Transmission of IP over Ethernet over IEEE 802.16 Networks
draft-ietf-16ng-ip-over-ethernet-over-802.16-04.txt

Status of this Memo

Copyright Notice

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 between point-to-point radio links, which are provided by IEEE 802.16 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 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.

Table of Contents

1. Introduction

IEEE 802.16 [IEEE802.16] defines a PMP (Point-to-Multipoint) radio transmission system connecting a BS (Base Station) with multiple SSs (Subscriber Stations). IEEE 802.16e [IEEE802.16e] amends the base specification with PHY and MAC functions for supporting mobile terminals while adopting the same data link principles.

Ethernet is a widely deployed transmission technology. It provides unicast transport, handles broadcast, and multicast traffic efficiently and provides rich services such as Virtual LANs. However, Ethernet has been originally architected and designed for a shared medium without special considerations for advanced wireless transmission systems. The focus on wired systems requires additional support functions when Ethernet is employed in the IEEE 802.16.

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 between the point-to-point radio links, which are provided by IEEE 802.16 between the BS and its associated SSs.

With the resource constraints of radio transmission systems and the particularities of the IEEE 802.16 MAC 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. Those IP specific support functions are preferably realized in a centralized device containing also the bridge for aggregation of traffic from all the BSs to provide a more straightforward management solution and allow for effectively commoditized BSs, which are deployed in the IEEE 802.16 based access network in a large volume.

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 [I-D.ietf-16ng-ps-goals].

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 realize an individual Service Flow which is identified by a 16 bit CID number and has a defined QoS profile.

Multiple connections can be established between a BS and a 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.

While uplink connections from the SSs to the BS provide only unicast transmission capabilities, 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. The use of multicast CIDs for realizing multicast transmissions, however, is not addressed in this document due to the ongoing standardization efforts for the management of multicast CIDs, the reduced transmission efficiency of multicast CIDs for small multicast groups, the missing support by [IEEE802.1D] for uni-directional broadcast channels as well as additional security threats of broadcast channels in a power-conservative wireless system.

