                Transmission of IPv6 via the IPv6 Convergence Sublayer
                        over IEEE 802.16 Networks

Status of This Memo

   This document specifies an Internet standards track protocol for the
   Internet community, and requests discussion and suggestions for
   improvements. Please refer to the current edition of the "Internet
   Official Protocol Standards" (STD 1) for the standardization state
   and status of this protocol. Distribution of this memo is unlimited.

Abstract

   IEEE Std 802.16 is an air interface specification for fixed and
   mobile Broadband Wireless Access Systems. Service-specific
   convergence sublayers to which upper-layer protocols interface are a
   part of the IEEE 802.16 MAC (Medium Access Control). The Packet
   convergence sublayer (CS) is used for the transport of all packet-
   based protocols such as Internet Protocol (IP) and IEEE 802.3 LAN/MAN
   CSMA/CD Access Method (Ethernet). IPv6 packets can be sent and
   received via the IP-specific part of the Packet CS. This document
   specifies the addressing and operation of IPv6 over the IP-specific
   part of the Packet CS for hosts served by a network that utilizes the
   IEEE Std 802.16 air interface. It recommends the assignment of a
   unique prefix (or prefixes) to each host and allows the host to use
   multiple identifiers within that prefix, including support for
   randomly generated interface identifiers.

Table of Contents

1. Introduction

    IEEE 802.16e is an air interface for fixed and mobile broadband
    wireless access systems. The IEEE 802.16 [802.16] standard specifies
    the air interface, including the Medium Access Control (MAC) layer
    and multiple physical layer (PHY) specifications. It can be deployed
    in licensed as well as unlicensed spectrum. While the PHY and MAC
    are specified in IEEE 802.16, the details of IPv4 and IPv6 operation
    over the air interface are not included. This document specifies the
    operation of IPv6 over the IEEE 802.16 air interface.

    IPv6 packets can be carried over the IEEE Std 802.16 specified air
    interface via:

    1. the IP-specific part of the Packet CS or

    2. the 802.3[802.3]-specific part of the Packet CS

    The scope of this specification is limited to the operation of IPv6
    over IP CS only.

    The IEEE 802.16 specification includes the PHY and MAC details. The
    convergence sublayers are a part of the MAC. The packet convergence
    sublayer includes the IP-specific part that is used by the IPv6
    layer.

    The mobile station (MS)/host is attached to an access router via a
    base station (BS). The host and the BS are connected via the IEEE
    Std 802.16 air interface at the link and physical layers. The IPv6
    link from the MS terminates at an access router that may be a part of
    the BS or an entity beyond the BS. The base station is a layer 2
    entity (from the perspective of the IPv6 link between the MS and
    access router (AR)) and relays the IPv6 packets between the AR and
    the host via a point-to-point connection over the air interface.

2. Terminology

    The terminology in this document is based on the definitions in "IP
    over 802.16 Problem Statement and Goals" [PS-GOALS].

    o   IP CS - The IP-specific part of the Packet convergence sublayer is
        referred to as IP CS. IPv6 CS and IP CS are used interchangeably.

    o   Subscriber station (SS), Mobile Station (MS), Mobile Node (MN) -
        The terms subscriber station, mobile station, and mobile node are
        used interchangeably in this document and mean the same, i.e., an
        IP host.

3. 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 RFC 2119 [RFC2119].

4. IEEE 802.16 Convergence Sublayer Support for IPv6

   The IEEE 802.16 MAC specifies two main service-specific convergence
   sublayers:

   1. ATM convergence sublayer

   2. Packet convergence sublayer

   The Packet CS is used for the transport of packet-based protocols,
   which include:

   1. IEEE Std 802.3(Ethernet)

   2. Internet Protocol (IPv4 and IPv6)

   The service-specific CS resides on top of the MAC Common Part
   Sublayer (CPS) as shown in Figure 1. The service-specific CS is
   responsible for:

   o accepting packets (Protocol Data Units, PDUs) from the upper
      layer,

   o performing classification of the packet/PDU based on a set of
      defined classifiers that are service specific,

   o delivering the CS PDU to the appropriate service flow and
      transport connection, and

   o receiving PDUs from the peer entity.

   Payload header suppression (PHS) is also a function of the CS but is
   optional.

   The figure below shows the concept of the service-specific CS in
   relation to the MAC:

```
------------------------------\
| ATM CS      | Packet CS    | \
------------------------------  \
| MAC Common Part Sublayer   |   \
| (Ranging, scheduling, etc.)|    802.16 MAC
------------------------------   /
|          Security          |  /
|(Auth, encryption, key mgmt)| /
------------------------------/
|            PHY             |
------------------------------
```
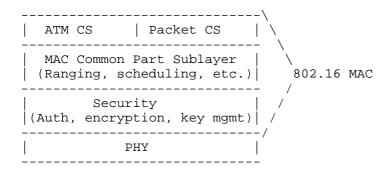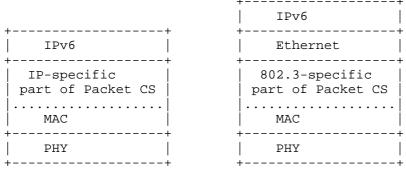
                 Figure 1: IEEE 802.16 MAC

Classifiers for each of the specific upper-layer protocols, i.e.,
Ethernet and IP, are defined in the IEEE 802.16 specification, which
enable the packets from the upper layer to be processed by the
appropriate service-specific part of the Packet CS. IPv6 can be
transported directly over the IP-specific part of the Packet CS (IP
CS). IPv4 packets also are transported over the IP-specific part of
the Packet CS. The classifiers used by IP CS enable the
differentiation of IPv4 and IPv6 packets and their mapping to
specific transport connections over the air interface.

