

I n t e r n a t i o n a l T e l e c o m m u n i c a t i o n U n i o n

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.170

(11/2005)

SERIES J: CABLE NETWORKS AND TRANSMISSION
OF TELEVISION, SOUND PROGRAMME AND OTHER
MULTIMEDIA SIGNALS

IPCablecom

IPCablecom security specification

Recommendation ITU-T J.170



Recommendation ITU-T J.170

IPCablecom security specification

Summary

This Recommendation defines authentication, access control, signalling and media content integrity, confidentiality, and non-repudiation security services for each of the network element interfaces.

IPCablecom security spans the entire IPCablecom architecture. The IPCablecom architectural framework Recommendation (ITU-T Rec. J.160) defines the overall IPCablecom architecture, as well as the system elements, interfaces, and functional requirements for the entire IPCablecom network.

Source

Recommendation ITU-T J.170 was approved on 29 November 2005 by ITU-T Study Group 9 (2005-2008) under Recommendation ITU-T A.8 procedures.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2010

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

| | Page |
|-----|--|
| 1 | Scope and introduction 1 |
| 1.1 | Scope 1 |
| 1.2 | Introduction 1 |
| 2 | References..... 2 |
| 2.1 | Normative references..... 2 |
| 2.2 | Informative references 3 |
| 3 | Terms and definitions 4 |
| 4 | Abbreviations, acronyms and conventions 5 |
| 4.1 | Abbreviations and acronyms 5 |
| 4.2 | Conventions 8 |
| 5 | Architectural overview of IPCablecom security..... 9 |
| 5.1 | IPCablecom reference architecture..... 9 |
| 5.2 | Threats 13 |
| 5.3 | Security architecture 19 |
| 6 | Security mechanisms 26 |
| 6.1 | IPsec 26 |
| 6.2 | Internet Key Exchange (IKE)..... 28 |
| 6.3 | SNMPv3 30 |
| 6.4 | Kerberos/PKINIT 31 |
| 6.5 | Kerberized key management 55 |
| 6.6 | End-to-end security for RTP 76 |
| 6.7 | End-to-end security for RTCP..... 76 |
| 6.8 | BPI+..... 77 |
| 6.9 | TLS 78 |
| 7 | Security profile 80 |
| 7.1 | Device and service provisioning 81 |
| 7.2 | Quality of Service (QoS) signalling 93 |
| 7.3 | Billing system interfaces 95 |
| 7.4 | Call signalling..... 97 |
| 7.5 | PSTN Gateway interface 104 |
| 7.6 | Media stream 106 |
| 7.7 | Audio Server services..... 126 |
| 7.8 | Electronic surveillance interfaces..... 129 |
| 7.9 | CMS provisioning 133 |
| 8 | IPCablecom certificates 134 |
| 8.1 | Generic structure 134 |
| 8.2 | Certificate trust hierarchy 135 |

| | Page |
|------|---|
| 9 | Cryptographic algorithms 145 |
| 9.1 | AES..... 145 |
| 9.2 | DES..... 145 |
| 9.3 | Block termination 146 |
| 9.4 | RSA signature..... 151 |
| 9.5 | HMAC-SHA-1 151 |
| 9.6 | Key derivation 151 |
| 9.7 | The MMH-MAC 152 |
| 9.8 | Random number generation 154 |
| 10 | Physical security 154 |
| 10.1 | Protection for MTA key storage..... 154 |
| 10.2 | MTA key encapsulation 157 |
| 11 | Secure software upgrade..... 157 |
| | Annex A – Oakley groups..... 158 |
| | Annex B – Kerberos Network Authentication Service 159 |
| | Annex C – PKINIT specification..... 159 |
| | Appendix I – IPCablecom administration guidelines and best practices..... 160 |
| | I.1 Routine CMS service key refresh..... 160 |
| | Appendix II – Example of MMH algorithm implementation 161 |

Recommendation ITU-T J.170

IPcablecom security specification

1 Scope and introduction

1.1 Scope

Authentication, access control, signalling and media content integrity, confidentiality, and non-repudiation security services must be provided as defined herein for each of the network element interfaces.

IPcablecom security spans the entire IPcablecom architecture. The IPcablecom architectural framework Recommendation (ITU-T Rec. J.160) defines the overall IPcablecom architecture, as well as the system elements, interfaces, and functional requirements for the entire IPcablecom network.

1.2 Introduction

1.2.1 Goals

This Recommendation describes the security relationships between the elements on the IPcablecom network. The general goals of this IPcablecom network security Recommendation and any implementations that encompass the requirements defined herein should be:

- **Secure network communications:** The IPcablecom network security must define a security architecture, methods, algorithms and protocols that meet the stated security service requirement. All media packets and all sensitive signalling communication across the network must be safe from eavesdropping. Unauthorized message modification, insertion, deletion and replays anywhere in the network must be easily detectable and must not affect proper network operation.
- **Reasonable cost:** The IPcablecom network security must define security methods, algorithms and protocols that meet the stated security service requirements such that a reasonable implementation can be manifested with reasonable cost and implementation complexity.
- **Network element interoperability:** All of the security services for any of the IPcablecom network elements must interoperate with the security services for all of the other IPcablecom network elements. Multiple vendors may implement each of the IPcablecom network elements as well as multiple vendors for a single IPcablecom network element.
- **Extensibility:** The IPcablecom security architecture, methods, algorithms and protocols must provide a framework into which new security methods and algorithms may be incorporated as necessary.

1.2.2 Assumptions

The following assumptions are made relative to the current scope of the IPcablecom security Recommendation:

- Embedded Media Terminal Adaptors (MTAs) are within the current scope. Stand-alone MTAs will be addressed in later phases and security issues for stand-alone MTAs are thus for future study.
- Network Call Signalling (NCS) is the only call signalling method, on the access network, addressed in this Recommendation.
- This version of the IPcablecom security Recommendation specifies security for a single administrative domain and the communications between domains.

- Security for chained RADIUS servers is not currently in the scope.

This Recommendation also does not include requirements for associated security operational issues (e.g., site security), back-office or inter/intra back-office security, service authorization policies or secure database handling. Record Keeping Servers (RKS), Network Management Systems, File Transfer Protocol (FTP) servers and Dynamic Host Configuration Protocol (DHCP) servers are all considered to be unique to any service provider's implementation and are beyond the scope of this Recommendation.

2 References

2.1 Normative references

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- Recommendation ITU-T J.112 (1998), *Transmission systems for interactive cable television services*.
- Recommendation ITU-T J.122 (2002), *Second-generation transmission systems for interactive cable television services – IP cable modems*.
- Recommendation ITU-T J.125 (2004), *Link privacy for cable modem implementations*.
- Recommendation ITU-T J.126 (2004), *Embedded Cable Modem device specification*.
- Recommendation ITU-T J.162 (2005), *Network call signalling protocol for the delivery of time-critical services over cable television networks using cable modems*.
- Recommendation ITU-T J.166 (2005), *IPCablecom Management Information Base (MIB) framework*.
- Recommendation ITU-T J.167 (2005), *Media terminal adapter (MTA) device provisioning requirements for the delivery of real-time services over cable television networks using cable modems*.
- Recommendation ITU-T X.509 (2000) | ISO/IEC 9594-8:2001, *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks*.
- Recommendation ITU-T X.690 (2002) | ISO/IEC 8825-1:2002, *Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)*.
- IETF RFC 1889 (1996), *RTP: A Transport Protocol for Real-Time Applications*.
- IETF RFC 1890 (1996), *RTP Profile for Audio and Video Conferences with Minimal Control*.
- IETF RFC 2104 (1997), *HMAC: Keyed-Hashing for Message Authentication*.
- IETF RFC 2246 (1999), *The TLS Protocol Version 1.0*.
- IETF RFC 2367 (1998), *PF_KEY Key Management API, Version 2*.
- IETF RFC 2401 (1998), *Security Architecture for the Internet Protocol*.

- IETF RFC 2403 (1998), *The Use of HMAC-MD5-96 within ESP and AH.*
- IETF RFC 2404 (1998), *The Use of HMAC-SHA-1-96 within ESP and AH.*
- IETF RFC 2406 (1998), *IP Encapsulating Security Payload (ESP).*
- IETF RFC 2407 (1998), *The Internet IP Security Domain of Interpretation for ISAKMP.*
- IETF RFC 2409 (1998), *The Internet Key Exchange (IKE).*
- IETF RFC 2437 (1998), *PKCS#1: RSA Cryptography Specification Version 2.0.*
- IETF RFC 2451 (1998), *The ESP CBC-Mode Cipher Algorithms.*
- IETF RFC 2459 (1999), *Internet X.509 Public Key Infrastructure Certificate and CRL Profile.*
- IETF RFC 2630 (1999), *Cryptographic Message Syntax.*
- IETF RFC 3261 (2002), *SIP: Session Initiation Protocol.*
- IETF RFC 3412 (2002), *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP).*
- IETF RFC 3414 (2002), *User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3).*
- IETF RFC 4120 (2005), *The Kerberos Network Authentication Service (V5).*
- IETF RFC 4556 (2006), *Public Key Cryptography for Initial Authentication in Kerberos (PKINIT).*
- FIPS PUB 81 (1980), *DES Modes of Operation.*
- FIPS PUB 180-1 (1995), *Secure Hash Algorithm (SHS).*
- FIPS PUB 197 (2001), *Advanced Encryption Standard (AES).*

2.2 Informative references

- Recommendation ITU-T J.160 (2005), *Architectural framework for the delivery of time-critical services over cable television networks using cable modems.*
- Recommendation ITU-T J.161 (2001), *Audio codec requirements for the provision of bidirectional audio service over cable television networks using cable modems.*
- Recommendation ITU-T J.163 (2005), *Dynamic quality of service for the provision of real-time services over cable television networks using cable modems.*
- Recommendation ITU-T J.164 (2005), *Event message requirements for the support of real-time services over cable television networks using cable modems.*
- Recommendation ITU-T J.171.x (2005), *IPCablecom Trunking Gateway Control Protocol (TGCP).*
- Recommendation ITU-T J.175 (2002), *Audio server protocol.*
- Recommendation ITU-T J.178 (2005), *IPCablecom CMS to CMS signalling.*
- FIPS PUB 140-1 (1994), *Security Requirements for Cryptographic Modules.*
- IETF RFC 1750 (1994), *Randomness Recommendations for Security.*
- IETF RFC 2327 (1998), *SDP: Session Description Protocol.*
- IETF RFC 2782 (2000), *A DNS RR for specifying the location of services (DNS SRV).*

- IETF RFC 3268 (2002), *Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS)*.
- HALEVI [S.], KRAWCZYK [H.]: MMH: Software Message Authentication in the Gbit/Second Rates, *Proceedings of the 4th Workshop on Fast Software Encryption*, Vol. 1267 Springer-Verloag, pp. 172-189, 1970.
- KILIAN [J.], ROGAWAY [P.]: How to Protect DES Against Exhaustive Key Search, *Edited version presented at Proceedings of Crypto '96*, July 1997.
- SCHNEIER [B.]: *Applied Cryptography*, John Wiley & Sons, Inc., Second edition, 1996.

3 Terms and definitions

This Recommendation defines the following terms:

- 3.1 access control:** Limiting the flow of information from the resources of a system only to authorized persons, programs, processes or other system resources on a network.
- 3.2 audio server:** An Audio Server plays informational announcements in IPCablecom network. Media announcements are needed for communications that do not complete and to provide enhanced information services to the user. The component parts of Audio Server services are Media Players and Media Player Controllers.
- 3.3 authentication:** The process of verifying the claimed identity of an entity to another entity.
- 3.4 authenticity:** The ability to ensure that the given information is without modification or forgery and was in fact produced by the entity that claims to have given the information.
- 3.5 authorization:** The act of giving access to a service or device if one has the permission to have the access.
- 3.6 cipher:** An algorithm that transforms data between plaintext and ciphertext.
- 3.7 ciphersuite:** A set, which must contain both an encryption algorithm and a message authentication algorithm (e.g., a MAC or an HMAC). In general, it may also contain a key management algorithm, which does not apply in the context of IPCablecom.
- 3.8 confidentiality:** A way to ensure that information is not disclosed to any one other than the intended parties. Information is encrypted to provide confidentiality. Also known as privacy.
- 3.9 cryptanalysis:** The process of recovering the plaintext of a message or the encryption key without access to the key.
- 3.10 downstream:** The direction from the head-end toward the subscriber location.
- 3.11 encryption:** A method used to translate information in plaintext into ciphertext.
- 3.12 endpoint:** A Terminal, Gateway or MCU.
- 3.13 event message:** Message capturing a single portion of a connection.
- 3.14 gateway:** Devices bridging between the IPCablecom Voice Communication world and the PSTN. Examples are the Media Gateway which provides the bearer circuit interfaces to the PSTN and transcodes the media stream, and the Signalling Gateway which sends and receives circuit switched network signalling to the edge of the IPCablecom network.
- 3.15 header:** Protocol control information located at the beginning of a protocol data unit.
- 3.16 integrity:** A way to ensure that information is not modified except by those who are authorized to do so.
- 3.17 Kerberos:** A secret-key network authentication protocol that uses a choice of cryptographic algorithms for encryption and a centralized key database for authentication.

- 3.18 key:** A mathematical value input into the selected cryptographic algorithm.
- 3.19 key exchange:** The swapping of public keys between entities to be used to encrypt communication between the entities.
- 3.20 key management:** The process of distributing shared symmetric keys needed to run a security protocol.
- 3.21 non-repudiation:** The ability to prevent a sender from denying later that he or she sent a message or performed an action.
- 3.22 privacy:** A way to ensure that information is not disclosed to any one other than the intended parties. Information is usually encrypted to provide confidentiality. Also known as "confidentiality".
- 3.23 private key:** The key used in public key cryptography that belongs to an individual entity and must be kept secret.
- 3.24 proxy:** A facility that indirectly provides some service or acts as a representative in delivering information thereby eliminating the need for a host to support the service.
- 3.25 public key:** The key used in public key cryptography that belongs to an individual entity and is distributed publicly. Other entities use this key to encrypt data to be sent to the owner of the key.
- 3.26 public key certificate:** A binding between an entity's public key and one or more attributes relating to its identity, also known as a digital certificate.
- 3.27 public key cryptography:** A procedure that uses a pair of keys, a public key and a private key for encryption and decryption, also known as an asymmetric algorithm. A user's public key is publicly available for others to use to send a message to the owner of the key. A user's private key is kept secret and is the only key that can decrypt messages sent encrypted by the user's public key.
- 3.28 root private key:** The private signing key of the highest-level Certification Authority. It is normally used to sign public key certificates for lower-level Certification Authorities or other entities.
- 3.29 X.509 certificate:** A public key certificate specification developed as part of the ITU-T Rec. X.500 standards directory.

4 Abbreviations, acronyms and conventions

4.1 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

| | |
|------|---|
| AES | Advanced Encryption Standard |
| AH | Authentication Header is an IPsec security protocol that provides message integrity for complete IP packets, including the IP header. |
| ASD | Application-Specific Data. An application-specific field in the IPsec header that along with the destination IP address provides a unique number for each SA. |
| BPI+ | Baseline Privacy Interface Plus is the security portion of ITU-T Rec. J.112 that runs on the MAC layer. |
| CBC | Cipher-block Chaining mode is an option in block ciphers that combine (XOR) the previous block of ciphertext with the current block of plaintext before encrypting that block of the message. |

| | |
|---------|--|
| CA | Certification Authority. A trusted organization that accepts certificate applications from entities, authenticates applications, issues certificates and maintains status information about certificates. |
| CA | Call Agent. The part of the CMS that maintains the communication state, and controls the line side of the communication. |
| CM | Cable Modem |
| CMS | Cryptographic Message Syntax |
| CMS | Call Management Server. Controls the audio connections. Also called a Call Agent in MGCP/SGCP terminology. This is one example of an Application Server. |
| CMTS | Cable Modem Termination System |
| CNAME | Canonical Name |
| COPS | Common Open Policy Service |
| CRL | Certificate Revocation List |
| CSR | Customer Service Record |
| DES | Data Encryption Standard |
| DF | Delivery Function |
| DH | Diffie-Hellman |
| DHCP | Dynamic Host Configuration Protocol |
| DNS | Domain Name Server |
| DNS SRV | A DNS RR for specifying the location of services |
| DOCSIS | Data-Over-Cable Service Interface Specification |
| DQoS | Dynamic Quality of Service |
| DSCP | DiffServ Code Point. A field in every IP packet that identifies the DiffServ Per-Hop Behavior. In IP version 4, the TOS byte is redefined to be the DSCP. In IP version 6, the Traffic Class octet is used as the DSCP. See IETF RFC 4556. |
| DTMF | Dual-Tone Multifrequency (tones) |
| ESP | IPsec Encapsulating Security |
| FQDN | Fully Qualified Domain Name. Refer to IETF RFC 821 for details. |
| GW | Gateway |
| HFC | Hybrid-Fibre/Coax |
| HMAC | Hashed Message Authentication Code. A message authentication algorithm, based on either SHA-1 or MD5 hash and defined in IETF RFC 2104. |
| IKE | Internet Key Exchange is a key management mechanism used to negotiate and derive keys for SAs in IPsec. |
| IKE | A notation defined to refer to the use of IKE with pre-shared keys for authentication. |
| INA | Interactive Network Adapter |
| IPsec | Internet Protocol Security |
| ISAKMP | Internet Security Association and Key Management Protocol |
| ISTP | Internet Signalling Transport Protocol |

| | |
|---------|---|
| IVR | Interactive Voice Response |
| KDC | Key Distribution Centre |
| MAC | Message Authentication Code. A fixed-length data item that is sent together with a message to ensure integrity; also known as a MIC. |
| MAC | Media Access Control. It is a sublayer of the Data Link Layer. It normally runs directly over the physical layer. |
| MD5 | Message Digest 5 |
| MG | Media Gateway |
| MGC | Media Gateway Controller |
| MGCP | Media Gateway Control Protocol |
| MIB | Management Information Base |
| MMH | Multilinear Modular Hash |
| MSB | Most Significant Bit |
| MTA | Media Terminal Adapter |
| NCS | Network Call Signalling |
| NVRAM | Non-Volatile Random Access Memory |
| OID | Object Identification |
| OSS | Operations Support Systems. The back-office software used for configuration, performance, fault, accounting, and security management. |
| PKCS | Public Key Cryptography Standards |
| PKI | Public Key Infrastructure. A process for issuing public key certificates, which includes standards, Certification Authorities, communication between authorities and protocols for managing certification processes. |
| PKCROSS | Utilizes PKINIT for establishing the inter-realm keys and associated inter-realm policies to be applied in issuing cross-realm service tickets between realms and domains in support of Intradomain and Interdomain CMS-to-CMS signalling (CMSS). |
| PSTN | Public Switched Telephone Network |
| QoS | Quality of Service |
| RADIUS | Remote Authentication Dial-In User Service |
| RC4 | A variable key length stream cipher offered in the ciphersuite, used to encrypt the media traffic in IPCablecom. |
| RFI | Radio Frequency Interface |
| RKS | Record Keeping Server. The device which collects and correlates the various Event Messages. |
| RSIP | Realm Specific IP |
| RSVP | Resource Reservation Protocol |
| RTCP | Real-Time Control Protocol |
| RTO | Retransmission Timeout |
| RTP | Real-Time Protocol |

| | |
|------|---|
| SA | Security Association |
| SDP | Session Description Protocol |
| SG | Signalling Gateway. A SG is a signalling agent that receives/sends SCN native signalling at the edge of the IP network. In particular, the SS7 SG function translates variants ISUP and TCAP in an SS7-Internet Gateway to a common version of ISUP and TCAP. |
| SIP | Session Initiation Protocol. An application-layer control (signalling) protocol for creating, modifying, and terminating sessions with one or more participants. |
| SIP+ | Session Initiation Protocol Plus. An extension to SIP. |
| SNMP | Simple Network Management Protocol |
| SRV | Server |
| SS7 | Signalling System No. 7. An architecture and set of protocols for performing out-of-band call signalling with a telephone network. |
| SSL | Secure Sockets Layer |
| TCAP | Transaction Capabilities Application Protocol. A protocol within the SS7 stack that is used for performing remote database transactions with a Signalling Control Point. |
| TD | Timeout for Disconnect |
| TFTP | Trivial File Transfer Protocol |
| TGS | Ticket Granting Server. A sub-system of the KDC used to grant Kerberos tickets. |
| TLS | Transport Layer Security |
| UDP | User Datagram Protocol |
| UNI | User-Network Interface |
| VCI | Virtual Channel Identifier |
| VPI | Virtual Path Identifier |

4.2 Conventions

If this Recommendation is implemented, the keywords "MUST" and "SHALL" as well as "REQUIRED" are to be interpreted as indicating a mandatory aspect of this Recommendation.

The keywords indicating a certain level of significance of a particular requirements that are used throughout this Recommendation are summarized below:

| | |
|--------------|---|
| "MUST" | This word or the adjective "REQUIRED" means that the item is an absolute requirement of this Recommendation. |
| "MUST NOT" | This phrase means that the item is an absolute prohibition of this Recommendation. |
| "SHOULD" | This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course. |
| "SHOULD NOT" | This phrase means that there may exist valid reasons in particular circumstances when the listed behaviour is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behaviour described with this label. |

"MAY"

This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

5 Architectural overview of IPCablecom security

5.1 IPCablecom reference architecture

Security requirements have been defined for every signalling and media link within the IPCablecom network. In order to understand the security requirements and specifications for IPCablecom, one must first understand the overall architecture. This clause presents a brief overview of the IPCablecom architecture. For a more detailed specification, refer to the IPCablecom architecture Recommendation (ITU-T Rec. J.160).

5.1.1 HFC network

In Figure 1, the access network between the MTAs and the CMTS is an HFC network, which employs J.112/J.122 physical layer and MAC layer protocols. J.125 enhanced security (see 6.8) and QoS protocols are enabled over this link.

