Transmission of IP over Ethernet over IEEE 802.16 Networks

Abstract

   This document describes the transmission of IPv4 over Ethernet, as
   well as IPv6 over Ethernet, in an access network deploying the IEEE
   802.16 cellular radio transmission technology. The Ethernet on top
   of IEEE 802.16 is realized by bridging connections that IEEE 802.16
   provides between a base station and its associated subscriber
   stations. Due to the resource constraints of radio transmission
   systems and the limitations of the IEEE 802.16 Media Access Control
   (MAC) functionality for the realization of an Ethernet, the
   transmission of IP over Ethernet over IEEE 802.16 may considerably
   benefit by adding IP-specific support functions in the Ethernet over
   IEEE 802.16 while maintaining full compatibility with standard IP
   over Ethernet behavior.

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.

Copyright Notice

Table of Contents

1. Introduction

   IEEE 802.16 [802.16] specifies a fixed-to-mobile, broadband wireless
   access system.

   The IEEE 802.16 standard defines a packet CS (Convergence Sublayer)
   for interfacing with specific packet-based protocols as well as a
   generic packet CS (GPCS) to provide an upper-layer, protocol-
   independent interface. This document describes transmission of IPv4
   and IPv6 over Ethernet via the Ethernet-specific part of the packet
   CS as well as of the GPCS in the access network based on IEEE 802.16.

   Ethernet has been originally architected and designed for a shared
   medium while the IEEE 802.16 uses a point-to-multipoint architecture
   like other cellular radio transmission systems. Hence, Ethernet on
   top of IEEE 802.16 is realized by bridging between IEEE 802.16 radio
   connections that connect a BS (Base Station) and its associated SSs
   (Subscriber Stations).

   Under the resource constraints of radio transmission systems and the
   particularities of the IEEE 802.16 for the realization of Ethernet,
   it makes sense to add IP-specific support functions in the Ethernet
   layer above IEEE 802.16 while maintaining full compatibility with
   standard IP over Ethernet behavior.

2. Requirements

   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 [RFC2119].

3. Terminology

   The terminology in this document is based on the definitions in "IP
   over 802.16 Problem Statement and Goals" [RFC5154].

4. The IEEE 802.16 Link Model

4.1. Connection-Oriented Air Interface

   The IEEE 802.16 MAC establishes connections between a BS and its
   associated SSs for the transfer of user data over the air. Each of
   these connections realizes an individual service flow, which is
   identified by a 16-bit Connection Identifier (CID) number and has a
   defined Quality of Service (QoS) profile.

   Multiple connections can be established between a BS and an SS, each
   with its particular QoS class and direction. Although the BS and all

the SSs are associated with unique 48-bit MAC addresses, packets
going over the air are only identified in the IEEE 802.16 MAC header
by the CID number of the particular connection. The connections are
established by MAC management messages between the BS and the SS
during network entry or also later on demand.

```
            [Subscriber  Side]                  [Network Side]


         |                   |                   |    +
         |                   |                   |    +
      +--+--+             +--+--+             +--+-+-+--+
      | MAC |             | MAC |             |   MAC   |
      +-----+             +-----+             +---------+
      | PHY |             | PHY |             |   PHY   |
      +-+-+-+             +-+-+-+             +-+-+-+-+-+
       + +                 | |                 | | + +
       + +                 | +-----CID#w------+ | + +
       + +                 +-------CID#x--------+ + +
       + +++++++++++++++++CID#y+++++++++++++++++ +
       ++++++++++++++++++CID#z++++++++++++++++++++
       SS#1                SS#2                BS
```
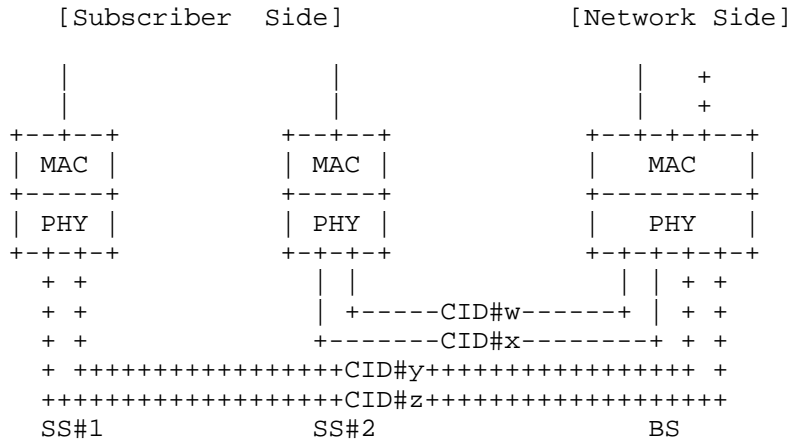
                  Figure 1: Basic IEEE 802.16 Link Model

4.2. MAC Addressing in IEEE 802.16

   Each SS has a unique 48-bit MAC address; the 48-bit MAC address is
   used during the initial ranging process for the identification of the
   SS and may be verified by the succeeding PKM (Privacy Key Management)
   authentication phase. Out of the successful authentication, the BS
   establishes and maintains the list of attached SSs based on their MAC
   addresses, purely for MAC management purposes.