Appendix A provides more background information about the issues arising with multicast CIDs in IEEE 802.16 systems.

```
             [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. Feeding User Data into the Appropriate Connections

Assignment of higher layer packets to particular Service Flows and related CIDs is performed by the convergence sublayer within the IEEE 802.16 MAC. It classifies the packets depending on the values in the defined set of the payload packet header fields and assigns the packets to the appropriate Service Flow.

To enable the transmission of different kind of payloads over IEEE 802.16, multiple convergence sublayers are defined, each specific for one kind of payload packet type, like Ethernet, IPv4, IPv6 or even for encapsulated payload, like IPv4 over Ethernet or IPv6 over Ethernet.

4.3. MAC addressing in IEEE 802.16

The 48-bit unique MAC address of a SS is used during the initial ranging process for the identification of a SS and is verified by the succeeding PKMv2 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 SSs as well as to the BS, the forwarding of packets over the air is performed only on base of the CID value. 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 a SS and its BS. This is beneficial when Ethernet frames with arbitrary MAC addresses have to be forwarded to a SS in the case that the SS is interconnected 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 connections belonging to a particular SS into a single link on the network side as shown in Figure 1 so that it provides plain layer 2 forwarding behavior between the radio link toward the subscriber side and its associated wired link on the network side.

5. Ethernet Network Model for IEEE 802.16

5.1. IEEE 802.16 Ethernet Link Model

According to [RFC4861], a link is defined as a communication facility or medium over which IP devices can communicate at the link layer, i.e. the layer immediately below IP. Ethernet fully satisfies the definition of the link. IEEE 802.16, however, has limitations on its transitive connectivity. IEEE 802.16 provides point-to-point connections between SSs and the BS but does not enable any direct SS to SS connectivity. Hence, it is required to interconnect each point-to-point connections between SSs and the BS so that Ethernet is realized over IEEE 802.16 access network.

This document defines an IEEE 802.16 Ethernet link model to provide above the interconnection functionality. The IEEE 802.16 Ethernet link model SHALL interconnect each point-to-point connections assigned to SSs at a centralized point, a.k.a. network-side bridge, as shown in Figure 2. This is equivalent to today's switched Ethernet with twisted pair wires connecting the hosts to a bridge ("Switch"). The single and centralized network-side bridge allows best control of the broadcasting forwarding behavior and prevents potential security threats coming up with cascaded bridges. Appendix B explains the drawbacks and the potential security threats of an architecture where a bridge interconnects BSs integrated with bridging function.

The BS SHALL forward all the Service Flows belonging to one SS to one port of the network-side bridge. No more than one SS SHALL be connected to one port of the network-side bridge. Separation method for multiple links on the connection between the BS and the network-side bridge is out of scope for this document. One implementation is to deploy flow identifiers (e.g. VLAN-IDs or GRE KEYS) on the wired path. Section 6 discusses the network-side bridge in detail.

If the SS is connected to another network consisting of multiple hosts behind the SS (i.e. SS#4 in the below figure) then the SS SHOULD support bridging according to [IEEE802.1D] and its amendment [IEEE802.16k], 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
       |         |                 |                 |       |   |
       |         |       +---------+---------+       |     +-+---+-+
       |         |       |     Net-Bridge    |       |     |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 [IEEE802.16][IEEE802.16e] does not define broadcast and multicast of IP packets. Also, MBS (Multicast and Broadcast Service as specified in 6.3.23 of [IEEE802.16e]) does not cover IP broadcast or multicast data because MBS is invisible to the IP layer. Hence IP broadcast and multicast packets SHOULD be replicated and then carried via unicast transport connections as IEEE 802.16 access link. The network-side bridge performs the replication and forwarding as specified in Section 6.3.

5.3. Deployment Scenarios for IP over Ethernet over IEEE 802.16

This section discusses two possible deployment scenarios based on the IEEE 802.16 Ethernet link model: Public Access scenario and Enterprise LAN scenario.

In both scenarios, an AR is connected to a network-side bridge, which acts as an aggregation point for all the connections from SSs. Multiple ARs can exist on a link and the link may consist of multiple hosts, either being co-located with a SS or behind a SS integrated with a subscriber-side bridge. The network behind a SS can support various access network technologies, e.g. 802.3, 802.11 or 802.15.

There is a big difference between the scenarios in terms of the service provider policies. The difference is also reflected in Section 6.1, Section 6.4, and Section 8.

5.3.1. Public Access Scenario

In the Public Access scenario, direct communication between nodes is restricted because of security and accounting issues. Figure 3 depicts the public access scenario.

In the 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 a NAS (Network Access Server).

```
                 +--+
                 |SS|
                 +--+*
                       *
                         *   +----+
               +--+        * |    +-------+
               |SS|* * * * **| BS +------+ \
               +--+        * |    +-----+ \ \  +---+
                   +--+  *   +----+      \ \ +-+  B|
                   |SS|*                  \ +--+N r|
                   +--+                    +---+e i|   +----+
        +----+                                 |t d+---+ AR +--Internet
        |Host+                              +--+  g|   +----+
        +----+\                  +----+    / +-+  e|
        +----+ +------+--+       |    +---+ /  +---+
        |Host+-+Bridge|SS|* * * *| BS |    /
        +----+ +------+--+    *  |    +---+
        +----+/             *    +----+
        |Host+      +--+  *
        +----+      |SS|*
                    +--+
```

Figure 3. Public Access Scenario

5.3.2. Enterprise LAN Scenario

The enterprise LAN scenario assumes trust relationship between all hosts and thus enables hosts to directly communicate with each other without detouring. There can be multiple ARs, which may reside on both the subscriber side and network side as shown in Figure 4.

