjacent infrastructure links, but reachability across the full routing fabric of the infrastructure network will only be possible if there is some routing protocol present on the infrastructure network. Addressing this problem is out of scope for this document.

[6.3.](#) Making service advertisements available on non-adjacent infrastructure

In order for service advertisements to be available on non-adjacent infrastructure, the infrastructure must provide SRP service for constrained stub networks, and must advertise the availability of such service so that stub routers can forward SRP updates to that SRP service, rather than providing SRP as a local service. This SRP service can be discovered using DNS-SD, using the `_dnssd-srp-tls` service type. If the stub network requires UDP-based SRP rather than tls-based SRP, the stub router **MUST** act as a proxy to deliver SRP updates over the tcp+tls transport.

For stub networks that use multicast DNS, stub routers must provide a discovery proxy service, and must advertise that service to the infrastructure. In turn, the infrastructure must configure that service to be discoverable by devices on the infrastructure, as described in [\[RFC8766\]](#), [Section 6](#).

[6.4.](#) Making service advertisements available on the internet

The mechanism described previously for making service advertisements available to non-adjacent infrastructure also scales to the internet, since it uses DNS. Indeed, the question an operator should ask before enabling such discovery is, do they want their stub network devices to be discoverable on the internet. If it becomes possible

to configure service advertising automatically, behavior similar to

that specified in [\[RFC6092\]](#), [Section 3.2](#) and 3.3, would be advised: do not automatically advertise stub network devices on the Internet.

[6.5](#). Distinction between non-adjacent infrastructure and global internet connectivity

Stub routers may be mobile, or fixed. That is, they may move from location to location along with some or all of their connected devices, attaching to whatever infrastructure is available. Or they may be fixed devices that are only ever expected to exist in one particular location.

For devices that are intended to be in a fixed location, the distinction between infrastructure links and the internet as a whole is meaningful; for mobile nodes it most likely is not, unless such a node is only going to ever attach to trusted infrastructure as it moves from location to location-not a common scenario.

For fixed links, the infrastructure may be trusted, in which case the distinction between infrastructure and internet can be expected to be managed by the infrastructure, and therefore only visible to the stub router in the sense that some non-adjacent destinations may be reachable (infrastructure destinations, for example) while others are not.

The reason for mentioning this here is to point out that the stub router can't be expected to manage this interface: it is up to the infrastructure network to do so, either implicitly or explicitly. [\[RFC7084\]](#) provides a set of default behaviors for home routers that may be adequate for automatically managing this interface, but further work in this area may be warranted.

[7](#). Normative References

- [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>>.
- [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>>.

- [RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", [RFC 6092](#), DOI 10.17487/RFC6092, January 2011, <<https://www.rfc-editor.org/info/rfc6092>>.

Lemon

Expires 27 October 2022

[Page 17]

Internet-Draft

Connecting Stub Networks

April 2022

- [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>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.
- [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>>.
- [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-12](#), 24 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-dnssd-srp-12>>.
- [I-D.sctl-advertising-proxy]
Cheshire, S. and T. Lemon, "Advertising Proxy for DNS-SD Service Registration Protocol", Work in Progress,

Internet-Draft, [draft-sctl-advertising-proxy-02](https://datatracker.ietf.org/doc/html/draft-sctl-advertising-proxy-02), 12 July 2021, <<https://datatracker.ietf.org/doc/html/draft-sctl-advertising-proxy-02>>.

Author's Address

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