                        Unique IPv6 Prefix Per Host
                draft-ietf-v6ops-unique-ipv6-prefix-per-host-10

Abstract

   This document outlines an approach utilising existing IPv6 protocols
   to allow hosts to be assigned a unique IPv6 prefix (instead of a
   unique IPv6 address from a shared IPv6 prefix). Benefits of unique
   IPv6 prefix over a unique service provider IPv6 address include
   improved host isolation and enhanced subscriber management on shared
   network segments.

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 March 22, 2018.

Table of Contents

1. Introduction

The concepts in this document are originally developed as part of a
large scale, production deployment of IPv6 support for a provider
managed shared access network service.

A shared network service, is a service offering where a particular L2
access network (e.g. wifi) is shared and used by multiple visiting
devices (i.e. subscribers). Many service providers offering shared
access network services, have legal requirements, or find it good
practice, to provide isolation between the connected visitor devices
to control potential abuse of the shared access network.

A network implementing a unique IPv6 prefix per host, can simply
ensure that devices cannot send packets to each other except through
the first-hop router. This will automatically provide robust
protection against attacks between devices that rely on link-local
ICMPv6 packets, such as DAD reply spoofing, ND cache exhaustion,
malicious redirects, and rogue RAs. This form of protection is much
more scalable and robust than alternative mechanisms such as DAD
proxying, forced forwarding, or ND snooping.

In this document IPv6 support does not preclude support for IPv4;
however, the primary objectives for this work was to make it so that
user equipment (UE) were capable of an IPv6 only experience from a
network operators perspective. In the context of this document, UE
can be 'regular' end-user-equipment, as well as a server in a
datacenter, assuming a shared network (wired or wireless).

Details of IPv4 support are out of scope for this document. This
document will also, in general, outline the requirements that must be
satisfied by UE to allow for an IPv6 only experience.

In most current deployments, User Equipment (UE) IPv6 address
assignment is commonly done using either IPv6 SLAAC RFC4862 [RFC4862]
and/or DHCP IA_NA (Identity Association - Non-temporary Address)
RFC3315 [RFC3315]. During the time when this approach was developed
and subsequently deployed, it has been observed that some operating
systems do not support the use of DHCPv6 for the acquisition of IA_NA
per RFC7934 [RFC7934]. To not exclude any known IPv6
implementations, IPv6 SLAAC based subscriber and address management
is the recommended technology to reach highest percentage of
connected IPv6 devices on a provider managed shared network service.
In addition an IA_NA-only network is not recommended per RFC 7934
RFC7934 [RFC7934] section 8. This document will detail the mechanics
involved for IPv6 SLAAC based address and subscriber management
coupled with stateless DHCPv6, where beneficial.

This document focuses upon the process for UEs to obtain a unique
IPv6 prefix.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in RFC 2119 [RFC2119].

2. Motivation and Scope of Applicability

The motivation for this work falls into the following categories:

o Deployment advice for IPv6 that will allow stable and secure IPv6
   only experience, even if IPv4 support is present

o Ensure support for IPv6 is efficient and does not impact the
   performance of the underlying network and in turn the customer
   experience

o Allow for the greatest flexibility across host implementation to
   allow for the widest range of addressing and configuration
   mechanisms to be employed. The goal here is to ensure that the
   widest population of UE implementations can leverage the
   availability of IPv6

o Lay the technological foundation for future work related to the
   use of IPv6 over shared media requiring optimized subscriber
   management

o Two devices (subscriber/hosts), both attached to the same provider
   managed shared network should only be able to communicate through
   the provider managed First Hop Router

o Provide guidelines regarding best common practices around IPv6
   neighborship discovery RFC4861 [RFC4861] and IPv6 address managent
   settings between the First Hop router and directly connected
   hosts/subscribers.

3. Design Principles

The First Hop router discussed in this document is the L3-Edge router
responsible for the communication with the devices (hosts and
subscribers) directly connected to a provider managed shared network,
and to transport traffic between the directly connected devices and
between directly connected devices and remote devices.

The work detailed in this document is focused on providing details
regarding best common practices of the IPv6 neighbor discovery and
related IPv6 address management settings between the First Hop router
and directly connected hosts/subscribers. The documented Best
Current Practice helps a service provider to better manage the shared
provider managed network on behalf of the connected devices.

This document recommends providing a unique IPv6 prefix to devices
connected to the managed shared network. Each unique IPv6 prefix can
function as control-plane anchor point to make sure that each device
receives expected subscriber policy and service levels (throughput,
QoS, security, parental-control, subscriber mobility management,
etc.).