```
                      +--+
                      |SS|
                      +--+*                                 +-----+
                            *   +-----+                     +Host|
                  +--+        * |    +--------+            /+-----+
                  |SS|* * * * **| BS +-------+ \          / +-----+
                  +--+        * |    +------+ \ \  +---+/ ++Host|
                      +--+  *   +-----+      \ \ +-+  B+ / +-----+
                      |SS|*                   \ +--+N r++
         +-----+      +--+                     +---+e i|   +-----+
       --+ AR ++                                   |t d+----+ AR +---
         +-----+ \                                +--+  g|   +-----+
                  \                 +-----+    / +-+  e|
           +-----+ +------+--+      |    +---+ /  +---+
           |Host+-+Bridge|SS|* * * *| BS |    /
           +-----+ +------+--+    * |    +---+
           +-----+/            *    +-----+
           |Host+      +--+  *
           +-----+      |SS|*
                        +--+
```

Figure 4. Enterprise LAN Scenario

6. Network-side Bridge Considerations

Network-side bridge is based on [IEEE802.1D] to interconnect point-to-point connections assigned to each SSs and pass Ethernet frames between the point-to-point connections. However, applying the IEEE 802.16 Ethernet link model and avoiding multicast/broadcast flooding require additional IP specific functionalities on the network-side bridge as well [IEEE802.1D].

Following sections discuss the additional functions of the network-side bridge based on Figure 5.

