Project	IEEE 802.16 Broadband Wireless Access Working Group <**http://ieee802.org/16**>
Title	Resolving ambiguities in the definition of CS PDU Formats
Date Submitted	2005-03-14
Source(s)	David A Castelow, Mark Thomas, Eyal Verbin Airspan Communications, Cambridge House, Oxford Road, Uxbridge, UK Voice: +44 1895 467281 Fax: +44 1895 467202 mailto:dcastelow@airspan.com Baraa Al Dabagh, Yigal Eliaspur Intel baraa.al.dabagh@intel.com
Re:	Supporting document for Comment to 802.16maint.
Abstract	There is ambiguity in the Ethernet and VLAN CS as to the data that is transported. This document provides resolution to these ambiguities.
Purpose	The document is intended for consideration within the comments resolution process.
Notice	
Release	
Patent Policy and Procedures	

Resolving ambiguities in the definition of CS PDU Formats

David Castelow, Mark Thomas, Eyal Verbin, Airspan Networks
Baraa Al Dabagh, Yigal Eliaspur, Intel
March 2005

References

[1] IEEE, "IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems," IEEE Std 802.16-2004.
[2] IEEE, "IEEE Draft Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems," IEEE P802.16-REVd/D5-2004.
[3] IEEE, "Corrigendum to IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems," 80216maint-04/10.

Introduction

The changes proposed in this document are to resolve ambiguity in the definition of the convergence sublayer PDU as described in IEEE 802.16-2004 [1, 2].

Description of problem

IEEE 802.16-2004 does not satisfactorily specify if the PDU exchanged between peer convergence sublayers contains the 32-bit Ethernet FCS for cases where the convergence sublayer type is Ethernet, VLAN, IP over Ethernet or IP over VLAN. If we fail to resolve this ambiguity, then the receiving higher-layer entity cannot determine if the last four octets in the service data unit (SDU) at the convergence sublayer service access point (CS SAP) represents FCS or data, and so BS and SS based on different interpretations will not be interoperable. It is certainly not satisfactory for inclusion of the FCS to be optional unless there exists an associated mechanism so that BS and SS can advertise their respective capabilities.

In addition to technical changes, this contribution includes editorial changes because the 802.1Q standard referred to in the body of 802.16-2004 is out of date and does not match the document in the list of references.

Proposed solution

There is a strong argument, based on analogy with other protocols in the 802 family, that the FCS must *not* be included. As a general rule, FCS must be calculated and appended in an 802 MAC layer when transmitting a frame, and checked in the peer MAC layer when receiving a frame. Received frames with incorrect FCS must be dropped in the MAC layer. The exact procedure for calculating FCS is specified independently for some of the MAC-layer protocols in the 802 family. We consider that the optional CRC applied in the MAC common part sublayer (CPS) of IEEE 802.16-2004 performs function analogous to the FCS used in other protocols in the 802 family.

Analogy with other protocols in the 802 family

We note that Section 6.4 of IEEE 802.1D-2004 specifies the general form of the SDU for the MAC SDU in an internal sublayer service for MAC relay in an 802 bridge. Here, M_UNITDATA is defined as (frame_type, destination_address, source_address, mac_service_data_unit, user_ priority, frame_check_sequence). This definition of the SDU may appear to be contrary to our argument as presented above. However, Section 6.5 of the same standard describes procedures for individual MAC protocols, and these show that the FCS is provided at the MAC SAP as a convenience for some protocols, and is irrelevant for others. Of course, 802.16 is not addressed in 802.1D, and so we have no definitive answer for the most interesting case.

Section 6.5.1 describes 802.3; the 802.3 MAC layer is allowed to re-use an FCS submitted in the SDU. If the FCS is missing, the 802.3 MAC layer must calculate a new FCS. We conclude that calculation of the FCS is naturally a function of the MAC layer, but that re-use of an FCS from the internal sublayer in a bridge is an optimisation to avoid unnecessary computation. Annex F offers an interesting (but only informative) perspective on minimising additional error in the case where an FCS can be re-used.

Section 6.5.4 describes 802.11; here the FCS need not be included in the SDU at the MAC SAP, because FCS is always calculated in the MAC layer using a method peculiar to (and defined in) 802.11. We consider that 802.16 should be treated like 802.11 in the sense that it defines a specialised method of calculating CRC/FCS within its own MAC layer definition.

Technical arguments demonstrating that the additional FCS is unnecessary

The case for including FCS is to ensure that the reconstructed frame suffers from an undetected error rate that is no greater than that of Ethernet. It is clear that to achieve this then a similar level of error protection is required. The CRCs used in 802.16 are the same as those used in Ethernet and therefore the protection offered is comparable. To provide service meeting the requirements of 802.1D, an equipment supplier needs to enable CRC calculation on those service flows. We do not see the need to mandate this behaviour in the standard, though we would strongly recommend its use for Ethernet and IP CS.