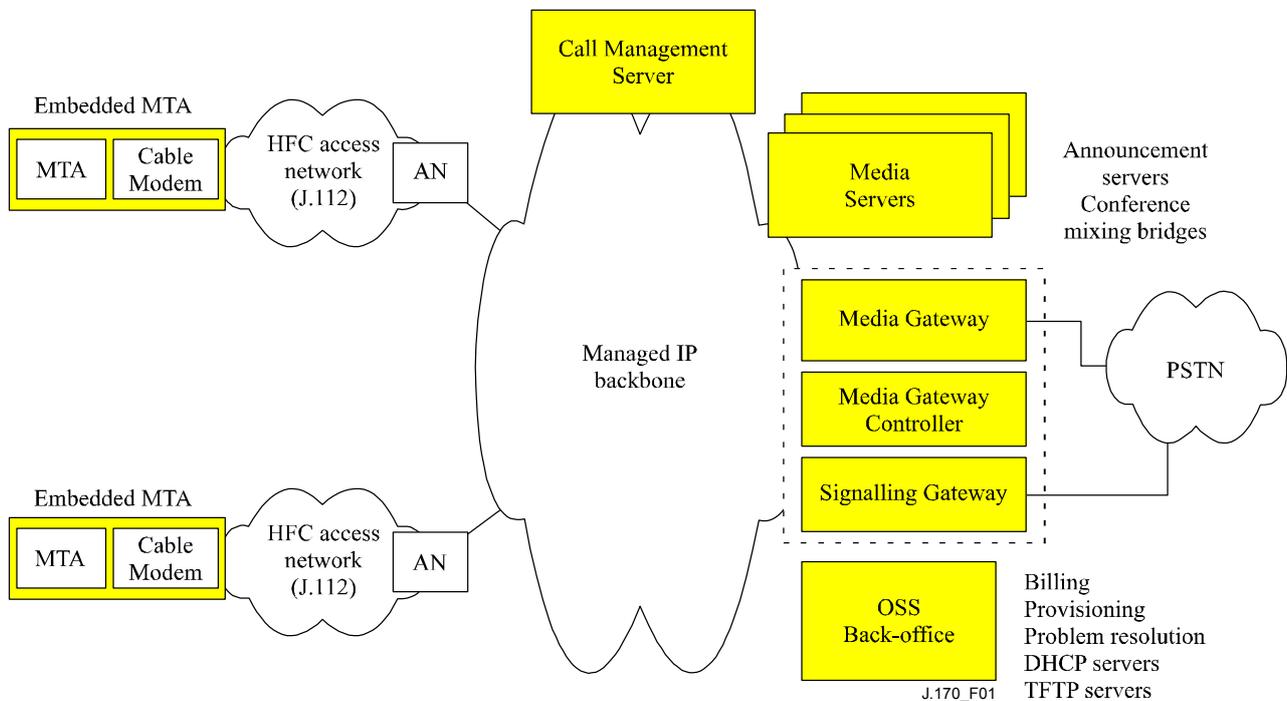


Figure 1 – IPCablecom single zone architecture

5.1.2 Call Management Server

In the context of voice communications applications, a central component of the system is the Call Management Server (CMS). It is involved in both call signalling and the establishment of Dynamic Quality of Service (DQoS). The CMS also performs queries at the PSTN Gateway for LNP (Local Number Portability) and other services necessary for voice communications, including interfacing with the PSTN.

As described in the IPCablecom Architecture Framework, the CMS is divided into the following functional components:

- Call Agent (CA) – The Call Agent maintains network intelligence and call state and controls the media gateway. Most of the time Call Agent is synonymous for Call Management Server.
- Gate Controller (GC) – The Gate Controller is a logical QoS management component that is typically part of the CMS. The GC coordinates all quality of service authorization and control on behalf of the application service – e.g., voice communications.
- Media Player Controller (MPC) – The MPC initiates and manages all announcement services provided by the Media Player. The MPC accepts requests from the CMS and arranges for the MP to provide the announcement in the appropriate stream so that the user hears the announcement.
- Media Gateway Controller (MGC) – The Media Gateway Controller maintains the gateway's portion of call state for communications traversing the Gateway.

A particular CMS can contain any subset of the above listed functional components.

5.1.3 Functional categories

The IPCablecom Architecture Framework identifies the following functional categories within the architecture:

- MTA device provisioning;
- quality of Service (HFC access network and managed IP Backbone);
- billing interface security;
- security (specified herein);
- network call signalling (NCS);
- PSTN interconnectivity;
- CODEC functionality and media stream mapping;
- Audio Server services;
- electronic surveillance (DF interfaces).

In most cases, each functional category corresponds to a particular IPCablecom Recommendation.

5.1.3.1 Device and service provisioning

During MTA provisioning, the MTA gets its configuration with the help of the DHCP and TFTP servers, as well as the OSS.

Provisioning interfaces need to be secured and have to configure the MTA with the appropriate security parameters (e.g., customer X.509 certificate signed by the Service Provider). This Recommendation specifies the steps in MTA provisioning, but provides detailed specifications only for the security parameters. Refer to ITU-T Rec. J.167 for a full specification on MTA provisioning and customer enrolment.

5.1.3.2 Dynamic Quality of Service

IPCablecom provides guaranteed Quality of Service (QoS) for each voice communication within a single zone with Dynamic QoS (ITU-T Rec. J.163).

DQoS is controlled by the Gate Controller function within the CMS and can guarantee Quality of Service within a single administrative domain. The Gate Controller utilizes the Common Open Policy Service (COPS) protocol to download QoS policy into the CMTS. After that, the QoS reservation is established via J.112/J.122 QoS messaging between the MTA and the CMTS on both sides of the connection.

5.1.3.3 Billing system interfaces

The CMS, CMTS and the PSTN Gateway are all required to send out billing event messages to the Record Keeping Server (RKS). This interface is specified to be RADIUS. Billing information should be checked for integrity and authenticity as well as kept private. This Recommendation defines security requirements and specifications for the communication with RKS.

5.1.3.4 Call signalling

The call signalling architecture defined within IPCablecom is network-based call signalling (NCS) per ITU-T Rec. J.162. The CMS is used to control call set-up, termination and most other call signalling functions. In the NCS architecture, the Call Agent function within the CMS is used in call signalling and utilizes the MGCP protocol.

5.1.3.5 PSTN interconnectivity

The PSTN interface to the voice communications capabilities of the IPCablecom network is through the Signalling and Media Gateways (SG and MG). Both of these gateways are controlled with the MGC (Media Gateway Controller). The MGC may be standalone or combined with a CMS. For further detail on PSTN Gateways, refer to ITU-T Rec. J.171.x-series.

All communications between the MGC and the SG and MG may be over the same-shared IP network and is subject to similar threats (e.g., privacy, masquerade, denial-of-service) that are encountered in other links in the same network. This Recommendation defines the security requirements and specifications for the PSTN Gateway links.

When communications from an MTA to a PSTN phone are made, bearer channel traffic is passed directly between an MTA and an MG. The protocols used in this case are RTP and RTCP, as in the MTA-to-MTA case. Both security requirements and specifications are very similar to the MTA-to-MTA bearer requirements and are fully defined in this Recommendation. After a voice communication enters the PSTN, the security requirements as well as specifications are the responsibility of the PSTN.

5.1.3.6 CODEC functionality and media stream mapping

The media stream between two MTAs or between an MTA and a PSTN Gateway utilizes the RTP protocol. Although ITU-T J.125 provides for privacy over the HFC network, the potential threats within the rest of the voice communications network require that the RTP packets be encrypted end-to-end¹.

In addition to RTP, there is an accompanying RTCP protocol, primarily used for reporting of RTCP statistics. In addition, RTCP packets may carry CNAME – a unique identifier of the sender of RTP packets. RTCP also defines a BYE message² that can be used to terminate an RTP session. These two additional RTCP functions raise privacy and denial-of-service threats. Due to these threats, RTCP security requirements are the same as the requirements for all other end-to-end (SIP+) signalling and are addressed in the same manner.

¹ In general, it is possible for an MTA-to-MTA or MTA-to-PSTN connection to cross the networks of several different Service Providers. In the process, this path may cross a PSTN network. This is an exception to the rule, where all RTP packets are encrypted end-to-end. The media traffic inside a PSTN network does not utilize RTP and has its own security requirements. Thus, in this case the encryption would not be end-to-end and would terminate at the PSTN Gateway on both sides of the intermediate PSTN network.

² The RTCP BYE message should not be confused with the SIP+ BYE message that is also used to indicate the end of a voice communication within the network.

In addition to MTAs and PSTN Gateways, Media Servers may also participate in the media stream flows. Media Servers are network-based components that operate on media flows to support various voice communications service options. Media Servers perform audio bridging, play terminating announcements, provide interactive voice response services, and so on. Both media stream and signalling interfaces to a Media Server are the same as the interfaces to an MTA. For more information on Codec functionality, see ITU-T Rec. J.161.

5.1.3.7 Audio Server services

Audio Server interfaces provide a suite of signalling protocols for providing announcement and audio services in an IPCablecom network.

5.1.3.7.1 Media Player Controller (MPC)

The Media Player Controller (MPC) initiates and manages all announcement services provided by the Media Player. The MPC accepts requests from the CMS and arranges for the MP to provide the announcement in the appropriate stream so that the user hears the announcement. The MPC also serves as the termination for certain calls routed to it for IVR services. When the MP collects information from the end-user, the MPC is responsible for interpreting this information and managing the IVR session accordingly. The MPC manages call state.

5.1.3.7.2 Media Player (MP)

The Media Player (MP) is a media resource server. It is responsible for receiving and interpreting commands from the MPC and for delivering the appropriate announcement(s) to the MTA. The MP provides the media stream with the announcement contents. The MP also is responsible for accepting and reporting user inputs (e.g., DTMF tones). The MP functions under the control of the MPC.

5.1.3.8 Electronic surveillance

The event interface between the CMS and the DF provides descriptions of calls, necessary to perform wiretapping. This information includes the media stream encryption key and the corresponding encryption algorithm. This event interface uses RADIUS and is similar to the CMS-RKS interface.

The COPS interface between the CMS and the CMTS is used to signal the CMTS to start/stop duplicating media packets to the DF for a particular call. This is the same COPS interface that is used for (DQoS) Gate Authorization messages.

The following clause summarizes general threats and the corresponding attacks that are relevant in the context of IP voice communications. This list of threats is not based on the knowledge of the specific protocols or security mechanisms employed in the network. A more specific summary of threats that are based on the functionality of each network element is listed in 5.2.6.

Some of the outlined threats cannot be addressed purely by cryptographic means – physical security and/or fraud management should also be used. These threats may be important, but cannot be fully addressed within the scope of IPCablecom. How vendors and cable operators implement fraud management and physical security will differ and in this case a standard is not required for interoperability.

5.2.1 Theft of network services

In the context of voice communications, the main services that may be stolen are:

- long-distance service;
- local (subscription) voice communications service;
- video conferencing;
- network-based three-way calling;
- Quality of Service.

5.2.1.1 MTA clones

One or more MTAs can masquerade as another MTA by duplicating its permanent identity and keys. The secret cryptographic keys may be obtained by either breaking the physical security of the MTA or by employing cryptanalysis.

When an MTA is broken into, the perpetrator can steal voice communications service and charge it all to the original owner. The feasibility of such an attack depends on where an MTA is located. This attack must be seriously considered in the cases when an MTA is located in an office or apartment building, or on a street corner.

An owner might break into his or her own MTA in at least one instance – after a false account with the cable operator providing the voice communications service had been set up. The customer name, address, and Social Security number may all be invalid or belong to someone else. The provided credit card number may be stolen. In that case, the owner of the MTA would not mind giving out the MTA cryptographic identity to others – he or she would not have to pay for service anyway.

In addition to cloning of the permanent cryptographic keys, temporary (usually symmetric) keys may also be cloned. Such an attack is more complex, since the temporary keys expire more often and have to be frequently redistributed. The only reason why someone would attempt this attack is if the permanent cryptographic keys are protected much better than the temporary ones, or if the temporary keys are particularly easy to steal or discover with cryptanalysis.

5.2.1.2 Other clones

It is conceivable that the cryptographic identity of another network element, such as a CMTS or a CMS, may be cloned. Such an attack is most likely to be mounted by an insider such as a corrupt or disgruntled employee.

5.2.1.3 Subscription fraud

A customer sets up an account under false information.

5.2.1.4 Non-payment for voice communications services

A customer stops paying his or her bill, but continues to use the MTA for voice communications service. This can happen if the network does not have an automated method to revoke the customer's access to the network.

5.2.1.5 Protocol attacks against an MTA

A weakness in the protocol can be manipulated to allow an MTA to authenticate to a network server with a false identity or hijack an existing voice communication. This includes replay and man-in-the-middle attacks.

5.2.1.6 Protocol attacks against other network elements

A perpetrator might employ similar protocol attacks to masquerade as a different network element, such as a CMTS or a CMS. Such an attack may be used in collaboration with cooperating MTAs to steal service.

5.2.1.7 Theft of services provided by the MTA

Services such as the support for multiple MTA ports, 3-way calling and call waiting may be implemented entirely in the MTA, without any required interaction with the network.

5.2.1.7.1 Attacks

MTA code to support these services may be downloaded illegally by an MTA clone, in which case the clone has to interact with the network to get the download. In that case, this threat is no different from the network service theft described in the previous clause.

Alternatively, downloading an illegal code image using some illegal out-of-band means can also enable these services. Such service theft is much harder to prevent (a secure software environment within the MTA may be required). On the other hand, in order for an adversary to go through this trouble, the price for these MTA-based services has to make the theft worthwhile.

An implication of this threat is that valuable services cannot be implemented entirely inside the MTA without a secure software environment in addition to tamper-proof protection for the cryptographic keys. (While a secure software environment within an MTA adds significant complexity, it is an achievable task.)

5.2.1.8 MTA moved to another network

A leased MTA may be reconfigured and registered with another network, contrary to the intent and property rights of the leasing company.

5.2.2 Bearer channel information threats

This class of threats is concerned with the breaking of privacy of voice communications over the IP bearer channel. Threats against non-VoIP communications are not considered here and assumed to require additional security at the application layer.

5.2.2.1 Attacks

Clones of MTAs and other network elements, as well as protocol manipulation attacks, also apply in the case of bearer channel information threats. These attacks are already described under the Service Theft threats.

MTA cloning attacks mounted by the actual owner of the MTA are less likely in this case, but not inconceivable. An owner of an MTA may distribute clones to unsuspecting victims, so that he or she can later spy on them.

5.2.2.1.1 Off-line cryptanalysis

Bearer channel information may be recorded and then analysed over a period of time, until the encryption keys are discovered through cryptanalysis. The discovered information may be of value even after a relatively long time has passed.

5.2.3 Signalling channel information threats

Signalling information, such as the caller identity and the services to which each customer subscribes, may be collected for marketing purposes. The caller identity may also be used illegally to locate a customer that wishes to keep his or her location private.

5.2.3.1 Attacks

Clones of MTAs and other network elements, as well as protocol manipulation attacks, also apply in the case of the signalling channel information threats. These attacks are already described in the clause on Service Theft threats.

MTA cloning attacks mounted by the actual owner of the MTA is theoretically possible in this case. An owner of an MTA may distribute clones to the unsuspecting victims, so that he or she can monitor their signalling messages (e.g., for information with marketing value). The potential benefits of such an attack seem unjustified, however.

5.2.3.1.1 Caller ID

A number of a party initiating a voice communication is revealed, even though a number is not generally available (i.e., is "unlisted") and the owner of that number enabled ID blocking.

5.2.3.1.2 Information with marketing value

Dialled numbers and the type of service customers use may be gathered for marketing purposes by other corporations.

5.2.4 Service disruption threats

This class of threats is aimed at disrupting the normal operation of voice communications. The motives for denial-of-service attacks may be malicious intent against a particular individual or against the service provider; or, perhaps a competitor wishes to degrade the performance of another service provider and use the resulting problems in an advertising campaign.

5.2.4.1 Attacks

5.2.4.1.1 Remote interference

A perpetrator is able to manipulate the protocol to close down ongoing voice communications. This might be achieved by masquerading as an MTA involved in such an ongoing communication. The same effect may be achieved if the perpetrator impersonates another network element, such as a gate controller or an edge router during either call set-up or voice packet routing.

Depending on the signalling protocol security, it might be possible for the perpetrator to mount this attack from the MTA, in the privacy of his or her own home.

Clones of MTAs and other network elements, as well as protocol manipulation attacks, also apply in the case of the service disruption threats. These attacks are described under Service Theft threats.

MTA cloning attacks mounted by the actual owner of the MTA can theoretically be used in service disruption against unsuspecting clone owners. However, since there are so many other ways to cause service disruption, such an attack cannot be taken seriously in this context.

5.2.5 Repudiation

In a network where masquerading (using the above-mentioned cloning and protocol manipulation techniques) is common or easily achievable, a customer may repudiate a particular communication (and, thus deny responsibility for paying for it) on that basis.

In addition, unless public key-based digital signatures are employed on each message, the source of each message cannot be absolutely proven. If a signature over a message that originated at an MTA is based on a symmetric key that is shared between that MTA and a network server (e.g., the CMS), it is unclear if the owner of the MTA can claim that the service provider somehow falsified the message.

However, even if each message were to carry a public key-based digital signature and if each MTA were to employ stringent physical security, the customer can still claim in court that someone else initiated that communication without his or her knowledge, just as a customer of a telecommunications carrier on the PSTN can claim, e.g., that particular long distance calls made from the customer's telephone were not authorized by the customer. Such telecommunications carriers commonly address this situation by establishing contractual and/or tariffed relationships with customers in which customers assume liability for unauthorized use of the customer's service. These same contractual principles are typically implemented in service contracts between information services providers such as ISPs and their subscribers. For these reasons, the benefits of non-repudiation seem dubious at best and do not appear to justify the performance penalty of carrying a public key-based digital signature on every message.

5.2.6 Threat summary

This clause provides a summary of the above threats and attacks and a brief assessment of their relative importance.

5.2.6.1 Primary threats

– **Theft of Service:** Attacks are:

- Subscription fraud: This attack is prevalent in today's telephony systems (i.e., the PSTN) and requires little economic investment. It can only be addressed with a Fraud Management system.
- Non-payment for services: Within the PSTN, telecommunications carriers usually do not prosecute the offenders, but simply shut down their accounts. Because prosecution is expensive and not always successful, it is a poor counter to this attack. Methods such as debit-based billing and device authorization (pay as you play), increasingly common in the wireless sector of the PSTN, might be a possible solution for this attack in the IP-Cablecom context. This threat can also be minimized with effective Fraud Management systems.
- MTA clones: This threat requires more technical knowledge than the previous two threats. A technically-knowledgeable adversary or underground organization might offer cloning services for profit. This threat is most effective when combined with subscription fraud, where an MTA registered under a fraudulent account is cloned. This threat can be addressed with both Fraud Management and physical security inside the MTA, or a combination of both.
- Impersonation of a network server: With proper cryptographic mechanisms, authorization and procedural security in place, this attack is unlikely, but has the potential for great damage.
- Protocol manipulation: Can occur only when security protocols are flawed or when not enough cryptographic strength is in place.

- **Bearer Channel Information Disclosure:** Attacks are:
 - Simple snooping: This would happen if voice packets were sent in the clear over some segment of the network. Even if that segment appears to be protected, an insider may still compromise it. This is the only major attack on privacy. The bearer channel privacy attacks listed below are possible but are all of secondary importance.
 - MTA clones: Again, this threat requires more technical knowledge but can be offered as a service by an underground organization. A most likely variation of this attack is when a publicly accessible MTA (e.g., in an office or apartment building) is cloned.
 - Protocol manipulation: A flawed protocol may somehow be exploited to discover bearer channel encryption keys.
 - Off-line cryptanalysis: Even when media packets are protected with encryption, they can be stored and analysed for long periods of time, until the decryption key is finally discovered. Such an attack is not likely to be prevalent, since it is justified only for particularly valuable customer-provided information (IP-Cablecom security is not required to protect data). This attack is more difficult to perform on voice packets (as opposed to data). Still, customers are very sensitive to this threat and it can serve as the basis for a negative publicity campaign by competitors.
- **Signalling Information Disclosure:** This threat is listed as primary only due to potential for bad publicity and customer sensitivity to keeping their numbers and location private. All of the attacks listed below are similar to those for bearer channel privacy and are not described here:
 - Simple snooping;
 - MTA clones;
 - Protocol manipulation;
 - Off-line cryptanalysis;
 - Service disruption.

5.2.6.2 Secondary threats

- **Theft of MTA-based services:** Based on the voice communications services that are planned for the near future, this threat does not appear to have potential for significant economic damage. This could possibly change with the introduction of new value-added services in the future.
- **Illegally registering a leased MTA with a different Service Provider:** Leased MTAs can normally be tracked. Most likely, this threat is combined with the actual theft of a leased MTA. Thus, this threat does not appear to have potential for widespread damage.

In Figure 3, each interface label is of the form:

<label>: <protocol> { <security protocol>/<key management protocol> }

If the key management protocol is missing, it is not needed for that interface. IPCablecom interfaces that do not require security are not shown in Figure 3.

Table 1 briefly describes each of the interfaces shown in Figure 3.