```
                      [Lower Side]                [Upper Side]
    +--+          +--+             +------------+
    |SS+----------+  +-------------*            |             +-----+
    +--+          |BS|             |Network-side*-------------+ AR |
    +--+          |  |             |Bridge      |             +-----+
    |SS+==========+  +=============*            |
    +--+          +--+             +------------+
```

Figure 5. Network-side Bridge

6.1. IEEE 802.16 Ethernet Link Model Considerations

In the IEEE 802.16 Ethernet link model, the network-side bridge SHOULD create a new lower side port whenever a new SS attaches to any of the BSs of the network or SHOULD remove a lower side port when an associated SS detaches from the BSs. Method for managing the port on the network-side bridge may depend on approaches to build multiple links on the connection between the BS and the network-side bridge. The port managing method is out of scope for this document.

6.1.1. Public Access Scenario Case

The network-side bridge SHOULD forward all packets received from any lower side ports to all upper side ports being in the forwarding state. Peer-to-peer communication SHOULD NOT be supported by the network-side bridge and all packets originated from a SS SHOULD be delivered to an AR.

While the network-side bridge forces all traffic from hosts to reach the AR, it still performs Learning Process and maintains Filtering Database as specified in [IEEE802.1D] and then forwards IP unicast packets from the AR based on the Filtering Database. However, IP broadcast and multicast packets SHOULD be treated with special rules as stated in Section 6.3.

6.1.2. Enterprise LAN Scenario Case

IP unicast packets from either SSs or AR SHALL be forwarded by [IEEE802.1D] based bridging. IP broadcast and multicast packets SHOULD be processed in the bridge according to the rules presented in Section 6.3.

6.2. Segmenting the Ethernet into VLAN

It is possible to restrict the size and coverage of an IP link by segmenting the Ethernet and grouping subsets of hosts into VLANs. Therefore, the network-side bridge MAY be enabled to support VLANs according to [IEEE802.1Q] by assigning and handling the VLAN-IDs of the virtual bridge ports.

If a SS itself contains a VLAN enabled bridge or is directly connected to a subscriber-side bridge supporting VLANs, the port associated with such a SS MAY be enabled as trunk port. On trunk ports, Ethernet frames are forwarded in the [IEEE802.1Q] frame format.

6.3. Multicast and Broadcast Packet Processing

All multicast and multicast control messages SHOULD be processed in the network-side bridge according to [RFC4605]. Broadcasting messages to all lower side ports of the network-side bridge SHOULD be prevented.

Further information on prevention of multicasting or broadcasting messages to all lower side ports are given in the following sections.

6.3.1. Multicast Transmission Considerations

Usually, bridges replicate the IP multicast packets and forward them into all of its available ports with the exception of the incoming port, like IP broadcast packets. As a result, the IP multicast packets would be transmitted even to SSs which do not participate in the corresponding multicast group. To allow bridges to handle IP multicast more efficiently, the IP multicast membership should be propagated between bridges.

IGMP/MLD proxying in [RFC4605] is a simple mechanism for multicast packets forwarding based on the Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) membership information, which works only in a basic tree topology. An IGMP/MLD proxy device does learning and proxying group membership information, and then forwards the IP multicast packets based on the membership information. Typically, the proxy device is located at an aggregation point, which has a single upstream interface and multiple downstream interfaces.

The IEEE 802.16 Ethernet link model in Section 5.1 has a simple tree topology and [RFC4541] provides an application guide how bridge, non-IP device, to examine and learn group membership. Hence, [RFC4605] can be applied to the IEEE 802.16 Ethernet link model to reduce the multicast packet flooding.

The network-side bridge in the IEEE 802.16 Ethernet link model SHOULD play a role in proxying IGMP/MLD messages as specified in [RFC4605]. The network-side bridge SHOULD perform the host portion of IGMP/MLD process on its upper side port and the router portion of IGMP/MLD process on its all lower side ports. Note that the network-side bridge SHOULD perform IGMP/MLD Querier on only lower side ports, which are already subscribed with received IGMP/MLD membership report messages. This is due to the reduction of flooding of IGMP/MLD Query messages. The network-side bridge SHOULD maintain subscription information on each lower side port with received IGMP/MLD membership report messages and forward multicast packets from a upper side port to lower side ports based on the subscription information. In case of multicast packets from lower side ports, the network-side bridge SHOULD forward the packets to an upper side port as well as lower side ports, except the incoming lower side port, based on the subscription information.

6.3.2. Broadcast Transmission Considerations

The typical 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 IEEE 802.16 based access network.

The network-side bridge in the IEEE 802.16 Ethernet link model SHOULD discard all IP broadcast packets except ARP, DHCP (DHCPv4 and DHCPv6), IGMP, and MLD related traffic. The ARP, DHCP, IGMP and MLD related packets SHOULD be forwarded with special rules specified in this specification. Note that packets destined for permanently assigned multicast addresses such as 224.0.0/24 in IPv4 or FF02::1 in IPv6 are also regarded as broadcast data.

6.4. Proxy ARP

Proxy ARP provides a process where a device on the link between hosts answers ARP Requests instead of the remote host. In this specification, the Proxy ARP mechanism is used to force all traffic from hosts to the access router and to avoid broadcasting ARP Requests over the air depending on the particular deployment scenario. The Proxy ARP process is usually co-located with the network-side bridge.

6.4.1. Public Access Scenario Case

