Network Working Group
Request for Comments: 5458
Category: Informational

H. Cruickshank
University of Surrey
P. Pillai
University of Bradford
M. Noisternig
University of Salzburg
S. Iyengar
Logica
March 2009

Security Requirements for the Unidirectional Lightweight Encapsulation (ULE) Protocol

Status of This Memo

Copyright Notice

Abstract

The MPEG-2 standard defined by ISO 13818-1 supports a range of transmission methods for a variety of services. This document provides a threat analysis and derives the security requirements when using the Transport Stream, TS, to support an Internet network-layer using Unidirectional Lightweight Encapsulation (ULE) defined in RFC 4326. The document also provides the motivation for link-layer security for a ULE Stream. A ULE Stream may be used to send IPv4 packets, IPv6 packets, and other Protocol Data Units (PDUs) to an arbitrarily large number of Receivers supporting unicast and/or multicast transmission.

The analysis also describes applicability to the Generic Stream Encapsulation (GSE) defined by the Digital Video Broadcasting (DVB) Project.

Table of Contents

1. Introduction

The MPEG-2 Transport Stream (TS) has been widely accepted not only for providing digital TV services, but also as a subnetwork technology for building IP networks. RFC 4326 [RFC4326] describes the Unidirectional Lightweight Encapsulation (ULE) mechanism for the transport of IPv4 and IPv6 Datagrams and other network protocol packets directly over the ISO MPEG-2 Transport Stream as TS Private Data. ULE specifies a base encapsulation format and supports an Extension Header format that allows it to carry additional header information to assist in network/Receiver processing. The encapsulation satisfies the design and architectural requirement for a lightweight encapsulation defined in RFC 4259 [RFC4259].

Section 3.1 of RFC 4259 presents several topological scenarios for MPEG-2 Transmission Networks. A summary of these scenarios is presented below:

A. Broadcast TV and Radio Delivery. This is not within the scope of this document.

B. Broadcast Networks used as an ISP. This resembles scenario A, but includes IP services to access the public Internet.

C. Unidirectional Star IP Scenario. This provides a data network delivering a common bit stream to typically medium-sized groups of Receivers.

D. Datacast Overlay. This employs MPEG-2 physical and link layers to provide additional connectivity such as unidirectional multicast to supplement an existing IP-based Internet service.

E. Point-to-Point Links. This connectivity may be provided using a pair of transmit and receive interfaces.

F. Two-Way IP Networks.

RFC 4259 states that ULE must be robust to errors and security threats. Security must also consider both unidirectional (A, B, C, and D) as well as bidirectional (E and F) links for the scenarios mentioned above.

An initial analysis of the security requirements in MPEG-2 transmission networks is presented in the "Security Considerations" section of RFC 4259. For example, when such networks are not using a wireline network, the normal security issues relating to the use of wireless links for transport of Internet traffic should be considered [RFC3819].

The security considerations of RFC 4259 recommend that any new encapsulation defined by the IETF should allow Transport Stream encryption and should also support optional link-layer authentication of the Subnetwork Data Unit (SNDU) payload. In ULE [RFC4326], it is suggested that this may be provided in a flexible way using Extension Headers. This requires the definition of a mandatory Extension Header, but has the advantage that it decouples specification of the security functions from the encapsulation functions.

This document extends the above analysis and derives in detail the security requirements for ULE in MPEG-2 transmission networks.

A security framework for deployment of secure ULE networks describing the different building blocks and the interface definitions is presented in Appendix A.

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Other terms used in this document are defined below:

ATSC: Advanced Television Systems Committee. A framework and a set of associated standards for the transmission of video, audio, and data using the ISO MPEG-2 Standard.

DVB: Digital Video Broadcast. A framework and set of associated standards published by the European Telecommunications Standards Institute (ETSI) for the transmission of video, audio, and data using the ISO MPEG-2 Standard [ISO-MPEG2].

Encapsulator: A network device that receives Protocol Data Units (PDUs) and formats these into Payload Units (known here as SNDUs) for output as a stream of TS Packets.

GCKS: Group Controller and Key Server. A server that authenticates and provides the policy and keying material to members of a secure group.

LLC: Logical Link Control [ISO-8802], [IEEE-802]. A link-layer protocol defined by the IEEE 802 standard, which follows the Ethernet Medium Access Control Header.

MAC: Message Authentication Code.

MPE: Multiprotocol Encapsulation [ETSI-DAT]. A scheme that encapsulates PDUs, forming a Digital Storage Media Command and Control (DSM-CC) Table Section. Each Section is sent in a series of TS Packets using a single TS Logical Channel.

MPEG-2: A set of standards specified by the Motion Picture Experts Group (MPEG) and standardised by the International Standards Organisation (ISO/IEC 13818-1) [ISO-MPEG2], and ITU-T (in H.222 [ITU-H222]).

NPA: Network Point of Attachment. In this document, refers to a 6-byte destination address (resembling an IEEE Medium Access Control address) within the MPEG-2 transmission network that is used to identify individual Receivers or groups of Receivers.

PDU: Protocol Data Unit. Examples of a PDU include Ethernet frames, IPv4 or IPv6 Datagrams, and other network packets.