Table 1 – IPCablecom security interfaces table

| Interface | Components | Description |
|------------------|-------------------|---|
| pkt-s0 | MTA-PS/OSS | Immediately after the DHCP sequence in the Secure Provisioning Flow, the MTA performs Kerberos-based key management with the Provisioning Server to establish SNMPv3 keys. The MTA bypasses Kerberized SNMPv3 and uses SNMPv2c in the Basic and Hybrid Flows. |
| pkt-s1 | MTA-TFTP | MTA Configuration file download. When the Provisioning Server in the Secure Provisioning Flow sends an SNMP Set command to the MTA, it includes both the configuration name and the hash of the file. Later, when the MTA downloads the file, it authenticates the configuration file using the hash value. The configuration file may be optionally encrypted. |
| pkt-s2 | CM-CMTS | J.112: This interface should be secured with BPI+ using BPI key management. BPI+ privacy layer on the HFC link. |
| pkt-s3 | MTA-MTA MTA-MG | RTP: End-to-end media packets between two MTAs, or between MTA and MG. RTP packets are encrypted directly with the chosen cipher. Message integrity is optionally provided by an HMAC (Hashed Message Authentication Code). Keys are randomly generated, and exchanged by the two endpoints inside the signalling messages via the CMS or other application server. |
| pkt-s4 | MTA-MTA MTA-MG | RTCP: RTCP control protocol for RTP. Message integrity and encrypted by selected cipher. The RTCP keys are derived using the same secret negotiated during the RTP key management. No additional key management messages are needed or utilized. |
| pkt-s5 | MTA-CMS | NCS: Message integrity and privacy via IPsec. Key management is with Kerberos with PKINIT (public key initial authentication) extension. |
| pkt-s6 | RKS-CMS | RADIUS: IPsec is used for both message integrity, as well as privacy. Key management is IKE or Kerberos. |
| pkt-s7 | RKS-CMTS | RADIUS: IPsec is used for both message integrity, as well as privacy. Key management is IKE or Kerberos. |
| pkt-s8 | CMS-CMTS | COPS: COPS protocol between the GC and the CMTS, used to download QoS authorization to the CMTS. Security is provided with IPsec for message integrity, as well as privacy. Key management is IKE or Kerberos. |
| pkt-s9 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s10 | MGC-MG | TGCP: IPCablecom interface to the PSTN Media Gateway. IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |
| pkt-s11 | | This interface has been removed from the IPCablecom architecture. |

Table 1 – IPCablecom security interfaces table

| Interface | Components | Description |
|------------------|---|---|
| pkt-s12 | MTA-MSO KDC | PKINIT: An AS-REQ message is sent to the KDC with public-key cryptography used for authentication. The KDC verifies the certificate and issues either a service ticket or a ticket granting ticket (TGT), depending on the contents of the AS Request. The AS Reply returned by the KDC contains a certificate chain and a digital signature that are used by the MTA to authenticate this message. In the case that the KDC returns a TGT, the MTA then sends a TGS Request to the KDC to which the KDC replies with a TGS Reply containing a service ticket. The TGS Request/Reply messages are authenticated using a symmetric session key inside the TGT. |
| pkt-s13 | MTA-telephony KDC | PKINIT: See pkt-s12 above. |
| pkt-s14 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s15 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s16 | CMS-CMS CMS-MGC CMS-EBP EBP-EBP | SIP: TLS is used for both message integrity and privacy. Certificates are used for mutual authentication during the TLS handshake. |
| pkt-s17 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s18 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s19 | | This interface has been removed from the IPCablecom architecture. |
| pkt-s20 | MPC-MP | ASP: IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |
| pkt-s21 | DF-CMS | RADIUS: IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |
| pkt-s22 | DF-CMTS | RADIUS: IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |
| pkt-s23 | DF-MGC | RADIUS: IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |
| pkt-s24 | DF-DF | RADIUS: IPsec is used for both message integrity and privacy. Key management is IKE+. |
| pkt-s25 | RKS-MGC | RADIUS: IPsec is used for both message integrity, as well as privacy. Key management is IKE or Kerberos. |
| pkt-s26 | OSS/Prov Serv – MSO KDC OSS/Prov Serv – Telephony KDC | The KDC uses Kerberos to map the MTA's MAC address to its FQDN for the purpose of authenticating the MTA before issuing it a ticket. |
| pkt-s27 | CMS-PS/OSS | HTTP: IPsec is used for both message integrity and privacy. Key management is IKE or Kerberos. |

5.3.2 Security assumptions

5.3.2.1 CMTS downstream messages are trusted

As mentioned previously, it is assumed that CMTS downstream messages cannot be easily modified in transit and a CMTS can be impersonated only at great expense.

Most messages secured in this Recommendation either move over the shared IP network in addition to the J.112 path, or do not go over J.112 at all.

In one case – the case of J.112 QoS messages exchanged between the CMTS and the CM – this assumption does not apply. Although J.112 QoS messages (both upstream and downstream) include an integrity check, the corresponding (BPI+) key management does not authenticate the identity of the CMTS. The CM is unable to cryptographically know that the network element it has connected to is the true CMTS for that network. However, even if a CMTS could be impersonated, it would allow only limited denial-of-service attacks. This vulnerability is not considered to be worth the effort and the expense of impersonating a CMTS.

5.3.2.2 Non-repudiation not supported

Non-repudiation, in this Recommendation, means that an originator of a message cannot deny that he or she sent that message. In this voice communications architecture, non-repudiation is not supported for most messages, with the exception of the top key management layer. This decision was based on the performance penalty incurred with each public key operation. The most important use for non-repudiation would have been during communications set up – to prove that a particular party had initiated that particular communication. However, due to very strict requirements on the set-up time, it is not possible to perform public key operations for each communication.

5.3.2.3 Root Private Key compromise protection

The cryptographic mechanisms defined in this Recommendation are based on a Public Key Infrastructure (PKI). As is the case with most other architectures that are based on a PKI, there is no automated recovery path from a compromise of a Root Private Key. However, with proper safeguards, the probability of this happening is very low, to the point that the risk of a Root Private Key compromise occurring is outweighed by the benefits of this architecture.

The corresponding Root Public Key is stored as a read-only parameter in many components of this architecture. Once the Root Private Key has been compromised, each manufacturer's certificate would have to be manually reconfigured.

Due to this limitation of a PKI, the Root Private Key must be very carefully guarded with procedural and physical security. And, it must be sufficiently long so that its value cannot be discovered with cryptographic attacks within the expected lifetime of the system.

5.3.2.4 Limited prevention of denial-of-service attacks

This Recommendation does not attempt to address all or even most denial-of-service attacks. The cryptographic mechanisms defined in this Recommendation prevent some denial-of-service attacks that are particularly easy to mount and are hard to detect. For example, they will prevent a compromised MTA from masquerading as other MTAs in the same upstream HFC segment and interrupting ongoing communications with illicit HANGUP messages.

This Recommendation will also prevent more serious denial-of-service attacks, such as an MTA masquerading as a CMS in a different network domain that causes all communications set-up requests to fail.

On the other hand, denial-of-service attacks where a router is taken out of service or is bombarded with bad IP packets are not addressed. In general, denial-of-service attacks that are based on damaging one of the network components can only be solved with procedural and physical security, which is out of the scope of this Recommendation.

Denial-of-service attacks where network traffic is overburdened with bad packets cannot be prevented in a large network (although procedural and physical security helps), but can usually be detected. Detection of such an attack and of its cause is out of scope of this Recommendation.

For example, denial-of-service attacks where a router is taken out of order or is bombarded with bogus IP packets cannot be prevented.

5.3.3 Susceptibility of network elements to attack

This clause describes the amount and the type of trust that can be assumed for each element of the voice communications network. It also describes the specific threats that are possible if each network component is compromised. These threats are based on the functionality specified for each component. The general categories of threats are described in 5.2.

Both the trust and the specific threats are described with the assumption that no cryptographic or physical security has been employed in the system, with the exception of the security per ITU-T Rec. J.125 that is on the HFC J.112 links. The goal of this security Recommendation is to address threats that are relevant to this voice communications system.

5.3.3.1 Managed IP network

It is assumed that the same IP network may be shared between multiple, possibly competing service providers. It is also assumed that the service provider may provide multiple services on the same IP network, e.g., Internet connectivity. No assumptions can be made about the physical security of each link in this IP network. An intruder can pop up at any location with the ability to monitor traffic, perform message modification and to reroute messages.

5.3.3.2 MTA

The MTA is considered to be an untrusted network element. It is operating inside customer premises, considered to be a hostile environment. It is assumed that a hostile adversary has the ability to open up the MTA and make software and even hardware modifications to fit his or her needs. This would be done in the privacy of the customer's home.

The MTA communicates with the CMTS over the shared J.112 path and has access to downstream and upstream messages from other MTAs within the same HFC segment.

An MTA is responsible for:

- initiating and receiving communications to/from another MTA or the PSTN;
- negotiating QoS.

A compromise of an MTA can result in:

- MTA clones that are capable of:
 - accessing basic service and any enhanced features in the name of another user's account;
 - violating privacy of the owner of the compromised MTA that does not know that the keys were stolen;
 - identity fraud;
- an MTA running a bad code image that disrupts communications made by other MTAs or degrades network performance.

5.3.3.3 CMTS

The CMTS communicates both over the J.112 path and over the shared IP network. When the CMTS sends downstream messages over the J.112 path, it is assumed that a perpetrator cannot modify them or impersonate the CMTS. Implementing ITU-T Rec. J.125 over that path provides for privacy.

However, when the CMTS is communicating over the shared IP network (e.g., with the CMS or another CMTS), no such assumptions can be made.

While the CMTS, as well as voice communications network servers, are more trusted than the MTAs, they cannot be trusted completely. There is always a possibility of an insider attack.

Insider attacks at the CMTS should be addressed by cryptographic authentication and authorization of the CMTS operators, as well as by physical and procedural security, which are all out of the scope of the IPCablecom Recommendations.

A CMTS is responsible for:

- reporting billing-related statistics to the RKS;
- QoS allocation for MTAs over the J.112 path;
- implementation of BPI+ (MAC layer security) and corresponding key management.

A compromise of a CMTS may result in:

- service theft by reporting invalid information to the RKS;
- unauthorized levels of QoS;
- loss of privacy, since the CMTS holds J.112 keys. This may not happen if additional encryption is provided above the MAC layer;
- degraded performance of some or all MTAs in that HFC segment;
- some or all of the MTAs in one HFC segment completely taken out of service.

5.3.3.4 Voice communications network servers are untrusted network elements

Application servers used for voice communications (e.g., CMS, RKS, Provisioning, OSS, DHCP and TFTP servers) reside on the network and can potentially be impersonated or subjected to insider attacks. The main difference would be in the damage that can be incurred in the case a particular server is impersonated or compromised.

Threats that are associated with each network element are discussed in the following subclauses. To summarize those threats, a compromise or impersonation of each of these servers can result in a wide-scale service theft, loss of privacy, and in highly damaging denial-of-service attacks.

In addition to authentication of all messages to and from these servers (specified in this Recommendation), care should be taken to minimize the likelihood of insider attacks. They should be addressed by cryptographic authentication and authorization of the operators, as well as by stringent physical and procedural security, which are all out of scope of the IPCablecom Recommendations.

5.3.3.4.1 CMS

The Call Management Server is responsible for:

- authorizing individual voice communications by subscribers;
- QoS allocation;
- initializing the billing information in the CMTS;

- distributing per communication keys for MTA-MTA signalling, bearer channel, and DQoS messages on the MTA-CMTS and CMTS-CMTS links;
- interface to PTSN gateway.

A compromised CMS can result in:

- free voice communications service to all of the MTAs that are located in the same network domain (up to 100 000). This may be accomplished by:
 - allowing unauthorized MTAs to create communications;
 - uploading invalid or wrong billing information to the CMTS;
 - combination of both of the above;
- loss of privacy, since the CMS distributes bearer channel keys;
- unauthorized allocation of QoS;
- unauthorized disclosure of customer identity, location (e.g., IP address), communication patterns, and a list of services to which the customer subscribes.

5.3.3.4.2 RKS

The RKS is responsible for collecting billing events and reporting them to the billing system. A compromised RKS may result in:

- free or reduced-rate service due to improper reporting of statistics;
- billing to a wrong account;
- billing customers for communications that were never made, i.e., fabricating communications;
- unauthorized disclosure of customer identity, personal information, service usage patterns, and a list of services to which the customer subscribes.

5.3.3.4.3 OSS, DHCP & TFTP servers

The OSS system is responsible for:

- MTA and service provisioning;
- MTA code downloads and upgrades;
- handling service change requests and dynamic reconfiguration of MTAs.

A compromise of the OSS, DHCP or TFTP server can result in:

- MTAs running illegal code, which may:
 - intentionally introduce bugs or render the MTA completely inoperable;
 - degrade voice communications performance on the IPCablecom or HFC network;
 - configure the MTA with features to which the customer is not entitled;
- MTAs configured with an identity and keys of another customer;
- MTAs configured with service options for which the customer did not pay;
- MTAs provisioned with a bad set of parameters that would make them perform badly or not perform at all.

5.3.3.5 PSTN Gateways

5.3.3.5.1 Media Gateway

The MG is responsible for:

- passing media packets between the IPCablecom network and the PSTN;

- reporting statistics to the RKS.

A compromise of the MG may result in:

- service theft by reporting invalid information to the RKS;
- loss of privacy on communications to/from the PSTN.

5.3.3.5.2 Signalling Gateway

The SG is responsible for translating call signalling between the IPCablecom network and the PSTN.

A compromise of the SG may result in:

- incorrect MTA identity reported to the PSTN;
- unauthorized services enabled within the PSTN;
- loss of PSTN connectivity;
- unauthorized disclosure of customer identity, location (e.g., IP address), usage patterns and a list of services to which the customer subscribes.

6 Security mechanisms

Unless explicitly stated otherwise, the following requirements apply to messages described by this Recommendation:

- ASN.1 encoded messages and objects MUST conform to the Distinguished Encoding Rules per ITU-T Rec. X.690.
- FQDNs used as components of principal names and principal identifiers MUST be rendered in lower case.
- FQDNs MUST NOT include the root domain (i.e., they MUST NOT include a trailing dot). All Kerberos messages in IPCablecom MUST utilize only UDP/IP.

6.1 IPsec

6.1.1 Overview

IPsec provides network-layer security that runs immediately above the IP layer in the protocol stack. It provides security for the TCP or UDP layer and above. It consists of two protocols, IPsec ESP and IPsec AH, as specified in RFC 2401.

IPsec ESP provides confidentiality and message integrity, IP header not included. IPsec AH provides only message integrity, but that includes most of the IP header (with the exception of some IP header parameters that can change with each hop). IPCablecom utilizes only the IPsec ESP protocol per RFC 2406, since authentication of the IP header does not significantly improve security within the IPCablecom architecture.

Each protocol supports two modes of use: transport mode and tunnel mode. IPCablecom only utilizes IPsec ESP transport mode. For more detail on IPsec and these two modes, refer to RFC 2401. Note that in RFC 2401, all implementations of ESP are required to support the concept of Security Associations (SAs). RFC 2401 also provides a general model for processing IP traffic relative to SAs. Although particular IPsec implementations need not follow the details of this general model, the external behavior of any IPsec implementation must match the external behavior of the general model. This ensures that components do not accept traffic from unknown addresses and do not send or accept traffic without security (when security is required). IPCablecom components that implement IPsec are expected to provide behavior that matches the general model described in RFC 2401.

6.1.2 IPCablecom profile for IPsec ESP (Transport mode)

6.1.2.1 IPsec ESP Transform Identifiers

IPsec Transform Identifier (1 byte) is used by IKE to negotiate an encryption algorithm that is used by IPsec. A list of available IPsec Transform Identifiers is specified in RFC 2407. Within IPCablecom, the same Transform Identifiers are used by all IPsec key management protocols: IKE, Kerberos and application layer (embedded in IP signalling messages).

Table 2 describes the IPsec Transform Identifiers (all of which use the CBC mode specified in RFC 2451) supported by IPCablecom.

Table 2 – IPsec ESP Transform Identifiers

| Transform ID | Value (Hex) | Key size (in bits) | MUST support | Description |
|--------------|-------------|--------------------|--------------|---|
| ESP_3DES | 0x03 | 192 | Yes | 3-DES in CBC mode |
| ESP_RC5 | 0x04 | 128 | No | RC5 in CBC mode |
| ESP_IDEA | 0x05 | 128 | No | IDEA in CBC mode |
| ESP_CAST | 0x06 | 128 | No | CAST in CBC mode |
| ESP_BLOWFISH | 0x07 | 128 | No | BLOWFISH in CBC mode |
| ESP_AES | 0x08 | 128 | No | AES-128 in CBC mode with 128-bit block size |
| ESP_NULL | 0x0B | 0 | Yes | Encryption turned off |

The ESP_3DES and ESP_NULL Transform IDs MUST be supported. ESP_AES is included as an optional encryption algorithm. For all of the above transforms, the CBC Initialization Vector (IV) is carried in the clear inside each ESP packet payload per RFC 2451. AES-128 (FIPS PUB 197) MUST be used in CBC mode with a 128-bit block size and a randomly generated Initialization Vector (IV). AES-128 requires 10 rounds of cryptographic operations.

IKE allows negotiation of the encryption key size. Other IPsec key management protocols used by IPCablecom do not allow key size negotiation, and so for consistency a single key size is listed for each Transform ID. If in the future it is desired to increase the key size for one of the above algorithms, IKE will use the built-in key-size negotiation, while other key management protocols will utilize a new Transform ID for the larger key size.

6.1.2.2 IPsec ESP authentication algorithms

The IPsec authentication algorithm (1 byte) is used by IKE to negotiate a packet-authentication algorithm that is used by IPsec. A list of available IPsec authentication algorithms is specified in RFC 2406. Within IPCablecom, the same authentication algorithms are used by all IPsec key management protocols: IKE, Kerberos, and application-layer (embedded in IP signalling messages).

IPCablecom supports the IPsec authentication algorithms in Table 3.

Table 3 – IPsec authentication algorithms

| Authentication algorithm | Value (Hex) | Key size (in bits) | MUST support | Description |
|--------------------------|-------------|--------------------|---------------------------------|----------------------------------|
| HMAC-MD5 | 0x01 | 128 | Yes (also required by RFC 2407) | First 12 bytes of the MD5 HMAC |
| HMAC-SHA | 0x02 | 160 | Yes | First 12 bytes of the SHA-1 HMAC |

The HMAC-MD5-96 and HMAC-SHA-1-96 authentication algorithms **MUST** be supported.

6.1.2.3 Replay protection

In general, IPsec provides an optional replay-protection service (anti-replay service). An IPsec sequence number outside of the current anti-replay window is flagged as a replay and the packet is rejected. When the anti-replay service is turned on, an IPsec sequence number cannot overflow and roll over to 0. Before that happens, a new Security Association must be created as specified in RFC 2406.

Within IPCablecom Security Specification, the IPsec anti-replay service **MUST** be turned on at all times. This is regardless of which key management mechanism is used with the particular IPsec interface.

6.1.2.4 Key management requirements

Within IPCablecom, IPsec is used on a number of different interfaces with different security and performance requirements. Because of this, several different key management protocols have been chosen for different IPCablecom interfaces. On some interfaces it is IKE (see 6.2), on other interfaces it is Kerberos/PKINIT (see 6.4.3.1).

When IKE is not used for key management, an alternative key management protocol needs an interface to the IPsec layer in order to create/update/delete IPsec Security Associations. IPsec Security Associations (SAs) **MUST** be automatically established or re-established as required. This implies that the IPsec layer also needs a way to signal a key management application when a new Security Association needs to be set up (e.g., the old SA is about to expire or there is no SA on a particular interface).

In addition, some network elements are required to run multiple key management protocols. In particular, the Application Server (such as a CMS) and the MTA **MUST** support multiple key management protocols. The MTA **MUST** support Kerberos/PKINIT on the MTA-CMS signalling interface. IKE **MUST** be supported on the CMS-CMTS and CMS-RKS interfaces.

The PF_KEY interface (see RFC 2367) **SHOULD** be used for IPsec key management within IPCablecom and would satisfy the above-listed requirements. For example, PF_KEY permits multiple key management applications to register for rekeying events. When the IPsec layer detects a missing Security Association, it signals the event to all registered key management applications. Based on the Identity Extension associated with that Security Association, each key management application decides if it should handle the event.

6.2 Internet Key Exchange (IKE)

6.2.1 Overview

IPCablecom utilizes RFC 2409 (IKE) as one of the key management protocols for IPsec. It is utilized on interfaces where:

- there is not a very large number of connections;
- the endpoints on each connection know about each other's identity in advance.

Within IPCablecom, IKE key management is completely asynchronous to call signalling messages and does not contribute to any delays during communications set-up. The only exception would be some unexpected error, where a Security Association is unexpectedly lost by one of the endpoints.

IKE is a peer-to-peer key management protocol. It consists of two phases. In the first phase, a shared secret is negotiated via a Diffie-Hellman key exchange. It is then used to authenticate the second IKE phase. The second phase negotiates another secret, used to derive keys for the IPsec ESP protocol.

6.2.2 IPCablecom profile for IKE

6.2.2.1 First IKE phase

There are several modes defined for authentication during the first IKE phase.

6.2.2.1.1 IKE authentication with signatures

In this mode, both peers **MUST** be authenticated with X.509 certificates and digital signatures. IPCablecom utilizes this IKE authentication mode on some IPsec interfaces. Whenever this mode is utilized, both sides **MUST** exchange X.509 certificates (although this is optional in RFC 2409).

6.2.2.1.2 IKE authentication with public-key encryption

IPCablecom **MUST NOT** utilize this IKE authentication with public key encryption. In order to perform this mode of IKE authentication, the initiator must already have the responder's public key, which is not supported by IPCablecom.

6.2.2.1.3 IKE authentication with pre-shared keys

A key derived by some out-of-band (e.g., manual) mechanism is used to authenticate the exchange. IPCablecom utilizes this IKE authentication mode on some IPsec interfaces. IPCablecom does not specify the out-of-band method for deriving pre-shared keys.