   While MAC addresses are assigned to all the BSs as well as the SSs,
   the forwarding of packets over the air is only based on the CID value
   of the particular connection in the IEEE 802.16 MAC header. Not
   relying on the MAC addresses in the payload for reception of a radio
   frame allows for the transport of arbitrary source and destination
   MAC addresses in Ethernet frames between an SS and its BS. This is
   required for bridging Ethernet frames toward an SS that is attached
   to a bridge connected to another network.

   Due to the managed assignment of the service flows and associated CID
   values to individual SSs, the BS is able to bundle all unicast
   connections belonging to a particular SS into a single link on the
   network side, as shown in Figure 1, so that it provides a single
   layer-2 link between the SS and its associated wired link on the
   network side.

4.3. Unidirectional Broadcast and Multicast Support

   Current IEEE 802.16 [802.16] does not support bidirectional native
   broadcast and multicast for IP packets. While downlink connections
   can be used for multicast transmission to a group of SSs as well as
   unicast transmission from the BS to a single SS, uplink connections
   from the SSs to the BS provide only unicast transmission
   capabilities. Furthermore, the use of multicast CIDs for realizing
   downlink multicast transmissions is not necessarily preferable due to
   the reduced transmission efficiency of multicast CIDs for small
   multicast groups. Appendix A provides more background information
   about the issues arising with multicast CIDs in IEEE 802.16 systems.

   MBS (Multicast and Broadcast Service), as specified in IEEE 802.16,
   also does not cover IP broadcast or multicast data because MBS is
   invisible to the IP layer.

4.4. IEEE 802.16 Convergence Sublayer for IP over Ethernet

   IEEE 802.16 provides two solutions to transfer Ethernet frames over
   IEEE 802.16 MAC connections.

   The packet CS is defined for handling packet-based protocols by
   classifying higher-layer packets depending on the values in the
   packet header fields and assigning the packets to the related service
   flow. The packet CS comprises multiple protocol-specific parts to
   enable the transmission of different kinds of packets over IEEE
   802.16. The Ethernet-specific part of the packet CS supports the
   transmission of Ethernet by defining classification rules based on
   Ethernet header information.

   The GPCS (Generic Packet Convergence Sublayer) may be used as an
   alternative to transfer Ethernet frames over IEEE 802.16. The GPCS
   does not define classification rules for each kind of payload but
   relies on higher-layer functionality outside of the scope of IEEE
   802.16 to provide the assignment of packets to particular service
   flows.

5. Ethernet Network Model for IEEE 802.16

   Like in today's wired Ethernet networks, bridging is required to
   implement connectivity between more than two devices. In IEEE
   802.16, the point-to-point connections between SSs and the BS can be
   bridged so that Ethernet is realized over the IEEE 802.16 access
   network.

5.1. IEEE 802.16 Ethernet Link Model

   To realize Ethernet on top of IEEE 802.16, all the point-to-point
   connections belonging to an SS MUST be connected to a network-side
   bridging function, as shown in Figure 2. This is equivalent to
   today's switched Ethernet with twisted pair wires or fibres
   connecting the hosts to a bridge ("Switch").

   The network-side bridging function can be realized either by a single
   centralized network-side bridge or by multiple interconnected
   bridges, preferably arranged in hierarchical order. The single
   centralized network-side bridge allows best control of the
   broadcasting and forwarding behavior of the Ethernet over IEEE
   802.16. Appendix B explains the issues of a distributed bridging
   architecture when no assumptions about the location of the access
   router can be made.

   The BS MUST forward all the service flows belonging to one SS to one
   port of the network-side bridging function. No more than one SS MUST
   be connected to one port of the network-side bridging function. The
   separation method for multiple links on the connection between the BS
   and the network-side bridging function is out of scope for this
   document. Either layer-2 transport or layer-3 tunneling may be used.