PID: Packet Identifier [ISO-MPEG2]. A 13-bit field carried in the header of TS Packets. This is used to identify the TS Logical Channel to which a TS Packet belongs [ISO-MPEG2]. The TS Packets forming the parts of a Table Section, Packetised Elementary Stream (PES), or other Payload Unit must all carry the same PID value. The all-zeros PID 0x0000 as well as other PID values are reserved for specific PSI/SI Tables [ISO-MPEG2]. The all-ones PID value 0x1FFF indicates a Null TS Packet introduced to maintain a constant bit rate of a TS Multiplex. There is no required relationship between the PID values used for TS Logical Channels transmitted using different TS Multiplexes.

Receiver: Equipment that processes the signal from a TS Multiplex and performs filtering and forwarding of encapsulated PDUs to the network-layer service (or bridging module when operating at the link layer).

SI Table: Service Information Table [ISO-MPEG2]. In this document, this term describes a table that is defined by another standards body to convey information about the services carried in a TS Multiplex. A Table may consist of one or more Table Sections; however, all sections of a particular SI Table must be carried over a single TS Logical Channel [ISO-MPEG2].

SNDU: SubNetwork Data Unit. An encapsulated PDU sent as an MPEG-2 Payload Unit.

TS: Transport Stream [ISO-MPEG2]. A method of transmission at the MPEG-2 layer using TS Packets; it represents Layer 2 of the ISO/OSI reference model. See also TS Logical Channel and TS Multiplex.

TS Multiplex: In this document, this term defines a set of MPEG-2 TS Logical Channels sent over a single lower-layer connection. This may be a common physical link (i.e., a transmission at a specified symbol rate, Forward Error Correction (FEC) setting, and transmission frequency) or an encapsulation provided by another protocol layer (e.g., Ethernet, or RTP over IP). The same TS Logical Channel may be repeated over more than one TS Multiplex (possibly associated with a different PID value) [RFC4259]; for example, to redistribute the same multicast content to two terrestrial TV transmission cells.

TS Packet: A fixed-length 188-byte unit of data sent over a TS Multiplex [ISO-MPEG2]. Each TS Packet carries a 4-byte header, plus optional overhead including an Adaptation Field, encryption details, and time stamp information to synchronise a set of related TS Logical Channels.

ULE Stream: An MPEG-2 TS Logical Channel that carries only ULE encapsulated PDUs. ULE Streams may be identified by definition of a stream_type in SI/PSI [ISO-MPEG2].

3. Threat Analysis

3.1. System Components

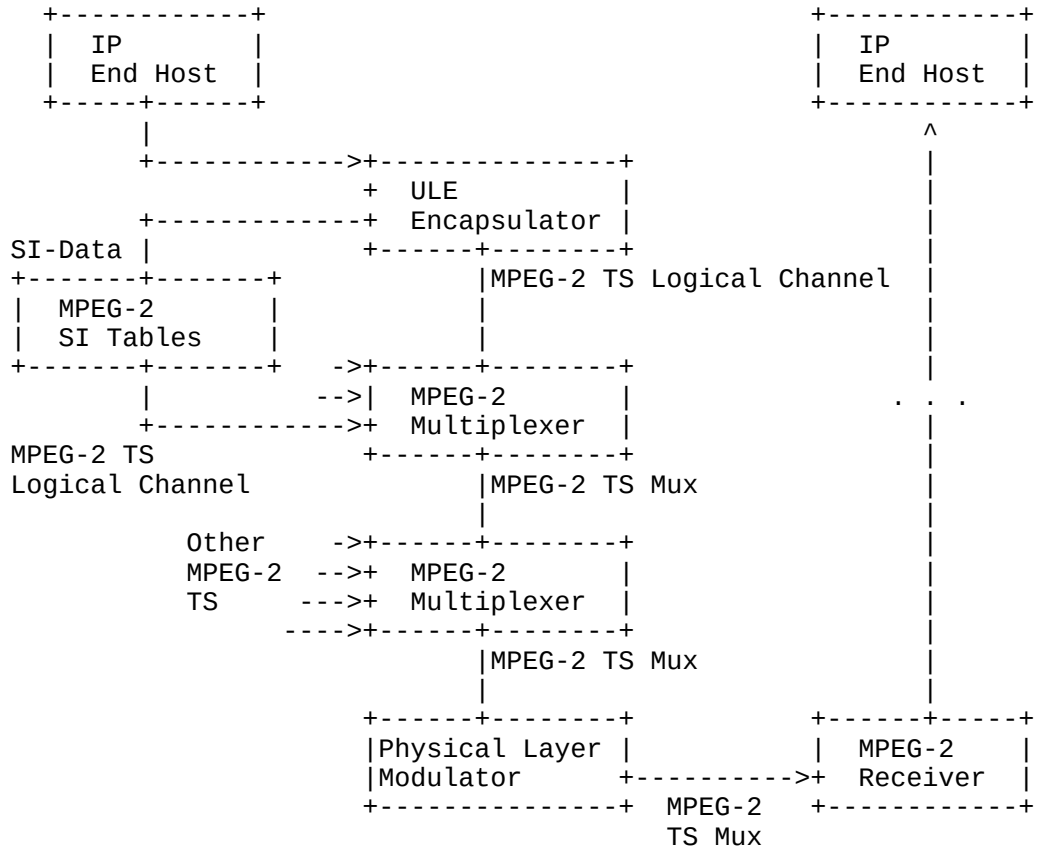

Figure 1: An example configuration for a unidirectional service for IP transport over MPEG-2 (adapted from [RFC4259])

As shown in Figure 1 above (from Section 3.3 of [RFC4259]), there are several entities within the MPEG-2 transmission network architecture. These include:

- ULE Encapsulation Gateways (the ULE Encapsulator)
- SI-Table signalling generator (input to the multiplexer)
- Receivers (the endpoints for ULE Streams)
- TS multiplexers (including re-multiplexers)

o Modulators

The TS Packets are carried to the Receiver over a physical layer that usually includes Forward Error Correction (FEC) coding that interleaves the bytes of several consecutive, but unrelated, TS Packets. FEC-coding and synchronisation processing makes injection of single TS Packets very difficult. Replacement of a sequence of packets is also difficult, but possible (see Section 3.2).