The figure below shows the options for IPv6 transport over the packet
CS of IEEE 802.16:

```
                                  +------------------+
                                  |      IPv6        |
       +------------------+       +------------------+
       |      IPv6        |       |    Ethernet      |
       +------------------+       +------------------+
       |  IP-specific     |       |  802.3-specific  |
       | part of Packet CS|       | part of Packet CS|
       |..................|       |..................|
       |      MAC         |       |      MAC         |
       +------------------+       +------------------+
       |      PHY         |       |      PHY         |
       +------------------+       +------------------+

       (1) IPv6 over            (2) IPv6 over
           IP-specific part         802.3/Ethernet-
           of Packet CS             specific part
                                    of Packet CS
```

   Figure 2: IPv6 over IP- and 802.3-specific parts of the Packet CS

The figure above shows that while there are multiple methods by which
IPv6 can be transmitted over an 802.16 air interface, the scope of
this document is limited to IPv6 operation over IP CS only.
Transmission of IP over Ethernet is specified in [IPoE-over-802.16].
Transmission of IPv4 over IP CS is specified in [IPv4-over-IPCS].

It should be noted that immediately after ranging (802.16 air
interface procedure) and exchange of SBC-REQ/RSP messages (802.16
specific), the MS and BS exchange their capabilities via REG-REQ
(Registration Request) and REG-RSP (Registration Response) 802.16 MAC
messages. These management frames negotiate parameters such as the
Convergence Sublayer supported by the MS and BS. By default, Packet,
IPv4, and 802.3/Ethernet are supported. IPv6 via the IP CS is
supported by the MS and the BS only when the IPv6 support bit in the
capability negotiation messages (REG-REQ and REG-RSP) implying such
support is indicated in the parameter "Classification/PHS options and
SDU (Service Data Unit) encapsulation support" (refer to [802.16]).
Additionally, during the establishment of the transport connection
for transporting IPv6 packets, the DSA-REQ (Dynamic Service Addition)
and DSA-RSP messages between the BS and MS indicate via the CS-
Specification TLV the CS that the connection being set up shall use.
When the IPv6 packet is preceded by the IEEE 802.16 6-byte MAC
header, there is no specific indication in the MAC header itself
about the payload type. The processing of the packet is based
entirely on the classifiers. Based on the classification rules, the
MAC layer selects an appropriate transport connection for the
transmission of the packet. An IPv6 packet is transported over a
transport connection that is specifically established for carrying
such packets.

Transmission of IPv6 as explained above is possible via multiple
methods, i.e., via IP CS or via Ethernet interfaces. Every Internet
host connected via an 802.16 link:

1.   MUST be able to send and receive IPv6 packets via IP CS when the
     MS and BS indicate IPv6 protocol support over IP CS

2.   MUST be able to send and receive IPv6 packets over the Ethernet
     (802.3)-specific part of the Packet CS when the MS and BS
     indicate IPv6 protocol support over Ethernet CS. However, when
     the MS and BS indicate IPv6 protocol support over both IP CS and
     Ethernet CS, the MS and BS MUST use IP CS for sending and
     receiving IPv6 packets.

When the MS and BS support IPv6 over IP CS, it MUST be used as the
default mode for transporting IPv6 packets over IEEE 802.16 and the
recommendations in this document that are followed. Inability to
negotiate a common convergence sublayer for IPv6 transport between

the MS and BS will result in failure to set up the transport
connection and thereby render the host unable to send and receive
IPv6 packets. In the case of a host that implements more than one
method of transporting IPv6 packets, the default choice of which
method to use (i.e., IPv6 over the IP CS or IPv6 over 802.3) is IPv6
over IP CS when the BS also supports such capability.

In any case, the MS and BS MUST negotiate at most one convergence
sublayer for IPv6 transport on a given link.

In addition, to ensure interoperability between devices that support
different encapsulations, it is REQUIRED that BS implementations
support all standards-track encapsulations defined for 802.16 by the
IETF. At the time of writing this specification, this is the only
encapsulation, but additional specifications are being worked on. It
is, however, not required that the BS implementations use all the
encapsulations they support; some modes of operation may be off by
configuration.

