v6ops
Internet-Draft
Intended status: Best Current Practice
Expires: April 21, 2016

J. Brzozowski
Comcast Cable
G. Van De Velde
Alcatel-Lucent

October 19, 2015

Unique IPv6 Prefix Per Host
draft-ietf-v6ops-unique-ipv6-prefix-per-host-00

Abstract

In some IPv6 environments the need has arisen for hosts to be able to utilise a unique IPv6 prefix even though the link or media may be shared. Typically hosts (subscribers) on a shared network, like Wi-Fi or Ethernet, will acquire unique IPv6 addresses from a common IPv6 prefix that is allocated or assigned for use on a specific link. Benefits of a unique IPv6 prefix compared to a unique IPv6 address from the service provider are going from enhanced subscriber management to improved isolation between subscribers.

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).

Status of This Memo

This Internet-Draft will expire on April 21, 2016.

Copyright Notice

Table of Contents

1. Introduction

The concepts in this document were originally developed as part of a large scale, production deployment of IPv6 support for a community Wi-Fi 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. Details of IPv4 support are out of scope for this document. This document will also, in general, outline the requirements that must be satified by UE to allow for an IPv6 only experience.

In most deployments today User Equipment (UE) IPv6 address assignment is commonly done using either IPv6 SLAAC RFC4862 [RFC4862] and/or DHCP IA_NA RFC3315 [RFC3315]. However, at current time there is a non-trivial UE/subscriber base not supporting DHCPv6 IA_NA, making IPv6 SLAAC based subscriber and address management for community Wi-Fi services the technology of choice as it does not exclude any known IPv6 implementation. This document will detail the mechanics involved for IPv6 SLAAC based address and subscriber management coupled with stateless DHCPv6, where beneficial.

A community Wi-Fi service is an environment to allow subscribers (hosts) to connect to a shared network providing Internet and/or closed network services. Often Service providers use community Wi-Fi networks to provide enhanced subscriber connectivity experiences. Additionally retail owners frequently provide community Wi-Fi services to improve their customers retail experience.

Upon further exploration the approach documented here has applicability in other environments including corporate, enterprise, or university settings where IPv6 support is desired over a shared media. Where applicable details related to the same will be provided.

2. Motivation and Scope of Applicability

The motivation for this work falls into the following categories:

- Deploy support for IPv6 that will allow for an IPv6 only experience, even if IPv4 support is present

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

- 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 the ensure that the widest population of UE implementations can leverage the availability of IPv6.

- Lay the technological foundation for future work related to the use of IPv6 over shared media like Wi-Fi

While this work was originally conceived in the context of large scale Wi-Fi networks, the scope of applicability is much broader. The techniques and concepts or subsets of the same may also be applicable in residential or SOHO networking environments.

3. Design Princinples

The Wireless LAN Gateway (WLAN-GW) discussed in this document is the L3-Edge router responsible for the communication with the Wi-Fi subscribers (hosts) and to aggregate the traffic from the Wi-Fi subscribers and the Wireless LAN network towards the community Wi-Fi provider.

The goal of a WLAN-GW is to provide sufficient data-plane throughput capacity to aggregate all Wi-Fi subscriber traffic, while at the same time it is functioning as control-plane anchor point to make sure that each subscriber is receiving the expected subscriber policy and service levels (throughput, QoS, security, parental-control, subscriber mobility management, etc.).

The work detailed in this document intends to provide details regarding the WLAN-GW Wi-Fi subscriber/host addressing methodology. Evolved WLAN-GW capabilities regarding fixed/mobile convergence, traffic steering, etc. are not the main focus and are outside the scope of this document.

4. Behaviour

This section outlines the essential components of the described system and interaction amongst the same.

4.1. Community Wi-Fi Network Topology Description

The topology and design referenced in this document is a generalized description of functional components currently deployed in a large scale subscriber oriented network.

```
                                         +-----+
                                         | AAA |
                                         +-----+
                                            /
                                         Radius
                                          /
                                    +----+
                                    | CP |
                +----+              +---------+
    UE--802.11--| AP |---Soft_GRE---| WLAN-GW |----Internet/WAN access
                +----+              +---------+
                                         |
                                      IP/HTTP
                                         |
                              +---------------------+
                              | HTTP Captive Portal |
                              +---------------------+
```

Figure 1

- UE: User Equipment.

- 802.11: Wireless Network

- AP: Access Point.

- Soft-GRE: Stateless GRE tunnel

- WLAN-GW: Wireless LAN Gateway

- CP: Control Plane component of the WLAN-GW

- AAA: Accounting, Authorisation and Authentication

- HTTP Captive Portal: Captive portal used to redirect traffic towards during subscriber onboarding process

While there are many ways for UE to associate to a Wi-Fi network (e.g. EAP-SIM, EAP-AKA, WPA2-PSK, etc.), community Wi-Fi predominantly leverages an HTTP Captive Portal. The key function for the Captive Portal is to identify the UE/subscriber and create on the WLAN-GW the corresponding UE/subscriber context for policy and accounting.

