

Recommendations for IPv6 in
Third Generation Partnership Project (3GPP) Standards

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This document contains recommendations from the Internet Engineering Task Force (IETF) IPv6 Working Group to the Third Generation Partnership Project (3GPP) community regarding the use of IPv6 in the 3GPP standards. Specifically, this document recommends that the 3GPP specify that multiple prefixes may be assigned to each primary PDP context, require that a given prefix must not be assigned to more than one primary PDP context, and allow 3GPP nodes to use multiple identifiers within those prefixes, including randomly generated identifiers.

The IPv6 Working Group supports the use of IPv6 within 3GPP and offers these recommendations in a spirit of open cooperation between the IPv6 Working Group and the 3GPP community. Since the original publication of this document as an Internet-Draft, the 3GPP has adopted the primary recommendations of this document.

Conventions Used In This Document

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 BCP 14, RFC 2119 [KEYWORD].

Table of Contents

1	Introduction.....	2
1.1	What is the 3GPP?.....	3
1.2	What is the IETF?.....	4
1.3	Terminology.....	4
1.3.1	3GPP Terminology.....	4
1.3.2	IETF Terminology.....	5
1.4	Overview of the IPv6 Addressing Architecture.....	6
1.5	An IP-Centric View of the 3GPP System.....	7
1.5.1	Overview of the UMTS Architecture.....	7
1.5.2	The PDP Context.....	10
1.5.3	IPv6 Address Autoconfiguration in GPRS.....	11
2	Recommendations to the 3GPP.....	13
2.1	Limitations of 3GPP Address Assignment.....	13
2.2	Advertising Multiple Prefixes.....	14
2.3	Assigning a Prefix to Only One Primary PDP Context.....	14
2.3.1	Is a /64 per PDP Context Too Much?.....	15
2.3.2	Prefix Information in the SGSN.....	16
2.4	Multiple Identifiers per PDP Context.....	16
3	Additional IPv6 Work Items.....	16
4	Security Considerations.....	17
	Appendix A: Analysis of Findings.....	18
	Address Assignment Solutions.....	18
	References.....	19
	Authors and Acknowledgements.....	22
	Editor's Address.....	22
	Full Copyright Statement.....	23

1. Introduction

In May 2001, the IPv6 Working Group (WG) held an interim meeting in Redmond, WA to discuss the use of IPv6 within the 3GPP standards. The first day of the meeting was a joint discussion with 3GPP, during which an architectural overview of 3GPP's usage of IPv6 was presented, and there was much discussion regarding particular aspects of IPv6 usage within 3GPP. At that meeting, a decision was made to form a design team to write a document offering advice from the IPv6 WG to the 3GPP community, regarding their use of IPv6. This document is the result of that effort.

This document offers recommendations to the 3GPP community from the IETF IPv6 Working Group. It is organized into three main sections:

1. An introduction (this section) that provides background information regarding the IETF IPv6 WG and the 3GPP and includes a high-level overview of the technologies discussed in this document.

2. Recommendations from the IPv6 WG to the 3GPP community. These can be found in section 2.
3. Further work items that should be considered by the IPv6 WG. These items are discussed in section 3.

It is the purpose of this document to provide advice from the IPv6 Working Group to the 3GPP community. We have limited the contents of this document to items that are directly related to the use of IPv6 within 3GPP. This document defines no standards, and it is not a definitive source of information regarding IPv6 or 3GPP. We have not chosen to explore 3GPP-related issues with other IETF protocols (i.e., SIP, IPv4, etc.), as they are outside the scope of the IPv6 Working Group.

The IPv6 Working Group fully supports the use of IPv6 within 3GPP, and we encourage 3GPP implementers and operators to participate in the IETF process. We are offering these suggestions in a spirit of open cooperation between the IPv6 Working Group and the 3GPP community, and we hope that our ongoing cooperation will help to strengthen both sets of standards.

The 3GPP address allocation information in this document is based on the 3GPP document TS 23.060 version 4.1.0 [OLD-TS23060]. At the 3GPP plenary meeting TSG #15 in March 2002, the 3GPP adopted the two primary recommendations contained in this document, allocating a unique prefix to each primary PDP context when IPv6 stateless address autoconfiguration is used, and allowing the terminals to use multiple interface identifiers. These changes were retroactively applied from 3GPP release 99 onwards, in TS23.060 versions 3.11.0, 4.4.0 and 5.1.0 [NEW-TS23060].