A Receiver in an MPEG-2 TS transmission network needs to identify a TS Logical Channel (or MPEG-2 Elementary Stream) to reassemble the fragments of PDUs sent by an L2 source [RFC4259]. In an MPEG-2 TS, this association is made via the Packet Identifier, PID [ISO-MPEG2]. At the sender, each source associates a locally unique set of PID values with each stream it originates. However, there is no required relationship between the PID value used at the sender and that received at the Receiver. Network devices may re-number the PID values associated with one or more TS Logical Channels (e.g., ULE Streams) to prevent clashes at a multiplexer between input streams with the same PID carried on different input multiplexes (updating entries in the PMT [ISO-MPEG2], and other SI tables that reference the PID value). A device may also modify and/or insert new SI data into the control plane (also sent as TS Packets identified by their PID value). However, there is only one valid source of data for each MPEG-2 Elementary Stream, bound to a PID value. (This observation could simplify the requirement for authentication of the source of a ULE Stream.)

In an MPEG-2 network, a set of signalling messages [RFC4947] may need to be broadcast (e.g., by an Encapsulation Gateway or other device) to form the L2 control plane. Examples of signalling messages include the Program Association Table (PAT), Program Map Table (PMT), and Network Information Table (NIT). In existing MPEG-2 transmission networks, these messages are broadcast in the clear (no encryption or integrity checks). The integrity as well as authenticity of these messages is important for correct working of the ULE network, i.e., supporting its security objectives in the area of availability, in addition to confidentiality and integrity. One method recently proposed [RFC5163] encapsulates these messages using ULE. In such cases all the security requirements of this document apply in securing these signalling messages.

ULE Stream security only concerns the security between the ULE Encapsulation Gateway (ULE Encapsulator) and the Receiver. In many deployment scenarios the user of a ULE Stream has to secure communications beyond the link since other network links are utilised in addition to the ULE link. Therefore, if authentication of the endpoints, i.e., the IP Sources, is required, or users are concerned

about loss of confidentiality, integrity, or authenticity of their communication data, they will have to employ end-to-end network security mechanisms, e.g., IPsec or Transport Layer Security (TLS). Governmental users may be forced by regulations to employ specific approved implementations of those mechanisms. Hence, for such cases, the requirements for confidentiality and integrity of the user data will be met by the end-to-end security mechanism and the ULE security measures would focus on providing traffic flow confidentiality either for user data that has already been encrypted or for users who choose not to implement end-to-end security mechanisms.

ULE links may also be used for communications where the two IP endpoints are not under central control (e.g., when browsing a public web site). In these cases, it may be impossible to enforce any end-to-end security mechanisms. Yet, a common objective is that users may make the same security assumptions as for wired links [RFC3819]. ULE security could achieve this by protecting the vulnerable (in terms of passive attacks) ULE Stream.

In contrast to the above, a ULE Stream can be used to link networks such as branch offices to a central office. ULE link-layer security could be the sole provider of confidentiality and integrity. In this scenario, users requiring high assurance of security (e.g., government use) will need to employ approved cryptographic equipment (e.g., at the network layer). An implementation of ULE Link Security equipment could also be certified for use by specific user communities.

3.2. Threats

The simplest type of network threat is a passive threat. This includes eavesdropping or monitoring of transmissions, with a goal to obtain information that is being transmitted. In broadcast networks (especially those utilising widely available low-cost physical layer interfaces, such as DVB), the passive threats are the major threats. One example is an intruder monitoring the MPEG-2 transmission broadcast and then extracting the data carried within the link. Another example is an intruder trying to determine the identity of the communicating parties and the volume of their traffic by sniffing (L2) addresses. This is a well-known issue in the security field; however, it is more of a problem in the case of broadcast networks such as MPEG-2 transmission networks because of the easy availability of Receiver hardware and the wide geographical span of the networks.

Active threats (or attacks) are, in general, more difficult to implement successfully than passive threats, and usually require more sophisticated resources and may require access to the transmitter. Within the context of MPEG-2 transmission networks, examples of active attacks are:

- Masquerading: An entity pretends to be a different entity. This includes masquerading other users and subnetwork control plane messages.

- Modification of messages in an unauthorised manner.

- Replay attacks: When an intruder sends some old (authentic) messages to the Receiver. In the case of a broadcast link, access to previous broadcast data is easy.

- Denial-of-Service (DoS) attacks: When an entity fails to perform its proper function or acts in a way that prevents other entities from performing their proper functions.

The active threats mentioned above are major security concerns for the Internet community [BELLOVIN]. Masquerading and modification of IP packets are comparatively easy in an Internet environment, whereas such attacks are in fact much harder for MPEG-2 broadcast links. This could, for instance, motivate the mandatory use of sequence numbers in IPsec, but not for synchronous links. This is further reflected in the security requirements for Case 2 and 3 in Section 4 below.