The Soft-GRE session is a stateless GRE tunnel between AP and the WLAN-GW. The AP is configured with the IP address or FQDN of the tunnel concentrator or aggregation point and initiates the GRE

tunnel, over IPv6 preferrably, by encapsulating packets towards the WLAN-GW. The WLAN-GW is configured as a GRE tunnel head-end server and accepts these GRE packets, while at the same time creating correct tunnel context to identify the AP. Soft-GRE is a very well established pragmatic technology. The use of GRE over IPv4 only furthers an operators dependence on IPv4 and should be deprecated by using GRE over IPv6 only.

The AP has, as seen in the illustration, an interface attached to the Wi-Fi network and will bridge traffic received on this Wi-Fi interface over the Soft-GRE tunnel to the WLAN-GW. This will include traffic from newly attached UE/subscribers which have not been identified or authorized on the Wi-Fi network. At the same time the AP implements split-horizon for BUM (broadcast, unknown and multicast) traffic, making sure that there is no undesired leakage of traffic between UE/subscribers attached to the Wi-Fi network.

The Control Plane (CP) of the WLAN-GW is a key component used during onboarding of UE/subscribers to identify the UE/subscriber and to exchange IP address related details. For that purpose it can make usage of DHCP, ARP, DHCPv6, ICMPv6 (RS/RA/NS/NA), Radius, Diameter, etc.

4.2. Wi-Fi Subscriber Onboarding Procedures

This section provides detail about Best Practice operational steps to onboard a UE/subscriber and the key architectural technology used to create the WLAN-GW UE/subscriber policy and IP addressing context.

The flow chart pictured below is providing a sequential overview of the operational steps performed to onboard a UE onto a community Wi-Fi network.

```
   UE               WLAN-GW              AAA         Captive-Portal DNS
   |                   |                  |                |         |
   |                   |                  |                |         |
   |--------RS-------->|                  |                |         |
   |                   |---Access-Req---->|                |         |
   |                   |<--Access-Acc-----|                |         |
   |                   |(=Radius UE info) |                |         |
   |<-------RA---------|                  |                |         |
   |(/64; M,L=0; O,A=1)|                  |                |         |
   |                   |                  |                |         |
   |------NS(DAD)----->|                  |                |         |
   |                   |                  |                |         |
   |-DHCPv6(info Req)->|                  |                |         |
   | (Ask DNS IP addr.)|                  |                |         |
   |                   |                  |                |         |
   |<---DHCPv6 Reply---|                  |                |         |
   |                   |                  |                |         |
   |--------------------------------DNS----------------------------->|
   |                   |                  |                |         |
   |------HTTP GET---->|                  |                |         |
   |<--HTTP Redirect---|                  |                |         |
   |                   |                  |                |         |
   |--------------------------------DNS----------------------------->|
   |                   |                  |                |         |
   |<------------------------HTTP Login------------------->|         |
   |                   |                  |                |         |
   |                   |                  |<-UE Identified-|         |
   |                   |<--Radius CoA-----|                |         |
   |                   |(remove HTTP Red.)|                |         |
   |                   |                  |                |         |
   |<--------UE/Subscriber connected to Internet/WAN------------------->
   |                   |                  |                |         |
   |                   |                  |                |         |
```

Figure 2

Note that the Wireless Access Point (AP) is not pictured in the flow chart above. This is because the AP is from architectural perspective functioning as a L2 bridge between the UE and WLAN-GW. For Wi-Fi community service the AP is configured to setup a Soft-GRE tunnel towards the WLAN-GW and to bridge relevant Wi-Fi traffic upon the Soft-GRE tunnel. The AP is also configured for split-horizon towards the Wi-Fi interface for subscriber isolation and security purpose. The AP will for the remainder of this document be silently inserted between UE and WLAN-GW to bridge traffic between the WLAN-GW and AP and vice versa

When a new UE connects to the community Wi-Fi it connects to the Wi-Fi network by attaching to the relevant 'open' SSID advertised for use as part of the community Wi-Fi offering. Once the UE/subscriber is attached to the Wi-Fi SSID it will initiate IP configuration. The focus of this document is to share IPv6 address assignment Best Practices, and hence will focus around those topics, eventhough there are many more aspects to deploy a quality community Wi-Fi service offering successfully.