4.1. IPv6 Encapsulation over the IP CS of the MAC

The IPv6 payload when carried over the IP-specific part of the Packet
CS is encapsulated by the 6-byte IEEE 802.16 generic MAC header. The
format of the IPv6 packet encapsulated by the generic MAC header is
shown in the figure below. The format of the 6-byte MAC header is
described in the [802.16] specification. The CRC (cyclic redundancy
check) is optional. It should be noted that the actual MAC address
is not included in the MAC header.

```
           ---------/ /-----------
          |    MAC SDU            |
           --------/ /------------
                    ||
                    ||
MSB                 \/                                        LSB
          ----------------------------------------------------------
         | Generic MAC header|  IPv6 Payload            | CRC  |
          ----------------------------------------------------------
```

Figure 3: IPv6 encapsulation

For transmission of IPv6 packets via the IP CS over IEEE 802.16, the
IPv6 layer interfaces with the 802.16 MAC directly. The IPv6 layer
delivers the IPv6 packet to the Packet CS of the IEEE 802.16 MAC.
The Packet CS defines a set of classifiers that are used to determine
how to handle the packet. The IP classifiers that are used at the
MAC operate on the fields of the IP header and the transport
protocol, and these include the IP Traffic class, Next header field,

Masked IP source and destination addresses, and Protocol source and
destination port ranges. Next header in this case refers to the last
header of the IP header chain. Parsing these classifiers, the MAC
maps an upper-layer packet to a specific service flow and transport
connection to be used. The MAC encapsulates the IPv6 packet in the
6-byte MAC header (MAC SDU) and transmits it. The figure below shows
the operation on the downlink, i.e., the transmission from the BS to
the host. The reverse is applicable for the uplink transmission.

```
     -----------                        ----------
    | IPv6 Pkt|                         |IPv6 Pkt|
     -----------                        ----------
       | |                                 /|\
       | |                                  |
   --[SAP]--------------------    ---------[SAP]--------
   ||-| |----------|        |    |          /|\        |
   || \ /        0---->[CID1] |    |    --- |-------   |
   || Downlink  0\/-->[CID2] |    |    |Reconstruct| |
   || classifiers0/\-->[....] |    |    | (undo PHS)| |
   ||         0---->[CIDn] |    |    ---   -------   |
   ||-------------|        |    |         /|\        |
   |                       |    |          |         |
   |   {SDU, CID,..}       |    |    {SDU, CID,..}   |
   |       |               |    |         /|\        |
   |       v               |    |          |         |
   ------[SAP]---------------    |-------[SAP]---------
   |     802.16 MAC CPS    |------>|  802.16 MAC CPS   |
   ---------------------------    ----------------------
           BS                             MS
```
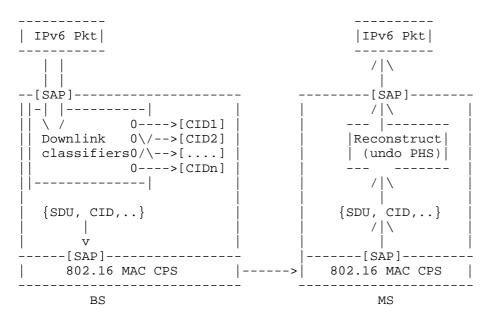
Figure 4: IPv6 packet transmission: Downlink

5. Generic Network Architecture Using the 802.16 Air Interface

   In a network that utilizes the 802.16 air interface, the host/MS is
   attached to an IPv6 access router (AR) in the network. The BS is a
   layer 2 entity only. The AR can be an integral part of the BS or the
   AR could be an entity beyond the BS within the access network. An AR
   may be attached to multiple BSs in a network. IPv6 packets between
   the MS and BS are carried over a point-to-point transport connection
   which is identified by a unique Connection Identifier (CID). The
   transport connection is a MAC layer link between the MS and the BS.
   The figures below describe the possible network architectures and are
   generic in nature. More esoteric architectures are possible but not
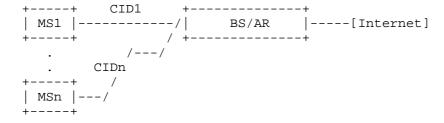   considered in the scope of this document.

   Option A:

```
         +-----+    CID1       +--------------+
         | MS1 |------------/|     BS/AR     |-----[Internet]
         +-----+          / +--------------+
           .           /---/
           .      CIDn
         +-----+     /
         | MSn |---/
         +-----+
```

              Figure 5: IPv6 AR as an integral part of the BS

   Option B:

```
         +-----+    CID1      +-----+           +-----------+
         | MS1 |----------/| BS1 |----------|     AR      |-----[Internet]
         +-----+        / +-----+           +-----------+
           .          /            _____
           .      CIDn /          ()_____()
         +-----+      /             L2 Tunnel
         | MSn |-----/
         +-----+
```
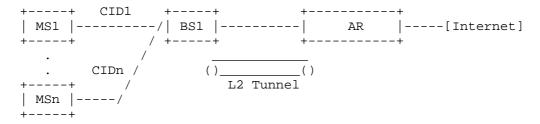
                Figure 6: IPv6 AR is separate from the BS

   The above network models serve as examples and are shown to
   illustrate the point-to-point link between the MS and the AR.