As explained in Section 3.1, the PID associated with an Elementary Stream can be modified (e.g., in some systems by reception of an updated SI table, or in other systems until the next announcement/discovery data is received). An attacker that is able to modify the content of the received multiplex (e.g., replay data and/or control information) could inject data locally into the received stream with an arbitrary PID value.

3.3. Threat Cases

Analysing the topological scenarios for MPEG-2 Transmission Networks in Section 1, the security threats can be abstracted into three cases:

- Case 1: Monitoring (passive threat). Here the intruder monitors the ULE broadcasts to gain information about the ULE data and/or tracking the communicating parties identities (by monitoring the destination NPA address). In this scenario, measures must be taken to protect the ULE payload data and the identity of ULE Receivers.

- Case 2: Locally conducting active attacks on the MPEG-TS multiplex. Here an intruder is assumed to be sufficiently sophisticated to override the original transmission from the ULE Encapsulation Gateway and deliver a modified version of the MPEG-TS transmission to a single ULE Receiver or a small group of Receivers (e.g., in a single company site). The MPEG-2 transmission network operator might not be aware of such attacks. Measures must be taken to ensure ULE data integrity and authenticity and preventing replay of old messages.

- Case 3: Globally conducting active attacks on the MPEG-TS multiplex. This assumes a sophisticated intruder able to override the whole MPEG-2 transmission multiplex. The requirements are similar to case 2. The MPEG-2 transmission network operator can usually identify such attacks and provide corrective action to restore the original transmission.

For both Cases 2 and 3, there can be two sub-cases:

- Insider attacks, i.e., active attacks from adversaries within the network with knowledge of the secret material.

- Outsider attacks, i.e., active attacks from adversaries without knowledge of the secret material.

In terms of priority, Case 1 is considered the major threat in MPEG-2 transmission systems. Case 2 is considered a lesser threat, appropriate to specific network configurations, especially when vulnerable to insider attacks. Case 3 is less likely to be found in an operational network, and is expected to be noticed by the MPEG-2 transmission operator. It will require restoration of the original transmission. The assumption being that physical access to the network components (multiplexers, etc.) and/or connecting physical media is secure. Therefore, Case 3 is not considered further in this document.

4. Security Requirements for IP over MPEG-2 TS

From the threat analysis in Section 3, the following security requirements can be derived:

Req 1. Data confidentiality MUST be provided by a link that supports ULE Stream Security to prevent passive attacks and reduce the risk of active threats.

Req 2. Protection of L2 NPA address is OPTIONAL. In broadcast networks, this protection can be used to prevent an intruder tracking the identity of ULE Receivers and the volume of their traffic.

Req 3. Integrity protection and source authentication of ULE Stream data are OPTIONAL. These can be used to prevent the active attacks described in Section 3.2.

Req 4. Protection against replay attacks is OPTIONAL. This is used to counter the active attacks described in Section 3.2.

Req 5. L2 ULE Source and Receiver authentication is OPTIONAL. This can be performed during the initial key exchange and authentication phase, before the ULE Receiver can join a secure session with the ULE Encapsulator (ULE source). This could be either unidirectional or bidirectional authentication based on the underlying key management protocol.

Other general requirements for all threat cases for link-layer security are:

GReq (a) ULE key management functions MUST be decoupled from ULE security services such as encryption and source authentication. This allows the independent development of both systems.

GReq (b) Support SHOULD be provided for automated as well as manual insertion of keys and policy into the relevant databases.

GReq (c) Algorithm agility MUST be supported. It should be possible to update the crypto algorithms and hashes when they become obsolete without affecting the overall security of the system.

GReq (d) The security extension header MUST be compatible with other ULE extension headers. The method must allow other extension headers (either mandatory or optional) to be used in combination with a security extension. It is RECOMMENDED that these are placed after the security extension header. This permits full protection for all headers. It also avoids situations where the SNDU has to be discarded on processing the security extension header, while preceding headers have already been evaluated. One exception is the Timestamp extension that SHOULD precede the security extension header [RFC5163]. In this case, the timestamp will be unaffected by security services such as data confidentiality and can be decoded without the need for key material.

Examining the threat cases in Section 3.3, the security requirements for each case can be summarised as:

- Case 1: Data confidentiality (Req 1) MUST be provided to prevent monitoring of the ULE data (such as user information and IP addresses). Protection of NPA addresses (Req 2) MAY be provided to prevent tracking ULE Receivers and their communications.

- Case 2: In addition to Case 1 requirements, new measures MAY be implemented such as authentication schemes using Message Authentication Codes, digital signatures, or Timed Efficient Stream Loss-Tolerant Authentication (TESLA) [RFC4082] in order to provide integrity protection and source authentication (Reqs 3 and 5). In addition, sequence numbers (Req 4) MAY be used to protect against replay attacks. In terms of outsider attacks, group authentication using Message Authentication Codes can provide the required level of security (Reqs 3 and 5). This will significantly reduce the ability of intruders to successfully inject their own data into the MPEG-TS stream. However, scenario 2 threats apply only in specific service cases, and therefore authentication and protection against replay attacks are OPTIONAL. Such measures incur additional transmission as well as processing overheads. Moreover, intrusion detection systems may also be needed by the MPEG-2 network operator. These should best be coupled with perimeter security policy to monitor common DoS attacks.

