

v6ops
Internet-Draft
Intended status: Best Current Practice
Expires: November 18, 2017

J. Brzozowski
Comcast Cable
G. Van De Velde
Nokia
May 17, 2017

Unique IPv6 Prefix Per Host
draft-ietf-v6ops-unique-ipv6-prefix-per-host-03

Abstract

In some IPv6 environments, the need has arisen for hosts to be able to utilize a unique IPv6 prefix, even though the link or media may be shared. Typically hosts (subscribers) on a shared network, either wired or wireless, such as Ethernet, WiFi, etc., will acquire unique IPv6 addresses from a common IPv6 prefix that is allocated or assigned for use on a specific link.

In most deployments today, IPv6 address assignment from a single IPv6 prefix on a shared network is done by either using IPv6 stateless address auto-configuration (SLAAC) and/or stateful DHCPv6. While this is still viable and operates as designed, there are some large scale environments where this concept introduces significant performance challenges and implications, specifically related to IPv6 router and neighbor discovery.

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 IPv6 address from the service provider include improved subscriber isolation and enhanced subscriber management.

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 <http://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."

Internet-Draft

Unique IPv6 Prefix Per Host

May 2017

This Internet-Draft will expire on November 18, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Motivation and Scope of Applicability	3
3.	Design Principles	4
4.	IPv6 Unique Prefix Assignment	4
5.	IPv6 Neighbor Discovery Best Practices	5
6.	IANA Considerations	7
7.	Security Considerations	7
8.	Acknowledgements	7
9.	Normative References	7
	Authors' Addresses	8

[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 network service. 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 datacentre, 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 [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](#)]. As such the use of IPv6 SLAAC based subscriber and address management for provider managed shared network services is the recommended technology of choice, as it does not exclude any known IPv6 implementation. 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 will focus upon the process for UEs to obtain a unique IPv6 prefix.

[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 neighborhood discovery and IPv6 address management 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.

The Best Current Practice documented in this note is to provide a unique IPv6 prefix to hosts/subscribers devices connected to the provider managed shared network. Each unique IPv6 prefix can function as control-plane anchor point to make sure that each subscriber is receiving 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

[RFC6106](#) [[RFC6106](#)], and the remaining information required to establish globally routable IPv6 connectivity. For that purpose, the UE/subscriber sends a RS (Router Solicitation) message.

The First Hop Router receives this UE/subscriber RS message and starts the process to compose the response to the UE/subscriber originated RS message. The First Hop Provider Router will answer using a unicast RA (Router Advertisement) to the UE/subscriber. This RA contains two important parameters for the UE/subscriber to consume: (1) a /64 prefix and (2) flags. The /64 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 UE/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 (UE/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 UE/subscriber can configure itself using SLAAC)
- o L-flag = 0 (The UE/subscriber is off-link, which means that the UE/subscriber will ALWAYS send packets to its default gateway, even if the destination is within the range of the /64 prefix)

The use of a unique IPv6 prefix per UE 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 provider router. Further, the flag combinations documented above maximize 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 UE/subscriber gets its own unique /64 IPv6 prefix for which it can use SLAAC or any other method to select its /128 unique address. In addition it will use stateless DHCPv6 to get the IPv6 address of the DNS server, however it SHOULD NOT use stateful DHCPv6 to receive a service provider managed IPv6 address. If the UE/subscriber desires to send anything external including other UE/subscriber devices (assuming device to device communications is enabled and supported), then, due to the L-bit set, it SHOULD send this traffic to the First Hop Provider Router.

After the UE/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 UE/subscriber device itself, it will need to verify that the address is unique on the shared network. The UE/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 UE/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 UE/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 UE/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 IPv6 Router Advertisement Interval = 300s
- 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 UE/subscriber IPv6 SLAAC connectivity model provides a consideration to make regarding resource consumption (i.e. memory, neighbor state) on the First Hop Router. To reduce undesired resource consumption on the First Hop Router the desire is to remove UE/subscriber context in the case of non-permanent UE, such as in the case of WiFi hotspots as quickly as possible. A possible solution is to use a subscriber inactivity timer which, 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 utilize [RFC 4941](#) RFC4941 [RFC4941] temporary 'private' addresses.

Similarly, when using this technology in a datacentre, the UE server may need to use several addresses from the same /64, 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 UE/subscriber using Radius attributes.

[6.](#) IANA Considerations

No IANA considerations are defined at this time.

[7.](#) Security Considerations

No additional security considerations are made in this document.

8. Acknowledgements

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

Tim Chown, Lorenzo Colitti, Killian Desmedt, Brad Hilgenfeld, Wim Henderickx, Erik Kline, Warren Kumari, Thomas Lynn, Jordi Palet, Phil Sanderson, Colleen Szymanik, Eric Vyncke, Sanjay Wadhwa

9. Normative References

- [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, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<http://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, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", [RFC 6106](#), DOI 10.17487/RFC6106, November 2010, <<http://www.rfc-editor.org/info/rfc6106>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", [BCP 204](#), [RFC 7934](#), DOI 10.17487/RFC7934, July 2016, <<http://www.rfc-editor.org/info/rfc7934>>.

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