When using pre-shared keys, the strength of the system is dependent upon the strength of the shared secret. The goal is to keep the shared secret from being the weak link in the chain of security. This implies that the shared secret needs to contain as much entropy (randomness) as the cipher being used. In other words, the shared secret should have at least 128-160 bits of entropy. This means if the shared secret is just a string of random 8-bit bytes, then the key can be 16-20 bytes. If the shared secret is derived from a passphrase that is a string of random alpha-nums (a-zA-Z0-9/+), then it should be at least 22-27 characters. This is because there are only 64 characters (6 bits) instead of 256 characters (8 bits) per 8-bit byte, which implies an expansion of 4/3 the length for the same amount of entropy. Both random 8-bit bytes and random 6-bit bytes assume truly random numbers. If there is any structure in the password/passphrase, like deriving from English, then even longer passphrases are necessary. A passphrase composed of English would need on the order of 60-100 characters, depending on mixing of case. Using English (or any language, for that matter) passphrases creates the problem that, if an attacker knows the language of the passphrase then they have less space to search. It is less random. This implies fewer bits of entropy per character, so a longer passphrase is required to maintain the same level of entropy.

6.2.2.2 Second IKE phase

In the second IKE phase, an IPsec ESP SA is established, including the IPsec ESP keys and ciphersuites. It is possible to establish multiple Security Associations with a single second-phase IKE exchange.

First, a shared second-phase secret is established, and then all the IPsec keying material is derived from it using the one-way function specified in RFC 2409.

The second-phase secret is built from encrypted nonces that are exchanged by the two parties. Another Diffie-Hellman exchange may be used in addition to the encrypted nonces. Within IPCablecom, IKE **MUST NOT** perform a Diffie-Hellman exchange in the second IKE phase in order to avoid the associated performance penalties.

The second IKE phase is authenticated using a shared secret that was established in the first phase. Supported authentication algorithms are the same as those specified for IPsec in 6.1.2.2.

6.2.2.3 Encryption algorithms for IKE exchanges

Both phase 1 and phase 2 IKE exchanges include some symmetrically-encrypted messages. The encryption algorithms supported as part of the IPCablecom profile for IKE MUST be the same algorithms identified in the IPCablecom profile for IPsec ESP in Table 2.

6.2.2.4 Diffie-Hellman groups

IKE defines specific sets of Diffie-Hellman parameters (i.e., prime and generator) that may be used for the phase 1 IKE exchanges. These are called groups in RFC 2409. The use of Diffie-Hellman groups within IPCablecom IKE is identical to that specified in RFC 2409. Note that this is different from the requirements pertaining to the IPCablecom use of groups in PKINIT described in 6.4.2.1.1. Annex A provides details of the first and second Oakley groups.

6.3 SNMPv3

Any mention of SNMP in this Recommendation without a specific reference to the SNMP protocol version must be interpreted as SNMPv3.

IPCablecom supports use of SNMPv2c coexistence for network management operations for devices provisioned under the Basic Flow or the Hybrid Flow. It also supports the SNMPv3/v2c coexistence for network management operations when the device is provisioned under the Secure Flow. Refer to the provisioning specification (ITU-T Rec. J.167) for the use of SNMP coexistence in IPCablecom. For any interface within the IPCablecom architecture utilizing SNMPv3 authentication MUST be turned on at all times and SNMPv3 privacy MAY also be utilized.

In order to establish SNMPv3 keys, all IPCablecom SNMP interfaces SHOULD utilize Kerberized SNMPv3 key management (as specified in 6.5.4). In addition, SNMPv3 key management techniques specified in RFC 3414 MAY also be used.

6.3.1 SNMPv3 Transform Identifiers

The SNMPv3 Transform Identifier (1 byte) is used by Kerberized key management to negotiate an encryption algorithm for use by SNMPv3.

For IPCablecom, the SNMPv3 Transform Identifiers in Table 4 MUST be supported.

Table 4 – SNMPv3 Transform Identifiers

| Transform ID | Value (Hex) | Key size (in bits) | MUST be supported | Description |
|--------------|-------------|--------------------|-------------------|--|
| SNMPv3_DES | 0x21 | 128 | Yes | DES in CBC mode The first 64 bits are used as the DES Key and the remaining 64 are used as the pre-IV. |
| SNMPv3_NULL | 0x20 | 0 | Yes | Encryption turned off |

The SNMPv3_DES and the SNMPv3_NULL Transform IDs MUST be supported. The DES encryption transform for SNMPv3 is specified in RFC 3414. Note that DES encryption does not provide strong privacy but is currently the only encryption algorithm specified by the SNMPv3 standard.

6.3.2 SNMPv3 authentication algorithms

SNMPv3 authentication algorithm (1 byte) is used by Kerberized key management to negotiate an SNMPv3 message authentication algorithm.

For IPCablecom, the SNMPv3 authentication algorithms in Table 5 are supported (both of which are specified in RFC 3414):

Table 5 – SNMPv3 authentication algorithms

| Authentication algorithm | Value (Hex) | Key size (in bits) | MUST be supported | Description |
|--------------------------|-------------|--------------------|---------------------------------|-------------|
| SNMPv3_HMAC-MD5 | 0x21 | 128 | Yes (also required by RFC 3414) | MD5 HMAC |
| SNMPv3_HMAC-SHA-1 | 0x22 | 160 | No (SHOULD be supported) | SHA-1 HMAC |

The SNMPv3_HMAC-MD5 authentication algorithm MUST be supported. The SNMPv3_HMAC-SHA-1 authentication algorithm SHOULD be supported.

6.4 Kerberos/PKINIT

6.4.1 Overview

IPCablecom utilizes the concept of Kerberized IPsec for signalling between an Application Server, such as the CMS, and the MTA. This refers to the ability to create IPsec security associations using keys derived from the subkeys exchanged using the Kerberos AP Request/AP Reply messages. On this interface, Kerberos (see IETF RFC 4120) is utilized with the PKINIT public key extension (also see IETF RFC 4556).

Kerberized IPsec consists of three distinct phases:

- 1) A client SHOULD obtain a TGT (Ticket Granting Ticket) from the KDC (Key Distribution Centre). Once the client obtains the TGT, it MUST use the TGT in the subsequent phase to authenticate to the KDC and obtain a ticket for the specific Application Server, e.g., a CMS.

In Kerberos, tickets are symmetric authentication tokens encrypted with a particular server's key. (For a TGT, the server is the KDC.) Tickets are used to authenticate a client to a server. A PKI equivalent of a ticket would be an X.509 certificate. In addition to authentication, a ticket is used to establish a session key between a client and a server, where the session key is contained in the ticket.

The logical function within the KDC that is responsible for issuing TGTs is referred to as an Authentication Server or AS.

- 2) A client obtains a ticket from the KDC for a specific Application Server. In this phase, a client can authenticate with a TGT obtained in the previous phase. A client can also authenticate to the KDC directly using a digital certificate or a password-derived key, bypassing phase 1.

The logical function within the KDC that is responsible for issuing Application Server tickets based on a TGT is referred to as the Ticket Granting Server (TGS). When the TGT is bypassed, it is the Authentication Server that issues the Application Server tickets.

- 3) A client utilizes the ticket obtained in the previous phase to establish a pair of Security Parameters (one to send and one to receive) with the server. This is the only key management phase that is not already specified in an IETF standard. The previous two phases are part of standard Kerberos, while this phase defines new messages that tie together Kerberos key management and IPsec.

Figure 4 illustrates the three phases of Kerberos-based key management for IPsec.

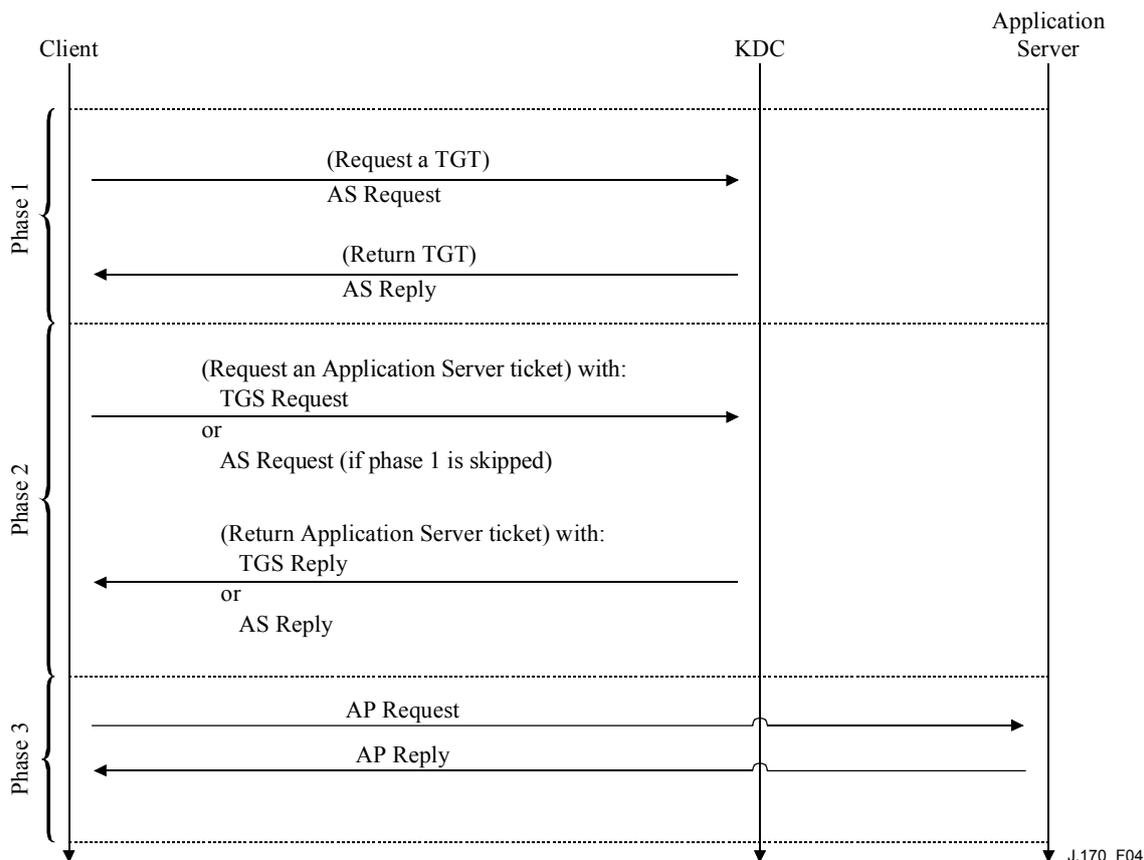


Figure 4 – Kerberos-based key management for IPsec

During the AS Request/AS Reply exchange (that can occur in either phase 1 or phase 2), the client and the KDC perform mutual authentication. In standard Kerberos, a client key that is shared with the KDC is used for this authentication (see 6.4.2.2). The same AS Request/AS Reply exchange may also be authenticated with digital signatures and certificates when the PKINIT public key extension is used (see 6.4.2). Both the TGT and the Application Server tickets used within IPCablecom have a relatively long lifetime (days or weeks). This is acceptable as 3-DES, a reasonably strong symmetric algorithm, is required by IPCablecom.

IPCablecom utilizes the concept of a TGT (Ticket Granting Ticket), used to authenticate subsequent requests for Application Server tickets. The use of a TGT has two main advantages:

- It limits the exposure of the relatively long-term client key (that is in some cases reused as the service key). This consideration does not apply to clients that use PKINIT.
- It reduces the number of public key operations that are required for PKINIT clients.

The Application Server ticket contains a symmetric session key, which MUST be used in phase 3 to establish a set of keys for the IPsec ESP protocol. The keys used by IPsec MUST expire after a configurable time-out period (e.g., 10 minutes). Normally, the same Application Server ticket SHOULD be used to automatically establish a new IPsec SA. However, there are instances where it is desirable to drop IPsec sessions after a Security Association time out and establish them on-demand later. This allows for improved system scalability, since an application server (e.g., CMS) does not need to maintain a SA for every client (e.g., MTA) that it controls. It also is possible that a group of application servers (e.g., CMS clusters) may control the same subset of clients (e.g., MTAs) for load balancing. In this case, the MTA is not required to maintain an SA with each CMS in that group. This clause provides specifications for how to automatically establish

a new IPsec SA right before an expiration of the old one and how to establish IPsec SAs on-demand, when a signalling message needs to be sent.

IPCablecom also utilizes the Kerberos protocol to establish SNMPv3 keys between the MTAs and the Provisioning Server. Kerberized SNMPv3 key management is very similar to the Kerberized IPsec key management and consists of the same phases that were explained above for Kerberized IPsec. Each MTA again utilizes the PKINIT extension to Kerberos to authenticate itself to the KDC with X.509 certificates.

Once an MTA obtains its service ticket for the Provisioning Server, it utilizes the same protocol that is used for Kerberized IPsec to authenticate itself to the Provisioning Server and to generate SNMPv3 keys. The key management protocol is specified to allow application-specific data that has different profiles for SNMPv3 and IPsec. The only exception is the Rekey exchange that is specified for IPsec in order to optimize the MTA hand-off between the members of a CMS cluster. The Rekey exchange is not utilized for SNMPv3 key management.

A recipient of any Kerberos message that does not fully comply with the IPCablecom requirements MUST reject the message.

6.4.1.1 Kerberos ticket storage

Kerberos clients that store tickets in persistent storage will be able to re-use the same Kerberos ticket after a reboot. In the event that PKINIT is used, this avoids the need to perform public key operations.

A Kerberos client MUST NOT obtain a new TGT upon reboot if it possesses a valid service ticket.

An MTA MUST store the Provisioning Server service ticket in persistent storage. An MTA MUST be capable of storing a minimum of $\text{pkcMtaDevEndPntCount}+1$ CMS service tickets in persistent storage, where $\text{pkcMtaDevEndPntCount}$ is the MIB object specifying the number of physical endpoints on the MTA. An MTA MUST store all CMS service tickets that correspond to active endpoints. This means that an MTA that reaches the maximum number of CMS service tickets that can be stored in persistent storage will not over-write CMS service tickets that correspond to active endpoints.

Kerberos clients other than MTAs SHOULD retain service tickets in persistent storage.

Note that Kerberos clients will need to store additional information in order to use and validate the ticket, such as the session key information, the client IP address, and the ticket validity period. Refer to 7.1 for additional information on re-using stored tickets.

6.4.2 PKINIT exchange

Figure 5 illustrates how a client may use PKINIT to either obtain a TGT (phase 1) or a Kerberos ticket for an Application Server (phase 2).

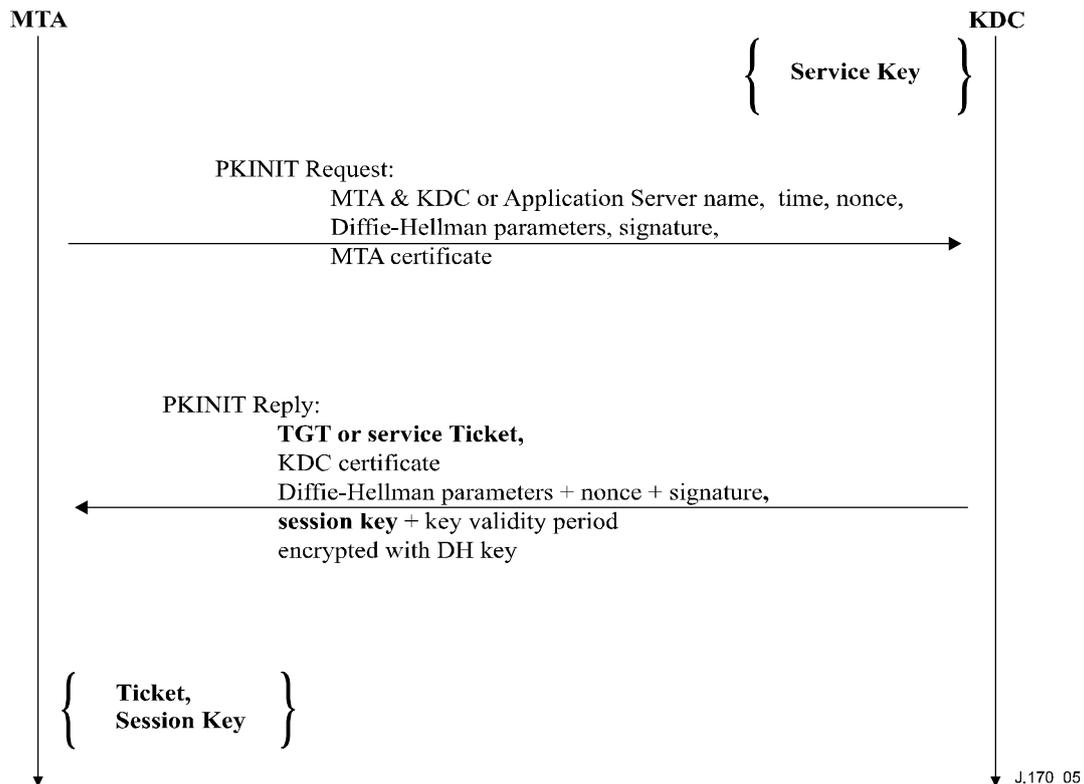


Figure 5 – PKINIT exchange

The PKINIT Request is carried as a Kerberos pre-authenticator field inside an AS Request and the PKINIT Reply is a pre-authenticator inside the AS Reply. The syntax of the Kerberos AS Request/Reply messages and how pre-authenticators plug in is specified in IETF RFC 4120.

In this clause, the PKINIT client is referred to as an MTA, as it is currently the only IPCablecom element that authenticates itself to the KDC with the PKINIT protocol. If in the future other IPCablecom elements will also utilize the PKINIT protocol, the same specifications will apply. IPCablecom use of the AS Request/AS Reply exchange without PKINIT is covered in 6.4.3.

Figure 5 lists several important parameters in the PKINIT Request and Reply messages. These parameters are:

PKINIT Request

- MTA (Kerberos principal) name – found in the KDC-REQ-BODY Kerberos structure (see IETF RFC 4120). For the format used in IPCablecom, see 6.4.7.
- KDC or Application Server (Kerberos principal) name – found in the KDC-REQ-BODY Kerberos structure (see IETF RFC 4120). For the format used in IPCablecom, see 6.4.6.
- Time – found in the PKAuthenticator structure, specified by PKINIT (IETF RFC 4556).
- Nonce – found in the PKAuthenticator structure, specified by PKINIT (IETF RFC 4556). There is also a second nonce in the KDC-REQ-BODY Kerberos structure.
- Diffie-Hellman parameters, signature and MTA certificate – these are all specified by PKINIT (IETF RFC 4556) and their use in IPCablecom is specified in 6.4.2.1.1.

PKINIT Reply

- TGT or Application Server Ticket – found in the KDC-REP Kerberos structure (see IETF RFC 4120).
- KDC Certificate, Diffie-Hellman parameters, signature – these are all specified by PKINIT (see IETF RFC 4556) and their use in IPCablecom is specified in 6.4.2.1.2.

- Nonce – found in the KdcDHKeyInfo structure, specified by PKINIT (IETF RFC 4556). This nonce must be the same as the one found in the PKAuthenticator structure of the PKINIT Request. There is another nonce in EncKDCRepPart Kerberos structure (see IETF RFC 4120). This nonce must be the same as the one found in the KDC-REQ-BODY of the PKINIT Request.
- Session key, key validity period – found in the EncKDCRepPart Kerberos structure (see IETF RFC 4120).

In Figure 5, the PKINIT exchange is performed at long intervals, in order to obtain an (intermediate) symmetric session key. This session key is shared between the MTA and the server via the server's ticket, where the application server may be the KDC (in which case the ticket is the TGT).

6.4.2.1 PKINIT profile for IPCablecom

A particular MTA implementation MUST utilize the PKINIT exchange to either obtain Application Server tickets directly, or obtain a TGT first and then use the TGT to obtain Application Server tickets. An MTA implementation MAY also support both uses of PKINIT, where the decision to get a TGT first or not is local to the MTA and is dependent on a particular MTA implementation. On the other hand, the KDC MUST be capable of processing PKINIT requests for both a TGT and for Application Server tickets.

The PKINIT exchange occurs independent of the signalling protocol, based on the current Ticket Expiration Time ($\text{Ticket}_{\text{EXP}}$) and on the PKINIT Grace Period ($\text{PKINIT}_{\text{GP}}$). If the PKINIT client is an MTA and the ticket it currently possesses corresponds to the Provisioning Server in the MIB, a KDC for a REALM that currently exists in the REALM table, or a CMS that currently exists in the CMS table, the MTA MUST initiate the PKINIT exchange at the time: $\text{Ticket}_{\text{EXP}} - \text{PKINIT}_{\text{GP}}$. If the PKINIT client is an MTA and the ticket it currently possesses does not correspond to the Provisioning Server in the MIB, a KDC for a REALM that currently exists in the REALM table, or a CMS that currently exists in the CMS table, the MTA MUST NOT initiate a PKINIT exchange. On the interfaces where $\text{PKINIT}_{\text{GP}}$ is not defined, the MTA SHOULD perform PKINIT exchanges on-demand.

In the case where PKINIT is used to obtain an Application Server ticket directly, the use of the grace period accounts for a possible clock skew between the MTA and the CMS or other application server. If the MTA is late with the PKINIT exchange, it still has until $\text{Ticket}_{\text{EXP}}$ before the Application Server starts rejecting the ticket. Similarly, if PKINIT is used to obtain a TGT, the grace period accounts for a possible clock skew between the MTA and the KDC.

The PKINIT exchange stops after the MTA obtains a new ticket, and therefore does not affect existing security parameters between the MTA and the CMS or other application server. Synchronizing the PKINIT exchange with the AP Request/Reply exchange is not required as long as the AS Request/Reply exchange results in a valid, non-expired Kerberos ticket.

The PKINIT Request/Reply messages contain public key certificates, which make them longer than a normal size of a UDP packet. In this case, large UDP packets MUST be sent using IP fragmentation.

A KDC server SHOULD be implemented on a separate host, independent of the Application Server. This would mean that frequent PKINIT operations from some MTAs will not affect the performance of any of the application servers or the performance of those MTAs that do not require frequent PKINIT exchanges.