Once the UE is connected to the Wi-Fi shared network, it will from an IPv6 perspective attempt to learn the default IPv6 gateway, the IPv6 prefix information, the DNS information, and the remaining information required to establish globally routable IPv6 connectivity. For that purpose the the UE/subscriber sends a RS (Router Solicitation) message. This RS is forwarded by the AP-bridge over the Soft-GRE interface, however due to the split-horizon configuration for BUM traffic it is not relayed to any other UE/Subscribers attached to the Wi-Fi network.

The WLAN-GW received this UE/subscriber RS message, and because it is the first time this UE/subscriber attaches to the Wi-Fi the UE/subscriber is by default not authorized. The WLAN-GW will now try to discover additional information about the subscriber information by querying the AAA server. This is done by sending a Radius Access-Request.

The Radius server receives this Access-Request, and performs a lookup in its policy database. If radius server discovers that the UE/subscriber is a fresh device trying to gain access onto the Wi-Fi network it will identify some parameters (e.g. IPv6 /64 prefix) to send back to the WLAN-GW together with a message to install a HTTP-redirect to a Captive Portal for further UE/subscriber identification. This will be sent from the AAA server to the WLAN-GW in a Radius Access-Ackowledge message.

The WLAN-GW will use the received Radius information to compose the response to the UE/subscriber originated RS message. The WLAN-GW will answer using a unicast RA (Router Advertisement) to the UE/subscriber. This RA contains a few important parameters for the EU/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 WLAN-GW or from a pool signalled by the Radius server in a radius attribute. The flags may 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 community Wi-Fi are:

- 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 over Wi-Fi is desired)

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

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

- L-flag = 0 (The UE/subscriber is off-link, which means that the UE/subscriber will send packets ALWAYS to his 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 WLAN-GW. 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 WLAN-GW.

Now that 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 Wi-Fi network.

At this stage the UE/subscriber has acquired a valid IPv6 address, however it may not have received one or more DNS server IPv6 address. The UE/subscriber can use stateless DHCPv6 exchange to identify a valid DNS server address(es). An alternative solution, albeit less supported by IPv6 hosts is to signal DNS server addresses is by

utilising RA extensions described in RNDSS RFC6106 [RFC6106] in which the router uses Router Advertisement options to advertise a list of DNS recursive server addresses and a DNS Search List to IPv6 UE/ subscribers. The use of RNDSS and stateless DHCPv6 for the configuration of hosts are not mutually exclusive. Both methods can and should be enabled simultaneously allowing for the widest range of hosts or UEs to learn and use DNS over IPv6. DNS server IPv6 address(es) sent via DHCPv6 and RDNSS must be identical.

At this moment the UE/subscriber has all information to be connected to the Internet, nevertheless the community Wi-Fi service provider has no idea about the identity or credentials of the UE/subscriber. For that purpose the Service provider has installed on the WLAN-GW a HTTP redirect for this particular UE/subscriber towards HTTP captive portal. First the subscriber utilises DNS to correlate the domain name with an IP address, next the HTTP GET is intercepted and an HTTP Redirect is issued to Redirect the HTTP session towards the Captive portal. The ultimate goal of this process is for the service provider to identify the UE/subscriber. From the moment the UE/ subscriber identified itself on the captive portal (login-ID/PW, PIN Challenge, etc.) then the captive portal informs the WLAN-GW about the correct policies (QoS, policing, etc.) and to remove the HTTP-redirect.

From now onwards the WLAN-GW has identified the UE/subscriber and installed all the subscriber context for identification, billing, traffic conditioning. The UE/subscriber can access the Internet/WAN within his agreed commuity Wi-Fi parameters.

4.3. UE IPv6 Addressing and Configuration

An over arching objective for any IPv6 deployment where subscriber endpoints or UEs are concerned must include an IPv6 only experience. Specifically, similar to residential broadband networks, Wi-Fi networks that support IPv6 must ensure there are no dependencies on IPv4. Due to fragmented support for various IPv6 address and configuration mechanisms network operators must effectively enable and support every combination of IPv6 address and configuration technique. Coordinating the configuraiton and values for the same is important to ensure proper UE behavior.

4.3.1. IPv6 Addressing

Stateless IPv6 address autoconfiguration is expected to be the primary mechanism for UEs to leverage when establing globally routable IPv6 connectivity. Stateful DHCPv6 is currently not utilized in this model for host addressing since stateful DHCPv6 is not universally supported for address acquisition. Stateful DHCPv6

may be considering in the future as part of enabling support for IPv6 prefix delegation [RFC3633].

4.3.2. IPv6 Configuration

In order to make an IPv6 only experience possible Wi-Fi network operators must ensure that UEs are able to reach all critical network services over IPv6. Today, many host operating systems still prefer querying DNS over IPv4. Additionally, widely deployed hosts do not truly leverage a single common approach for IPv6 configuration. As such the following should be expected to be available to support a proper IPv6 only configuration:

- RDNSS [RFC6106] is enabled by default and is expected to contain one or more globally routable IPv6 addresses

- Stateless DHCPv6 [RFC3315] is enabled by default and will minimally transmit one or more DNS server IPv6 addresses. To ensure the desired behavior is triggered IPv6 router advertisements transmitted by the WLAN-GW will set the M flag to 0 and the O flag to 1.