- Case 3: As stated in Section 3.3, the requirements here are similar to Case 2, but since the MPEG-2 transmission network operator can usually identify such attacks, the constraints on intrusion detections are less than in Case 2.

Table 1 below shows the threats that are applicable to ULE networks, and the relevant security mechanisms to mitigate those threats.

	Security Mechanism					
Threat	Data Privacy	Data freshness	Source Authentication	Data Integrity	Intrusion Detection	Identity Protection
Monitoring	X	-	-	-	-	X
Masquerading	X	-	X	X	-	X
Replay Attacks	-	X	X	X	X	-
DoS Attacks	-	X	X	X	X	-
Modification of Messages	-	-	X	X	X	-

Table 1: Security techniques to mitigate network threats in ULE Networks

5. Design Recommendations for ULE Security Extension Header

Table 1 may assist in selecting fields within a ULE Security Extension Header framework.

Security services may be grouped into profiles based on security requirements, e.g., a base profile (with payload encryption and identity protection) and a second profile that extends this to also provide source authentication and protection against replay attacks. Although the use of specific security techniques is optional, it is RECOMMENDED that receiver devices should implement all the techniques in Reqs 2-5 of Section 4 to ensure interoperability of all profiles.

A modular design of ULE security may allow it to use and benefit from existing key management protocols, such as the Group Secure Association Key Management Protocol (GSAKMP) [RFC4535] and the Group Domain of Interpretation (GDOI) [RFC3547] defined by the IETF Multicast Security (MSEC) working group. This does not preclude the use of other key management methods in scenarios where this is more appropriate.

IPsec [RFC4301] and TLS [RFC5246] also provide a proven security architecture defining key exchange mechanisms and the ability to use a range of cryptographic algorithms. ULE security can make use of these established mechanisms and algorithms. See Appendix A for more details.

6. Compatibility with Generic Stream Encapsulation

RFC 5163 [RFC5163] describes three new Extension Headers that may be used with Unidirectional Link Encapsulation, ULE, [RFC4326] and the Generic Stream Encapsulation (GSE) that has been designed for the Generic Mode (also known as the Generic Stream (GS)), offered by second-generation DVB physical layers [GSE].

The security threats and requirements presented in this document are applicable to ULE and GSE encapsulations.

7. Summary

This document analyses a set of threats and security requirements. It defines the requirements for ULE security and states the motivation for link security as a part of the Encapsulation layer.

ULE security must provide link-layer encryption and ULE Receiver identity protection. The framework must support the optional ability to provide for link-layer authentication and integrity assurance, as well as protection against insertion of old (duplicated) data into the ULE Stream (i.e., replay protection). This set of features is optional to reduce encapsulation overhead when not required.

ULE Stream security between a ULE Encapsulation Gateway and the corresponding Receiver(s) is considered an additional security mechanism to IPsec, TLS, and application layer end-to-end security, and not as a replacement. It allows a network operator to provide similar functions to that of IPsec, but in addition provides MPEG-2 transmission link confidentiality and protection of ULE Receiver identity (NPA address).

Appendix A describes a set of building blocks that may be used to realise a framework that provides ULE security functions.

8. Security Considerations

Link-layer (L2) encryption of IP traffic is commonly used in broadcast/radio links to supplement end-to-end security (e.g., provided by TLS [RFC5246], SSH [RFC4251], IPsec [RFC4301]).

A common objective is to provide the same level of privacy as wired links. It is recommended that an ISP or user provide end-to-end security services based on well-known mechanisms such as IPsec or TLS.

This document provides a threat analysis and derives the security requirements to provide link encryption and optional link-layer integrity/authentication of the SNDU payload.

There are some security issues that were raised in RFC 4326 [RFC4326] that are not addressed in this document (i.e., are out of scope), e.g.:

- The security issue with un-initialised stuffing bytes. In ULE, these bytes are set to 0xFF (normal practice in MPEG-2).

- Integrity issues related to the removal of the LAN FCS in a bridged networking environment. The removal of bridged frames exposes the traffic to potentially undetected corruption while being processed by the Encapsulator and/or Receiver.

- There is a potential security issue when a Receiver receives a PDU with two Length fields. The Receiver would need to validate the actual length and the Length field and ensure that inconsistent values are not propagated by the network.

9. Acknowledgments

The authors acknowledge the help and advice from Gorry Fairhurst (University of Aberdeen). The authors also acknowledge contributions from Laurence Duquerroy and Stephane Coombes (ESA), and Yim Fun Hu (University of Bradford).

10. References

10.1. Normative References

[ISO-MPEG2] "Information technology -- generic coding of moving pictures and associated audio information systems, Part I", ISO 13818-1, International Standards Organisation (ISO), 2000.

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

[RFC4326] Fairhurst, G. and B. Collini-Nocker, "Unidirectional Lightweight Encapsulation (ULE) for Transmission of IP Datagrams over an MPEG-2 Transport Stream (TS)", RFC 4326, December 2005.

10.2. Informative References

[BELLOVIN] S. Bellovin, "Security Problems in the TCP/IP Protocol Suite", Computer Communications Review 2:19, pp. 32-48, April 1989. http://www.cs.columbia.edu/~smb/

[ETSI-DAT] EN 301 192, "Digital Video Broadcasting (DVB); DVB Specifications for Data Broadcasting", European Telecommunications Standards Institute (ETSI).

[GSE] TS 102 606, "Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE) Protocol, "European Telecommunication Standards, Institute (ETSI), 2007.