Kerberos tickets MUST NOT be issued for a period of time that is longer than 7 days. The MTA clock MUST NOT drift more than 2.5 minutes within that period (7 days). The PKINIT Grace Period ($\text{PKINIT}_{\text{GP}}$) MUST be at least 15 minutes.

6.4.2.1.1 PKINIT Request

The PKINIT Request message (PA-PK-AS-REQ) in IETF RFC 4556 is defined as:

```
PA-PK-AS-REQ ::= SEQUENCE {
    signedAuthPack      [0] ContentInfo
    trustedCertifiers   [1] SEQUENCE OF TrustedCas OPTIONAL,
    kdcCert              [2] IssuerAndSerialNumber OPTIONAL
    encryptionCert      [3] IssuerAndSerialNumber OPTIONAL
}
```

The following fields **MUST** be present in PA-PK-AS-REQ for IPCablecom (and all other fields **MUST NOT** be present):

- signedAuthPack – a signed authenticator field, needed to authenticate the client. It is defined in Cryptographic Message Syntax identified by the SignedData OID:

```
{iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs7(7) 2}.
```

SignedData is defined as:

```
SignedData ::= SEQUENCE {
    version              CMSVersion,
    digestAlgorithms     DigestAlgorithmIdentifiers,
    encapContentInfo     EncapsulatedContentInfo,
    certificates          [0] IMPLICIT CertificateSet OPTIONAL,
    crls                 [1] IMPLICIT CertificateRevocationLists OPTIONAL,
    signerInfos          SignerInfos
}
```

- digestAlgorithms – for now **MUST** contain an algorithm identifier for SHA-1. Other digest algorithms may optionally be supported in the future.
- encapContentInfo – is of type EncapsulatedContentInfo that is defined by Cryptographic Message Syntax as:

```
EncapsulatedContentInfo ::= SEQUENCE {
    eContentTypeContent Type,
    eContent [0] EXPLICIT OCTET STRING OPTIONAL
}
```

Here eContentType indicates the type of data and for PKINIT must be set to:

```
{iso(1) org(3) dod(6) internet(1) security(5) kerberosv5(2) pkinit(3)
    pkauthdata(1)}
```

eContent is a data structure of type AuthPack encoded inside an OCTET STRING:

```
AuthPack ::= SEQUENCE {
    pkAuthenticator      [0] PKAuthenticator,
    clientPublicValue    [1] SubjectPublicKeyInfo OPTIONAL
}
```

The optional clientPublicValue parameter inside the AuthPack **MUST** always be present for IPCablecom. (This parameter specifies the client's Diffie-Hellman public value.)

```
PKAuthenticator ::= SEQUENCE {
    cusec [0] INTEGER,
    -- for replay prevention as in RFC 4120
    ctime [1] KerberosTime,
    -- for replay prevention as in RFC 4120
    nonce [2] INTEGER,
    -- zero only if client will accept
    -- cached DH parameters from KDC;
    -- must be non-zero otherwise
}
```

```

    pachecksum      [3]  Checksum
                      -- Checksum over KDC-REQ-BODY
                      -- Defined by Kerberos spec
}

```

The pachecksum field MUST use the Kerberos checksum type rsa-md5, a plain MD5 checksum over the KDC-REQ-BODY.

The nonce field MUST be non-zero, indicating that the client does not support the caching of Diffie-Hellman values and their expiration.

- certificates – required by IPCablecom. This field MUST contain an MTA Device Certificate and an MTA Manufacturer Certificate. This field MUST NOT contain any other certificates. All IPCablecom certificates are X.509 certificates for RSA Public keys as specified in clause 8.
- crls – MUST NOT be filled in by the MTA.
- signerInfos – MUST be a set with exactly one member that holds the MTA signature. This signature is a part of a SignerInfo data structure defined within the Cryptographic Message Syntax. All optional fields in this data structure MUST NOT be used in IPCablecom. The digestAlgorithm MUST be set to SHA-1:

```
iso(1) identified-organization(3) oiw(14) secsig(3) algorithm(2) 26
```

and the signatureAlgorithm MUST be set to rsaEncryption:

```
iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs-1(1) 1
```

PKINIT allows an Ephemeral-Ephemeral Diffie-Hellman exchange as part of the PKINIT Request/Reply sequence. (Ephemeral-Ephemeral means that both parties during each exchange randomly generate the Diffie-Hellman private exponents.) The Kerberos session key is returned to the MTA in the PKINIT Reply, encrypted with a secret that is derived from the Diffie-Hellman exchange. Within IPCablecom, the Ephemeral-Ephemeral Diffie-Hellman MUST be supported.

The IKE specification in RFC 2409 defines Diffie-Hellman parameters as Oakley groups. Within the IPCablecom PKINIT profile the 2nd Oakley group MUST be supported and the 1st Oakley group MAY also be supported. Annex A provides details of the first and second Oakley groups.

When generating Diffie-Hellman private keys, a device MUST generate a key of length at least 144 bits when the first Oakley group is used and MUST generate a key of length at least 164 bits when the second Oakley group is used.

For further details of PKINIT, please refer to IETF RFC 4556.

Additionally, PKINIT supports a Static-Ephemeral Diffie-Hellman exchange, where the client is required to possess a Diffie-Hellman certificate in addition to an RSA certificate. This mode MUST NOT be used within IPCablecom.

PKINIT also allows a single client RSA key to be used both for digital signatures and for encryption – wrapping the Kerberos session key in the PKINIT Reply. This mode MUST NOT be used within IPCablecom.

PKINIT has an additional option for a client to use two separate RSA keys – one for digital signatures and one for encryption. This mode MUST NOT be used within IPCablecom.

Upon receipt of a PA-PK-AS-REQ, the KDC MUST:

- 1) check the validity of the certificate chain (MTA Device Certificate, MTA Manufacturer Certificate, MTA Root Certificate);
- 2) check the validity of the signature in the (single) SignerInfo field;
- 3) check the validity of the checksum in the PKAuthenticator.

6.4.2.1.2 PKINIT Reply

The PKINIT Reply message (PA-PK-AS-REP) in IETF RFC 4556 is defined as follows:

```
PA-PK-AS-REP ::= CHOICE {  
    dhSignedData    [0] ContentInfo,  
    encKeyPack      [1] ContentInfo,  
}
```

IPCablecom MUST use only the dhSignedData choice, which is needed for a Diffie-Hellman exchange.

The value of the Kerberos session key is not present in PA-PK-AS-REP. It is found in the encrypted portion of the AS Reply message that is specified in IETF RFC 4556. The AS Reply MUST be encrypted with 3-DES CBC, where the corresponding Kerberos etype value MUST be des3-cbc-md5. Other encryption types may be supported in the future.

The client MUST use PA-PK-AS-REP to determine the encryption key used on the AS Reply. This PKINIT Reply contains the KDC's Diffie-Hellman public value that is used to generate a shared secret (part of the key agreement). This shared secret is used to encrypt/decrypt the private part of the AS Reply.

- dhSignedData – dhSignedData is identified by the SignedData oid:

```
{iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs7(7) 2}
```

Within SignedData (specified in 6.4.2.1.1):

- digestAlgorithms for now MUST contain an algorithm identifier for SHA-1. Other digest algorithms may optionally be supported in the future.
- encapContentInfo is of type pkdhkeydata, where eContentType contains the following OID value:

```
{iso(1) org(3) dod(6) internet(1) security(5) kerberosv5(2) pkinit(3)  
    pkdhkeydata(2)}
```

eContent is of type KdcDHKeyInfo (encoded inside an OCTET STRING):

```
KdcDHKeyInfo ::= SEQUENCE {  
    subjectPublicKey [0] BIT STRING,  
    -- used only when utilizing Diffie-Hellman  
    -- Equals public exponent (g^a mod p)  
    -- INTEGER encoded as payload of  
    -- BIT STRING  
    nonce [1] INTEGER,  
    -- Binds response to the request  
    -- Exception: Set to zero when KDC  
    -- is using a cached DH value  
    dhKeyExpiration [2] KerberosTime OPTIONAL  
    -- Expiration time for KDC's cached  
    -- DH value
```

The nonce MUST be the same nonce that was passed in by the client in the PKINIT Request.

- The subjectPublicKey MUST be the Diffie-Hellman public value generated by the KDC. The Diffie-Hellman-derived key is used to directly encrypt part of the AS Reply. The requirements on the length of the Diffie-Hellman private exponent are as defined in 6.4.2.1.1.
- The dhKeyExpiration MUST not be present as caching of Diffie-Hellman values is not permitted.

- certificates – required by IPCablecom. This field MUST contain a KDC certificate. If a Local System CA issued the KDC certificate, then the corresponding Local System CA Certificate MUST also be present. The Service Provider CA Certificate MUST also be present in this field. This field MAY contain the Service Provider Root CA certificate (refer to 8.2.1 for validating the Service Provider Root CA certificate if it is included in the PKINIT Reply). This field MUST NOT contain any other certificates. If the MTA is configured with a specific service provider name, it MUST verify that the Service Provider name is identical to the value of the OrganizationName attribute in the subjectName of the Service Provider certificate. If the Local System Certificate is present, then the MTA MUST verify that the Service Provider name is identical to the value of the OrganizationName attribute in the subjectName of the Local System Certificate. In addition to standard certificate verification rules specified in RFC 2459, an MTA MUST verify that the KDC certificate includes a subjectAltName extension in the format specified in 8.2.3.4.1. The MTA MUST verify that the extension contains a valid KDC principal name and that the KDC realm in this extension is identical to the server realm name in the encrypted portion of the AS Reply message (EncKDCRepPart).
- crls – this optional field MAY be filled in by the KDC.
- signerInfos – MUST be a set with exactly one member that holds the KDC signature. This signature is a part of a SignerInfo data structure defined within the Cryptographic Message Syntax. All optional fields in this data structure MUST NOT be used in IPCablecom. The digestAlgorithm MUST be set to SHA-1:

```
iso(1) identified-organization(3) oiw(14) secsig(3) algorithm(2) 26
```

The signatureAlgorithm MUST be set to rsaEncryption:

```
iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs-1(1) 1
```

Upon receipt of a PA-PK-AS-REP, the client MUST:

- 1) check the value of the nonce in the eContent field;
- 2) check the validity of the KDC certificate;
- 3) check the validity of the signature in the SignerInfo field.

6.4.2.1.2.1 PKINIT error messages

In the case that a PKINIT Request is rejected, instead of a PKINIT Reply the KDC MUST return a Kerberos error message of type KRB_ERROR, as defined in IETF RFC 4556. Any error code that is defined in IETF RFC 4556 for PKINIT MAY be returned.

- The KRB_ERROR MUST use typed-data of REQ-NONCE to bind the error message to the nonce from the KDC-REQ-BODY portion of the AS-REQ message. This error message MUST NOT include the optional e-cksum member that would contain a keyed checksum of the error reply. The use of this field is not possible during the PKINIT exchange, since the client and the KDC do not share a symmetric key.

When a client receives an error message from the KDC, in some cases this Recommendation calls for the client to take some recovery steps and then send a new AS Request or TGS Request. When a client is responding to an error message, it is not a retry and MUST NOT be considered to be part of the client's back-off and retry procedure specified in 6.4.8. The client MUST reset its timers accordingly, to reflect that the new request in response to an error message is not a retry.

Although this Recommendation calls for a KDC to return some specific error codes under certain error conditions, in the case when a KDC is repeatedly getting the same error from the same client IP address, it MAY at some point choose to stop sending back any further replies (errors or otherwise) to this client.

6.4.2.1.2.1.1 Clock skew error

When the KDC clock and the client clock are off by more than the limit for a clock skew, an error code `KRB_AP_ERR_SKEW` MUST be returned. The value for the maximum clock skew allowed by the KDC MUST NOT exceed 5 minutes. The optional client's time in the `KRB_ERROR` MUST be filled out, and the client MUST compute the difference (in seconds) between the two clocks based upon the client and server time contained in the `KRB_ERROR` message. The client SHOULD store this clock difference in non-volatile memory and MUST use it to adjust Kerberos timestamps in subsequent KDC request messages (AS Request and TGS Request) by adding the clock skew to its local clock value each time. The client MUST maintain a separate clock skew value for each realm. The clock skew values are intended for uses only within the Kerberos protocol and SHOULD NOT otherwise affect the value of the local clock (since a clock skew is likely to vary from realm to realm).

In the case that a KDC request fails due to a clock skew error, a client MUST immediately retry after adjusting the Kerberos timestamp inside the KDC Request message.

In addition, the MTA MUST validate the time offset returned in the clock skew error, to make sure that it does not exceed a maximum allowable amount. This maximum time offset MUST NOT exceed 1 hour. This MTA check against a maximum time offset protects against an attack in which a rogue KDC attempts to fool an MTA into accepting an expired KDC certificate.

6.4.2.1.3 Pre-authenticator for Provisioning Server Location

An AS Request sent by the MTA MUST include this `PROV-SRV-LOCATION` pre-authenticator that the KDC can use to locate the Provisioning Server.

The pre-authenticator type MUST be `-1` (according to IETF RFC 4120, the negative type is used for application-specific pre-authenticators). Its ASN.1 encoding is specified as:

```
PROV-SRV-LOCATION ::= GeneralString
                    -- Provisioning Server's FQDN
```

6.4.2.2 Profile for the Kerberos AS Request/AS Reply messages

As mentioned earlier, the PKINIT Request and Reply are pre-authenticator fields embedded into the AS Request/AS Reply messages. The IPCablecom-specific `PROV-SRV-LOCATION` pre-authenticator MUST be used in combination with PKINIT. All other pre-authenticators MUST NOT be used in combination with PKINIT.

The optional fields `enc-authorization-data`, `additional-tickets` and `rtime` in the `KDC-REQ-BODY` MUST NOT be present in the AS Request. All other optional fields in the AS Request MAY be present for IPCablecom. The client MUST NOT set any of the `KDCOptions` in the `AS-REQUEST`, except that the `DISABLE-TRANSITED-CHECK` option MAY be set.

The MTA MUST include its IP address in the optional `addresses` field of the `KDC-REQ-BODY`. The KDC MUST verify that the `addresses` field in the `KDC-REQ-BODY` contains exactly one IP address and that it is identical to the IP address in the IP header of the AS Request. After the KDC validates the `addresses` field, it MUST include it in the `caddr` fields of the issued ticket and the AS Reply. The KDC MUST reject an AS Request that does not include the MTA's IP address. In this case the KDC MUST return a `KDC_ERR_POLICY` error code.

If a KDC receives an `AS-REQ` message in which any of the `KDCOptions` are set, except for the `DISABLE-TRANSITED-CHECK` option, the KDC MUST return an error with the error code `KDC_ERR_POLICY`.

In the AS Reply, `key-expiration`, `starttime`, and `renew-till` optional fields MUST NOT be present. The session key contained in the `AS-REPLY` (which MUST be identical to the session key in the ticket) MUST be etype `des3-cbc-md5`.

The encrypted part of the AS Reply is of the type EncryptedData. The ASN.1 definition of EncryptedData that is used inside multiple Kerberos objects is missing from the Kerberos revisions IETF RFC 4120. In all cases, EncryptedData MUST be DER-encoded with EXPLICIT tags as the following ASN.1 structure:

```
EncryptedData ::= SEQUENCE {
    etype    [0] INTEGER,          -- EncryptionType
    kvno     [1] INTEGER OPTIONAL, -- service key
                                         -- version number
    cipher   [2] OCTET STRING     -- ciphertext
}
```

When EncryptedData contains ciphertext that is encrypted with a service key, the 'kvno' element MUST be present and MUST identify the version of the service key that was used to encrypt the data. When EncryptedData contains ciphertext that is encrypted with a Kerberos session key or with a reply key derived from a PKINIT pre-authenticator, the 'kvno' element MUST NOT be present. This is the case for the encrypted portion of the AS Reply.

The encryption type for an encrypted portion of the AS Reply MUST be set to des3-cbc-md5. In order to generate the value of the 'cipher' element of the EncryptedData, the following data MUST be concatenated and processed in the following sequence before being encrypted with 3-DES CBC, IV=0:

- 8-byte random byte sequence, called a confounder;
- an MD5 checksum, which is the MD5 hash of the concatenation of the three quantities (the confounder + sixteen NULL octets + the text to be encrypted [not including any padding]);
- AS Reply part that is to be encrypted;
- random padding up to a multiple of 8.

Upon receipt of an AS-REPLY, the client MUST check the validity of the checksum in the encrypted portion of the AS-REPLY.

6.4.2.3 Profile for Kerberos tickets

In Kerberos tickets, authorization-data, starttime and renew-till optional fields MUST NOT be present. The optional caddr field MUST be present when requested in an AS-REQUEST or when present in a TGT of a TGS Request (see IETF RFC 4120). The only ticket flags that are currently supported within IPCablecom are the INITIAL, PRE-AUTHENT and TRANSITED-POLICY-CHECKED flags. If the KDC receives any request that would otherwise cause it to set any other flag, it MUST return an error with the error code KDC_ERR_POLICY. The KDC MUST NOT generate tickets with any other flags set. The session key contained in the ticket (which MUST be identical to the session key in the AS-REPLY) MUST be etype des3-cbc-md5. Since the transited encoding information normally required by PKINIT is not used in IPCablecom, a KDC MAY choose to leave as a null string the 'contents' field of the TransitedEncoding portion of a ticket issued in response to a PKINIT request.

The encrypted part of the Kerberos ticket MUST be encrypted with the encryption type set to des3-cbc-md5, using the same procedure as described in 6.4.2.2.

Upon receipt of a ticket for a service, the server MUST:

- 1) check the validity of the checksum in the encrypted portion of the ticket;
- 2) check that the ticket has not expired.

Currently, all the service keys are pre-shared using an out-of-band mechanism between the KDC and the device providing the service. In the future, IPCablecom may support a method that does not require these keys to be pre-shared.

6.4.3 Symmetric Key AS Request/AS Reply exchange

In IPCablecom, a Kerberos client MAY use standard symmetric-key authentication (with a client key) during the AS Request/AS Reply exchange. Also, in IPCablecom, a client not utilizing PKINIT is, at the same time, an Application Server for which other clients might obtain tickets. This means that an IPCablecom entity may utilize the same symmetric key for both client authentication and for decrypting its service tickets.

The Kerberos AS Request/AS Reply exchange, in general, is allowed to occur with no client authentication. The client, in those cases, would authenticate itself later by proving that it is able to decrypt the AS Reply with its symmetric key and make use of the session key.

Such use of Kerberos is not acceptable within IPCablecom. This approach would allow a rogue client to continuously generate AS Requests on behalf of other clients and receive the corresponding AS Replies. Although this rogue client would be unable to decrypt each AS Reply, it will know some of the fields that it should contain. This, and the availability of the matching encrypted AS Replies, would aid an attacker in the discovery of another client's key with cryptanalysis.

Therefore, IPCablecom requires that whenever an AS Request is not using a PKINIT pre-authenticator, it MUST instead use a different pre-authenticator, of type PA-ENC-TS-ENC. This pre-authenticator is specified as:

```
PA-ENC-TS-ENC ::= SEQUENCE {
    patimestamp      [0] KerberosTime,
                    -- client's time
    pausec           [1] INTEGER OPTIONAL
    pachecksum       [2] CheckSum OPTIONAL
                    -- keyed checksum of
                    -- KDC-REQ-BODY
}
```

The PA-ENC-TS-ENC pre-authenticator MUST be encrypted with the client key using the encryption type des3-cbc-md5, as described in 6.4.2.2. All optional fields inside PA-ENC-TS-ENC MUST be present for IPCablecom. The pachecksum field MUST be a keyed checksum of type des3-cbc-md5 and MUST be validated by the KDC. The encrypted timestamp is used by the KDC to authenticate the client. At the same time, the timestamp inside this pre-authenticator is used to prevent replays. The KDC checks for replays upon the receipt of this pre-authenticator; this is similar to the checking performed by an Application Server upon receipt of an AP Request message.

If the timestamp in the PA-ENC-TS-ENC pre-authenticator differs from the current KDC time by more than `pktcKdcToMtaMaxClockSkew` then KDC MUST reply with a clock skew error message and the client MUST respond to this error message as specified in 6.4.2.1.2.1.1.

If the realm, target server name (e.g., the name of the KDC), along with the client name, time and microsecond fields from the PA-ENC-TS-ENC pre-authenticator match any recently-seen such tuples, the `KRB_AP_ERR_REPEAT` error MUST be returned. The KDC MUST remember any such pre-authenticator presented within acceptable clock skew period, so that a replay attempt is guaranteed to fail.

If the Application Server loses track of any authenticator presented within acceptable clock skew period, it MUST reject all requests until the acceptable clock skew interval has passed.

Symmetric-key AS Request/AS Reply exchange is illustrated in Figure 6.

