Internet Engineering Task Force (IETF)
Request for Comments: 5948
Category: Standards Track
ISSN: 2070-1721

S. Madanapalli
iRam Technologies
S. Park
Samsung Electronics
S. Chakrabarti
IP Infusion
G. Montenegro
Microsoft Corporation

August 2010

Transmission of IPv4 Packets over the IP Convergence Sublayer of IEEE 802.16

Abstract

IEEE 802.16 is an air interface specification for wireless broadband access. IEEE 802.16 has specified multiple service-specific Convergence Sublayers for transmitting upper-layer protocols. The Packet CS (Packet Convergence Sublayer) is used for the transport of all packet-based protocols such as the Internet Protocol (IP) and IEEE 802.3 (Ethernet). The IP-specific part of the Packet CS enables the transport of IPv4 packets directly over the IEEE 802.16 Media Access Control (MAC) layer.

This document specifies the frame format, the Maximum Transmission Unit (MTU), and the address assignment procedures for transmitting IPv4 packets over the IP-specific part of the Packet Convergence Sublayer of IEEE 802.16.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc5948.

Table of Contents

1. Introduction

IEEE 802.16 [IEEE802_16] is a connection-oriented access technology for the last mile. The IEEE 802.16 specification includes the Physical (PHY) and Media Access Control (MAC) layers. The MAC layer includes various Convergence Sublayers (CSs) for transmitting higher-layer packets, including IPv4 packets [IEEE802_16].

The scope of this specification is limited to the operation of IPv4 over the IP-specific part of the Packet CS (referred to as "IPv4 CS") for hosts served by a network that utilizes the IEEE Std 802.16 air interface.

This document specifies a method for encapsulating and transmitting IPv4 [RFC0791] packets over the IPv4 CS of IEEE 802.16. This document also specifies the MTU and address assignment method for hosts using IPv4 CS.

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].

2. Terminology

- Mobile Station (MS) -- The term "MS" is used to refer to an IP host. This usage is more informal than that in IEEE 802.16, in which "MS" refers to the interface implementing the IEEE 802.16 MAC and PHY layers and not to the entire host.

- Last mile -- The term "last mile" is used to refer to the final leg of delivering connectivity from a communications provider to a customer.

Other terminology in this document is based on the definitions in [RFC5154].

3. Typical Network Architecture for IPv4 over IEEE 802.16

The network architecture follows what is described in [RFC5154] and [RFC5121]. Namely, each MS is attached to an Access Router (AR) through a Base Station (BS), a Layer 2 (L2) entity (from the perspective of the IPv4 link between the MS and the AR).

For further information on the typical network architecture, see [RFC5121], Section 5.

3.1. IEEE 802.16 IPv4 Convergence Sublayer Support

As described in [IEEE802_16], the IP-specific part of the Packet CS allows the transmission of either IPv4 or IPv6 payloads. In this document, we are focusing on IPv4 over the Packet Convergence Sublayer.

For further information on the IEEE 802.16 Convergence Sublayer and encapsulation of IP packets, see Section 4 of [RFC5121] and [IEEE802_16].

4. IPv4 CS Link in 802.16 Networks

In 802.16, the transport connection between an MS and a BS is used to transport user data, i.e., IPv4 packets in this case. A transport connection is represented by a service flow, and multiple transport connections can exist between an MS and a BS.

When an AR and a BS are co-located, the collection of transport connections to an MS is defined as a single IPv4 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 IPv4 link.

Each host belongs to a different IPv4 link and is assigned a unique IPv4 address, similar to the recommendations discussed in "Analysis of IPv6 Link Models for IEEE 802.16 Based Networks" ([RFC4968]).

4.1. IPv4 CS Link Establishment

In order to enable the sending and receiving of IPv4 packets between the MS and the AR, the link between the MS and the AR via the BS needs to be established. This section explains the link establishment procedure, as described in Section 6.2 of [RFC5121]. Steps 1-4 are the same as those indicated in Section 6.2 of [RFC5121]. In step 5, support for IPv4 is indicated. In step 6, a service flow is created that can be used for exchanging IP-layer signaling messages, e.g., address assignment procedures using DHCP.