   If the Ethernet over IEEE 802.16 is extended to multiple end stations
   behind the SS (i.e., SS#4 in the figure below), then the SS SHOULD
   support bridging according to [802.1D] and its amendment [802.16k],
   a.k.a. subscriber-side bridge, between all its subscriber-side ports
   and the IEEE 802.16 air link.

```
                --------------------- IP Link ------------------------

        [Subscriber Side]         [Network Side]         [Subscriber Side]
          |       |                   |                    |      |   |
         ETH     ETH                 ETH                  ETH    ETH ETH
          |       |                   |                    |      |   |
          |       |         +---------+---------+          |     +-+---+-+
          |       |         | Bridging Function |          |     |Bridge |
          |       |         +--+-+---------+-+--+          |     +---+---+
          |       |            | +         + |              |         |
        +--+--+  +--+--+     +--+-+-+     +--+-+-+       +--+--+    +--+--+
        | MAC |  | MAC |     | MAC  |     | MAC  |       | MAC |    | MAC |
        +-----+  +-----+     +------+     +------+       +-----+    +-----+
        | PHY |  | PHY |     | PHY  |     | PHY  |       | PHY |    | PHY |
        +-+-+-+  +-+-+-+     +-+-+-+-+    +-+-+-+-+      +-+-+-+    +-+-+-+
          +       | |          | | +       + | |          | |         +
          +       | +--CID#u-+ | +       + | +-CID#x--+ |          +
          +       +----CID#v---+ +       + +---CID#y----+          +
          +++++++++++++++CID#w++++++      ++++++CID#z+++++++++++++++

          SS#1     SS#2      BS#1         BS#2       SS#3      SS#4
```
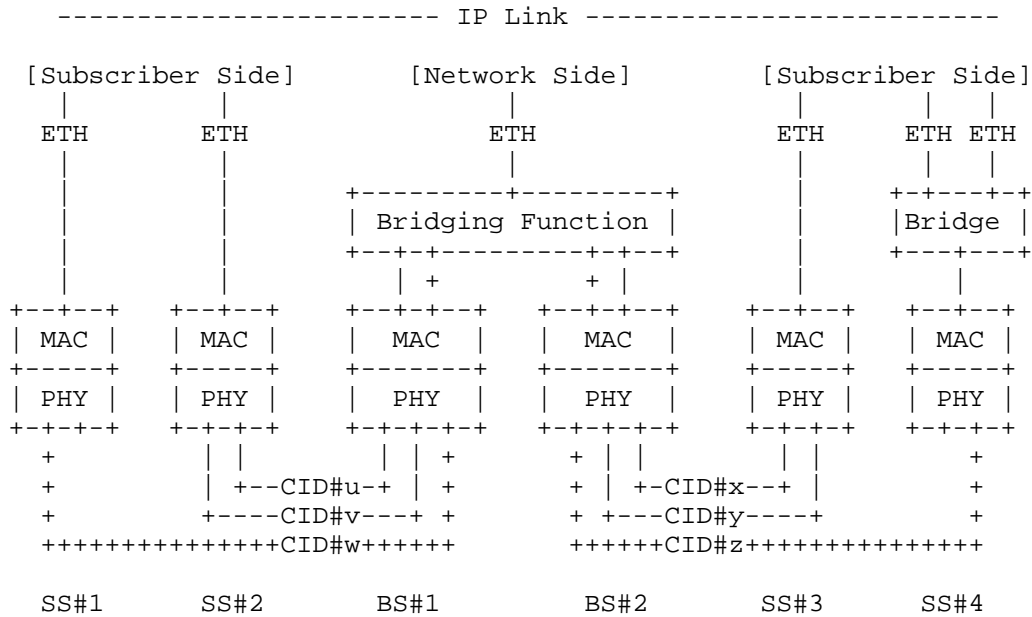
Figure 2: IEEE 802.16 Ethernet Link Model

5.2. Ethernet without Native Broadcast and Multicast Support

   Current IEEE 802.16 does not define broadcast and multicast of
   Ethernet frames. Hence, Ethernet frames that are broadcast or
   multicast SHOULD be replicated and then carried via unicast transport
   connections on the IEEE 802.16 access link. The network-side
   bridging function performs the replication and forwarding for
   Ethernet broadcast and multicast over the IEEE 802.16 radio links.

5.3. Network-Side Bridging Function

   The network-side bridging function MUST create a new radio-side port
   whenever a new SS attaches to any of the BSs of the network, or it
   MUST remove a radio-side port when an associated SS detaches from the
   BSs. The method for managing the port on the network-side bridging
   function may depend on the protocol used for establishing multiple
   links on the connection between the BS and the network-side bridge.
   The port-managing method is out of scope for this document.

   The network-side bridging function MUST be based on [802.1D] and its
   amendment [802.16k] to interconnect the attached SSs and pass
   Ethernet frames between the point-to-point connections associated
   with the attached SSs. However, to enhance the IEEE 802.16 Ethernet
   link model by avoiding broadcast or multicast packet flooding,

   additional IP-specific functionalities MAY be provided by the
   network-side bridging function in addition to the mandatory
   functions, according to Section 5.1 of [802.1D].

5.4. Segmenting the Ethernet into VLANs

   It is possible to restrict the size and coverage of the broadcast
   domain by segmenting the Ethernet over IEEE 802.16 into VLANs and
   grouping subsets of hosts into particular VLANs with each VLAN
   representing an IP link. Therefore, the network-side bridging
   function MAY be enabled to support VLANs according to [802.1Q] by
   assigning and handling the VLAN-IDs on the virtual bridge ports.

   If an SS is directly connected to a subscriber-side bridge supporting
   VLANs, the port associated with such an SS MAY be enabled as trunk
   port. On trunk ports, Ethernet frames are forwarded in the [802.1Q]
   frame format.