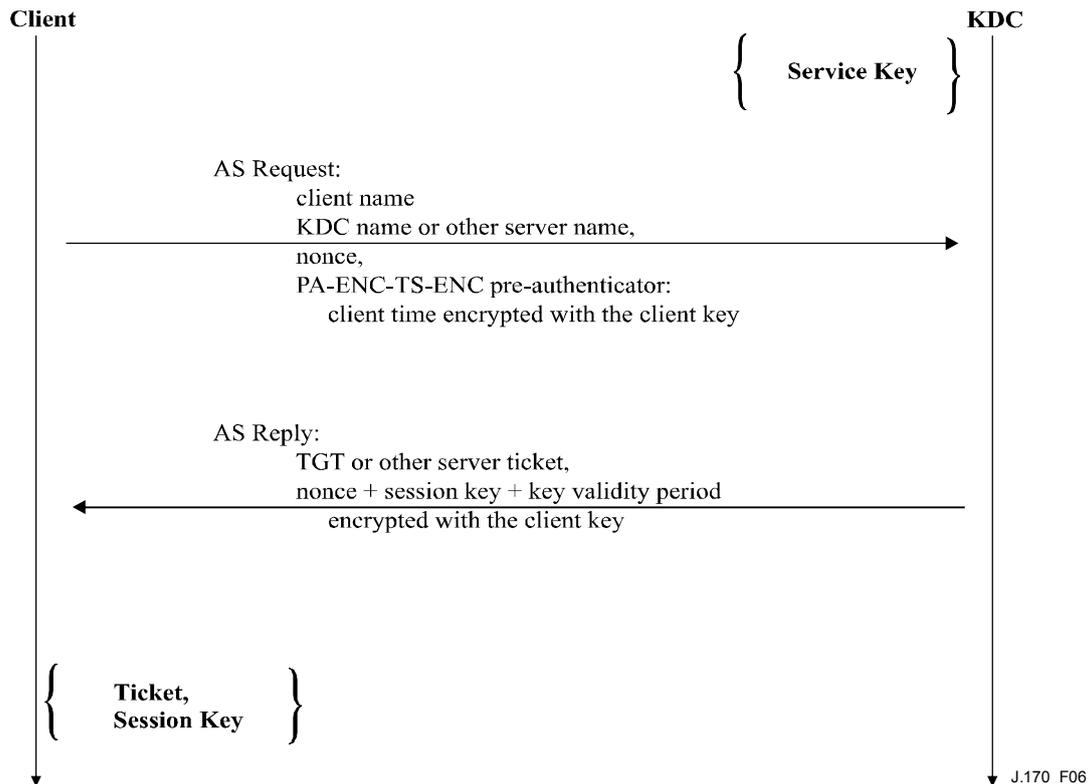


Figure 6 – Symmetric-key AS Request/AS Reply exchange

6.4.3.1 Profile for the Symmetric-key AS Request/AS Reply exchanges

The content of the AS Request/AS Reply messages is the same as in the case of the PKINIT pre-authentication (see 6.4.2.2) with the exception of the type of the pre-authenticator that is used.

In general, clients using a symmetric-key form of the AS Request/AS Reply exchange are not required to always possess a valid TGT or a valid Application Server ticket. A client MAY obtain both a TGT and Application Server tickets on demand, as they are needed for the key management with the Application Server.

However, there may be cases where a client is required to quickly switch between servers for load balancing and the additional symmetric-key exchanges with the KDC are undesirable. In those cases, a client MAY be optimized to obtain tickets in advance, so that the key management would take only a single round trip (AP Request/AP Reply exchange).

In the case that the KDC rejects the AS Request, it returns a KRB_ERROR message instead of the AS Reply, as specified in IETF RFC 4120. The KRB_ERROR MUST use typed-data of REQ-NONCE to bind the error message to the nonce from the AS-REQ message. This error message MUST include the optional e-cksum member that would contain an rsa-md5-des3 keyed checksum of the error reply, unless pre-authentication failed to prove knowledge of the shared symmetric key in which case the e-cksum MUST NOT be used.

The rsa-md5-des3 checksum MUST be computed as follows:

- 1) Prepend the message with an 8-byte random byte sequence, called a confounder.
- 2) Take an MD5 hash of the result of step 1.
- 3) Prepend the hash with the same 8-byte confounder.
- 4) Take the 3DES session key from the ticket and XOR each byte with F0.

- 5) Use 3DES in CBC mode to encrypt the result of step 3, using the key in step 4 and with IV (initialization vector) = 0.

Once a client receives an AS Reply, it SHOULD save both the obtained ticket and the session key information (found in the enc-part member of the reply) in non-volatile memory. Thus, the client will be able to reuse the same Kerberos ticket after a reboot, avoiding the need to perform the AS Request again.

Kerberos tickets MUST NOT be issued for a period of time that is longer than 7 days (same as for PKINIT exchanges).

Upon receipt of a KRB_ERROR that contains an e-cksum field, the recipient MUST verify the validity of the checksum.

6.4.4 Kerberos TGS Request/TGS Reply exchange

In the cases where a client obtained a TGT, that TGT is then used in the TGS Request/TGS Reply exchange to obtain a specific Application Server ticket. This is part of the Kerberos standard, as specified in IETF RFC 4120.

A TGS Request includes a KRB_AP_REQ data structure (the same structure used in an AP Request: see 6.4.4.1). This data structure contains the TGT as well as an authenticator that is used by the client to prove the possession of the corresponding session key. The TGS Reply has the same format as an AS Reply, except that it is encrypted using a different key – the session key from the TGT.

Figure 7 illustrates the TGS Request/TGS Reply exchange:

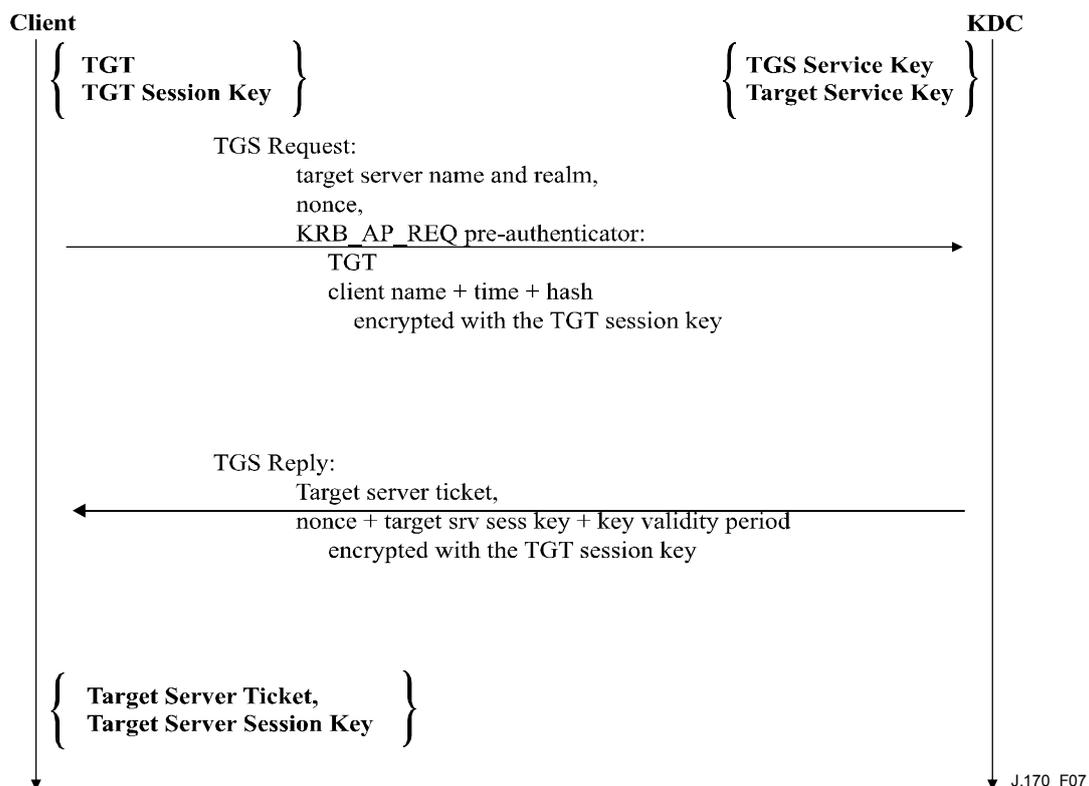


Figure 7 – Kerberos TGS Request/TGS Reply exchange

Figure 7 lists several important parameters in the TGS Request and Reply messages. These parameters are:

TGS Request

- Target server (principal) name and realm, nonce – found in the KDC-REQ-BODY Kerberos structure (see IETF RFC 4120).
- TGS pre-authenticator – found in the KDC-REQ Kerberos structure, inside the padata field (see IETF RFC 4120). The pre-authenticator type in this case is PA-TGS-REQ.
- KRB_AP_REQ – the value of the pre-authenticator of type PA-TGS-REQ.
- TGT – inside the KRB_AP_REQ.
- Client name, time – inside the Kerberos Authenticator structure, which is embedded in an encrypted form in the KRB_AP_REQ.

TGS Reply

- Target server ticket – found in the KDC-REP Kerberos structure (see IETF RFC 4120).
- Target server session key, nonce, key validity period – found in the EncKDCRepPart Kerberos structure (see IETF RFC 4120).

In general, the TGS Request/Reply exchange may be performed on-demand – whenever an Application Server ticket is needed to establish Security Parameters. If the client is an MTA and a ticket it currently possesses corresponds to the Provisioning Server in the MIB or a CMS that currently exists in the CMS table, it MUST initiate the TGS Request/Reply exchange at the time: $\text{Ticket}_{\text{EXP}} - \text{TGS}_{\text{GP}}$. Here, $\text{Ticket}_{\text{EXP}}$ is the expiration time of the current Application Server ticket and TGS_{GP} is the TGS Grace Period. If the client is an MTA and the ticket it currently possesses does not correspond to the Provisioning Server in the MIB or a CMS that currently exists in the CMS table, the MTA MUST NOT initiate a PKINIT exchange.

The validity of the Application Server tickets MUST NOT extend beyond the expiration time of the TGT that was used to obtain the server ticket.

6.4.4.1 TGS Request profile

The optional padata element in the KDC-REQ data structure MUST consist of exactly one element – a pre-authenticator of type PA-TGS-REQ. The value of this pre-authenticator is the KRB_AP_REQ data structure. Within KRB_AP_REQ:

- 1) options in the ap-options field MUST NOT be present;
- 2) the ticket is the TGT;
- 3) the encrypted authenticator MUST contain the checksum field – an MD5 checksum of the ASN.1 encoding of the KDC-REQ-BODY data structure. It MUST NOT contain any other optional fields;
- 4) the authenticator MUST be encrypted using 3-DES CBC with the following Kerberos etype: value des3-cbc-md5 as specified in 6.4.2.2.

The optional fields from enc-authorization-data, additional-tickets and rtime in the KDC-REQ-BODY MUST NOT be present in the TGS Request. The optional field cname SHOULD NOT be present. All other optional fields in the TGS Request MAY be present for IPCablecom. The KDC MUST reject a TGT that has any ticket flags set, apart from the flags INITIAL, PRE-AUTHENT or TRANSITED-POLICY-CHECKED. If the KDC receives any request that would otherwise cause it to set any flag in the service ticket, apart from the PRE-AUTHENT and TRANSIT-POLICY-CHECKED flags, it MUST return an error with the error code KDC_ERR_POLICY. The KDC MUST NOT generate TGT-based service tickets with any other flags set.

If the TGT contains a caddr field, the KDC MUST verify that it is a single IP address and that it is identical to the IP address in the IP header of the TGS Request. The KDC MUST reject TGS Requests from an MTA with a TGT that does not include the MTA's IP address, returning a KDC_ERR_POLICY error code (refer to 6.4.4.3)

Upon receipt of a TGS Request, the KDC MUST:

- 1) check the validity of the TGT;
- 2) check the validity of the checksum in the authenticator.

6.4.4.2 TGS Reply profile

In the TGS Reply, key-expiration, starttime, and renew-till optional fields MUST NOT be present. The encrypted part of the TGS Reply MUST be encrypted with the encryption type set to des3-cbc-md5, using the same procedure as described in 6.4.2.2.

Upon receipt of a TGS Reply, the client MUST:

- 1) use the value of the nonce to bind the reply to the corresponding TGS Request;
- 2) check the validity of the checksum in the encrypted portion of the TGS Reply.

6.4.4.3 Error reply

If the KDC is able to successfully parse the TGS Request and the TGT that is inside of it, but the TGS Request is rejected, it MUST return a Kerberos error message of type KRB_ERROR, as defined in IETF RFC 4120. The error message MUST include the optional e-cksum member, which is the keyed hash over the KRB_ERROR message. The checksum type MUST be rsa-md5-des3, calculated using the procedure described in 6.4.3.1.

The KRB_ERROR MUST also include typed-data of REQ-NONCE to bind the error message to the nonce from the TGS-REQ message.

Upon receipt of a KRB_ERROR, the client MUST check the validity of the checksum.

6.4.5 Kerberos server locations and naming conventions

6.4.5.1 Kerberos realms

A realm name MAY use the same syntax as a domain name, however Kerberos Realms MUST be in all capitals. For a full specification of Kerberos realms, refer to IETF RFC 4120.

6.4.5.2 KDC

Kerberos principal identifier for the local KDC when it is in a role of issuing tickets is always: krbtgt/<realm>@<realm>, where <realm> is the Kerberos realm corresponding to the particular IPCablecom zone. This is the service name listed inside a TGT.

A Kerberos client MUST query KDC FQDNs for a particular realm name using DNS SRV records, as specified in RFC 2782 and as shown below:

```
<Service Name>.<Protocol>.<Name>  TTL Class  SRV  Priority  Weight  Port  Target
```

where:

- the Service Name for Kerberos in IPCablecom MUST be "_kerberos";
- the Protocol for Kerberos in IPCablecom MUST be "_udp";
- the Name MUST be the Kerberos realm name that this record corresponds to;
- TTL, Class, SRV, Priority, Weight, Port, and Target have the standard meaning as defined in RFC 2782.

For example, assume the presence of a realm, PACKETCABLE.COM, with two KDCs: kdc1.packetcable.com and kdc2.packetcable.com. These KDCs have different priorities. The DNS SRV records in this case would be:

```
_kerberos._udp.PACKETCABLE.COM.      86400      IN      SRV      0      0      88
kdc1.packetcable.com.
_kerberos._udp.PACKETCABLE.COM.      86400      IN      SRV      1      0      88
kdc2.packetcable.com.
```

To obtain records pertaining to the realm PACKETCABLE.COM, the MTA would send a DNS SRV request for:

```
kerberos._udp.PACKETCABLE.COM
```

The client, upon receiving a response for a DNS SRV request, MUST consider the priority/weight as described in the algorithm in RFC 2782 and contact the servers in that order. A client MUST contact the next server based on priority/weight and so on, till all possible server FQDNs and the corresponding IPs are exhausted, if it fails to get a suitable response from the first server listed (refer to 6.4.8 for timeout procedures).

For example, after the above DNS SRV records are retrieved, the client will try kdc1.packetcable.com first, based on its priority. (Priority for kdc1.packetcable.com is 0, while priority for kdc2.packetcable.com is 1: a lower priority number means a higher priority.)

When an IPcablecom KDC is requesting information from a Provisioning Server (e.g., the mapping of an MTA MAC address to its corresponding FQDN) it MUST use a principal name of type NT-PRINCIPAL (1) with a single component "kdcquery" (without quotes).

In an ASCII representation, the principal identifier is as follows:

```
kdcquery@<realm>
```

where <realm> is the Kerberos realm of the KDC.

6.4.5.3 CMS

A CMS Kerberos principal identifier MUST be constructed from the CMS FQDN as follows:

```
cms/<FQDN>@<realm>
```

where <FQDN> is the CMS's FQDN (in lower case) and <realm> is its Kerberos realm.

For example, a CMS with an FQDN 'iptel-cms1.company1.com' and with a realm name 'COMPANY1.COM' would have the principal identifier:

```
'cms/iptel-cms1.company1.com@COMPANY1.COM'
```

The Kerberos PrincipalName data structure (inside the Kerberos messages) is defined as follows:

```
PrincipalName ::= SEQUENCE {
    name-type      [0]  INTEGER,
    name-string    [1]  SEQUENCE OF GeneralString
}
```

Within this data structure, name-type MUST be NT-SRV-HST (which has the value of 3 according to the Kerberos specification). The name-string element of the data structure MUST have exactly two components, where the first component has the string value "cms" (without the quotes) and the second component is the CMS's FQDN in lower case.

For the full syntax of Kerberos principal names, refer to IETF RFC 4120.

For the purpose of setting up an IPsec connection between the CMS and RKS, the first component of the CMS principal name MUST be of the form "cms:<ElementID", where the <ElementID> is described in 6.4.5.5.

In the case of a combined network element that integrates the functions of multiple logical elements within the PacketCable reference architecture (e.g., a single network element that provides both CMS and MGC functionality), the principal name may include all server functions as specified in 6.4.5.5.

6.4.5.4 Provisioning Server

When an IPCablecom MTA Provisioning Server is acting in the role of an SNMP manager, it MUST use a principal name of type NT-SRV-HST (3) with the following two components:

- 1) "mtaprovsrvr" (without quotes);
- 2) the FQDN of the Provisioning Server (in lower case).

In ASCII representation, the Provisioning Server's principal identifier MUST be as follows:

```
mtaprovsrvr /<Prov Server FQDN>@<realm>
```

where <realm> is the Kerberos realm of the Provisioning Server.

When an IPCablecom Provisioning Server is providing a service (to the KDC) that maps each MTA MAC address to its corresponding FQDN, it MUST use a principal name of type NT-SRV-HST (3) with the following two components:

- 1) "mtafqdnmap" (without quotes);
- 2) the FQDN of the Provisioning Server (in lower case).

In ASCII representation, the principal identifier MUST be as follows:

```
mtafqdnmap/<Prov Server FQDN>@<realm>
```

where <realm> is the Kerberos realm of the Provisioning Server.

6.4.5.5 Names of other Kerberized services

All Kerberized services within IPCablecom except for the KDC krbtgt service (see 6.4.5.2), MUST be assigned a service principal name of type KRB_NT_SRV_HST (Value=3), which has the following form according to the Kerberos specification: <service name>/<FQDN>.

This means that the first component of the service principal name is the service name in lower case, and the last is either an FQDN in lower case or an IP address of the corresponding host. If a specific host has an assigned FQDN, its principal name includes an FQDN and not an IP address. When a KDC receives a ticket request for a service on this host with an IP address instead of an FQDN as the second component of the service principal name, the KDC MUST reject such a request.

When a KDC database contains a service with a principal name that has an IP address as the second component, all ticket requests for this service MUST use the same service principal name with the same IP address as the second component. When a KDC receives a ticket request for this service with an FQDN as the second component of the service principal name, the KDC MUST reject such a request. (This scenario could happen if a service principal is defined in the KDC database at the time when the corresponding host does not have an FQDN, and then later an FQDN for this host is defined as well.)

When an IP address is used, it MUST be formatted as follows:

```
[A.B.C.D]
```

where A, B, C and D are components of an IPv4 address expressed as decimal numbers. The components of an IP address MUST be separated by a period '.' and the IP address MUST be surrounded by square brackets.

The following is an example of a principal name based on an IP address:

```
df/[192.35.65.4]
```

Figure 3 shows a number of interfaces for which the necessary security is provided by IPsec. In addition to supplying the required key management using IKE with pre-shared keys, some vendors may choose to implement, and operators to deploy, a Kerberized key management scheme for these interfaces.

This Recommendation requires that the RKS verifies billing event messages by ensuring that the Element ID contained in the message matches correctly the IP address at the far end of the IPsec Security Associations. In order to ensure that the RKS is able to maintain this mapping when Kerberized key management is used to generate the Security Associations, devices that communicate with the RKS include their Element ID in their principal name. This information is then passed to the RKS in the cname field of the ticket that the KDC issues; this ticket is passed to the RKS in the AP-REQ that is used to initiate the IPsec Security Associations.

The first component of the principal name for the various IPCablecom devices **MUST** be as follows:

- 1) BP: bp[:<ElementID>]
- 2) CMTS: cmts[:<ElementID>]
- 3) DF: df[:<ElementID>]
- 4) MG: mg[:<ElementID>]
- 5) MGC: mgc[:<ElementID>]
- 6) MP: mp[:<ElementID>]
- 7) MPC: mpc[:<ElementID>]
- 8) RKS: rks[:<ElementID>]
- 9) SG: sg[:<ElementID>]

where:

<ElementID> is the identifier that appears in billing event messages and it **MUST** be included in a principal name of every server that is capable of generating event messages.

Element ID is defined as an 5-octet right-justified, space-padded ASCII-encoded numerical string (J.164). When converting the Element ID for use in a principal name, any spaces **MUST** be converted to ASCII zeros (0x48).

For example, a CMTS that has the Element ID "311" will have a principal name whose first component is "cmts:00311". Similarly, a DF with no Element ID will have a principal name whose first component is "df".

Components that contain combined elements (such as a CMS with an integrated MGC) **MUST** indicate this in the principal name by including all component names, joined with the character "&", in the first component of the principal name. The following is an example of a principal name for a combined CMS and MGC with a single IP address:

```
cms:00210&mgc:00211/[192.35.65.4]
```

If the combined component uses a single ElementID, the principal name would be:

```
cms:00210&mgc:00210/[192.35.65.4]
```

6.4.6 MTA principal names

An MTA principal name **MUST** be of type NT-SRV-HST with exactly two components, where the first component **MUST** be the string "mta" (not including the quotes) and the second component **MUST** be the FQDN of the MTA: mta /<MTA FQDN>, where <MTA FQDN> is the FQDN of the MTA in lower case.

For example, if an MTA FQDN is "mta12345.mso1.com" and its realm is "MSO1.COM", the principal identifier would be: mta/mta12345.mso1.com@MSO1.COM.

6.4.7 Mapping of MTA MAC address to MTA FQDN

The MTA authenticates itself with the MTA Device Certificate in the AS Request, where the certificate contains the MTA MAC address but not its FQDN. In order to authenticate the MTA principal name (containing the FQDN), the KDC MUST map the MTA MAC address (from the MTA Device certificate) to the MTA FQDN, in order to verify the principal name in the AS Request.

The protocol for retrieving the MTA FQDNs is Kerberos-based. The Provisioning Server MUST listen for the request on UDP port 2246 and MUST return the response to the UDP port from which the request was transmitted on the client:

- 1) MTA FQDN Request – sent from the KDC to the Provisioning Server, containing the MTA MAC address and the hash of the MTA public key. This message consists of the Kerberos KRB_AP_REQ concatenated with KRB_SAFE.
- 2) MTA FQDN Reply – a reply to the KDC by the Provisioning Server, containing the MTA FQDN. This message consists of the Kerberos KRB_AP_REP concatenated with KRB_SAFE.
- 3) MTA FQDN Error Reply – an error reply in response to the MTA FQDN Request. This message is the Kerberos KRB_ERROR.

The format of each of these messages is specified in the clauses below.