4.2. Frame Format for IPv4 Packets

IPv4 packets are transmitted in Generic IEEE 802.16 MAC frames in the data payloads of the 802.16 PDU (see Section 3.2 of [RFC5154]).

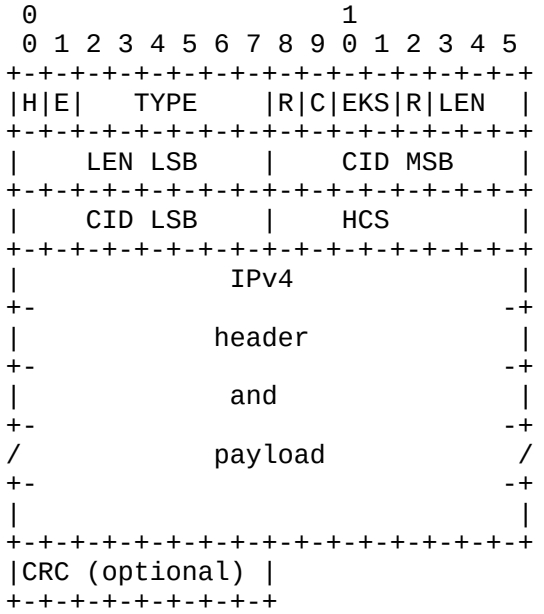

Figure 1. IEEE 802.16 MAC Frame Format for IPv4 Packets

Here, "MSB" means "most significant byte", and "LSB" means "least significant byte".

H: Header Type (1 bit). Shall be set to zero, indicating that it is a Generic MAC PDU.

E: Encryption Control. 0 = Payload is not encrypted; 1 = Payload is encrypted.

R: Reserved. Shall be set to zero.

C: Cyclic Redundancy Check (CRC) Indicator. 1 = CRC is included; 0 = No CRC is included.

EKS: Encryption Key Sequence.

LEN: The Length, in bytes, of the MAC PDU, including the MAC header and the CRC, if present (11 bits).

CID: Connection Identifier (16 bits).

HCS: Header Check Sequence (8 bits).

CRC: An optional 8-bit field. The CRC is appended to the PDU after encryption.

TYPE: This field indicates the subheaders (Mesh subheader, Fragmentation subheader, Packing subheader, etc.) and special payload types (e.g., Automatic Repeat reQuest (ARQ)) present in the message payload.

4.3. Maximum Transmission Unit

The MTU value for IPv4 packets on an IEEE 802.16 link is configurable (e.g., see the end of this section for some possible mechanisms). The default MTU for IPv4 packets over an IEEE 802.16 link SHOULD be 1500 octets. Given the possibility for "in-the-network" tunneling, supporting this MTU at the end hosts has implications on the underlying network, for example, as discussed in [RFC4459].

Per [RFC5121], Section 6.3, the IP MTU can vary to be larger or smaller than 1500 octets.

If an MS transmits 1500-octet packets in a deployment with a smaller MTU, packets from the MS may be dropped at the link layer silently. Unlike IPv6, in which departures from the default MTU are readily advertised via the MTU option in Neighbor Discovery (via router advertisement), there is no similarly reliable mechanism in IPv4, as

the legacy IPv4 client implementations do not determine the link MTU by default before sending packets. Even though there is a DHCP option to accomplish this, DHCP servers are required to provide the MTU information only when requested.

Discovery and configuration of the proper link MTU value ensures adequate usage of the network bandwidth and resources. Accordingly, deployments should avoid packet loss due to a mismatch between the default MTU and the configured link MTUs.

Some of the mechanisms available for the IPv4 CS host to find out the link's MTU value and mitigate MTU-related issues are:

- Recent revision of 802.16 by the IEEE (see IEEE 802.16-2009 [IEEE802_16]) to (among other things) allow the provision of the Service Data Unit or MAC MTU in the IEEE 802.16 SBC-REQ/SBC-RSP phase, such that clients that are compliant with IEEE 802.16 can infer and configure the negotiated MTU size for the IPv4 CS link. However, the implementation must communicate the negotiated MTU value to the IP layer to adjust the IP Maximum Payload Size for proper handling of fragmentation. Note that this method is useful only when the MS is directly connected to the BS.