6. IPv6 Link

   "Neighbor Discovery for IP Version 6 (IPv6)" [RFC4861] defines link
   as a communication facility or medium over which nodes can
   communicate at the link layer, i.e., the layer immediately below IP.
   A link is bounded by routers that decrement the Hop limit field in
   the IPv6 header. When an MS moves within a link, it can keep using
   its IP addresses. This is a layer 3 definition, and note that the
   definition is not identical with the definition of the term '(L2)
   link' in IEEE 802 standards.

6.1. IPv6 Link in 802.16

   In 802.16, the transport connection between an MS and a BS is used to
   transport user data, i.e., IPv6 packets in this case. A transport
   connection is represented by a CID, and multiple transport
   connections can exist between an MS and a BS.

When an AR and a BS are colocated, the collection of transport
connections to an MS is defined as a single link. When an AR and a
BS are separated, it is recommended that a tunnel be established
between the AR and a BS whose granularity is no greater than 'per MS'
or 'per service flow' (An MS can have multiple service flows which
are identified by a service flow ID). Then the tunnel(s) for an MS,
in combination with the MS's transport connections, forms a single
point-to-point link.

The collection of service flows (tunnels) to an MS is defined as a
single link. Each link that uses the same higher-layer protocol has
only an MS and an AR. Each MS belongs to a different link. A
different prefix should be assigned to each unique link. This link
is fully consistent with a standard IP link, without exception, and
conforms with the definition of a point-to-point link in neighbor
discovery for IPv6 [RFC4861]. Hence, the point-to-point link model
for IPv6 operation over the IP-specific part of the Packet CS in
802.16 SHOULD be used. A unique IPv6 prefix(es) per link (MS/host)
MUST be assigned.

6.2. IPv6 Link Establishment in 802.16

In order to enable the sending and receiving of IPv6 packets between
the MS and the AR, the link between the MS and the AR via the BS
needs to be established. This section illustrates the link
establishment procedure.

The MS goes through the network entry procedure as specified by
802.16. A high-level description of the network entry procedure is
as follows:

1. The MS performs initial ranging with the BS. Ranging is a
    process by which an MS becomes time aligned with the BS. The MS
    is synchronized with the BS at the successful completion of
    ranging and is ready to set up a connection.

2. The MS and BS exchange basic capabilities that are necessary for
    effective communication during the initialization using SBC-REQ/
    RSP (802.16 specific) messages.

3. The MS progresses to an authentication phase. Authentication is
    based on Privacy Key Management version 2 (PKMv2) as defined in
    the IEEE Std 802.16 specification.

4. On successful completion of authentication, the MS performs
    802.16 registration with the network.

5. The MS and BS perform capability exchange as per 802.16
    procedures. Protocol support is indicated in this exchange. The
    CS capability parameter indicates which classification/PHS
    options and SDU encapsulation the MS supports. By default,
    Packet, IPv4, and 802.3/Ethernet shall be supported; thus,
    absence of this parameter in REG-REQ (802.16 message) means that
    named options are supported by the MS/SS. Support for IPv6 over
    the IP-specific part of the Packet CS is indicated by Bit #2 of
    the CS capability parameter (refer to [802.16]).

6. The MS MUST request the establishment of a service flow for IPv6
    packets over IP CS if the MS and BS have confirmed capability for
    supporting IPv6 over IP CS. The service flow MAY also be
    triggered by the network as a result of pre-provisioning. The
    service flow establishes a link between the MS and the AR over
    which IPv6 packets can be sent and received.

7. The AR and MS SHOULD send router advertisements and solicitations
    as specified in neighbor discovery [RFC4861].

The above flow does not show the actual 802.16 messages that are used
for ranging, capability exchange, or service flow establishment.
Details of these are in [802.16].

6.3. Maximum Transmission Unit in 802.16

The MTU value for IPv6 packets on an 802.16 link is configurable.
The default MTU for IPv6 packets over an 802.16 link SHOULD be 1500
octets.

The 802.16 MAC PDU is composed of a 6-byte header followed by an
optional payload and an optional CRC covering the header and the
payload. The length of the PDU is indicated by the Len parameter in
the Generic MAC header. The Len parameter has a size of 11 bits.
Hence, the total MAC PDU size is 2048 bytes. The IPv6 payload size
can vary. In certain deployment scenarios, the MTU value can be
greater than the default. Neighbor discovery for IPv6 [RFC4861]
defines an MTU option that an AR MUST advertise, via router
advertisement (RA), if a value different from 1500 is used. The MN
processes this option as defined in [RFC4861]. Nodes that implement
Path MTU Discovery [RFC1981] MAY use the mechanism to determine the
MTU for the IPv6 packets.