There is a subtlety, as Ethernet frames may (and frequently will be) fragmented and transmitted, and the 802.16 CRC is applied to the fragment (plus its header) rather than on the frame. Fragments are transmitted with a Fragment Sequence Number (FSN) that is either 3 bits (modulo 8) or 11 bits (modulo 2048). Within a fragmented frame, the first and last fragments are specially marked, and intermediate fragments are marked as “middle”. This gives rise to the possibility that 8 or 2048 intermediate fragments may be received in error and discarded, resulting in a reconstructed frame that is erroneous. The probability of this event depends on the probability of fragment error: a graph of this is shown in Figure 1. This plot is based on an assumption of separate fragments being independent, uncorrelated events. It can be seen that for fragment error rates below about 4% the probability of error is lower than that of the overall FCS, irrespective of the FSN scheme adopted. In fact the probabilities are lower than this, because, in order that the frame error is undetected, the frame length in the first fragment must match the received and reconstructed frame, and this length change must be undetected, automatically reducing the probability of the undetected frame error rate by *at least* the probability of an undetected fragment.

An additional case still needs considering, and that is that the loss of fragments is a result of an outage and that therefore the probabilities are correlated. A simple model is that the duration of an outage is uniformly distributed. The probability of an undetected frame error depends on the system loosing exactly 8 or 2048 fragments, so on this basis the probability is 1/8 or 1/2048 times the probability of the outage event itself. However, the outage must also result in an undetected change in length in the first fragment that matches the reconstructed frame length. The probability of this event is therefore less than the probability of an undetected error, which, as has already been stated, is the same as the rate for FCS failure in Ethernet.

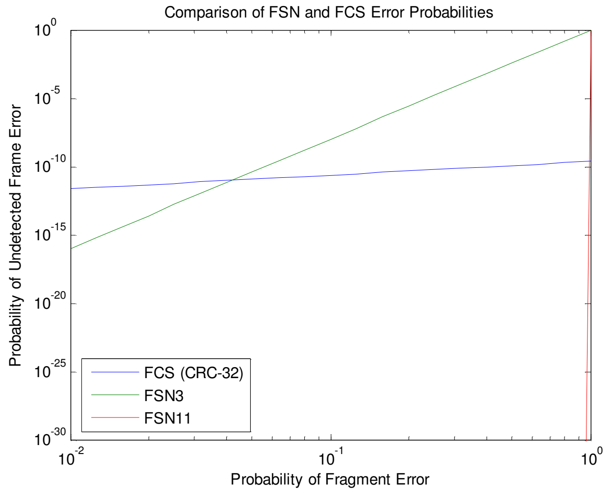

Figure 1. Probability of Undetected Frame Error as a function of the Probability of Fragment Error.

Technical Text Changes

At page 8, line 5, include the following:

[*Add the following paragraph at the end of section 5.2.4.1:*]

The IEEE Std. 802.3/Ethernet PDU consists of the following fields: Destination MAC address, source MAC address, length/type, data. The Ethernet frame check sequence (FCS) does not form part of the Ethernet PDU in the CS.

At page 8, line 25, include the following:

[*Add the following paragraph at the end of section 5.2.5.1:*]

The IEEE Std. 802.1Q/VLAN tagged frame consists of the following fields: Destination MAC address, source MAC address, length/type, tag control information, data. The Ethernet frame check sequence (FCS) does not form part of the VLAN tagged frame in the CS.

Editorial Text Changes

Replace page 8, line 14, with the following:

5.2.5 IEEE Std 802.1Q-~~1998~~2003 virtual local area network (VLAN) specific part

This CS shall be employed when IEEE Std 802.1Q-~~1998~~2003 tagged VLAN frames are to be carried over the IEEE Std 802.16 network.

Page 8, line 16, modify as follows:

5.2.5.1 IEEE Std 802.1Q-~~1998~~2003 VLAN CS PDU format

Page 8, line 21, modify as follows:

The format of the IEEE Std 802.1Q-~~1998~~2003 VLAN CS PDU shall be as shown in Figure 14 (when header suppression is enabled at the connection but not applied to the CS PDU) or Figure 15 (with header suppression). In the case PHS is not enabled, PHSI field shall be omitted.

Replace Page 8, line 26 to line 34 with the following:

5.2.5.2 IEEE Std 802.1Q-~~1998~~2003 CS classifiers

The following parameters are relevant for IEEE Std 802.1Q-~~1998~~2003 CS classifiers:

~~LC~~Ethernet header classification parameters—zero or more of the ~~LLC~~Ethernet header classification parameters (Destination MAC address, source MAC address, Ethertype/SAP).

IEEE Std 802.1D-~~1998~~2003Parameters—zero or more of the IEEE classification parameters (IEEE Std 802.1D-~~1998~~2003Priority Range, IEEE Std 802.1Q-~~1998~~2003VLAN ID).

For IP over IEEE Std 802.1Q-~~1998~~2003 VLAN, IP headers may be included in classification. In this case, the IP classification parameters (11.13.19.3.4.2—11.13.19.3.4.7) are allowed.

Additional information

Support for native IP

IEEE 802.16-2004 allows for a convergence sublayer type of IP, meaning that IP datagrams are carried over 802.16 without encapsulation in Ethernet headers. IP does not provide protection against transmission errors, because this is assumed to be the responsibility of layer 2. It follows from this that 802.16 must inherently be able to reliably detect transmission errors and discard errored packets without depending on the Ethernet FCS. If 802.16 L2 transmission is reliable, then the Ethernet FCS is redundant; if 802.16 L2 transmission is unreliable then this needs to be addressed in the 802.16 MAC layer.

Secondary management channel

If the secondary management channel carries Ethernet frames, then FCS must not be included in secondary management traffic.