6. Transmission of IP over Ethernet over IEEE 802.16 Link

6.1. Generic IP over Ethernet Network Scenario

   The generic IP over Ethernet network scenario assumes that all hosts
   are residing on the same link. It enables the hosts to directly
   communicate with each other without detouring. There can be multiple
   Access Routers (ARs) on the link, and these may reside both on the
   subscriber side as well as on the network side, as shown in Figure 3.

```
                      +--+--+
                   ---|AR|SS|
                      +--+--+*                            +----+
                         *   +----+                      +Host|
            +----+--+       * |    +-------+             /+----+
            |Host|SS|* * * **|  BS +------+ \           / +----+
            +----+--+       * |    +-----+ \ \         / ++Host|
               +----+--+  *   +----+       \ \ +-+--------+ / +----+
               |Host|SS|*                   \ +--+         ++
          +----+  +----+--+                  +---+Bridging|   +----+
         --+ AR ++                               |Function+---+ AR +---
          +----+ \                               +--+      |   +----+
                 \              +----+    / +-+--------+
            +----+ +------+--+      |    +---+ /
            |Host+-+Bridge|SS|* * * *|  BS |    /
            +----+ +------+--+     * |    +---+
            +----+/              *    +----+
            |Host+ +----+--+   *
            +----+ |Host|SS|*
                   +----+--+
```
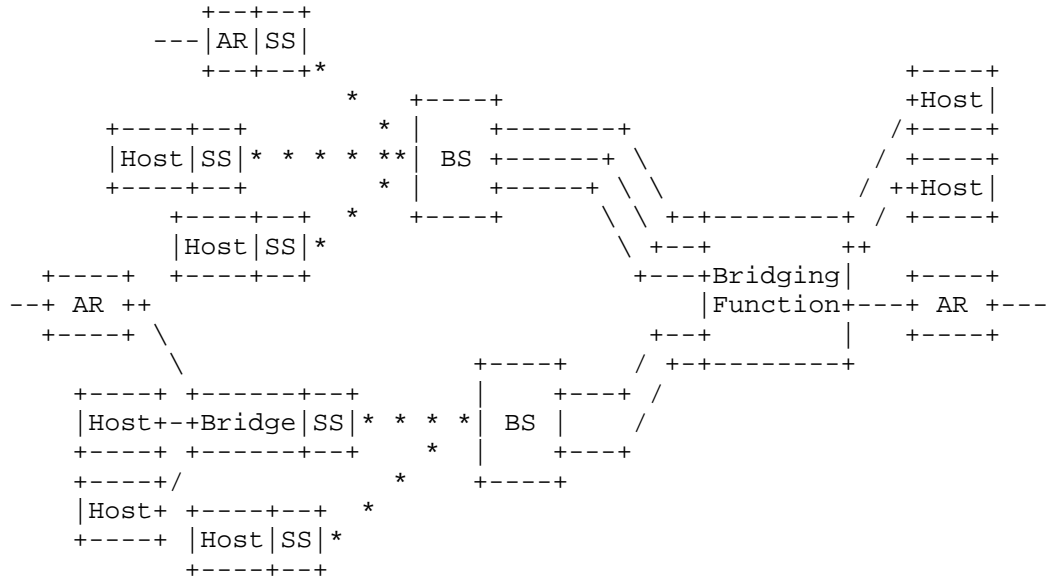
     Figure 3: Generic IP over Ethernet Network Scenario Using IEEE 802.16

6.2. Transmission of IP over Ethernet

6.2.1. IPv4-over-Ethernet Packet Transmission

   [RFC0894] defines the transmission of IPv4 packets over Ethernet
   networks. It contains the specification of the encapsulation of the
   IPv4 packets into Ethernet frames as well as rules for mapping IP
   addresses onto Ethernet MAC addresses. Hosts transmitting IPv4 over
   Ethernet packets over the IEEE 802.16 MUST follow the operations
   specified in [RFC0894].

6.2.1.1. Address Configuration

   IPv4 addresses can be configured manually or assigned dynamically
   from Dynamic Host Configuration Protocol for IPv4 (DHCPv4) servers
   [RFC2131].

6.2.1.2. Address Resolution

   The Address Resolution Protocol (ARP) [RFC0826] MUST be used for
   finding the destination Ethernet MAC address.

6.2.2. IPv6-over-Ethernet Packet Transmission

   [RFC2464] defines transmission of IPv6 packets over Ethernet
   networks, which includes an encapsulation of IPv6 packets into
   Ethernet frames; that document includes rules for mapping IPv6
   addresses to Ethernet addresses (i.e., MAC addresses). Hosts
   transmitting IPv6-over-Ethernet packets over IEEE 802.16 MUST follow
   the operations specified in [RFC2464].