1.1 What is the 3GPP?

The Third Generation Partnership Project (3GPP) is a global standardization partnership founded in late 1998. Its Organizational Partners have agreed to co-operate in the production of technical specifications for a Third Generation Mobile System, based on the evolved GSM core networks.

The 3GPP Organizational Partners consist of several different standardization organizations: ETSI from Europe, Standards Committee T1 Telecommunications (T1) in the USA, China Wireless Telecommunication Standard Group (CWTS), Korean Telecommunications Technology Association (TTA), the Association of Radio Industries and Businesses (ARIB), and the Telecommunication Technology Committee(TTC) in Japan.

The work is coordinated by a Project Co-ordination Group (PCG), and structured into Technical Specification Groups (TSGs). There are five TSGs: Core Network (TSG CN), Radio Access Networks (TSG RAN), Services and System Aspects (TSG SA), GSM/EDGE Radio Access Network (GERAN), and the Terminals (TSG T). The TSGs are further divided into Working Groups (WGs). The technical work is done in the working groups, and later approved in the TSGs.

3GPP working methods are different from IETF working methods. The major difference is where the majority of the work is done. In 3GPP, the work is done in face-to-face meetings, and the mailing list is used mainly for distributing contributions, and for handling documents that were not handled in the meeting, due to lack of time. Decisions are usually made by consensus, though voting does exist. However, it is rather rare to vote. 3GPP documents are public and can be accessed via the 3GPP web site [3GPP-URL].

1.2 What is the IETF?

The Internet Engineering Task Force (IETF) is a large, open, international community of network designers, operators, vendors, and researchers, concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF is also the primary standards body developing Internet protocols and standards. It is open to any interested individual. More information about the IETF can be found at the IETF web site [IETF-URL].

The actual technical work of the IETF is done in working groups, organized by topic into several areas (e.g., routing, transport, security, etc.). The IPv6 Working Group is chartered within the Internet area of the IETF. Much of the work is handled via mailing lists, and the IETF holds meetings three times per year.

1.3 Terminology

This section defines the 3GPP and IETF terminology used in this document. The 3GPP terms and their meanings have been taken from [TR21905].

1.3.1 3GPP Terminology

APN	Access Point Name. The APN is a logical name referring to a GGSN and an external network.
CS	Circuit Switched
GERAN	GSM/EDGE Radio Access Network

GGSN	Gateway GPRS Support Node. A router between the GPRS network and an external network (i.e., the Internet).
GPRS	General Packet Radio Services
GTP-U	General Tunneling Protocol - User Plane
MT	Mobile Termination. For example, a mobile phone handset.
PDP	Packet Data Protocol
PDP Context	A PDP connection between the UE and the GGSN.
PS	Packet Switched
SGSN	Serving GPRS Support Node
TE	Terminal Equipment. For example, a laptop attached through a 3GPP handset.
UE	User Equipment (TE + MT + USIM). An example would be a mobile handset with a USIM card inserted and a laptop attached.
UMTS	Universal Mobile Telecommunications System
USIM	Universal Subscriber Identity Module. Typically, a card that is inserted into a mobile phone handset.
UTRAN	Universal Terrestrial Radio Access Network

1.3.2 IETF Terminology

IPv6	Internet Protocol version 6 [RFC 2460]
NAS	Network Access Server
NAT	Network Address Translator
NAT-PT	Network Address Translation with Protocol Translation. An IPv6 transition mechanism. [NAT-PT]
PPP	Point-to-Point Protocol [PPP]
SIIT	Stateless IP/ICMP Transition Mechanism [SIIT]

1.4 Overview of the IPv6 Addressing Architecture

The recommendations in this document are primarily related to IPv6 address assignment. To fully understand the recommended changes, it is necessary to understand the IPv6 addressing architecture, and current IPv6 address assignment mechanisms.

The IPv6 addressing architecture represents a significant evolution from IPv4 addressing [ADDRARCH]. It is required that all IPv6 nodes be able to assemble their own addresses from interface identifiers and prefix information. This mechanism is called IPv6 Host Autoconfiguration [AUTOCONF], and it allows IPv6 nodes to configure themselves without the need for stateful configuration servers (i.e., DHCPv6) or statically configured addresses.