[IEEE-802] "Local and metropolitan area networks-Specific requirements Part 2: Logical Link Control", IEEE 802.2, IEEE Computer Society, (also ISO/IEC 8802-2), 1998.

[ISO-8802] ISO/IEC 8802.2, "Logical Link Control", International Standards Organisation (ISO), 1998.

[ITU-H222] H.222.0, "Information technology, Generic coding of moving pictures and associated audio information Systems", International Telecommunication Union, (ITU-T), 1995.

[RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, June 2001.

[RFC3547] Baugher, M., Weis, B., Hardjono, T., and H. Harney, "The Group Domain of Interpretation", RFC 3547, July 2003.

[RFC3715] Aboba, B. and W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements", RFC 3715, March 2004.

[RFC3819] Karn, P., Ed., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, July 2004.

[RFC4082] Perrig, A., Song, D., Canetti, R., Tygar, J., and B. Briscoe, "Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction", RFC 4082, June 2005.

[RFC4251] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Protocol Architecture", RFC 4251, January 2006.

[RFC4259] Montpetit, M.-J., Fairhurst, G., Clausen, H., Collini-Nocker, B., and H. Linder, "A Framework for Transmission of IP Datagrams over MPEG-2 Networks", RFC 4259, November 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4535] Harney, H., Meth, U., Colegrove, A., and G. Gross, "GSAKMP: Group Secure Association Key Management Protocol", RFC 4535, June 2006.

[RFC4947] Fairhurst, G. and M. Montpetit, "Address Resolution Mechanisms for IP Datagrams over MPEG-2 Networks", RFC 4947, July 2007.

[RFC5163] Fairhurst, G. and B. Collini-Nocker, "Extension Formats for Unidirectional Lightweight Encapsulation (ULE) and the Generic Stream Encapsulation (GSE)", RFC 5163, April 2008.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5374] Weis, B., Gross, G., and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", RFC 5374, November 2008.

Appendix A. ULE Security Framework

This section describes a security framework for the deployment of secure ULE networks.

A.1. Building Blocks

This ULE Security framework describes the following building blocks as shown in Figure 2 below:

- The Key Management Block
- The ULE Security Extension Header Block
- The ULE Databases Block

Within the Key Management Block, the communication between the Group Member entity and the Group Server entity happens in the control plane. The ULE Security Header Block applies security to the ULE SNDU and this happens in the ULE data plane. The ULE Security Databases Block acts as the interface between the Key Management Block (control plane) and the ULE Security Header Block (ULE data plane) as shown in Figure 2. The Security Databases Block exists in both the group member and server sides. However, it has been omitted from Figure 2 just for clarity.