6.2.2.1. Router Discovery, Prefix Discovery and Parameter Discovery

   Router Discovery, Prefix Discovery, and Parameter Discovery
   procedures are achieved by receiving Router Advertisement messages.
   However, periodic Router Advertisement messages can waste radio
   resource and disturb SSs in dormant mode in IEEE 802.16. Therefore,
   the AdvDefaultLifetime and MaxRtrAdvInterval SHOULD be overridden
   with high values specified in Section 8.3 in [RFC5121].

6.2.2.2. Address Configuration

   When stateful address autoconfiguration is required, the stateful
   address configuration according to [RFC3315] MUST be performed. In
   this case, an AR supports a Dynamic Host Configuration Protocol for
   IPv6 (DHCPv6) server or relay function.

   When stateless address autoconfiguration is required, the stateless
   address configuration according to [RFC4862] and [RFC4861] MUST be
   performed.

6.2.2.3. Address Resolution

   The Neighbor Discovery Protocol (NDP) [RFC4861] MUST be used for
   determining the destination Ethernet MAC address.

6.2.3. Maximum Transmission Unit

   [RFC2460] mandates 1280 bytes as a minimum Maximum Transmission Unit
   (MTU) size for the link layer and recommends at least 1500 bytes for
   IPv6 over Ethernet transmission. [RFC0894] also specifies 1500 bytes
   as a maximum length of IPv4 over Ethernet. Therefore, the default
   MTU of IPv6 packets and IPv4 packets on an Ethernet over IEEE 802.16
   link MUST be 1500 bytes.

6.2.4. Prefix Assignment

   As Ethernet over IEEE 802.16 may only build a part of a larger
   Ethernet of arbitrary structure, any kind of prefix assignment that
   is feasible for Ethernet is applicable for Ethernet over IEEE 802.16

as well. The same IPv4 prefix and the same set of IPv6 prefixes MAY
be assigned to all hosts attached to the Ethernet over IEEE 802.16 to
make best usage of Ethernet behavior. Sharing the prefix means
locating all hosts on the same subnetwork.

7. Operational Enhancements for IP over Ethernet over IEEE 802.16

   This section presents operational enhancements in order to improve
   network performance and radio resource efficiency for transmission of
   IP packets over Ethernet over IEEE 802.16 networks.

7.1. IP Multicast and Broadcast Packet Processing

   All multicast and multicast control messages can be processed in the
   network-side bridging function, according to [RFC4541]. Broadcasting
   messages to all radio-side side ports SHOULD be prevented.

   Further information on the prevention of multicasting or broadcasting
   messages to all radio-side ports is given in the following sections.

7.1.1. Multicast Transmission Considerations

   Usually, bridges replicate the IP multicast packets and forward them
   into all of its available ports except the incoming port. As a
   result, the IP multicast packets would be transmitted over the air --
   even to hosts that have not joined the corresponding multicast group.
   To allow bridges to handle IP multicast more efficiently, the IP
   multicast membership information should be propagated between
   bridges.

   In the IEEE 802.16 Ethernet link model in Section 5.1, the network-
   side bridging function can process all multicast data and multicast
   control messages according to [RFC4541] in order to maintain IP
   multicast membership states and forward IP multicast data to only
   ports suitable for the multicast group.

7.1.2. Broadcast Transmission Considerations

   The ordinary bridge floods the IP broadcast packets out of all
   connected ports except the port on which the packet was received.
   This behavior is not appropriate with scarce resources and dormant-
   mode hosts in a wireless network such as an access network based on
   IEEE 802.16.

   The network-side bridging function in the IEEE 802.16 Ethernet link
   model SHOULD flood all IP broadcast packets except ARP-, DHCPv4-, and
   Internet Group Management Protocol (IGMP)-related traffic.

IGMP-related broadcast packets can be forwarded according to the
[RFC4541]. ARP-related broadcast SHOULD be processed as specified in
Section 7.3.

7.2. DHCP Considerations

In the IPv4-over-Ethernet case, DHCPv4 clients may send DHCPDISCOVER
and DHCPREQUEST messages with the BROADCAST bit set to request the
DHCPv4 server to broadcast its DHCPOFFER and DHCPACK messages. The
network-side bridging function SHOULD filter these broadcast
DHCPOFFER and DHCPACK messages and forward the broadcast messages
only to the host defined by the client hardware address in the chaddr
information element.

Alternatively, the DHCP Relay Agent Information option (option 82)
[RFC3046] MAY be used to avoid DHCPv4 broadcast replies. Option 82
consists of two types of sub-options: Circuit ID and Remote ID. The
DHCPv4 Relay Agent is usually located on the network-side bridging
function as the Layer 2 DHCPv4 Relay Agent. The port number of the
network-side bridging function can be used as Circuit ID, and Remote
ID may be left unspecified. Note that using option 82 requires
DHCPv4 servers that are aware of option 82.

In the IPv6-over-Ethernet case, DHCPv6 clients use their link-local
addresses and the All_DHCP_Relay_Agents_and_Servers multicast address
to discover and communicate with DHCPv6 servers or Relay Agents on
their link. Hence, DHCPv6-related packets are unicasted or
multicasted. The network-side bridging function SHOULD handle the
DHCPv6-related unicast packets based on [802.1D] and SHOULD transmit
the DHCPv6-related multicast packets as specified in Section 7.1.1.