The network-side bridge SHOULD filter broadcast ARP Requests and SHOULD respond to all the ARP Requests from lower side port with the access router’s Ethernet MAC address so that all IPv4 packets from SSs are forwarded to the access router.

6.4.2. Enterprise LAN Scenario Case

The network-side bridge 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 network-side bridge is updating its Filtering Database according to [IEEE802.1D].

Upon receiving the ARP Requests from SSs, the network-side bridge SHOULD unicast ARP Replies back to SSs with Ethernet address of target host provided that the target address matches an entry in the ARP Cache. Otherwise, the network-side bridge MAY flood the ARP Requests. The network-side bridge SHOULD silently discard any received self-ARP Requests.

7. Access Router Considerations

In the public access scenario, all packets between SSs will always be relayed via access router. In this scenario, the access router SHOULD forward packets via the same interface on which the access router received the packets, if the source and destination addresses are reachable on the same interface. This would result in a Redirect message from the access router [RFC0792][RFC4861]. The Redirect message SHOULD be suppressed.

8. Prefix Assignment

8.1. Public Access Scenario Case

Unique IPv6 prefix per SS SHOULD be assigned because it results in layer 3 separation between SSs and it forces all packets from SSs to be transferred to an AR. The AR SHOULD assign the IPv6 prefixes with Prefix Information option as specified in [RFC4861].

One IPv4 prefix SHOULD be assigned to all SSs in a way that it benefits from high address assignment efficiency when concerning scarce IPv4 address space. The prefix can be manually configured or automatically with subnet mask option in DHCPv4 [RFC2132].

8.2. Enterprise LAN Scenario Case

The AR SHOULD assign all SSs one IPv4 prefix in IPv4 over Ethernet and one or more IPv6 prefixes in IPv6 over Ethernet so that all SSs under the same AR share the subnet prefix. Sharing the prefix means locating all SSs on the same subnetwork.

9. Transmission of IP over Ethernet

9.1. IPv4 over Ethernet

[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 of IP addresses onto Ethernet MAC addresses. IP over Ethernet over IEEE802.16 MUST follow the operations specified in [RFC0894].

9.1.1. Address Configuration

IPv4 addresses can be configured manually or assigned dynamically from DHCPv4 server [RFC2131].

DHCP clients may send DHCP DISCOVER and DHCP REQUEST messages with the BROADCAST bit set to request the DHCP server to broadcast its DHCP OFFER and DHCP ACK messages. The network-side bridge SHOULD filter these broadcast DHCP OFFER and DHCP ACK messages and forwards 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 DHCP broadcast replies. The option-82 consists of two type of sub-options; Circuit ID and Remote ID. DHCP Relay Agent is usually located on the network-side bridge as Layer 2 DHCP Relay Agent, like described in [TR101]. Port number of the network-side bridge is possible as Circuit ID and Remote ID may be left unspecified. Note that using option-82 requires option-82 aware DHCP servers.

9.1.2. Address Resolution

SSs MUST use Address Resolution Protocol (ARP) [RFC0826] for finding an Ethernet MAC address of destination.

9.2. IPv6 over Ethernet

[RFC2464] defines transmission of IPv6 Packets over Ethernet Networks. In this document, encapsulation of IPv6 packets into Ethernet frames and mapping rules for IPv6 address to Ethernet address (i.e. MAC address) MUST follow [RFC2464].

9.2.1. Router Discovery, Prefix Discovery and Parameter Discovery

Router Discovery, Prefix Discovery and Parameter Discovery procedures are achieved by receiving Router Advertisement messages. In this specification, SSs perform above the discovery process by solicited Router Advertisement messages because periodic Router Advertisement messages are discarded on the network-side bridge following the Broadcast Data Forwarding Rules in Section 6.1.2.

9.2.2. Address Configuration

9.2.2.1. Stateful Address Autoconfiguration

When the’M’ flag in the received RA is set, a SS SHOULD perform stateful address configuration according to [RFC3315]. In this case, an AR supports DHCPv6 server or relay function.

9.2.2.2. Stateless Address Autoconfiguration

SS SHOULD derive its global IPv6 addresses based on prefix and EUI-64-derived interface identifier and then confirmed through Duplicate Address Detection (DAD) as specified in [RFC4862] and [RFC4861].

9.2.3. Address Resolution

SS SHOULD send Neighbor Solicitation destined for solicited-node address corresponding to the target address in order to determine the MAC address of a neighbor and then resolve the MAC address by receiving Neighbor Advertisement as specified in [RFC4861].

9.3. Maximum Transmission Unit Consideration

[RFC2460] mandates 1280 bytes as a minimum Maximum Transmission Unit (MTU) size for 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 and encourages to support full-length packets. Therefore, IEEE 802.16 frame SHOULD support for carrying 1518 bytes payload that includes 18 bytes Ethernet header and 1500 bytes IP packet.

In the deployment scenarios of IP over Ethernet over IEEE 802.16, it is likely that the link between BS and network-side bridge is implemented by GRE or VLAN because the WiMAX Forum has chosen GRE for the mobile WiMAX architecture and VLAN works well with conventional Ethernet technologies.

In the case of GRE-based implementation, it does not introduce additional considerations for MTU size. GRE is able to carry any size of packet as IP is able to fragment and reassemble packets exceeding the MTU of the underlying transport.

However, when VLAN is implemented in the link between a BS and a network-side bridge, there may be restrictions on the supported packet size. The bridge adds VLAN tags to untagged Ethernet frame and increases the length of the original Ethernet frame by 4 bytes each VLAN tag, which may cause the Ethernet frame to be discarded in the link between the bridge and an AR. Therefore, the network operator should consider the size of stacked VLAN tags when implementing VLAN and setting the MTU of the link. In IPv6 case, the AR can advertise the MTU through router advertisement as defined in [RFC4861]. If MTU is advertised through router advertisement, the SS SHOULD use the MTU from the router advertisement.