4. IPv6 Unique Prefix Assignment

When a UE connects to the shared provider managed network and is
attached, it will initiate IP configuration phase. During this phase
the UE will, from an IPv6 perspective, attempt to learn the default
IPv6 gateway, the IPv6 prefix information, the DNS information
RFC8106 [RFC8106], and the remaining information required to
establish globally routable IPv6 connectivity. For that purpose, the
the subscriber sends a RS (Router Solicitation) message.

The First Hop Router receives this subscriber RS message and starts
the process to compose the response to the subscriber originated RS
message. The First Hop Provider Router will answer using a solicited
RA (Router Advertisement) to the subscriber.

When the First Hop Provider Router sends a solicited RA response, or
periodically sends unsolicited RAs, the RA MUST be sent only to the

subscriber that has been assigned the Unique IPv6 prefix contained in the RA. This is achieved by sending a solicited RA response or unsolicited RAs to the all-nodes group, as detailed in RFC4861 [RFC4861] section 6.2.4 and 6.2.6, but instead of using the link-layer multicast address associated with the all-nodes group, the link-layer unicast address of the subscriber that has been assigned the Unique IPv6 prefix contained in the RA MUST be used as the link-layer destination RFC6085 [RFC6085]. Or, optionally in some cases, a solicited RA response could be sent unicast to the link-local address of the subscriber as detailed in RFC4861 [RFC4861] section 6.2.6, nevertheless unsolicited RAs are always sent to the all-nodes group.

This solicited RA contains two important parameters for the subscriber to consume: a Unique IPv6 prefix (currently a /64 prefix) and some flags. The Unique IPv6 prefix can be derived from a locally managed pool or aggregate IPv6 block assigned to the First Hop Provider Router or from a centrally allocated pool. The flags indicate to the subscriber to use SLAAC and/or DHCPv6 for address assignment; it may indicate if the autoconfigured address is on/off-link and if 'Other' information (e.g. DNS server address) needs to be requested.

The IPv6 RA flags used for best common practice in IPv6 SLAAC based Provider managed shared networks are:

o M-flag = 0 (subscriber address is not managed through DHCPv6), this flag may be set to 1 in the future if/when DHCPv6 prefix delegation support is desired)

o O-flag = 1 (DHCPv6 is used to request configuration information i.e. DNS, NTP information, not for IPv6 addressing)

o A-flag = 1 (The subscriber can configure itself using SLAAC)

o L-flag = 0 (the prefix is not an on-link prefix, which means that the subscriber will never assume destination addresses that match the prefix are on-link and will always send packets to those addresses to the appropriate gateway according to route selection rules.)

The use of a unique IPv6 prefix per subscriber adds an additional level of protection and efficiency as it relates to how IPv6 Neighbor Discovery and Router Discovery processing. Since the UE has a unique IPv6 prefix all traffic by default will be directed to the First Hop router. Further, the flag combinations documented above maximise the IPv6 configurations that are available by hosts including the use of privacy IPv6 addressing.

The architected result of designing the RA as documented above is
that each subscriber gets its own unique IPv6 prefix. Each host can
consequently use SLAAC or any other method of choice to select its
/128 unique address. Either stateless DHCPv6 RFC3736 [RFC3736] or
IPv6 Router Advertisement Options for DNS Configuration RFC8106
[RFC8106] can be used to get the IPv6 address of the DNS server. If
the subscriber desires to send anything external including other
subscriber devices (assuming device to device communications is
enabled and supported), then, due to the L-bit being unset, it SHOULD
send this traffic to the First Hop Router.

After the subscriber received the RA, and the associated flags, it
will assign itself a 128 bit IPv6 address using SLAAC. Since the
address is composed by the subscriber device itself, it will need to
verify that the address is unique on the shared network. The
subscriber will for that purpose, perform Duplicate Address Detection
algorithm. This will occur for each address the UE attempts to
utilize on the shared provider managed network.

5. IPv6 Neighbor Discovery Best Practices

An operational consideration when using IPv6 address assignment using
IPv6 SLAAC is that after the onboarding procedure, the subscriber
will have a prefix with certain preferred and valid lifetimes. The
First Hop Provider Router extends these lifetimes by sending an
unsolicited RA, the applicable MaxRtrAdvInterval on the first hop
router MUST therefore be lower than the preferred lifetime. One
consequence of this process is that the First Hop Router never knows
when a subscriber stops using addresses from a prefix and additional
procedures are required to help the First Hop Router to gain this
information. When using stateful DHCPv6 IA_NA for IPv6 subscriber
address assignment, this uncertainty on the First Hop Router is not
of impact due to the stateful nature of DHCPv6 IA_NA address
assignment.