Interface identifiers can be globally unique, such as modified EUI-64 addresses [ADDRARCH], or non-unique, such as randomly generated identifiers. Hosts that have a globally unique identifier available may also choose to use randomly generated addresses for privacy [PRIVADDR] or for other reasons. IPv6 hosts are free to generate new identifiers at any time, and Duplicate Address Detection (DAD) is used to protect against the use of duplicate identifiers on a single link [IPV6ND].

A constant link-local prefix can be combined with any interface identifier to build an address for communication on a locally attached link. IPv6 routers may advertise additional prefixes (site-local and/or global prefixes)[IPV6ND]. Hosts can combine advertised prefixes with their own interface identifiers to create addresses for site-local and global communication.

IPv6 introduces architectural support for scoped unicast addressing [SCOPARCH]. A single interface will typically have multiple addresses for communication within different scopes: link-local, site-local and/or global [ADDRARCH]. Link-local addresses allow for local communication, even when an IPv6 router is not present. Some IPv6 protocols (i.e., routing protocols) require the use of link-local addresses. Site-local addressing allows communication to be administratively contained within a single site. Link-local or site-local connections may also survive changes to global prefix information (e.g., site renumbering).

IPv6 explicitly associates each address with an interface. Multiple-interface hosts may have interfaces on more than one link or in more than one site. Links and sites are internally identified using zone identifiers. Proper routing of non-global traffic and proper address selection are ensured by the IPv6 scoped addressing architecture [SCOPARCH].

IPv6 introduces the concept of privacy addresses [PRIVADDR]. These addresses are generated from an advertised global prefix and a randomly generated identifier, and are used for anonymous access to Internet services. Applications control the generation of privacy addresses, and new addresses can be generated at any time.

The IPv6 site renumbering specification [SITEREN] relies upon the fact that IPv6 nodes will generate new addresses when new prefixes are advertised on the link, and that they will deprecate addresses that use deprecated prefixes.

In the future, additional IPv6 specifications may rely upon the ability of IPv6 nodes to use multiple prefixes and/or multiple identifiers to dynamically create new addresses.

1.5 An IP-Centric View of the 3GPP System

The 3GPP specifications define a Third Generation Mobile System. An overview of the packet switched (PS) domain of the 3GPP Release 99 system is described in the following sections. The authors hope that this description is sufficient for the reader who is unfamiliar with the UMTS packet switched service, to understand how the UMTS system works, and how IPv6 is currently defined to be used within it.

1.5.1 Overview of the UMTS Architecture

The UMTS architecture can be divided into two main domains -- the packet switched (PS) domain, and the circuit switched (CS) domain. In this document, we will concentrate on the PS domain, or General Packet Radio Services (GPRS).




Figure 1: Simplified GPRS Architecture




Figure 2: GPRS Protocol Stacks




Figure 3: Laptop Attached to 3GPP Handset

The GPRS core network elements, shown in Figures 1 and 2, are the User Equipment (UE), Serving GPRS Support Node (SGSN), and Gateway GPRS Support Node (GGSN). The UTRAN comprises Radio Access Network Controllers (RNC) and the UTRAN base stations.

GGSN: A specialized router that functions as the gateway between the GPRS network and the external networks, e.g., Internet. It also gathers charging information about the connections. In many ways, the GGSN is similar to a Network Access Server (NAS).

SGSN: The SGSN's main functions include authentication, authorization, mobility management, and collection of billing information. The SGSN is connected to the SS7 network and through that, to the Home Location Register (HLR), so that it can perform user profile handling, authentication, and authorization.

GTP-U: A simple tunnelling protocol running over UDP/IP and used to route packets between RNC, SGSN and GGSN within the same, or between different, UMTS backbone(s). A GTP-U tunnel is identified at each end by a Tunnel Endpoint Identifier (TEID).

Only the most significant elements of the GPRS system are discussed in this document. More information about the GPRS system can be found in [OLD-TS23060].

1.5.2 The PDP Context

The most important 3GPP concept in this context is a PDP Context. A PDP Context is a connection between the UE and the GGSN, over which the packets are transferred. There are two kinds of PDP Contexts -- primary, and secondary.