- Configuration and negotiation of MTU size at the network layer by using the DHCP interface MTU option [RFC2132].

This document recommends that implementations of IPv4 and IPv4 CS clients SHOULD use the DHCP interface MTU option [RFC2132] in order to configure its interface MTU accordingly.

In the absence of DHCP MTU configuration, the client node (MS) has two alternatives: 1) use the default MTU (1500 bytes), or 2) determine the MTU by the methods described in IEEE 802.16-2009 [IEEE802_16].

Additionally, the clients are encouraged to run Path MTU (PMTU) Discovery [RFC1191] or Packetization Layer Path MTU Discovery (PLPMTUD) [RFC4821]. However, the PMTU mechanism has inherent problems of packet loss due to ICMP messages not reaching the sender and IPv4 routers not fragmenting the packets due to the Don't Fragment (DF) bit being set in the IP packet. The above-mentioned path MTU mechanisms will take care of the MTU size between the MS and its correspondent node across different flavors of convergence layers in the access networks.

5. Subnet Model and IPv4 Address Assignment

The subnet model recommended for IPv4 over IEEE 802.16 using IPv4 CS is based on the point-to-point link between the MS and the AR [RFC4968]; hence, each MS shall be assigned an address with a 32-bit prefix length or subnet mask. The point-to-point link between the MS and the AR is achieved using a set of IEEE 802.16 MAC connections (identified by service flows) and an L2 tunnel (e.g., a Generic Routing Encapsulation (GRE) tunnel) for each MS between the BS and the AR. If the AR is co-located with the BS, then the set of IEEE 802.16 MAC connections between the MS and the BS/AR represent the point-to-point connection.

The "next hop" IP address of the IPv4 CS MS is always the IP address of the AR, because the MS and the AR are attached via a point-to-point link.

5.1. IPv4 Unicast Address Assignment

DHCP [RFC2131] SHOULD be used for assigning an IPv4 address for the MS. DHCP messages are transported over the IEEE 802.16 MAC connection to and from the BS and relayed to the AR. In case the DHCP server does not reside in the AR, the AR SHOULD implement a DHCP relay agent [RFC1542].

5.2. Address Resolution Protocol

The IPv4 CS does not allow for transmission of Address Resolution Protocol (ARP) [RFC0826] packets. Furthermore, in a point-to-point link model, address resolution is not needed.

5.3. IP Broadcast and Multicast

Multicast or broadcast packets from the MS are delivered to the AR via the BS through the point-to-point link. This specification simply assumes that the broadcast and multicast services are provided. How these services are implemented in an IEEE 802.16 Packet CS deployment is out of scope of this document.

6. Security Considerations

This document specifies transmission of IPv4 packets over IEEE 802.16 networks with the IPv4 Convergence Sublayer and does not introduce any new vulnerabilities to IPv4 specifications or operation. The security of the IEEE 802.16 air interface is the subject of [IEEE802_16]. In addition, the security issues of the network

architecture spanning beyond the IEEE 802.16 Base Stations is the subject of the documents defining such architectures, such as the Worldwide Interoperability for Microwave Access (WiMAX) network architecture [WMF].

7. Acknowledgements

The authors would like to acknowledge the contributions of Bernard Aboba, Dave Thaler, Jari Arkko, Bachet Sarikaya, Basavaraj Patil, Paolo Narvaez, and Bruno Sousa for their review and comments. The working group members Burcak Beser, Wesley George, Max Riegel, and DJ Johnston helped shape the MTU discussion for the IPv4 CS link. Thanks to many other members of the 16ng Working Group who commented on this document to make it better.

8. References

8.1. Normative References

[IEEE802_16] "IEEE Std 802.16-2009, Draft Standard for Local and Metropolitan area networks, Part 16: Air Interface for Broadband Wireless Access Systems", May 2009.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[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.

[RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, October 1993.

[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.

8.2. Informative References

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

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC4459] Savola, P., "MTU and Fragmentation Issues with In-the-Network Tunneling", RFC 4459, April 2006.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.

[RFC4840] Aboba, B., Davies, E., and D. Thaler, "Multiple Encapsulation Methods Considered Harmful", RFC 4840, April 2007.