7. IPv6 Prefix Assignment

   The MS and the AR are connected via a point-to-point connection at
   the IPv6 layer. Hence, each MS can be considered to be on a separate
   subnet. A CPE (Customer Premise Equipment) type of device that
   serves multiple IPv6 hosts may be the end point of the connection.
   Hence, one or more /64 prefixes SHOULD be assigned to a link. The
   prefixes are advertised with the on-link (L-bit) flag set as
   specified in [RFC4861]. The size and number of the prefixes are a
   configuration issue. Also, Dynamic Host Configuration Protocol
   (DHCP) or Authentication, Authorization, and Accounting (AAA)-based
   prefix delegation MAY be used to provide one or more prefixes to MS
   for an AR connected over 802.16. The other properties of the
   prefixes are also dealt with via configuration.

8. Router Discovery

8.1. Router Solicitation

   On completion of the establishment of the IPv6 link, the MS may send
   a router solicitation message to solicit a router advertisement
   message from the AR to acquire necessary information as per the
   neighbor discovery for IPv6 specification [RFC4861]. An MS that is
   network attached may also send router solicitations at any time.
   Movement detection at the IP layer of an MS in many cases is based on
   receiving periodic router advertisements. An MS may also detect
   changes in its attachment via link triggers or other means. The MS
   can act on such triggers by sending router solicitations. The router
   solicitation is sent over the IPv6 link that has been previously
   established. The MS sends router solicitations to the all-routers
   multicast address. It is carried over the point-to-point link to the
   AR via the BS. The MS does not need to be aware of the link-local
   address of the AR in order to send a router solicitation at any time.
   The use of router advertisements as a means for movement detection is
   not recommended for MNs connected via 802.16 links as the frequency
   of periodic router advertisements would have to be high.

8.2. Router Advertisement

   The AR SHOULD send a number (configurable value) of router
   advertisements to the MS as soon as the IPv6 link is established.
   The AR sends unsolicited router advertisements periodically as per
   [RFC4861]. The interval between periodic router advertisements is
   however greater than the specification in neighbor discovery for
   IPv6, and is discussed in the following section.

8.3. Router Lifetime and Periodic Router Advertisements

   The router lifetime SHOULD be set to a large value, preferably in
   hours. This document overrides the specification for the value of
   the router lifetime in "Neighbor Discovery for IP Version 6 (IPv6)"
   [RFC4861]. The AdvDefaultLifetime in the router advertisement MUST
   be either zero or between MaxRtrAdvInterval and 43200 seconds. The
   default value is 2 * MaxRtrAdvInterval.

   802.16 hosts have the capability to transition to an idle mode, in
   which case, the radio link between the BS and MS is torn down.
   Paging is required in case the network needs to deliver packets to
   the MS. In order to avoid waking a mobile that is in idle mode and
   consuming resources on the air interface, the interval between
   periodic router advertisements SHOULD be set quite high. The
   MaxRtrAdvInterval value specified in this document overrides the
   recommendation in "Neighbor Discovery for IP Version 6
   (IPv6)"[RFC4861]. The MaxRtrAdvInterval MUST be no less than 4
   seconds and no greater than 21600 seconds. The default value for
   MaxRtrAdvInterval is 10800 seconds.

9. IPv6 Addressing for Hosts

   The addressing scheme for IPv6 hosts in 802.16 networks follows the
   IETF's recommendation for hosts specified in "IPv6 Node Requirements"
   [RFC4294]. The IPv6 node requirements [RFC4294] specify a set of
   RFCs that are applicable for addressing, and the same is applicable
   for hosts that use 802.16 as the link layer for transporting IPv6
   packets.

9.1. Interface Identifier

   The MS has a 48-bit globally unique MAC address as specified in
   802.16 [802.16]. This MAC address MUST be used to generate the
   modified EUI-64 format-based interface identifier as specified in "IP
   Version 6 Addressing Architecture" [RFC4291]. The modified EUI-64
   interface identifier is used in stateless address autoconfiguration.
   As in other links that support IPv6, EUI-64-based interface
   identifiers are not mandatory and other mechanisms, such as random
   interface identifiers, "Privacy Extensions for Stateless Address
   Autoconfiguration in IPv6" [RFC4941], MAY also be used.

9.2. Duplicate Address Detection

   DAD SHOULD be performed as per "Neighbor Discovery for IP Version 6
   (IPv6)", [RFC4861] and "IPv6 Stateless Address Autoconfiguration"
   [RFC4862]. The IPv6 link over 802.16 is specified in this document
   as a point-to-point link. Based on this criteria, it may be

redundant to perform DAD on a global unicast address that is
configured using the EUI-64 or generated as per RFC 4941 [RFC4941]
for the interface as part of the IPv6 Stateless Address
Autoconfiguration Protocol [RFC4862] as long as the following two
conditions are met:

1. The prefixes advertised through the router advertisement messages
    by the access router terminating the 802.16 IPv6 link are unique
    to that link.

2. The access router terminating the 802.16 IPv6 link does not
    autoconfigure any IPv6 global unicast addresses from the prefix
    that it advertises.

9.3. Stateless Address Autoconfiguration

   When stateless address autoconfiguration is performed, it MUST be
   performed as specified in [RFC4861] and [RFC4862].