The primary PDP Context initially defines the link to the GGSN. For instance, an IP address is assigned to each primary PDP Context. In addition, one or more secondary PDP Contexts can be added to a primary PDP Context, sharing the same IP address. These secondary PDP Contexts can have different Quality of Service characteristics than the primary PDP Context.

Together, a primary PDP Context and zero or more secondary PDP Contexts define, in IETF terms, a link. GPRS links are point-to-point. Once activated, all PDP contexts have equal status, meaning that a primary PDP context can be deleted while keeping the link between the UE and the GGSN, as long as there are other (secondary) PDP contexts active for the same IP address.

There are currently three PDP Types supported in GPRS -- IPv4, IPv6, and PPP. This document will only discuss the IPv6 PDP Type.

There are three basic actions that can be performed on a PDP Context: PDP Context Activation, Modification, and Deactivation. These actions are described in the following.

Activate PDP Context

Opens a new PDP Context to a GGSN. If a new primary PDP Context is activated, there is a new link created between a UE and a GGSN. A UE can open multiple primary PDP Contexts to one or more GGSNs.

Modify PDP Context

Changes the characteristics of a PDP Context, for example QoS attributes.

Deactivate PDP Context

Deactivates a PDP Context. If a primary PDP Context and all secondary PDP contexts associated with it are deactivated, a link between the UE and the GGSN is removed.

The APN is a name which is logically linked to a GGSN. The APN may identify a service or an external network. The syntax of the APN corresponds to a fully qualified domain name. At PDP context activation, the SGSN performs a DNS query to find out the GGSN(s) serving the APN requested by the terminal. The DNS response contains a list of GGSN addresses from which the SGSN selects one (in a round-robin fashion).




Figure 3: Correspondence of PDP Contexts to IPv6 Links

1.5.3 IPv6 Address Autoconfiguration in GPRS

GPRS supports static and dynamic address allocation. Two types of dynamic address allocation are supported -- stateless, and stateful. Stateful address configuration uses an external protocol to connect to a server that gives the IP address, e.g., DHCP.

The stateless IPv6 autoconfiguration works differently in GPRS than in Ethernet networks. GPRS nodes have no unique identifier, whereas Ethernet nodes can create an identifier from their EUI-48 address. Because GPRS networks are similar to dialup networks, the stateless address autoconfiguration in GPRS was based on PPPv6 [PPPv6].

3GPP address autoconfiguration has the following steps:

1. The Activate PDP Context message is sent to the SGSN (PDP Type=IPv6, PDP Address = 0, etc.).
2. The SGSN sends a Create PDP Context message to the GGSN with the above parameters.
3. GGSN chooses an interface identifier for the PDP Context and creates the link-local address. It answers the SGSN with a Create PDP Context response (PDP Address = link-local address).
4. The SGSN sends an Activate PDP Context accept message to the UE (PDP Address = link-local address).
5. The UE keeps the link-local address, and extracts the interface identifier for later use. The UE may send a Router Solicitation message to the GGSN (first hop router).
6. After the PDP Context Activation, the GGSN sends a Router Advertisement to the UE.
7. The UE should be configured not to send a Neighbor Solicitation message. However, if one is sent, the GGSN will silently discard it.
8. The GGSN updates the SGSN with the whole IPv6 address.

Each connected handset or laptop will create a primary PDP context for communication on the Internet. A handset may create many primary and/or secondary PDP contexts throughout the life of its connection with a GGSN.

Within 3GPP, the GGSN assigns a single 64-bit identifier to each primary PDP context. The GGSN also advertises a single /64 prefix to the handset, and these two items are assembled into a single IPv6 address. Later, the GGSN modifies the PDP context entry in the SGSN to include the whole IPv6 address, so that the SGSN can know the single address of each 3GPP node (e.g., for billing purposes). This address is also used in the GGSN to identify the PDP context associated with each packet. It is assumed that 3GPP nodes will not

generate any addresses, except for the single identifier/prefix combination assigned by the GGSN. DAD is not performed, as the GGSN will not assign the same address to multiple nodes.

2 Recommendations to the 3GPP

In the spirit of productive cooperation, the IPv6 Working Group recommends that the 3GPP consider three changes regarding the use of IPv6 within GPRS. Specifically, we recommend that the 3GPP:

1. Specify that multiple prefixes may be assigned to each primary PDP context,
2. Require that a given prefix must not be assigned to more than one primary PDP context, and
3. Allow 3GPP nodes to use multiple identifiers within those prefixes, including randomly generated identifiers.