[RFC4968] Madanapalli, S., "Analysis of IPv6 Link Models for 802.16 Based Networks", RFC 4968, August 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.

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

[WMF] "WiMAX End-to-End Network Systems Architecture Stage 2-3 Release 1.2, http://www.wimaxforum.org/", January 2008.

Appendix A. Multiple Convergence Layers -- Impact on Subnet Model

Two different MSs using two different Convergence Sublayers (e.g., an MS using Ethernet CS only and another MS using IPv4 CS only) cannot communicate at the data link layer and require interworking at the IP layer. For this reason, these two nodes must be configured to be on two different subnets. For more information, refer to [RFC4840].

Appendix B. Sending and Receiving IPv4 Packets

IEEE 802.16 MAC is a point-to-multipoint connection-oriented air interface, and the process of sending and receiving IPv4 packets is different from multicast-capable shared-medium technologies like Ethernet.

Before any packets are transmitted, an IEEE 802.16 transport connection must be established. This connection consists of an IEEE 802.16 MAC transport connection between the MS and the BS and an L2 tunnel between the BS and the AR (if these two are not co-located). This IEEE 802.16 transport connection provides a point-to-point link between the MS and the AR. All the packets originating at the MS always reach the AR before being transmitted to the final destination.

IPv4 packets are carried directly in the payload of IEEE 802.16 frames when the IPv4 CS is used. IPv4 CS classifies the packet based on upper-layer (IP and transport layers) header fields to place the packet on one of the available connections identified by the CID. The classifiers for the IPv4 CS are source and destination IPv4 addresses, source and destination ports, Type-of-Service, and IP Protocol field. The CS may employ Packet Header Suppression (PHS) after the classification.

The BS optionally reconstructs the payload header if PHS is in use. It then tunnels the packet that has been received on a particular MAC connection to the AR. Similarly, the packets received on a tunnel interface from the AR would be mapped to a particular CID using the IPv4 classification mechanism.

The AR performs normal routing for the packets that it receives, processing them per its forwarding table. However, the DHCP relay agent in the AR MUST maintain the tunnel interface on which it receives DHCP requests so that it can relay the DHCP responses to the correct MS. The particular method is out of scope of this specification as it need not depend on any particularities of IEEE 802.16.

Appendix C. WiMAX IPv4 CS MTU Size

The WiMAX (Worldwide Interoperability for Microwave Access) forum has defined a network architecture [WMF]. Furthermore, WiMAX has specified IPv4 CS support for transmission of IPv4 packets between the MS and the BS over the IEEE 802.16 link. The WiMAX IPv4 CS and this specification are similar. One significant difference, however, is that the WiMAX Forum [WMF] has specified the IP MTU as 1400 octets [WMF] as opposed to 1500 in this specification.

Hence, if an IPv4 CS MS configured with an MTU of 1500 octets enters a WiMAX network, some of the issues mentioned in this specification may arise. As mentioned in Section 4.3, the possible mechanisms are not guaranteed to work. Furthermore, an IPv4 CS client is not capable of doing ARP probing to find out the link MTU. On the other hand, it is imperative for an MS to know the link MTU size. In practice, an MS should be able to sense or deduce the fact that it is operating within a WiMAX network (e.g., given the WiMAX-specific particularities of the authentication and network entry procedures), and adjust its MTU size accordingly. Even though this method is not perfect, and the potential for conflict may remain, this document recommends a default MTU of 1500. This represents the WG's consensus (after much debate) to select the best value for IEEE 802.16 from the point of view of the IETF, in spite of the WiMAX Forum's deployment.

Authors' Addresses

Syam Madanapalli
iRam Technologies
#H304, Shriram Samruddhi, Thubarahalli
Bangalore - 560066
India

EMail: smadanapalli@gmail.com

Soohong Daniel Park
Samsung Electronics
416 Maetan-3dong, Yeongtong-gu
Suwon 442-742
Korea

EMail: soohong.park@samsung.com

Samita Chakrabarti
IP Infusion
1188 Arques Avenue
Sunnyvale, CA
USA

EMail: samitac@ipinfusion.com

Gabriel Montenegro
Microsoft Corporation
Redmond, WA
USA

EMail: gabriel.montenegro@microsoft.com