9.4. Stateful Address Autoconfiguration

   When stateful address autoconfiguration is performed, it MUST be
   performed as specified in [RFC4861] and [RFC3315].

10. Multicast Listener Discovery

   "Multicast Listener Discovery Version 2 (MLDv2) for IPv6" [RFC3810]
   SHOULD be supported as specified by the hosts and routers attached to
   each other via an 802.16 link. The access router that has hosts
   attached to it via a point-to-point link over an 802.16 SHOULD NOT
   send periodic queries if the host is in idle/dormant mode. The AR
   can obtain information about the state of a host from the paging
   controller in the network.

11. Security Considerations

   This document does not introduce any new vulnerabilities to IPv6
   specifications or operation. The security of the 802.16 air
   interface is the subject of [802.16]. It should be noted that 802.16
   provides capability to cipher the traffic carried over the transport
   connections. A traffic encryption key (TEK) is generated by the MS
   and BS on completion of successful authentication and is used to
   secure the traffic over the air interface. An MS may still use IPv6
   security mechanisms even in the presence of security over the 802.16
   link. In addition, the security issues of the network architecture
   spanning beyond the 802.16 base stations are the subject of the
   documents defining such architectures, such as WiMAX Network
   Architecture [WiMAXArch] in Sections 7.2 and 7.3 of Stage 2, Part 2.

12. Acknowledgments

   The authors would like to acknowledge the contributions of the 16NG
   working group chairs Soohong Daniel Park and Gabriel Montenegro as
   well as Jari Arkko, Jonne Soininen, Max Riegel, Prakash Iyer, DJ
   Johnston, Dave Thaler, Bruno Sousa, Alexandru Petrescu, Margaret
   Wasserman, and Pekka Savola for their review and comments. Review
   and comments by Phil Barber have also helped in improving the
   document quality.

13. References

13.1. Normative References

   [802.16]   "IEEE Std 802.16e: IEEE Standard for Local and
                       metropolitan area networks, Amendment for
                       Physical and Medium Access Control Layers for
                       Combined Fixed and Mobile Operation in Licensed
                       Bands", October 2005, <http://standards.ieee.org/
                       getieee802/download/802.16e-2005.pdf>.

   [RFC1981]   McCann, J., Deering, S., and J. Mogul, "Path MTU
                       Discovery for IP version 6", RFC 1981,
                       August 1996.

   [RFC2119]   Bradner, S., "Key words for use in RFCs to
                       Indicate Requirement Levels", BCP 14, RFC 2119,
                       March 1997.

   [RFC3810]   Vida, R. and L. Costa, "Multicast Listener
                       Discovery Version 2 (MLDv2) for IPv6", RFC 3810,
                       June 2004.

   [RFC4291]   Hinden, R. and S. Deering, "IP Version 6
                       Addressing Architecture", RFC 4291,
                       February 2006.

   [RFC4861]   Narten, T., Nordmark, E., Simpson, W., and H.
                       Soliman, "Neighbor Discovery for IP version 6
                       (IPv6)", RFC 4861, September 2007.

   [RFC4862]   Thomson, S., Narten, T., and T. Jinmei, "IPv6
                       Stateless Address Autoconfiguration", RFC 4862,
                       September 2007.

13.2. Informative References

   [802.3]   "IEEE Std 802.3-2005: IEEE Standard for
                       Information technology-Telecommunications and
                       information exchange between systems-Local and
                       metropolitan area networks--Specific requirements
                       Part 3: Carrier Sense Multiple Access with
                       Collision Detection (CSMA/CD) Access Method and
                       Physical Layer Specifications", December 2005,
                       <http://standards.ieee.org/getieee802/
                       802.3.html>.

   [IPoE-over-802.16] Jeon, H., Riegel, M., and S. Jeong, "Transmission
                       of IP over Ethernet over IEEE 802.16 Networks",
                       Work in Progress, January 2008.

   [IPv4-over-IPCS]   Madanapalli, S., Park, S., and S. Chakrabarti,
                       "Transmission of IPv4 packets over IEEE 802.16's
                       IP Convergence Sublayer", Work in Progress,
                       November 2007.

   [PS-GOALS]   Jee, J., Madanapalli, S., and J. Mandin, "IP over
                       802.16 Problem Statement and Goals", Work
                       in Progress, December 2007.

   [RFC3315]   Droms, R., Bound, J., Volz, B., Lemon, T.,
                       Perkins, C., and M. Carney, "Dynamic Host
                       Configuration Protocol for IPv6 (DHCPv6)",
                       RFC 3315, July 2003.

   [RFC4294]   Loughney, J., "IPv6 Node Requirements", RFC 4294,
                       April 2006.

   [RFC4941]   Narten, T., Draves, R., and S. Krishnan, "Privacy
                       Extensions for Stateless Address
                       Autoconfiguration in IPv6", RFC 4941,
                       September 2007.

   [WMF]   "WiMAX Forum", <http://www.wimaxforum.org>.

   [WiMAXArch]   "WiMAX End-to-End Network Systems Architecture",
                       September 2007.