Following is a reference table of the key IPv6 router discovery and
neighbor discovery timers for provider managed shared networks:

o Maximum IPv6 Router Advertisement Interval = 300s (or 600s when
   battery consumption is a concern according RFC7772 [RFC7772]).

o IPv6 Router LifeTime = 3600s

o Reachable time = 30s

o IPv6 Valid Lifetime = 3600s

o IPv6 Preferred Lifetime = 1800s

o Retransmit timer = 0s

The stateless nature of the subscriber IPv6 SLAAC connectivity model
introduces non-optimal resource consumption (i.e. memory, neighbor
state) on the First Hop Router. To reduce undesired resource
consumption, such as in the case of WiFi hotspots, is to remove
subscriber context as quickly as possible when a device or subscriber
becomes non-active. A possible solution is to use a subscriber or
device inactivity timer, that after tracking a pre-defined (currently
unspecified) number of minutes, deletes the subscriber context on the
First Hop Router.

When employing stateless IPv6 address assignment, a number of widely
deployed operating systems will attempt to utilise RFC 4941 RFC4941
[RFC4941] temporary 'private' addresses.

Similarly, when using this technology in a datacenter, the UE server
may need to use several addresses from the same Unique IPv6 Prefix,
for example because is using multiple virtual hosts, containers, etc.
in the bridged virtual switch. This can lead to the consequence that
a UE has multiple /128 addresses from the same IPv6 prefix. The
First Hop Provider Router MUST be able to handle the presence and use
of multiple globally routable IPv6 addresses.

For accounting purposes, the First Hop Provider Router must be able
to send usage statistics per subscriber using Radius attributes.

6. IANA Considerations

No IANA considerations are defined at this time.

7. Security Considerations

The mechanics of IPv6 privacy extensions RFC4941 [RFC4941] is
compatible with assignment of a unique IPv6 Prefix per Host.
However, when combining both IPv6 privacy extensions and a unique
IPv6 Prefix per Host a reduced privacy experience for the subscriber
is introduced, because a prefix may be associated with a subscriber,
even when the subscriber implemented IPv6 privacy extensions RFC4941
[RFC4941].

No other additional security considerations are made in this
document.

8. Acknowledgements

   The authors would like to thank the following, in alphabetical order,
   for their contributions:

   Ben Campbell, Brian Carpenter, Tim Chown, Lorenzo Colitti, Killian
   Desmedt, David Farmer, Brad Hilgenfeld, Wim Henderickx, Erik Kline,
   Suresh Krishnan, Warren Kumari, Thomas Lynn, Jordi Palet, Phil
   Sanderson, Colleen Szymanik, Jinmei Tatuya, Eric Vyncke, Sanjay
   Wadhwa

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>.

   [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins,
              C., and M. Carney, "Dynamic Host Configuration Protocol
              for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July
              2003, <https://www.rfc-editor.org/info/rfc3315>.

   [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol
              (DHCP) Service for IPv6", RFC 3736, DOI 10.17487/RFC3736,
              April 2004, <https://www.rfc-editor.org/info/rfc3736>.

   [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>.

   [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
              Address Autoconfiguration", RFC 4862,
              DOI 10.17487/RFC4862, September 2007,
              <https://www.rfc-editor.org/info/rfc4862>.

   [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy
              Extensions for Stateless Address Autoconfiguration in
              IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007,
              <https://www.rfc-editor.org/info/rfc4941>.

   [RFC6085] Gundavelli, S., Townsley, M., Troan, O., and W. Dec,
              "Address Mapping of IPv6 Multicast Packets on Ethernet",
              RFC 6085, DOI 10.17487/RFC6085, January 2011,
              <https://www.rfc-editor.org/info/rfc6085>.

   [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy
              Consumption of Router Advertisements", BCP 202, RFC 7772,
              DOI 10.17487/RFC7772, February 2016,
              <https://www.rfc-editor.org/info/rfc7772>.

   [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi,
              "Host Address Availability Recommendations", BCP 204,
              RFC 7934, DOI 10.17487/RFC7934, July 2016,
              <https://www.rfc-editor.org/info/rfc7934>.

   [RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli,
              "IPv6 Router Advertisement Options for DNS Configuration",
              RFC 8106, DOI 10.17487/RFC8106, March 2017,
              <https://www.rfc-editor.org/info/rfc8106>.

Authors' Addresses

   John Jason Brzozowski
   Comcast Cable
   1701 John F. Kennedy Blvd.
   Philadelphia, PA
   USA

   Email: john_brzozowski@cable.comcast.com

   Gunter Van De Velde
   Nokia
   Antwerp
   Belgium

   Email: gunter.van_de_velde@nokia.com