6.4.7.1 MTA FQDN Request

The KDC MUST first verify the digital signature and certificate chain in the PKINIT Request, before sending out an MTA FQDN Request message to determine the MTA MAC address to FQDN mapping.

In the case where the PKINIT Request and certificate signatures are all valid but the manufacturer certificate is revoked, the KDC MAY still proceed with the MTA FQDN Request. In this case, the KDC MUST provide the revocation time in the MTA FQDN Request.

The MTA FQDN Request MUST be formatted as in Table 6.

Table 6 – MTA FQDN Request format

| Field name | Length | Description |
|------------|----------|---|
| KRB_AP_REQ | Variable | DER-encoded. The length is in the ASN.1 header. |
| KRB_SAFE | Variable | DER-encoded |

In the KRB_AP_REQ, only the following option is supported:

- MUTUAL-REQUIRED – mutual authentication required. This option MUST always be set.
- All other options are not supported.

The encrypted authenticator in the KRB_AP_REQ MUST contain the following field, which is optional in Kerberos:

- seq-number: random value generated by the KDC

All other optional fields within the encrypted authenticator are not supported within IPCablecom. In 6.5.2.2 there is a requirement that the recipient of a KRB_AP_REQ accepts certain negative values of seq-number; that requirement does not apply when processing the KRB_AP_REQ embedded in a received MTA FQDN message. The authenticator itself MUST be encrypted using

3-DES CBC with the Kerberos etype value des3-cbc-md5 with the session key from the ticket that is contained in this KRB_AP_REQ object. The encryption method for des3-cbc-md5 is specified in 6.4.2.2.

KRB_SAFE MUST contain the following field, which is optional in Kerberos:

- seq-number: same value as in the KRB_AP_REQ, to tie KRB_SAFE to KRB_AP_REQ and avoid replay attacks.

All other optional fields within KRB_SAFE are not supported within IPCablecom. The keyed checksum within KRB_SAFE MUST be of type rsa-md5-des3 and MUST be computed with the session key in the accompanying KRB_AP_REQ. The method for computing an rsa-md5 des3 keyed checksum is specified in 6.4.3.1.

The data that is wrapped inside KRB_SAFE MUST be formatted as in Table 7:

Table 7 – KRB_SAFE format

| Field name | Length | Description |
|-----------------------------------|----------|---|
| Message Type | 1 byte | 1 = MTA FQDN Request |
| Enterprise Number | 4 bytes | Network byte order, MSB first. 1 = IPCablecom |
| Protocol Version | 1 byte | 2 for this version |
| MTA MAC Address | 6 bytes | MTA MAC Address |
| MTA Pub Key Hash | 20 bytes | SHA-1 hash of DER-encoded SubjectPublicKeyInfo. |
| Manufacturer Cert Revocation Time | 4 bytes | 0 = MTA Manufacturer cert not revoked Otherwise, this is UTC time, number of seconds since midnight of Jan 1, 1970, in network byte order. |

Once the KDC has sent an MTA FQDN Request, it MUST save the nonce value that was contained in the seq-number field in order to validate a matching MTA FQDN Reply.

If the KDC times out before getting a reply it MUST give up and simply drop the PKINIT request with no error code returned. The KDC MUST NOT retry in this case, since it would still have to handle retries of PKINIT Request from the MTA. At the same time, after a time out the KDC SHOULD increase its time out value on the next request to the same Provisioning Server using an exponential back-off algorithm.

The Provisioning Server receiving this message MUST validate the KRB_AP_REQ and verify that it is not a replay using the procedure specified in the Kerberos standard (see IETF RFC 4120), also described in 6.5.2. After the KRB_AP_REQ has been validated, the Provisioning Server MUST also verify the KRB_SAFE component: that the checksum keyed with the session key is valid and that the seq-number field matches the KRB_AP_REQ.

If the Manufacturer Cert Revocation Time field is 0 and the Provisioning Server supports the storage of MTA public key hashes, then it MUST update the MTA public key hash in its database. If the public key hash has changed or is saved for the first time, the Provisioning Server MUST also record the time this update (to the MTA public key hash) is performed.

- If the Manufacturer Cert Revocation Time field is non-zero, the Provisioning Server MUST validate that the public key hash has not changed from the previous update and that the revocation time is after the last update to the MTA public key hash. If not, the error code KRB_MTAMA_ERR_PUBKE_NOT_TRUSTED MUST be returned. If the Provisioning Server does not support storage of MTA public key hashes and the Manufacturer Cert Revocation Time field is non-zero, the same error code MUST be returned.

6.4.7.2 MTA FQDN reply

The MTA FQDN Reply MUST be formatted as in Table 8:

Table 8 – MTA FQDN format

| Field name | Length | Description |
|------------|----------|---|
| KRB_AP_REP | Variable | DER-encoded. The length is in the ASN.1 header. |
| KRB_SAFE | Variable | DER-encoded |

The encrypted part of the KRB_AP_REP MUST contain the following field, which is optional in Kerberos:

- seq-number: echoes the value in the KRB_AP_REQ

All other optional fields within the encrypted part of the KRB_AP_REP are not supported within IPCablecom. It MUST be encrypted using 3-DES CBC with the Kerberos etype value des3-cbc-md5 and MUST be computed with the session key from the preceding KRB_AP_REQ. The encryption method for des3-cbc-md5 is specified in 6.4.2.2.

KRB_SAFE MUST contain the following field, which is optional in Kerberos:

- seq-number: same value as in the KRB_AP_REP, to tie KRB_SAFE to KRB_AP_REP and avoid replay attacks.

All other optional fields within KRB_SAFE are not supported within IPCablecom. The keyed checksum within KRB_SAFE MUST be of type rsa-md5-des3 and MUST be computed with the session key from the preceding KRB_AP_REQ. The method for computing an rsa-md5-des3 keyed checksum is specified in 6.4.3.1.

The data that is wrapped inside KRB_SAFE MUST be formatted as in Table 9:

Table 9 – KRB_SAFE data format

| Field name | Length | Description |
|-------------------|----------|---|
| Message Type | 1 byte | 2 = MTA FQDN Reply |
| Enterprise Number | 4 bytes | Network byte order, MSB first 1 = IPCablecom |
| Protocol Version | 1 byte | 2 for this version |
| MTA FQDN | variable | MTA FQDN |
| MTA IP Address | 4 bytes | MTA-IP Address (MSB first) |

After the KDC receives this reply message, it MUST validate the integrity of both the KRB_AP_REP and KRB_SAFE objects (see IETF RFC 4120) and MUST also verify that the value of the seq-number field is the same for both. If this integrity check fails, the KDC MUST immediately discard the reply and proceed as if the message had never been received (e.g., if the KDC was waiting for a valid MTA FQDN Reply, it should continue to do so).

The Provisioning Server MAY set the MTA IP Address field of the MTA FQDN Reply to zero. If the KDC receives an MTA FQDN REPLY with a non-zero MTA IP Address field, it MUST compare it to the IP address contained in the AS Request. If this check fails, then the KDC MUST NOT respond to the AS Request.

6.4.7.3 MTA FQDN error

If the Provisioning Server is able to successfully parse the KRB_AP_REQ and the ticket that is inside of it, but the MTA FQDN Request is rejected, it MUST return an error message.

All errors MUST be returned as a KRB_ERROR message, as specified in IETF RFC 4120. It MUST include typed-data of REQ-SEQ to bind the error message to the sequence number from the authenticator in the KRB_AP_REQ. Also, the error message MUST include the optional e-cksum member, which is the keyed hash over the KRB_ERROR message. The checksum type MUST be rsa-md5-des3 and MUST be computed with the session key from the preceding KRB_AP_REQ as specified in 6.4.3.1.

In the case that the client time field inside KRB_AP_REQ differs from the Provisioning Server's clock by more than the maximum allowable clock skew, a clock skew error MUST be handled as specified in 6.5.2.3.2.

If the error is application-specific (not a Kerberos-related error), then KRB_ERROR MUST include typed-data of type TD-APP-DEFINED-ERROR (value 106). The value of this typed-data is specified in IETF RFC 4120 as follows:

```
AppSpecificTypedData ::= SEQUENCE {
    oid                [0]  OPTIONAL OBJECT IDENTIFIER,
                       -- identifies the application
    data-value        [1]  OCTET STRING
                       -- application-specific data
}
```

Inside AppSpecificTypedData, the oid field MUST be set to:

```
enterprises (1.3.6.1.4.1) cableLabs (4491) clabProjects (2) clabProjPacketCable
(2) pktcSecurity (4) errorCodes (1) FQDN (3)
```

The data-value field MUST correspond to the following typed-data value:

```
pktcKrbMtaMappingError ::= SEQUENCE {
    e-code[0] INTEGER,
    e-text[1] GeneralString OPTIONAL,
    e-data[2] OCTET STRING OPTIONAL
}
```

The e-code field MUST correspond to one of the following error code values:

| | | |
|-----------------------------------|---|-----------------------------------|
| KRB_MTAMAP_ERR_NOT_FOUND | 1 | MTA MAC Address not found |
| KRB_MTAMAP_ERR_PUBKEY_NOT_TRUSTED | 2 | MTA public key is not trusted |
| KRB_MTAMAP_VERSION_UNSUP | 3 | Unsupported Version Number |
| KRB_MTAMAP_MSGTYPE_UNKNOWN | 4 | Unrecognized Message Type |
| KRB_MTAMAP_ENTERPRISE_UNKNOWN | 5 | Unrecognized Enterprise Number |
| KRB_MTAMAP_NOT_YET_VALID | 6 | MTA not yet valid |
| KRB_MTAMAP_ERR_GENERIC | 7 | Generic MTA name mapping error |

The optional e-text field can be used for informational purposes (i.e., logging, network troubleshooting) and the optional e-data field is reserved for future use to transport any application data associated with a specific error.

Upon receipt of a KRB_ERROR from the Provisioning Server, the KDC MUST check the validity of the checksum. If the KRB_ERROR passes the validity check, the KDC MUST send a corresponding KRB_ERROR to the MTA (as specified in 6.4.2.1.2), in response to the PKINIT Request. The application specific MAC-FQDN error codes MUST be mapped to Kerberos error codes in the error reply to the MTA according to Table 10.

Table 10 – Mapping of KRB_MTAMAP_ERR to KRB_ERR

| | |
|-----------------------------------|-----------------------------|
| KRB_MTAMAP_ERR_NOT_FOUND | KDC_ERR_C_PRINCIPAL_UNKNOWN |
| KRB_MTAMAP_ERR_PUBKEY_NOT_TRUSTED | KDC_ERR_CLIENT_REVOKED |
| KRB_MTAMAP_VERSION_UNSUP | KRB_ERR_GENERIC |
| KRB_MTAMAP_MSGTYPE_UNKNOWN | KRB_ERR_GENERIC |
| KRB_MTAMAP_ENTERPRISE_UNKNOWN | KRB_ERR_GENERIC |
| KRB_MTAMAP_NOT_YET_VALID | KDC_ERR_CLIENT_NOTYET |
| KRB_MTAMAP_ERR_GENERIC | KRB_ERR_GENERIC |

6.4.8 Server key management time-out procedure

The Kerberos client MUST implement a retransmission strategy using exponential back-off with configurable initial and maximum retransmission timer values for any KDC or application server requests that have not been acknowledged by the server. The Kerberos client MUST update the client timestamp field with the current time-of-day reading for each such retry. During an exponential back-off, when a previous time out value was T_i , then the next time out value, value T_{i+1} , MUST satisfy the following criteria:

$$1.5 \times T_i \leq T_{i+1} \leq 2.5 \times T_i$$

After successfully processing an AS Request or TGS Request and generating a corresponding reply, the KDC MUST save:

- the AS Request or TGS Request (e.g., the full AS Request/TGS Request or a hash of the AS Request/TGS Request);
- the full KDC reply.

The KDC MUST maintain this information for all requests with the client time field that is within the time window $(T - \Delta T_{MAX}, T + \Delta T_{MAX})$, where T is the current time and ΔT_{MAX} is the maximum clock skew that is allowed by the KDC policy.

The KDC MAY also save:

- the client principal identifier;
- the information that uniquely identifies the client pre-authentication field in the AS Request (PKINIT or encrypted timestamp in the case of non-public key AS Request) or TGS Request (PA-TGS-REQ).

The KDC MAY maintain this information for all requests with the client time field that is within the time window $(T - \Delta T_{MAX}, T + \Delta T_{MAX})$, where T is the current time and ΔT_{MAX} is the maximum clock skew that is allowed by the KDC policy. If the AS Request or TGS Request is identical to the one previously received, the KDC MUST respond with the same reply message. If only the principal name and pre-authenticator (PKINIT, encrypted timestamp or PA-TGS-REQ) match, then the KDC MUST perform one of the following:

- If the received AS Request or TGS Request passes all other error checks, the KDC may reply with a cached reply message.
- Reject this message as a replay.

The MTA may have learned several IP addresses for a KDC or application server (refer to 6.4.5.2 for more information on obtaining IP addresses from Realm Names and forming a local list of IP addresses based on prioritization). If the number of retransmissions for a KDC IP address has reached its maximum configured value and there are more IP addresses for the same KDC that have not been tried, then the MTA MUST direct the retransmissions to the remaining alternate addresses in its local list. Each time that the MTA switches to a new KDC IP address for retransmissions, it MUST start a new exponential back-off procedure. If there are no more KDC IP addresses to try,

then the MTA SHOULD actively query the name server in order to detect the possible change of KDC network interfaces, regardless of the Time To Live (TTL) associated with the DNS record to see if any other IP addresses have become available. If there are new IP Addresses discovered, the MTA MUST go through the retransmission strategy again for the newly discovered IP Addresses.

For Kerberized key management with application servers, when an application layer is informed that key management with a particular IP address failed, it is normally up to the application layer to select the next IP address. The switch-over algorithm between multiple IP addresses mapped to the same FQDN is specified by each corresponding application protocol. For example, in the case of the Kerberized key management between the MTA and the CMS, refer to the NCS specification (ITU-T Rec. J.162) There are also cases when key management is performed independent of the application layer, e.g., to pre-establish security associations during MTA initialization. In those cases, it is up to a specific MTA implementation to decide if to failover and how to failover to another application server IP address.

An application server may not respond to application messages (e.g., NCS messages) from the MTA. This may occur if the MTA has valid security parameters with the application server, but the security parameters on the server have been lost or corrupted (e.g., the CMS rebooted and lost all IPsec Security Associations).

In the case of NCS signalling, an MTA MUST no longer use any previously established IPsec SAs with a particular CMS each time the NCS back-off and retry algorithm places an MTA endpoint controlled by that CMS into a DISCONNECTED state. After an MTA endpoint has moved to a DISCONNECT state, it will start sending RSIP/disconnect NCS messages which will need to be protected by newly established IPsec SAs.

6.4.9 Service key versioning

The service key that is shared between a KDC and an application server, to encrypt/decrypt service tickets, is a versioned key (refer to IETF RFC 4120). This key may be changed either due to a routine key refresh, or because it was compromised. When the service key is changed, the application server MUST retain the older key for a period of time that is at least as long as the ticket lifetime used when issuing service tickets (i.e., up to 7 days). In the case of a routine service key change, the application server MUST accept any ticket that is encrypted with an older key that it has retained and is still valid (not compromised). This key versioning on the application server will prevent against many MTAs from suddenly flooding a KDC with PKINIT Requests for new tickets.

If a service key is changed because it has been compromised, the application server MUST flag all older key versions it has retained as invalid and reject any AP Request that contains a ticket that is encrypted with one of these invalid keys. When rejecting the AP Request, the application server MUST respond as specified in IETF RFC 4120 with a KRB_AP_ERR_BADKEYVER error. The application server MUST still decrypt the rejected ticket, using the invalid service key, in order to extract the session key. This session key is needed to securely bind the KRB_ERROR reply message to the AP Request message using a keyed checksum (see 6.5.2.3.1). Note that this step is necessary in order to prevent denial-of-service attacks, which could otherwise occur if the MTA was unable to verify the authenticity of the KRB_ERROR message.

Upon receiving this error reply, the MTA MUST discard the service ticket which is no longer valid and fetch a new one from its KDC.

6.5 Kerberized key management

6.5.1 Overview

This clause specifies how Kerberos tickets are used to perform key management between a client and an Application Server, where a client is able to get a Kerberos ticket for the server but not the other way around.

The same protocol described here applies in a symmetric case – where both sides of a key management interface are able to get a ticket for each other, i.e., each side is both a client and a server. In the symmetric case, only the AP Request and AP Reply messages apply.

The Kerberos session key is used in the AP Request and AP Reply messages that are exchanged in order to re-establish security parameters. The sub-key from the AP-REQ and AP-REP are used to derive all of the secret keys used for both directions. The AP Request and AP Reply messages are small enough to fit into a standard UDP packet, not requiring fragmentation.

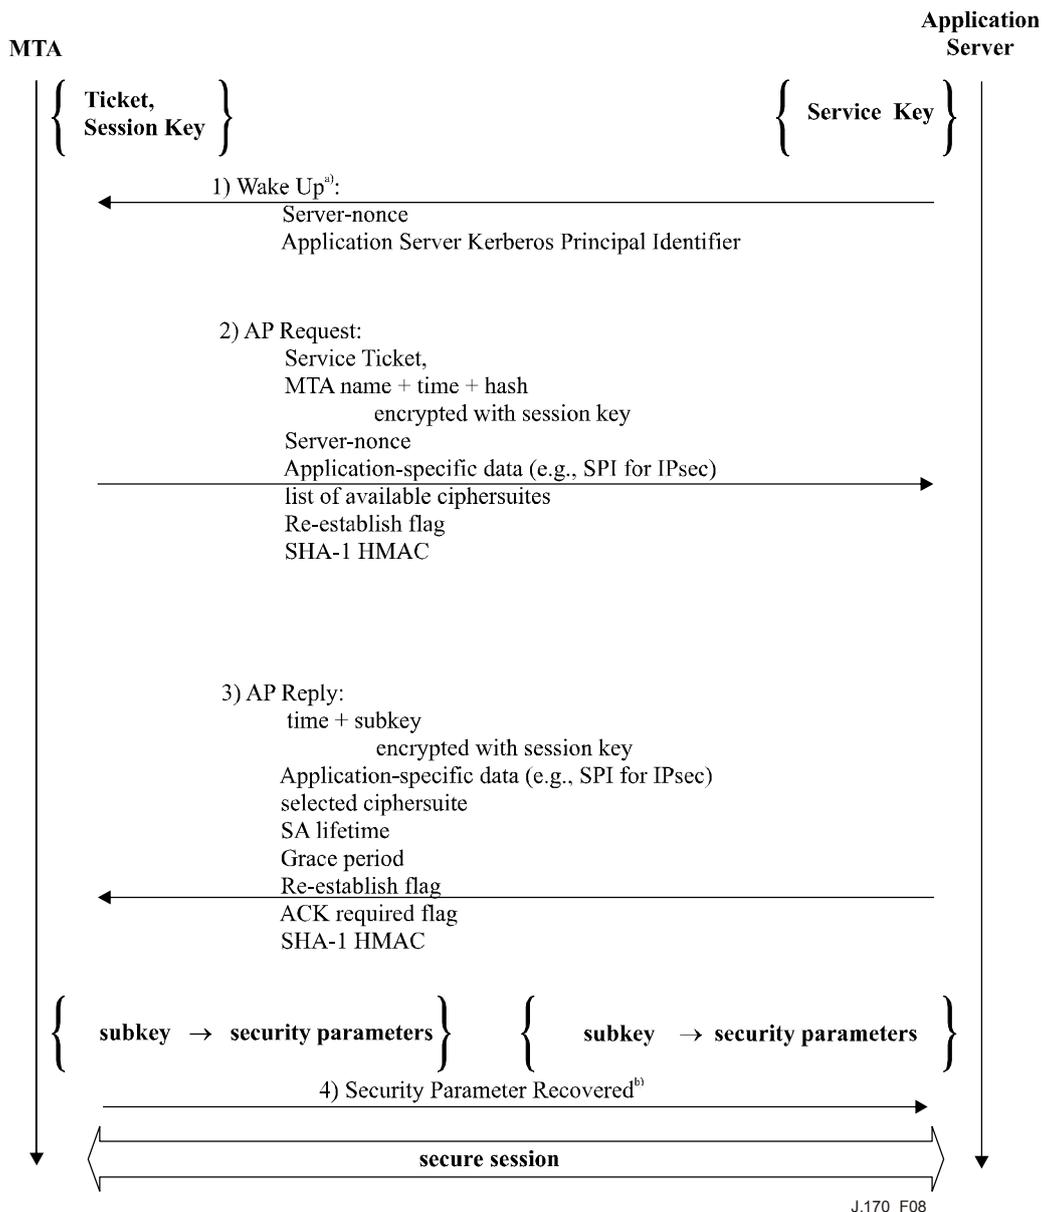
A Kerberos AP Request/Reply exchange MAY occur periodically, to ensure that there are always valid security parameters between the client and the Application Server. It MAY also occur on demand, where the security parameters are allowed to time out and are re-established the next time that application traffic needs to be sent over a secure link.

The UDP port used for all key management messages between the client and the Application Server MUST be 1293 (on both devices).

A recipient of any Kerberized Key Management message that does not fully comply with the IPCablecom requirements MUST reject the message.

6.5.2 Kerberized key management messages

Figure 8 illustrates an AP Request/AP Reply exchange:



- a) This message is sent whenever key management is initiated by an Application Server.
- b) This message is optional, sent whenever the ACK-required flag is set in the preceding AP Reply.

Figure 8 – Kerberos AP Request/AP Reply exchange

- 1) **Wake Up** – An Application Server sends this message when it initiates a new key management exchange.

To prevent denial-of-service attacks, this message includes a Server-nonce field – a random value generated by the Application Server. The client includes the exact value of this Server-nonce in the subsequent AP Request.