7.3. Address Resolution Considerations

In the IPv4-over-Ethernet case, ARP Requests are usually broadcasted
to all hosts on the same link in order to resolve an Ethernet MAC
address, which would disturb all hosts on the same link. Proxy ARP
provides the function in which a device on the same link as the hosts
answers ARP Requests instead of the remote host. When transmitting
IPv4 packets over the IEEE 802.16 Ethernet link, the Proxy ARP
mechanism is used by the network-side bridging function to avoid
broadcasting ARP Requests over the air.

The network-side bridging function SHOULD maintain an ARP cache large
enough to accommodate ARP entries for all its serving SSs. The ARP
cache SHOULD be updated by any packets including ARP Requests from
SSs in the same way the normal layer-2 bridging device is updating
its Filtering Database according to [802.1D].

Upon receiving an ARP Request from an SS, the network-side bridging
function SHOULD unicast an ARP Reply back to the SS with the Ethernet
address of the target host, provided that the target address matches
an entry in the ARP cache. However, in case of receiving an ARP
Request from a host behind a subscriber-side bridge, the network-side
bridging function SHOULD discard the request if the target host is
also behind the same subscriber-side bridge, i.e., on the same port
of the network-side bridge. Otherwise, the ARP Request MAY be
flooded. The network-side bridging function SHOULD silently discard
any received self-ARP Request.

In the IPv6-over-Ethernet case, Neighbor Solicitation messages are
multicasted to the solicited-node multicast address for the address
resolution, including a duplicate address detection. The solicited-
node multicast address facilitates the efficient querying of hosts
without disturbing all hosts on the same link. The network-side
bridging function SHOULD transmit the Neighbor Solicitation messages
specified in Section 7.1.1.

8. Public Access Recommendations

   In the public access scenario, direct communication between nodes is
   restricted because of security and accounting issues. Figure 4
   depicts the public access scenario.

   In this scenario, the AR is connected to a network-side bridge. The
   AR MAY perform security filtering, policing, and accounting of all
   traffic from hosts, e.g., like an NAS (Network Access Server).

   If the AR functions as the NAS, all the traffic from SSs SHOULD be
   forwarded to the AR, not bridged at the network-side bridging
   function -- even in the case of traffic between SSs served by the
   same AR. The bridge SHOULD forward upstream traffic from hosts
   toward the AR but MUST perform normal bridging operation for
   downstream traffic from the AR and MUST bridge SEcure Neighbor
   Discovery (SEND) [RFC3971] messages to allow applicability of
   security schemes.

   In the IPv4-over-Ethernet case, MAC-Forced Forwarding (MAC-FF)
   [RFC4562] can be used for the public access network to ensure that
   traffic from all hosts is always directed to the AR. The MAC-FF is
   performed in the network-side bridging function; thus, the bridge
   filters broadcast ARP Requests from all the hosts and responds to the
   ARP Requests with an Ethernet MAC address of the AR.

   In the IPv6-over-Ethernet case, unique IPv6 prefixes per SS can be
   assigned because doing so forces all IPv6 packets from SSs to be
   transferred to the AR and thus results in layer-3 separation between

SSs. Alternatively, common IPv6 prefixes can be assigned to all SSs
served by the same AR in order to exploit the efficient multicast
support of Ethernet link in the network side. In this case, a Prefix
Information Option (PIO) [RFC4861] carrying the common IPv6 prefixes
SHOULD be advertised with the On-link flag (L-Flag) reset so that it
is not assumed that the addresses matching the prefixes are available
on-link.