Making these changes would provide several advantages for 3GPP implementers and users:

Laptops that connect to 3GPP handsets will work without any software changes. Their implementation of the standard IPv6 over PPP, address assignment, and autoconfiguration mechanisms will work without any modification. This will eliminate the need for vendors and operators to build and test special 3GPP drivers and related software. As currently specified, the 3GPP standards will be incompatible with laptop implementations that generate their own identifiers for privacy or other purposes.

IPv6 software implementations could be used in 3GPP handsets without any modifications to the IPv6 protocol mechanisms. This will make it easier to build and test 3GPP handsets.

Applications in 3GPP handsets will be able to take advantage of different types of IPv6 addresses (e.g., static addresses, temporary addresses for privacy, site-scoped addresses for site only communication, etc.)

The GPRS system will be better positioned to take advantage of new IPv6 features that are built around the current addressing architecture.

2.1 Limitations of 3GPP Address Assignment

The current 3GPP address assignment mechanism has the following limitations:

The GGSN only advertises a single /64 prefix, rather than a set of prefixes. This will prevent the participation of 3GPP nodes (e.g., handsets or 3GPP-attached laptops) in IPv6 site renumbering, or in other mechanisms that expect IPv6 hosts to create addresses based on multiple advertised prefixes.

A 3GPP node is assigned a single identifier and is not allowed to generate additional identifiers. This will prevent the use of privacy addresses by 3GPP nodes. This also makes 3GPP mechanisms not fully compliant with the expected behavior of IPv6 nodes, which will result in incompatibility with popular laptop IPv6 stacks. For example, a laptop that uses privacy addresses for web browser connections could not currently establish a web browser connection over a 3GPP link.

These limitations could be avoided by enabling the standard IPv6 address allocation mechanisms in 3GPP nodes. The GGSN could advertise one or more prefixes for the local link in standard IPv6 Router Advertisements, and IPv6 addresses could be assembled, as needed, by the IPv6 stack on the handset or laptop. An interface identifier could still be assigned by the GGSN, as is currently specified in the 3GPP standards. However, the handset or laptop could generate additional identifiers, as needed for privacy or other reasons.

2.2 Advertising Multiple Prefixes

For compliance with current and future IPv6 standards, the IPv6 WG recommends that the 3GPP allow multiple prefixes to be advertised for each primary PDP context. This would have several advantages, including:

3GPP nodes could participate in site renumbering and future IPv6 mechanisms that rely on the use of multiple global prefixes on a single link.

Site-local prefixes could be advertised on 3GPP links, if desired, allowing for site-constrained communication that could survive changes to global prefix information (e.g., site renumbering).

2.3 Assigning a Prefix to Only One Primary PDP Context

The IPv6 WG recommends that the 3GPP treat a primary PDP context, along with its secondary PDP contexts, as a single IPv6 link, and that the GGSN view each primary PDP context as a single subnet. Accordingly, a given global (or site-local) prefix should not be assigned to more than one PDP context.

Because multiple IPv6 hosts may attach through a 3GPP handset, the IPv6 WG recommends that one or more /64 prefixes should be assigned to each primary PDP context. This will allow sufficient address space for a 3GPP-attached node to allocate privacy addresses and/or route to a multi-link subnet [MULTLINK], and will discourage the use of NAT within 3GPP-attached devices.

2.3.1 Is a /64 per PDP Context Too Much?

If an operator assigns a /64 per PDP context, can we be assured that there is enough address space for millions of mobile devices? This question can be answered in the positive using the Host Density (HD) Ratio for address assignment efficiency [HD]. This is a measure of the number of addresses that can practically and easily be assigned to hosts, taking into consideration the inefficiencies in usage resulting from the various address assignment processes. The HD ratio was empirically derived from actual telephone number and data network address assignment cases.

We can calculate the number of easily assignable /64's making the following assumptions:

An HD ratio of 0.8 (representing the efficiency that can be achieved with no particular difficulty).

Only addresses with the 3-bit prefix 001 (the Aggregatable Global Unicast Addresses defined by RFC 2373) are used, resulting in 61 bits of assignable address space.