10. IANA Considerations

This document has no actions for IANA.

11. Security Considerations

This document does not introduce new vulnerability to operations of IPv4 over Ethernet and IPv6 over Ethernet. [RFC3971] can be adopted for securing neighbor discovery processes.

12. Acknowledgments

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

13. References

13.1. Normative References

[IEEE802.16]
IEEE Std 802.16-2004, "IEEE Standard for Local and metropolitan area networks, Part 16: Air Interface for Fixed Broadband Wireless Access Systems", October 2004.

[IEEE802.16e]
IEEE Std 802.16e-2005, "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", December 2005.

[IEEE802.16k]
IEEE Std 802.16k-2007, "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges, Amendment5 for Bridging of IEEE 802.16", March 2007.

[IEEE802.1D]
IEEE Std 802.1D-2004, "IEEE Standard for Local and metropolitan area networks, Media Access Control (MAC) Bridges", June 2004.

[IEEE802.1Q]
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.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.

[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.

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

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, 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.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 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

[I-D.ietf-16ng-ps-goals]
Jee, J., Madanapalli, S., Mandin, J., and S. Park, "IP over 802.16 Problem Statement and Goals", draft-ietf-16ng-ps-goals-03 (work in progress), November 2007.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[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.

[TR101] DSL Forum, "Migration to Ethernet-Based DSL Aggregation", April 2006.

Appendix A. Multicast CID Deployment Considerations

IEEE 802.16 allows for downlink CIDs associated to multiple SSs to support efficient transport of multicast and broadcast data. Broadcast CIDs are used by IEEE 802.16 for MAC signaling messages like frame synchronization or messages describing the allocation of tones within a frame to particular CIDs. Such information is transferred over a broadcast connection and received by all associated subscriber stations in parallel. It is also possible to establish multicast connections by assigning a downlink CID to a number of subscriber stations. MAC messages sent to a CID with multiple subscribers are received and decoded in parallel by subscibed stations.

Multicast CIDs are highly efficient means to distribute the same information in parallel to a high number of subscribers under the same base station. The deployment of multicast CIDs for multicast and broadcast services requires a standardized mechanism for establishing and maintaining the multicast CIDs and the association of the multicast CIDs with multicast and broadcast services. Such a protocol is not yet available but under development by the Networking Working Group of the WiMAX Forum.

A drawback of multicast CIDs for Ethernet over IEEE802.16 is the unidirectional nature of multicast CIDs. While it is possible to multicast information downstream to a number of stations in parallel, there are no upstream multicast connections. In upstream direction unicast CIDs have to be used for sending multicast messages over the air to the basestation requiring a special multicast forwarding function in the BS 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 available. Unfortunately IEEE802.1D cant be applied 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 wake-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 services, it is harmful for link layer broadcast control messages aimed for a single station, like an ARP Request. All other stations are wasting scarce battery power for receiving, decoding and discarding the message. Low power consumption is an extremely important aspect in a wireless communication system and it is necessary to protect subscriber stations from denial of service attacks by wasting battery power due to malicious ARP requests.

Furthermore it should kept in mind that multicast CIDs are only efficient for a large number of subscribed stations 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 resource for transferring the same information as a unicast connections

Appendix B. Distributed Bridging Considerations

A large Ethernet link can be realized by cascading smaller bridges. This behavior would allow the network-side bridging function to be realized by a bridge connecting bridges integrated with the BSs. While this works for the plain Ethernet behavior, it introduces some drawbacks and even potential security threats for the transmission of IP over Ethernet over IEEE 802.16.

The Proxy ARP function described in Section 6.4 prevents that ARP broadcast messages have to be forwarded to each of the associated SSs, when the ARP proxy is aware of the existence of the queried IP address at one of the bridge ports. If the queried IP address is not known to ARP proxy the bridge has to flood all its ports with the ARP request.

Distributing the bridging function into the BSs would imply that the Proxy ARP function is split into multiple Proxy ARP functions each knowing only about the subset of the IP addresses, which are directly connected by the BS. IP addresses belonging to the same link but being connected to other BSs would not be known to the Proxy ARP functions and would cause that ARP requests for these IP addresses are broadcasted to all SSs. This causes a huge waste of radio resources for transmitting ARP requests and potentially more critical even, it would waste scarce battery power in the SSs.

A malicious user would be able to deploy this behavior for denial of service attacks by exhausting the batteries of SSs by just sending ARP Requests.

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

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

Phone: +49-89-636-75194
Email: maximilian.riegel@nsn.com

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

Phone: +82-42-860-1877
Email: sjjeong@gmail.com

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).