Appendix A. WiMAX Network Architecture and IPv6 Support

   The WiMAX (Worldwide Interoperability for Microwave Access) forum
   [WMF] has defined a network architecture in which the air interface
   is based on the IEEE 802.16 standard. The addressing and operation
   of IPv6 described in this document are applicable to the WiMAX
   network as well.

   WiMAX is an example architecture of a network that uses the 802.16
   specification for the air interface. WiMAX networks are also in the
   process of being deployed in various parts of the world, and the
   operation of IPv6 within a WiMAX network is explained in this
   appendix.

   The WiMAX network architecture consists of the Access Service Network
   (ASN) and the Connectivity Service Network (CSN). The ASN is the
   access network that includes the BS and the AR in addition to other
   functions such as AAA, mobile IP foreign agent, paging controller,
   location register, etc. The ASN is defined as a complete set of
   network functions needed to provide radio access to a WiMAX
   subscriber. The ASN is the access network to which the MS attaches.
   The IPv6 access router is an entity within the ASN. The term ASN is
   specific to the WiMAX network architecture. The CSN is the entity
   that provides connectivity to the Internet and includes functions
   such as mobile IP home agent and AAA. The figure below shows the
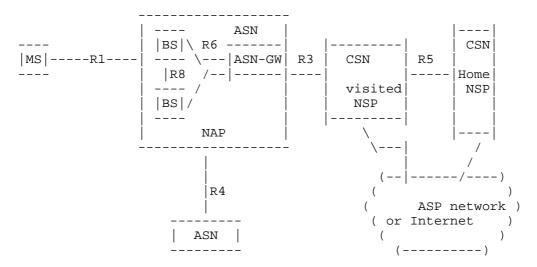   WiMAX reference model:

```
                         ------------------
                         | ----     ASN    |                     |----|
        ----             | |BS|\ R6 -------|     |---------|      | CSN|
       |MS|-----R1----|  | ---- \---|ASN-GW| R3  |  CSN    | R5   |    |
        ----             |  |R8  /--|------|----|           |-----|Home|
                         | ---- /           |    | visited |      | NSP|
                         | |BS|/            |    |  NSP    |      |    |
                         | ----             |    |---------|      |    |
                         |       NAP        |        \            |----|
                         ------------------          \---|         /
                             |                           |        /
                             |                        (--|------/----)
                             |R4                      (              )
                             |                        (    ASP network )
                         ---------                    ( or Internet    )
                         | ASN   |                    (                )
                         ---------                      (----------)
```

                 Figure 7: WiMAX network reference model

Three different types of ASN realizations called profiles are defined
by the architecture. ASNs of profile types A and C include BS' and
ASN-gateway(s) (ASN-GW), which are connected to each other via an R6
interface. An ASN of profile type B is one in which the
functionality of the BS and other ASN functions are merged together.
No ASN-GW is specifically defined in a profile B ASN. The absence of
the R6 interface is also a profile B specific characteristic. The MS
at the IPv6 layer is associated with the AR in the ASN. The AR may
be a function of the ASN-GW in the case of profiles A and C and is a
function in the ASN in the case of profile B. When the BS and the AR
are separate entities and linked via the R6 interface, IPv6 packets
between the BS and the AR are carried over a Generic Routing
Encapsulation (GRE) tunnel. The granularity of the GRE tunnel should
be on a per-MS basis or on a per-service-flow basis (an MS can have
multiple service flows, each of which is identified uniquely by a
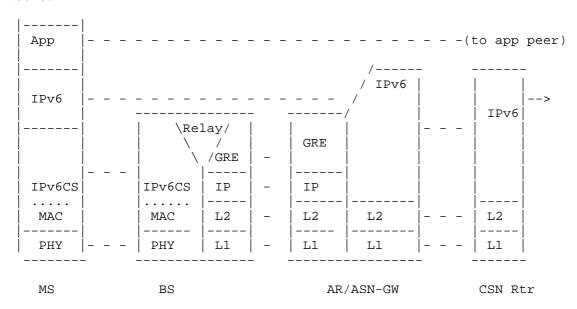service flow ID). The protocol stack in WiMAX for IPv6 is shown
below:

```
|-------|
| App   |- - - - - - - - - - - - - - - - - - - - - -(to app peer)
|       |
|-------|                                      /------     -------
|       |                                     / IPv6 |     |     |
| IPv6  |- - - - - - - - - - - - - - - -  /        |     |     |-->
|       |       ---------------   -------/          |     | IPv6|
|-------|       |   \Relay/   |   |      |          | - - |     |
|       |       |    \   /    |   | GRE  |          |     |     |
|       |       |     \ /GRE  | - |      |          |     |     |
|       |- - -  |     |-----|   |------|          |     |     |
| IPv6CS|       |IPv6CS| IP  | - | IP   |          |     |     |
| ..... |       |......|-----|   |------|--------|   |-----|
|  MAC  |       | MAC | L2  | - | L2   | L2     |- - -| L2  |
|-------|       |------|-----|   |----- |--------|   |-----|
| PHY   |- - -  | PHY | L1  | - | L1   | L1     |- - -| L1  |
 --------        ---------------   -----------------     -------

   MS              BS               AR/ASN-GW           CSN Rtr
```

Figure 8: WiMAX protocol stack

As can be seen from the protocol stack description, the IPv6 end-
points are constituted in the MS and the AR. The BS provides lower-
layer connectivity for the IPv6 link.