Using these assumptions, a total of 490 trillion (490×10^{12}) /64 prefixes can be assigned. This translates into around 80,000 PDP Contexts per person on the earth today. Even assuming that a majority of these IPv6 /64 prefixes will be used by non-3GPP networks, there is still clearly a sufficient number of /64 prefixes.

Given this, it can be safely concluded that the IPv6 address space will not be exhausted if /64 prefixes are allocated to primary PDP contexts.

For more information regarding policies for IPv6 address assignment, refer to the IAB/IESG recommendations regarding address assignment [IABAA], and the APNIC, ARIN and RIPE address allocation policy [AAPOL].

2.3.2 Prefix Information in the SGSN

Currently, the 3GPP standards allow only one prefix and one identifier for each PDP context. So, the GGSN can send a single IPv6 address to the SGSN, to be used for billing purposes, etc.

Instead of using the full IPv6 address to identify a PDP context, the IPv6 WG recommends that the SGSN be informed of each prefix that is currently assigned to a PDP context. By assigning a prefix to only one primary PDP context, the SGSN can associate a prefix list with each PDP context.

2.4 Multiple Identifiers per PDP Context

The IPv6 WG also recommends that the 3GPP standards be modified to allow multiple identifiers, including randomly generated identifiers, to be used within each assigned prefix. This would allow 3GPP nodes to generate and use privacy addresses, and would be compatible with future IPv6 standards that may depend on the ability of IPv6 nodes to generate new interface identifiers for communication.

This is a vital change, necessary to allow standards-compliant IPv6 nodes to connect to the Internet through 3GPP handsets, without modification. It is expected that most IPv6 nodes, including the most popular laptop stacks, will generate privacy addresses. The current 3GPP specifications will not be compatible with those implementations.

3 Additional IPv6 Work Items

During our work on this document, we have discovered several areas that could benefit from further informational or standards-track work within the IPv6 Working Group.

The IPv6 WG should work to define a point-to-point architecture and specify how the standard IPv6 address assignment mechanisms are applicable to IPv6 over point-to-point links. We should also review and clarify the IPv6 over PPP specification [PPP] to match the current IPv6 addressing architecture [ADDRARCH].

The IPv6 WG should consider publishing an "IPv6 over PDP Contexts" (or similar) document. This document would be useful for developers writing drivers for IPv6 stacks to work over 3GPP PDP Contexts.

The IPv6 working group should undertake an effort to define the minimal requirements for all IPv6 nodes.

4 Security Considerations

This document contains recommendations on the use of the IPv6 protocol in 3GPP standards. It does not specify a protocol, and it introduces no new security considerations.

Appendix A: Analysis of Findings

This section includes some analysis that may be useful to understanding why the IPv6 working group is making the above recommendations. It also includes some other options that were explored, and the reasons why those options were less suitable than the recommendations outlined above.

A.1 Address Assignment Solutions

In order to allow for the configuration and use of multiple IPv6 addresses per primary PDP Context having different interface identifiers, some modifications to the current 3GPP specifications would be required.

The solutions to achieve this were evaluated against the following factors:

- Scarcity and high cost of wireless spectrum
- Complexity of implementation and state maintenance
- Stability of the relevant IETF standards
- Impact on current 3GPP standards

Two solutions to allow autoconfiguration of multiple addresses on the same primary PDP Context were considered:

1. Assign one or more entire prefixes (/64s) to a PDP Context upon PDP Context activation and allow the autoconfiguration of multiple addresses.
 - a) The assignment may be performed by having the GGSN advertise one or more /64 prefixes to the mobile device.
 - b) The assignment may be performed by building "prefix delegation" functionality into the PDP Context messages or by using layer 3 mechanisms such as [PREFIXDELEGATION]. In this way, the prefix is not assigned to the link between the GGSN and the mobile device (as in 1a), but it is assigned to the mobile device itself. Note that [PREFIXDELEGATION] cannot be considered stable and has not, at this stage, been adopted by the IPv6 WG as a WG document.
2. Share the same prefix between multiple PDP Contexts connected to the same GGSN (and APN). Given that mobile devices may generate multiple addresses using more than one interface identifier, this would require DAD for the newly generated addresses over the air interface, and a proxy DAD, function which would increase the complexity and the amount of state to

be kept in the GGSN. Also, the GGSN would need to determine when the temporary addresses are no longer in use, which would be difficult. One possible solution could be using periodic unicast neighbor solicitations for the temporary addresses [IPv6ND].