The AR should relay packets between SSs within the same AR.

```
            +-+--+
            |H|SS|                   +- - - - - - - - - - +
            +-+--+*    +----+   | +------+
       +-+--+       *  |    +-----+        |            |
       |H|SS|* * * * *| BS +-----+Bridge+-+
       +-+--+       *  |    +-----+     | | +-----+ |
                     *  +----+   | +------+ | | B   |
           +-+--+ *            |          +-+ r  | | +-------+
           |H|SS|*             |          | i +---+AR(NAS)+--
    +---+    +-+--+            |          | d | | +-------+
    | H ++                    +-+ g  |
    +---+ \           +----+  | +------+ | | e  | |
    +---+  +--+--+      |    +----+     | | +-----+
    | H +--+Br|SS|* * * * | BS |  | |Bridge+-+          |
    +---+  +--+--+    *  |    +----+     |
    +---+ /          *  +----+  | +------+           |
    | H ++    +-+--+ *            | Bridging Function |
    +---+    |H|SS|*              +- - - - - - - - - - +
             +-+--+
```
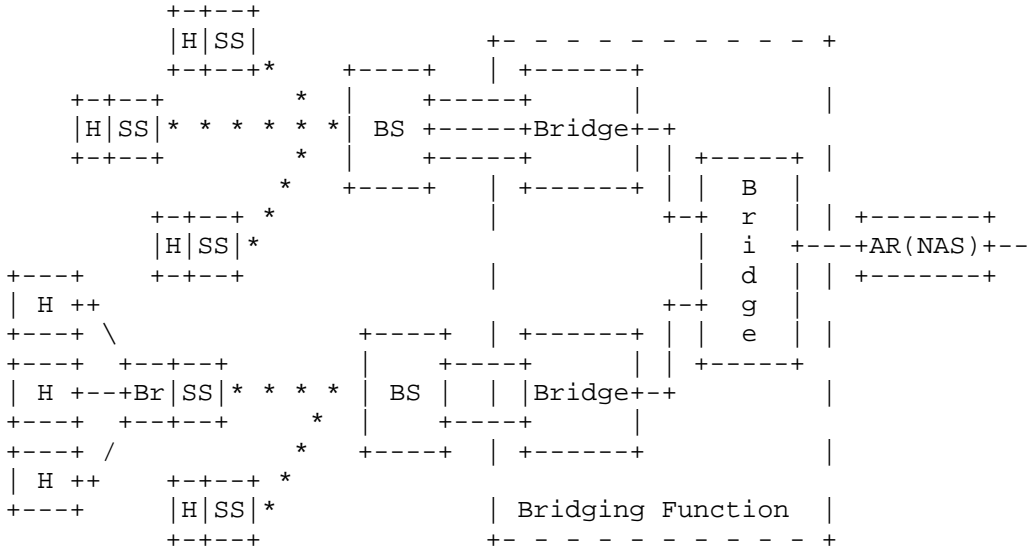
            Figure 4: Public Access Network Using IEEE 802.16

9. Security Considerations

   This recommendation does not introduce new vulnerabilities to IPv4
   and IPv6 specifications or operations. The security of the IEEE
   802.16 air interface between SSs and BS is the subject of [802.16],
   which provides the capabilities of admission control and ciphering of
   the traffic carried over the air interface. A Traffic Encryption Key
   (TEK) is generated by the SS and BS on completion of successful
   mutual authentication and is used to secure the air interface.

   The IEEE 802.16 Ethernet link model described in Section 5.1
   represents a bridged (switched) Ethernet architecture with point-to-
   point links between the SS and its bridge port. Even though the
   bridged Ethernet model prevents messaging between SSs on the same
   link without passing through the bridge, it is still vulnerable,
   e.g., by malicious reconfiguration of the address table of the bridge

in the learning process. This recommendation does not cause new
security issues beyond those that are already known for the bridged
Ethernet architecture. For example, link security mechanisms
according to [802.1AE] can be used on top of this recommendation to
resolve the security issues of the bridged Ethernet.

As the generic IP over Ethernet network using IEEE 802.16 emulates a
standard Ethernet link, existing IPv4 and IPv6 security mechanisms
over Ethernet can still be used. The public access network using
IEEE 802.16 can secure isolation of each of the upstream links
between hosts and AR by adopting SEcure Neighbor Discovery (SEND)
[RFC3971] for securing neighbor discovery processes.

10. Acknowledgments

   The authors would like to thank David Johnston, Dave Thaler, Jari
   Arkko, and others for their inputs to this work.

11. References

11.1. Normative References

   [802.16]   IEEE Std 802.16-2009, "IEEE Standard for Local and
               metropolitan area networks, Part 16: Air Interface for
               Fixed Broadband Wireless Access Systems", May 2009.

   [802.16k]   IEEE Std 802.16k-2007, "IEEE Standard for Local and
               metropolitan area networks, Media Access Control (MAC)
               Bridges, Amendment 5: Bridging of IEEE 802.16",
               March 2007.

   [802.1D]   IEEE Std 802.1D-2004, "IEEE Standard for Local and
               metropolitan area networks, Media Access Control (MAC)
               Bridges", June 2004.

   [802.1Q]   IEEE Std 802.1Q-2005, "IEEE Standard for Local and
               metropolitan area networks, Virtual Bridged Local Area
               Networks", May 2005.

   [RFC0826]   Plummer, D., "Ethernet Address Resolution Protocol: Or
               converting network protocol addresses to 48.bit Ethernet
               address for transmission on Ethernet hardware", STD 37,
               RFC 826, November 1982.

   [RFC0894]   Hornig, C., "Standard for the transmission of IP datagrams
               over Ethernet networks", STD 41, RFC 894, April 1984.

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

   [RFC2131] Droms, R., "Dynamic Host Configuration Protocol",
              RFC 2131, March 1997.

   [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6
              (IPv6) Specification", RFC 2460, December 1998.

   [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet
              Networks", RFC 2464, December 1998.

   [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.

   [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.

   [RFC5121] Patil, B., Xia, F., Sarikaya, B., Choi, JH., and S.
              Madanapalli, "Transmission of IPv6 via the IPv6
              Convergence Sublayer over IEEE 802.16 Networks", RFC 5121,
              February 2008.