Appendix B. IPv6 Link in WiMAX

   WiMAX is an example of a network based on the IEEE Std 802.16 air
   interface. This section describes the IPv6 link in the context of a
   WiMAX network. The MS and the AR are connected via a combination of:

   1. The transport connection that is identified by a Connection
       Identifier (CID) over the air interface, i.e., the MS and BS, and

   2. A GRE tunnel between the BS and AR that transports the IPv6
       packets

   From an IPv6 perspective, the MS and the AR are connected by a point-
   to-point link. The combination of transport connection over the air
   interface and the GRE tunnel between the BS and AR creates a (point-
   to-point) tunnel at the layer below IPv6.

   The collection of service flows (tunnels) to an MS is defined as a
   single link. Each link has only an MS and an AR. Each MS belongs to
   a different link. No two MSs belong to the same link. A different
   prefix should be assigned to each unique link. This link is fully
   consistent with a standard IP link, without exception, and conforms
   with the definition of a point-to-point link in [RFC4861].

Appendix C. IPv6 Link Establishment in WiMAX

   The mobile station performs initial network entry as specified in
   802.16. On successful completion of the network entry procedure, the
   ASN gateway/AR triggers the establishment of the initial service flow
   (ISF) for IPv6 towards the MS. The ISF is a GRE tunnel between the
   ASN-GW/AR and the BS. The BS in turn requests the MS to establish a
   transport connection over the air interface. The end result is a
   transport connection over the air interface for carrying IPv6 packets
   and a GRE tunnel between the BS and AR for relaying the IPv6 packets.
   On successful completion of the establishment of the ISF, IPv6
   packets can be sent and received between the MS and AR. The ISF
   enables the MS to communicate with the AR for host configuration
   procedures. After the establishment of the ISF, the AR can send a
   router advertisement to the MS. An MS can establish multiple service
   flows with different quality of service (QoS) characteristics. The
   ISF can be considered as the primary service flow. The ASN-GW/AR
   treats each ISF, along with the other service flows to the same MS,
   as a unique link that is managed as a (virtual) interface.

Appendix D. Maximum Transmission Unit in WiMAX

   The WiMAX forum [WMF] has specified the Max SDU size as 1522 octets.
   Hence, the IPv6 path MTU can be 1500 octets. However, because of the
   overhead of the GRE tunnel used to transport IPv6 packets between the
   BS and AR and the 6-byte MAC header over the air interface, using a
   value of 1500 would result in fragmentation of packets. It is
   recommended that the MTU for IPv6 be set to 1400 octets in WiMAX
   networks, and this value (different from the default) be communicated
   to the MS. Note that the 1522-octet specification is a WiMAX forum
   specification and not the size of the SDU that can be transmitted
   over 802.16, which has a higher limit.

Authors' Addresses

    Basavaraj Patil
    Nokia Siemens Networks
    6000 Connection Drive
    Irving, TX   75039
    USA

    EMail: basavaraj.patil@nsn.com

    Frank Xia
    Huawei USA
    1700 Alma Dr. Suite 500
    Plano, TX   75075
    USA

    EMail: xiayangsong@huawei.com

    Behcet Sarikaya
    Huawei USA
    1700 Alma Dr. Suite 500
    Plano, TX   75075
    USA

    EMail: sarikaya@ieee.org

    JinHyeock Choi
    Samsung AIT
    Networking Technology Lab
    P.O.Box 111
    Suwon, Korea   440-600

    EMail: jinchoe@samsung.com

    Syam Madanapalli
    Ordyn Technologies
    1st Floor, Creator Building, ITPL.
    Off Airport Road
    Bangalore, India   560066

    EMail: smadanapalli@gmail.com

Intellectual Property

   The IETF takes no position regarding the validity or scope of any
   Intellectual Property Rights or other rights that might be claimed to
   pertain to the implementation or use of the technology described in
   this document or the extent to which any license under such rights
   might or might not be available; nor does it represent that it has
   made any independent effort to identify any such rights. Information
   on the procedures with respect to rights in RFC documents can be
   found in BCP 78 and BCP 79.

   Copies of IPR disclosures made to the IETF Secretariat and any
   assurances of licenses to be made available, or the result of an
   attempt made to obtain a general license or permission for the use of
   such proprietary rights by implementers or users of this
   specification can be obtained from the IETF on-line IPR repository at
   http://www.ietf.org/ipr.

   The IETF invites any interested party to bring to its attention any
   copyrights, patents or patent applications, or other proprietary
   rights that may cover technology that may be required to implement
   this standard. Please address the information to the IETF at
   ietf-ipr@ietf.org.