Considering all the factors when evaluating the solutions, the recommendation is to use Solution 1a. This solution requires the least modification to the current 3GPP standards and maintains all the advantages of the other solutions.

Effectively, this would mean that each APN in a GGSN would have a certain number of /64 prefixes that can be handed out at PDP context Activation, through Router Advertisements. Therefore, instead of using the full IPv6 address to identify a primary PDP context, the IPv6 WG recommends that the GGSN use the entire prefix (together with other 3GPP specific information) and that the SGSN be informed of the prefixes that are assigned to a PDP context. By assigning a given prefix to only one primary PDP context, the GGSN and SGSN can associate a prefix list with each PDP context, as needed.

Note that the recommended solution does not imply or assume that the mobile device is a router. The MT is expected to use the /64 for itself and may also use this prefix for devices attached to it. However, this is not necessary if each device behind the MT is connected to a separate primary PDP Context and therefore can use a /64, which is not shared with other devices. The MT is also expected to handle DAD locally for devices attached to it (e.g., laptops) without forwarding Neighbor Solicitations over the air to the GGSN.

References

- [OLD-TS23060] TS 23.060, "General Packet Radio Service (GPRS); Service description; Stage 2", V4.1.0
- [NEW-TS23060] TS 23.060 version 3.11.0 (release 99), 4.4.0 (release 4) and 5.1.0 (release 5).
- [3GPP-URL] <http://www.3gpp.org>
- [IETF-URL] <http://www.ietf.org>
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996
- [KEYWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1999.

- [TR21905] 3GPP TR 21.905, "Vocabulary for 3GPP Specifications", V5.0.0
- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [NAT-PT] Tsirtsis, G. and P. Shrisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [PPP] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [SIIT] Nordmark, N., "Stateless IP/ICMP Translation Algorithm", RFC 2765, February 2000.
- [ADDRARCH] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [IPV6ND] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [AUTOCONF] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998
- [PRIVADDR] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [IPV6ETH] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [PPPV6] Haskin, D. and E. Allen, "IP Version 6 over PPP", RFC 2472, December 1998.
- [MULTLINK] C. Huitema, D. Thaler, "Multi-link Subnet Support in IPv6", Work in Progress.
- [SITEREN] C. Huitema, "IPv6 Site Renumbering", Work in Progress.
- [HD] Durand, A. and C. Huitema, "The Host-Density Ratio for Address Assignment Efficiency: An update on the H ratio", RFC 3194, November 2001.
- [IABAA] IAB, IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

- [AAPOL] APNIC, ARIN, RIPE-NCC, "IPv6 Address Allocation and Assignment Global Policy", Work in Progress.
- [SCOPARCH] S. Deering, et. al., "IPv6 Scoped Address Architecture", Work in Progress.
- [CELLREQ] J. Arkko, et. al., "Minimum IPv6 Functionality for a Cellular Host", Work in Progress.
- [PREFDEL] J. Martin, B. Haberman, "Automatic Prefix Delegation Protocol for Internet Protocol Version 6 (IPv6)", Work in Progress.

Authors and Acknowledgements

This document was written by the IPv6 3GPP design team:

Steve Deering, Cisco Systems
EMail: deering@cisco.com

Karim El-Malki, Ericsson Radio Systems
EMail: Karim.El-Malki@era.ericsson.se

Paul Francis, Tahoe Networks
EMail: francis@tahoenetworks.com

Bob Hinden, Nokia
EMail: hinden@iprg.nokia.com

Christian Huitema, Microsoft
EMail: huitema@windows.microsoft.com

Niall Richard Murphy, Hutchison 3G
EMail: niallm@enigma.ie

Markku Savela, Technical Research Centre of Finland
Email: Markku.Savela@vtt.fi

Jonne Soininen, Nokia
EMail: Jonne.Soininen@nokia.com

Margaret Wasserman, Wind River
EMail: mrw@windriver.com

Information was incorporated from a presentation co-authored by:

Juan-Antonio Ibanez, Ericsson Eurolab

Editor's Address

Comments or questions regarding this document should be sent to:

Margaret Wasserman
Wind River
10 Tara Blvd., Suite 330
Nashua, NH 03062 USA

Phone: (603) 897-2067
EMail: mrw@windriver.com

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