5. Operational Considerations

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 WLAN-GW extends these lifetimes by sending an unsolicited RA, the applicable MaxRtrAdvInterval on the WLAN-GW MUST therefore be lower than the preferred lifetime. As a consequence of this process is that the WLAN-GW never knows when a UE/subscriber stops using addresses from a prefix and additional procedures are required to help the WLAN-GW to gain this information. When using stateful DHCPv6 IA_NA for IPv6 UE/subscriber address assignment this uncertainty on the WLAN-GW is not of impact due to the stateful nature of DHCPv6 IA_NA address assignment.

Following is reference table of the key IPv6 router discovery and neighbor discovery timers:

- IPv6 Router Advertisement Interval = 300s

- IPv6 Router LifeTime = 3600s

- Reachable time = 30s

- IPv6 Valid Lifetime = 3600s

- IPv6 Preferred Lifetime = 1800s

- Retransmit timer = 0s

The stateless nature of the UE/subscriber IPv6 SLAAC connectivity model provides value to make sure that the UE/subscriber context is timely removed from the WLAN-GW to avoid ongoing resource depletion. A possible solution is to use a subscriber inactivity timer which after tracking a pre-defined (currently unspecified) # of minutes deletes the subscriber context on the WLAN-GW.

When using SLAAC the UE/subscriber the IP address assignment happens without a WLAN-GW controlled state machine, and as result there is no state-information on the WLAN-GW about actual IPv6 address usage. To accomodate this the WLAN-GW can periodically perform a Subscriber Host Connectivity Verification (i.e. periodically ping each IPv6 UE/subscriber from the WLAN-GW) to make sure that the subscriber table on the WLAN-GW is correct and that the inactive UE/subscribers are removed.

When employing stateless IPv6 address assignment a number of widely deployed operating systems will attempt to utilize RFC 4941 RFC4941 [RFC4941] temporary 'private' addresses. This can lead to the consequence that a UE has multiple /128 addresses from the same IPv6 prefix. The WLAN-GW MUST be able to handle the presence and use of multiple globally routable IPv6 addresses.

When geo-localisation is of importance the WLAN-GW needs to have information about the Access Point to which the UE/subscriber is connected. In an environment using DHCPv6 IA_NA for IPv6 address assignment this is achieved by having the AP insert an interface-id RFC3315 [RFC3315] in the UE/subscriber DHCPv6 Solicit message. The interface-id format expected is [ap-mac;ssid;[o-s]], e.g. [00:11:22:33:44:55;example;o] (o stands for open, s for secure). This way the service provider can learn both the AP-MAC (identifies location) and the SSID (identifies service). When a service provider uses SLAAC IPv6 address assignment it becomes harder for the service provider to rely on this type of information and alternate solutions have to be used to acquire the MAC address of the Access Point to which the UE/subscriber is connected. A solution could be for the WLAN-GW to support NSoGRE to harvest the Access-Point MAC address to which the UE/subscriber is connected.

For security purposes it will be important for the service provider to have the capability on the WLAN-GW to have supported mechanics for LI (Lawfull Intercept) and the installation of IPv6 filters per subscriber.

For accounting purposes the WLAN-GW must be able to send usage statistics per UE/subscriber using Radius attributes.

6. Future work

- Support for IPv6 prefix delegation over Wi-Fi

7. IANA Considerations

No IANA considerations are defined at this time.

8. Security Considerations

No Additional Security Considerations are made in this document.

9. Acknowledgements

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

Lorenzo Colitti, Killian Desmedt, Brad Hilgenfeld, Wim Henderickx, Erik Kline, Thomas Lynn, Phil Sanderson, Colleen Szymanik, Sanjay Wadhwa

10. References

10.1. 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>.

[RFC6180] Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", RFC 6180, DOI 10.17487/RFC6180, May 2011, <http://www.rfc-editor.org/info/rfc6180>.

10.2. Informative References

[I-D.ietf-v6ops-v4v6tran-framework]
Carpenter, B., Jiang, S., and V. Kuarsingh, "Framework for IP Version Transition Scenarios", draft-ietf-v6ops-v4v6tran-framework-02 (work in progress), July 2011.

[RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, DOI 10.17487/RFC6343, August 2011, <http://www.rfc-editor.org/info/rfc6343>.

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
Alcatel-Lucent
Antwerp
Belgium

Email: gunter.van_de_velde@alcatel-lucent.com