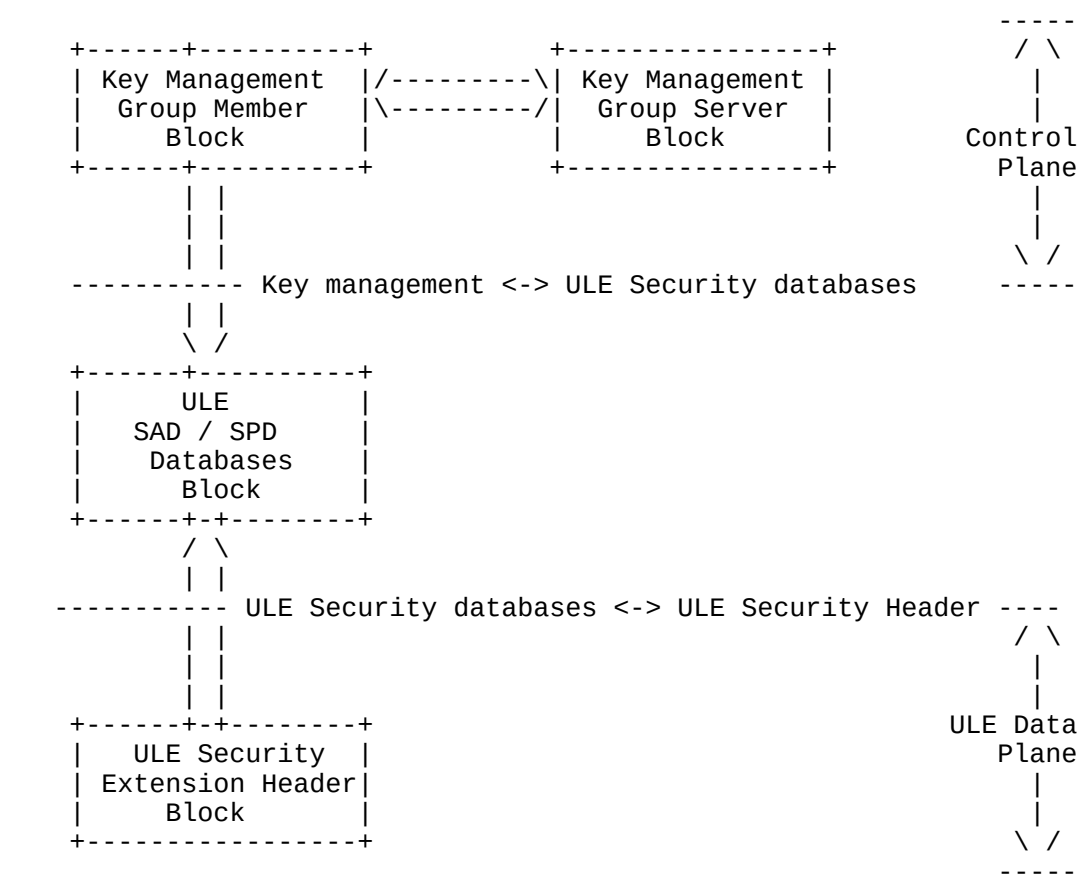

Figure 2: Secure ULE Framework Building Blocks

A.1.1. Key Management Block

A key management framework is required to provide security at the ULE level using extension headers. This key management framework is responsible for user authentication, access control, and Security Association negotiation (which include the negotiations of the security algorithms to be used and the generation of the different session keys as well as policy material). The key management framework can be either automated or manual. Hence, this key management client entity (shown as the Key Management Group Member Block in Figure 2) will be present in all ULE Receivers as well as at the ULE Encapsulators. The ULE Encapsulator could also be the Key Management Group Server Entity (shown as the Key Management Group Server Block in Figure 2).

This happens when the ULE Encapsulator also acts as the Key Management Group Server. Deployment may use either automated key management protocols (e.g., GSAKMP [RFC4535]) or manual insertion of keying material.

A.1.2. ULE Security Databases Block

There needs to be two databases, i.e., similar to the IPsec databases.

- o ULE-SAD: ULE Security Association Database contains all the Security Associations that are currently established with different ULE peers.

- o ULE-SPD: ULE Security Policy Database contains the policies as described by the system manager. These policies describe the security services that must be enforced.

While traditionally link-layer security has operated using simple policy mechanisms, it is envisaged that ULE security should provide flexibility comparable to IPsec. The above design is based on the two databases defined for IPsec [RFC4301]. These databases could be used to implement either simple policies (as in traditional link security services) or more complex policies (as in IPsec).

The exact details of the header patterns that the SPD and SAD will have to support for all use cases will be described in a separate document. This document only highlights the need for such interfaces between the ULE data plane and the Key Management control plane.

A.1.3. ULE Extension Header Block

A new security extension header for the ULE protocol is required to provide the security features of data confidentiality, identity protection, data integrity, data authentication, and mechanisms to prevent replay attacks. Security keying material will be used for the different security algorithms (for encryption/decryption, MAC generation, etc.), which are used to meet the security requirements, described in detail in Section 4 of this document.

This block will use the keying material and policy information from the ULE Security Database Block on the ULE payload to generate the secure ULE Extension Header or to decipher the secure ULE extension header to get the ULE payload. An example overview of the ULE Security extension header format along with the ULE header and payload is shown in Figure 3 below.

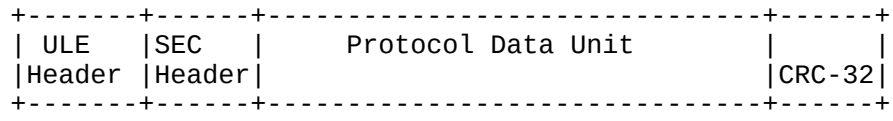

Figure 3: ULE Security Extension Header Placement

A.2. Interface Definition

Two new interfaces have to be defined between the blocks as shown in Figure 2 above. These interfaces are:

- Key Management Block <-> ULE Security Databases Block

- ULE Security Databases Block <-> ULE Security Header Block

While the first interface is used by the Key Management Block to insert keys, security associations, and policies into the ULE Database Block, the second interface is used by the ULE Security Extension Header Block to get the keys and policy material for generation of the security extension header.

A.2.1. Key Management <-> ULE Security Databases

This interface is between the Key Management Block of a group member (GM client) and the ULE Security Database Block (shown in Figure 2). The Key Management GM entity will communicate with the GCKS and then get the relevant security information (keys, cipher mode, security service, ULE_Security_ID, and other relevant keying material as well as policy) and insert this data into the ULE Security Database Block. The Key Management could be either automated (e.g., GSAKMP [RFC4535] or GDOI [RFC3547]), or security information could be manually inserted using this interface.

Examples of interface functions are:

- Insert_record_database (char * Database, char * record, char * Unique_ID);

- Update_record_database (char * Database, char * record, char * Unique_ID);

- Delete_record_database (char * Database, char * Unique_ID);

The definitions of the variables are as follows:

- Database - This is a pointer to the ULE Security databases

- o record - This is the rows of security attributes to be entered or modified in the above databases

- o Unique_ID - This is the primary key to look up records (rows of security attributes) in the above databases

A.2.2. ULE Security Databases <-> ULE Security Header

This interface is between the ULE Security Database and the ULE Security Extension Header Block as shown in Figure 2. When sending traffic, the ULE encapsulator uses the Destination Address, the PID, and possibly other information such as L3 source and destination addresses to locate the relevant security record within the ULE Security Database. It then uses the data in the record to create the ULE security extension header. For received traffic, the ULE decapsulator on receiving the ULE SNDU will use the Destination Address, the PID, and a ULE Security ID inserted by the ULE encapsulator into the security extension to retrieve the relevant record from the Security Database. It then uses this information to decrypt the ULE extension header. For both cases (either send or receive traffic) only one interface is needed since the main difference between the sender and receiver is the direction of the flow of traffic. An example of such an interface is as follows:

- o Get_record_database (char * Database, char * record, char * Unique_ID);

Appendix B. Motivation for ULE Link-Layer Security

Examination of the threat analysis and security requirements in Sections 3 and 4 has shown that there is a need to provide security in MPEG-2 transmission networks employing ULE. This section compares the placement of security functionalities in different layers.