11.2. Informative References

   [802.1AE] IEEE Std 802.1AE-2006, "IEEE Standard for Local and
              metropolitan area networks Media Access Control (MAC)
              Security", August 2006.

   [RFC3046] Patrick, M., "DHCP Relay Agent Information Option",
              RFC 3046, January 2001.

   [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure
              Neighbor Discovery (SEND)", RFC 3971, March 2005.

   [RFC4541] Christensen, M., Kimball, K., and F. Solensky,
              "Considerations for Internet Group Management Protocol
              (IGMP) and Multicast Listener Discovery (MLD) Snooping
              Switches", RFC 4541, May 2006.

   [RFC4562] Melsen, T. and S. Blake, "MAC-Forced Forwarding: A Method
              for Subscriber Separation on an Ethernet Access Network",
              RFC 4562, June 2006.

   [RFC5154] Jee, J., Madanapalli, S., and J. Mandin, "IP over IEEE
              802.16 Problem Statement and Goals", RFC 5154, April 2008.

Appendix A. Multicast CID Deployment Considerations

   Multicast CIDs are a highly efficient means to distribute the same
   information concurrently to multiple SSs under the same BS. However,
   the deployment of multicast CIDs for multicast or broadcast data
   services suffers from the following drawbacks.

   A drawback of multicast CIDs for Ethernet over IEEE 802.16 is the
   unidirectional nature of multicast CIDs. While it is possible to
   multicast information downstream to a number of SSs in parallel,
   there are no upstream multicast connections. In the upstream
   direction, unicast CIDs have to be used for sending multicast
   messages over the air to the BS, requiring a special multicast
   forwarding function for sending the information back to the other SSs
   on a multicast CID. While similar in nature to a bridging function,
   there is no appropriate forwarding model available. [802.1D] cannot
   take advantage of the multicast CIDs because it relies on unicast
   connections or bidirectional broadcast connections.

   A further drawback of deploying multicast CIDs for distributing
   broadcast control messages, like ARP Requests, is the inability to
   prevent the waking up of dormant-mode SSs by messages not aimed for
   them. Whenever a message is sent over a multicast CID, all
   associated stations have to power up and receive and process the
   message. While this behavior is desirable for multicast and
   broadcast traffic, it is harmful for link-layer broadcast control
   messages aimed for a single SS, like an ARP Request. All other SSs
   are wasting scarce battery power for receiving, decoding, and
   discarding the message. Low power consumption is an extremely
   important aspect in a wireless communication.

   Furthermore, it should be kept in mind that multicast CIDs are only
   efficient for a large number of subscribed SSs in a cell. Due to
   incompatibility with advanced radio-layer algorithms based on
   feedback information from the receiver side, multicast connections
   require much more radio resources for transferring the same
   information as unicast connections.

Appendix B. Centralized vs. Distributed Bridging

   This specification introduces a network-side bridging function, which
   can be realized either by a centralized device or by multiple
   interconnected bridges in a distributed manner. One common
   implementation of the distributed model is the scenario where a
   bridge is directly attached to the BS, such that the interface
   between BS and bridging function becomes a software interface within
   the operation system of the BS/bridge device.

The operational enhancements described in Section 7 of this document
are based on the availability of additional information about all the
hosts attached to the Ethernet. Flooding all ports of the bridge can
be avoided when a priori information is available to determine the
port to which an Ethernet frame has to be delivered.

Best performance can be reached by a centralized database containing
all information about the hosts attached to the Ethernet. A
centralized database can be established by either a centralized
bridge device or a hierarchical bridging structure with dedicated
uplink and downlink ports like in the public access case, where the
uppermost bridge is able to retrieve and maintain all the
information.

As the generic case of the IP over Ethernet over IEEE 802.16 link
model does not make any assumptions about the location of the AR (an
AR may eventually be attached to an SS), a centralized bridging
system is recommended for the generic case. In the centralized
system, every connection over the air of a link should be attached to
a single centralized bridge.

A distributed bridging model is appropriate, in particular, for the
public access mode, where Ethernet frames, which do not have entries
in the bridge behind the BS, are sent upstream until finally reaching
a bridge that has an entry for the destination MAC address.

Authors' Addresses

    Hongseok Jeon
    Electronics Telecommunications Research Institute
    161 Gajeong-dong, Yuseong-gu
    Daejeon,   305-350
    KOREA

    Phone: +82-42-860-3892
    EMail: hongseok.jeon@gmail.com

    Sangjin Jeong
    Electronics Telecommunications Research Institute
    161 Gajeong-dong, Yuseong-gu
    Daejeon,   305-350
    KOREA

    Phone: +82-42-860-1877
    EMail: sjjeong@etri.re.kr

    Max Riegel
    Nokia Siemens Networks
    St-Martin-Str 76
    Munich,   81541
    Germany

    Phone: +49-89-5159-32728
    EMail: maximilian.riegel@nsn.com