B.1. Security at the IP Layer (Using IPsec)

The security architecture for the Internet Protocol [RFC4301] describes security services for traffic at the IP layer. This architecture primarily defines services for the Internet Protocol (IP) unicast packets, as well as manually configured IP multicast packets.

It is possible to use IPsec to secure ULE Streams. The major advantage of IPsec is its wide implementation in IP routers and hosts. IPsec in transport mode can be used for end-to-end security transparently over MPEG-2 transmission links with little impact.

In the context of MPEG-2 transmission links, if IPsec is used to secure a ULE Stream, then the ULE Encapsulator and Receivers are equivalent to the security gateways in IPsec terminology. A security gateway implementation of IPsec uses tunnel mode. Such usage has the following disadvantages:

- o There is an extra transmission overhead associated with using IPsec in tunnel mode, i.e., the extra IP header (IPv4 or IPv6).

- o There is a need to protect the identity (NPA address) of ULE Receivers over the ULE broadcast medium; IPsec is not suitable for providing this service. In addition, the interfaces of these devices do not necessarily have IP addresses (they can be L2 devices).

- o Multicast is considered a major service over ULE links. The current IPsec specifications [RFC4301] only define a pairwise tunnel between two IPsec devices with manual keying. Work is in progress in defining the extra detail needed for multicast and to use the tunnel mode with address preservation to allow efficient multicasting. For further details refer to [RFC5374].

B.2. Link Security below the Encapsulation Layer

Link layer security can be provided at the MPEG-2 TS layer (below ULE). MPEG-2 TS encryption encrypts all TS Packets sent with a specific PID value. However, an MPEG-2 TS may typically multiplex several IP flows, belonging to different users, using a common PID. Therefore, all multiplexed traffic will share the same security keys.

This has the following advantages:

- o The bit stream sent on the broadcast network does not expose any L2 or L3 headers, specifically all addresses, type fields, and length fields are encrypted prior to transmission.

- o This method does not preclude the use of IPsec, TLS, or any other form of higher-layer security.

However it has the following disadvantages:

- o When a PID is shared between several users, each ULE Receiver needs to decrypt all MPEG-2 TS Packets with a matching PID, possibly including those that are not required to be forwarded. Therefore, it does not have the flexibility to separately secure individual IP flows.

- When a PID is shared between several users, the ULE Receivers will have access to private traffic destined to other ULE Receivers, since they share a common PID and key.

- IETF-based key management that is very flexible and secure is not used in existing MPEG-2 based systems. Existing access control mechanisms in such systems have limited flexibility in terms of controlling the use of keying and rekeying. Therefore, if the key is compromised, this will impact several ULE Receivers.

Currently, there are few deployed L2 security systems for MPEG-2 transmission networks. Conditional access for digital TV broadcasting is one example. However, this approach is optimised for TV services and is not well-suited to IP packet transmission. Some other systems are specified in standards such as MPE [ETSI-DAT], but there are currently no known implementations and these methods are not applicable to GSE.

B.3. Link Security as a Part of the Encapsulation Layer

Examining the threat analysis in Section 3 has shown that protection of ULE Stream from eavesdropping and ULE Receiver identity are major requirements.

There are several advantages in using ULE link-layer security:

- The protection of the complete ULE Protocol Data Unit (PDU) including IP addresses. The protection can be applied either per IP flow or per Receiver NPA address.

- Ability to protect the identity of the Receiver within the MPEG-2 transmission network at the IP layer and also at L2.

- Efficient protection of IP multicast over ULE links.

- Transparency to the use of Network Address Translation (NATs) [RFC3715] and TCP Performance Enhancing Proxies (PEP) [RFC3135], which require the ability to inspect and modify the packets sent over the ULE link.

This method does not preclude the use of IPsec at L3 (or TLS [RFC5246] at L4). IPsec and TLS provide strong authentication of the endpoints in the communication.

L3 end-to-end security would partially deny the advantage listed above (use of PEP, compression, etc.), since those techniques could only be applied to TCP packets bearing a TCP-encapsulated IPsec packet exchange, but not the TCP packets of the original applications, which in particular inhibits compression.

End-to-end security (IPsec, TLS, etc.) may be used independently to provide strong authentication of the endpoints in the communication. This authentication is desirable in many scenarios to ensure that the correct information is being exchanged between the trusted parties, whereas Layer 2 methods cannot provide this guarantee.

Authors' Addresses

Haitham Cruickshank
Centre for Communications System Research (CCSR)
University of Surrey
Guildford, Surrey, GU2 7XH
UK
EMail: h.cruickshank@surrey.ac.uk

Prashant Pillai
Mobile and Satellite Communications Research Centre (MSCRC)
School of Engineering, Design and Technology
University of Bradford
Richmond Road, Bradford BD7 1DP
UK
EMail: p.pillai@bradford.ac.uk

Michael Noisternig
Multimedia Comm. Group, Dpt. of Computer Sciences
University of Salzburg
Jakob-Haringer-Str. 2
5020 Salzburg
Austria
EMail: mnoist@cosy.sbg.ac.at

Sunil Iyengar
Space & Defence
Logica
Springfield Drive
Leatherhead
Surrey KT22 7LP
UK
EMail: sunil.iyengar@logica.com