This message also contains the Server Kerberos principal identifier, used by the client to find or to obtain a correct Kerberos ticket for that Application Server.

The Wake Up message MUST be formatted as the concatenation of the following fields:

- Key Management Message ID – 1-byte value. Always set to 0x01.
- Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established.

| Domain of Interpretation (DOI) values | |
|--|------------------------|
| Value | Target protocol |
| 1 | IPsec |
| 2 | SNMPv3 |

- Protocol Version – 1 byte. The high-order nibble is the major version number, and the lower-order nibble is the minor version number. For IPCablecom, the major number MUST be 1, and the minor number MUST be 0.
- Server-nonce: a 4-byte random binary string. Its value MUST NOT be all 0s.
- Server Kerberos principal identifier: a printable, null-terminated ASCII string, representing a fully qualified Kerberos Principal identifier of the Application Server, as defined in 6.4.5.

Once the Application Server has sent a Wake Up, it MUST save the Server-nonce. The Application Server MUST keep this nonce in order to validate a matching AP Request. In the case of a time-out, the Application Server MUST adhere to the exponential retry back-off procedure described in 6.4.8. The Application Server MUST begin each retry by re-sending a Wake Up message with a new server-nonce value. When the "Timeout Procedure" has completed without success, the Application Server MUST discard the server-nonce from the last retry, after which it will no longer accept a matching AP Request.

- 2) AP Request – MUST be sent by the client in order to establish a new set of security parameters. Any time the client receives a Wake Up message from a valid application server that is listed as part of client configuration data, it MUST respond with the AP Request message specified below. If a client receives a Wake Up message from an unknown application server, the client MUST NOT respond.

In addition, this Recommendation specifies the use of this message by the client to periodically establish a new set of security parameters with the Application Server – see 6.5.4.2. It also specifies the use of this message by the client to establish a new set of security parameters with the Application Server, when the client somehow loses the security parameters (e.g., after a reboot) – see 6.5.3.5.

The client starts out with a valid Kerberos ticket, previously obtained during a PKINIT exchange. The Application Server starts out with its Service Key that it can use to decrypt and validate Kerberos tickets.

The client sends an AP Request that includes a ticket and an authenticator, encrypted with the session key. The Application Server gets the session key out of the ticket and uses it to decrypt and then validate the authenticator.

The AP Request includes the Kerberos KRB_AP_REQ message along with some additional information, specific to IPCablecom. It MUST consist of the concatenation of the following fields:

- Key Management Message ID – 1-byte value. Always set to 0x02.
- Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established. See the table of DOI values above.

- Protocol Version – 1 byte. The high-order nibble is the major version number, and the lower-order nibble is the minor version number. For IPCablecom, the major number MUST be 1, and the minor number MUST be 0.
- KRB_AP_REQ – DER encoding of the KRB_AP_REQ Kerberos message, as specified in IETF RFC 4120.
- Server-nonce – a 4-byte random binary string. If this AP Request is in response to a Wake Up, then the value MUST be identical to that of the Server-nonce field in the Wake Up message. If this AP Request is in response to a Rekey (see 6.5.2.1), then the value MUST be identical to that of the Server-nonce field in the Rekey message. Otherwise, the value MUST be all 0s.
- Application-specific data – additional information that must be communicated by the client to the server, dependent on the target protocol for which security is being established (e.g., IPsec or SNMPv3).
- List of ciphersuites available at the client:
Number of entries in this list (1 byte)
Each entry has the following format:

| | |
|--|---|
| Authentication Algorithm (1 byte) | Encryption Transform ID (1 byte) |
|--|---|

The actual values of the authentication algorithms and encryption transform IDs are dependent on the target protocol.

- Re-establish flag – a 1-byte Boolean value. When the value is TRUE (1), the client is making an attempt to automatically establish a new set of Security Parameters before the old one expires. Otherwise, the value is FALSE (0).
- SHA-1 HMAC (20 bytes) over the contents of this message, not including this field. The 20-byte key for this HMAC is determined by taking a SHA-1 hash of the session key.
- Whenever the AP Request is received (by the Application Server), it MUST verify the value of this HMAC. If this integrity check fails, the Application Server MUST immediately discard the AP Request and proceed as if the message had never been received (e.g., if the Application Server was waiting for a valid AP Request, it should continue to do so).

Once the client has sent an AP Request, it MUST save the nonce value that was contained in the seq-number field (a different nonce from the server-nonce specified above) along with the server Kerberos principal identifier in order to validate a matching AP Reply. If the client generated this AP Request on its own, it MUST adhere to the exponential retry backoff procedure described in 6.4.8.

If the AP Request was generated in response to a message sent by the Application Server (Wake Up or Rekey), then the client MUST save the nonce and Server Kerberos Principal Identifier until the time specified by the appropriate Key Management MIB variables (pktcMtaDevProvSolicitedKeyTimeout for Prov Server, pktcMtaDevCmsSolicitedKeyTimeout for CMS). After the timeout has been exceeded or when the "Timeout Procedure" has completed without success, the client MUST discard this (nonce, server Kerberos principal identifier) pair, after which it will no longer accept a matching AP Reply.

If the MTA generated an AP Request on its own and has reached the maximum number of retries with a particular application server IP address failing to get an AP Reply, it must retry with alternate application server IP addresses as specified in 6.4.8.

In the case that the Server-nonce is 0 (not filled in), and the Application Server is currently waiting for a reply to a Wake Up or Rekey message from a client at this IP address, it MUST reject the AP Request and not reply to the client. If the Application Server is not waiting for a reply to a Wake Up or Rekey message, it MUST verify that this AP Request is not a replay using the procedure specified in the Kerberos standard (IETF RFC 4120):

- If the timestamp in the AP Request differs from the current Application Server time by more than the acceptable clock skew then the Application Server MUST reply with an error message specified in 6.5.2.3.2.
- If the realm, Application Server name, along with the client name, time and microsecond fields from the Kerberos Authenticator (in the AP Request) match any recently-seen such tuples, the KRB_AP_ERR_REPEAT error MAY be returned.

The Application Server MUST remember any authenticator presented within acceptable clock skew, so that a replay attempt is guaranteed to fail.

- If the Application Server loses track of any authenticator presented within `pktsrvrToMtaMaxClockSkew`, it MUST reject all requests until the clock skew interval has passed.

In the case that the Server-nonce is not 0, the Application Server MAY follow the above procedure in order to fully conform with the Kerberos specification (IETF RFC 4120). In this case, the above procedure is not required because matching the Server-nonce in the Wake Up or Rekey message against the Server-nonce in the AP Request also prevents replays.

3) AP Reply – Sent by the Application Server in response to AP Request.

The AP Reply MUST include a randomly generated sub-key (inside the Kerberos KRB_AP_REP message), encrypted with the same session key.

The AP Reply includes the Kerberos KRB_AP_REP message along with some additional information, specific to IPCablecom. It MUST consist of the concatenation of the following fields:

- Key Management Message ID – 1-byte value. Always set to 0x03.
- Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established. See the table of DOI values in this clause.
- Protocol Version – 1 byte. The high-order nibble is the major version number, and the lower-order nibble is the minor version number. For IPCablecom, the major number MUST be 1, and the minor number MUST be 0.
- KRB_AP_REP – DER encoding of the KRB_AP_REP Kerberos message, as specified in IETF RFC 4120.
- Application-specific data – additional information that must be communicated by the server to the client, dependent on the target protocol for which security is being established (e.g., IPsec or SNMPv3).
- Selected ciphersuite for the target protocol, using the same format as defined for AP Request. The number of entries in the list MUST be one.
- Security parameters lifetime – a 4-byte value, MSB first, indicating the number of seconds from now, when these security parameters are due to expire.
- Grace period – a 4-byte value in seconds, MSB first. This indicates to the client to start creating a new set of security parameters (with a new AP Request/AP Reply exchange) when the timer gets to within this period of their expiration time.

- Re-establish flag – a 1-byte Boolean value. When the value is TRUE (1), a new set of security parameters MUST be established before the old one expires. When the value is FALSE (0), the old set of security parameters MUST be allowed to expire.
- ACK-required flag – a 1-byte Boolean value. When the value is TRUE (1), the AP Reply message requires an acknowledgement, in the form of the Security Parameter Recovered message.
- SHA-1 HMAC (20 bytes) over the contents of this message, not including this field. The 20-byte key for this HMAC is determined by taking a SHA-1 hash of the session key.

Whenever the AP Reply is received (by the client), it MUST:

- verify the value of this HMAC. If this integrity check fails, the client MUST immediately discard the AP Reply;
- verify that the AP Reply Source IP Address matches the AP Request Destination IP Address in the list of outstanding AP Requests. The Client MUST immediately discard the AP Reply, which cannot be matched for the corresponding AP Request;
- verify that the nonce value contained in the seq-number field in AP Reply matches the one in the corresponding AP Request. The Client MUST immediately discard the AP Reply if seq-number field value in AP Reply does not match.

If the AP Reply is discarded, the Client MUST proceed as if the message had never been received (e.g., if the client was waiting for a valid AP Reply, it should continue to do so). Once the Application Server has sent an AP Reply with the ACK-required flag set, it MUST compute the expected value in the Security Parameter Recovered message and save it for an appropriate timeout period during which it will accept a matching Security Parameter Recovered Message. Once the appropriate timeout period is exceeded, the Application Server MUST discard the saved values and no longer accept a matching Security Parameter Recovered Message

Each time the Application Server times out waiting for the Security Parameter Recovered message, it MUST continue with the exponential back-off algorithm until all retries have been exhausted, as specified in clause 6.4.8. The Application Server MUST begin each retry by re-sending a Wake Up message with a new server-nonce value.

- 4) Security Parameter Recovered – Sent by the client to the Application Server to acknowledge that it received an AP Reply and successfully set up new Security Parameters. This message is only sent when ACK-required flag is set in the AP Reply.

This message MUST consist of the concatenation of the following:

- Key Management Message ID – 1-byte value. Always set to 0x04.
- Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established. See the table of DOI values in this clause.
- Protocol Version – 1 byte. The high-order nibble is the major version number, and the lower-order nibble is the minor version number. For IPCablecom, the major number MUST be 1, and the minor number MUST be 0.
- HMAC – a 20-byte SHA-1 HMAC of the preceding AP Reply message. The 20-byte key for this HMAC is determined by taking a SHA-1 hash of the sub-key from the AP Reply.

If the receiver (Application Server) gets a bad Security Parameter Recovered message that does not match an AP Reply, the Application Server MUST discard it and proceed as if this Security Parameter Recovered message was never received.

6.5.2.1 Rekey messages

The Rekey message replaces the Wake Up message and provides better performance, whenever a receiver (Application Server) wants to trigger the establishment of a Security Parameter with a specified client. The Rekey message requires the availability of the shared Server Authentication Key, which is not always available. Thus, support for the Wake Up message is still required.

The Rekey message was added specifically for use with the NCS-based clustered Call Agents, potentially consisting of multiple IP addresses and multiple hosts. Any IP address or host within one cluster needs the ability to quickly establish a new Security Parameter with a client, without a significant impact to the ongoing voice communication.

The use of the Rekey message eliminates the need for the AP Reply message, thus reducing the key management overhead to a single round trip. This is illustrated in Figure 9.

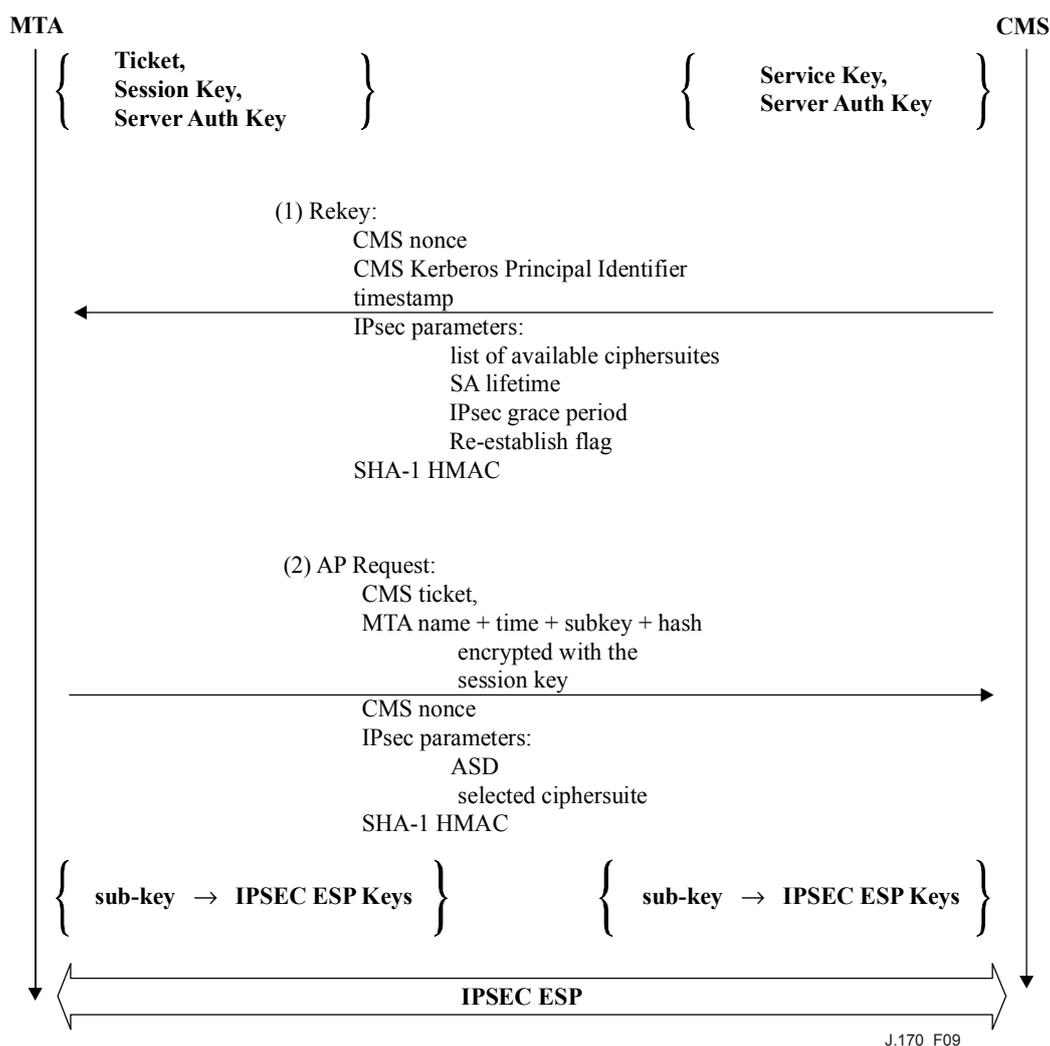


Figure 9 – Rekey message to establish a security parameter

The messages listed in Figure 9 are defined as follows:

- 1) **Rekey** – Sent by the Application Server to establish a new set of security parameters. It MUST be a concatenation of the following:
 - Key Management Message ID – 1-byte value. Always set to 0x05.
 - Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established. See the table of DOI values in this clause.

- Protocol Version – 1 byte. The high-order nibble is the major version number, and the lower-order nibble is the minor version number. For IPCablecom, the major number MUST be 1, and the minor number MUST be 0.
- Server-nonce – a 4-byte random binary string. Its value MUST NOT be all 0s.
- Server Kerberos principal identifier – a printable, null-terminated ASCII string, representing a fully qualified Kerberos principal identifier of the Application Server, as defined in 6.4.5. This allows the client to both find the right Server Authentication Key and to pick the right Kerberos ticket for the subsequent AP Request message.
- Timestamp – a string of the format YYMMDDhhmmssZ, representing UTC time. This string is not NULL-terminated.
- Application-specific data – additional information that must be communicated by the server to the client, dependent on the target protocol for which security is being established (e.g., IPsec).
- List of ciphersuites available at the server – see above specification for the AP Request message.
- Security parameters lifetime – a 4-byte value, MSB first. This indicates the number of seconds from now, when this set of security parameters is due to expire.
- Grace period – a 4-byte value in seconds, MSB first. This indicates to the client to start creating a new set of security parameters (with a new AP Request/AP Reply exchange) when the timer gets to within this period of their expiration time.
- Re-establish flag – a 1-byte Boolean value. When the value is TRUE (1), a new set of security parameters MUST be established before the old one expires. When the value is FALSE (0), the old set of security parameters MUST be allowed to expire.
- SHA-1 HMAC over the concatenation of all of the above-listed fields.

The Server Authentication Key used for this HMAC is uniquely identified by the following name pair (client principal name, server principal name). This key MUST be updated at the Application Server right after it sends an AP Reply message. It MUST be set to a (20-byte) SHA-1 hash of the Kerberos session key used in that AP Reply. The client MUST also update this key as soon as it receives the AP Reply. (Note that multiple AP Replies will continue using the same Kerberos session key, until it expires. That means that the derived Server Authentication Key may have the same value as the old one.)

It is possible that the Application Server sends a Rekey message as soon as it sends an AP Reply (from another IP address), and before the client is able to derive the new Server Authentication Key. In that case, the client will not authenticate the Rekey message and the Application Server will have to retry.

Similarly, after sending an AP Reply the Application Server might immediately send an IP packet using the just established Security Parameter, when the client is not yet ready to receive it. In this case, the client will reject the packet and the Application Server will have to retransmit.

Both of these error cases could be completely avoided with a 3-way handshake (a client acknowledging an AP Reply with a Security Association Recovered message).

Whenever the Rekey message is received (by the client), it MUST verify the value of this HMAC. If this integrity check fails, the client MUST immediately discard this message and proceed as if the message had never been received.

Once the Application Server has sent a Rekey, it MUST save the server-nonce in order to validate a matching AP Request. In the case of a time-out, the Application Server MUST adhere to the exponential retry backoff procedure described in 6.4.8. The Application Server MUST begin each retry by re-sending a Rekey message with a new server-nonce

value. When the "Time-out Procedure" has completed without success, the Application Server MUST discard the server-nonce from the last retry, after which it will no longer accept a matching AP Request.

When this Rekey message is received and validated by the client, all previously existing outgoing Security Parameters with this Application Server IP address MUST be removed at this time. If the client previously had a timer set for automatic refresh of Security Parameters with this Application Server IP address, that automatic refresh MUST be reset or disabled.

The client MUST verify that this Rekey message is not a replay using the procedure similar to the one for AP Request in the Kerberos standard (IETF RFC 4120):

- If $|T_{CMS} - (T_{MTA} + Skew)| > \text{The acceptable Clock Skew}$, then the client MUST drop the message. Here, T_{CMS} is the timestamp in the Rekey message and T_{MTA} is the reading of the MTA local clock. Skew is the saved difference between the Application Server and MTA clock. `PkctSrvrToMtaMaxClockSkew` is currently in the MTA MIB (see ITU-T Rec. J.166) as the variable `pkctMtaDevCmsMaxClockSkew`.
- If the Server-nonce, principal name and timestamp fields match any recently seen (within the `pkctSrvrToMtaMaxClockSkew`) Rekey messages, then the client MUST drop the message.

- 2) AP Request – MUST be sent by the client as a response to a Rekey message. Unlike the AP Request message described above, this one MUST also include the sub-key (inside `KRB_AP_REQ` ASN.1 structure). `KRB_AP_REQ` will have a Kerberos flag set, indicating that an AP Reply MUST NOT follow.

The format of the AP Request is as specified above in 6.5.2. The only difference is that the list of ciphersuites here MUST contain exactly one entry – the ciphersuite selected by the client from the list provided in the Rekey message.

Right before the client sends out this AP Request, it MUST establish the security parameters with the corresponding server IP address. If the corresponding Rekey message had the Re-establish flag set, the client MUST be prepared to automatically re-establish new security parameters, as specified in 6.5.

Once this AP Request is received and verified by the Application Server, the server MUST also establish the security parameters.

6.5.2.2 IPCablecom profile for `KRB_AP_REQ/KRB_AP_REP` messages

In the `KRB_AP_REQ`, only the following option is supported:

- `MUTUAL-REQUIRED` – mutual authentication required. When this option is set, the server MUST respond with an AP Reply message. When this option is not set, the AP Reply message MUST NOT be sent in reply.
- All other options MUST NOT be set. If an application server receives a request containing the unsupported option `USE-SESSION-KEY`, it MUST return an error with the error code `KRB_AP_ERR_METHOD`. If an application server receives a request containing any other unsupported options, it MUST return an error with the error code `KRB_ERR_GENERIC`. When `MUTUAL-REQUIRED` is set, the encrypted authenticator in the `KRB_AP_REQ` MUST contain the following field, which is optional in Kerberos:
 - `seq-number` – MUST contain a pseudo-random number generated by the client (to be used as a nonce).

The server MUST accept otherwise-valid `KRB-AP-REQ` messages that contain a `seq-number` in the range -2^{31} to -1 .

When MUTUAL-REQUIRED is not set, the encrypted authenticator MUST contain the following field that is optional in Kerberos:

- sub-key – used to generate security parameters for the target protocol. The sub-key type MUST be set to –1³. The actual sub-key length is dependent on the target protocol.

When MUTUAL-REQUIRED is set, the target protocol is IPsec and the client is an MTA, the client MAY include the subkey field; in the case that the target protocol is IPsec and the client is other than an MTA, the client SHOULD include the subkey field. For IPsec, the subkey, if present, MUST contain a pseudo-random number of length 46 octets generated by the client.

Other optional fields in the authenticator MUST NOT be present. If the authenticator contains the authorization-data field, the application server MUST return an error with the error code KRB_ERR_GENERIC. If the authenticator contains any other optional fields (apart from subkey and authorization-data), the application server MUST ignore those fields.

The negative key type is used to indicate that it is application-specific and not defined in the Kerberos specification. When the Kerberos specification is updated to include this key type, the IPCablecom spec will be updated accordingly.

The authenticator itself MUST be encrypted using 3-DES CBC with the Kerberos etype value: des3-cbc-md5 as specified in 6.4.2.2.

In the encrypted part of the KRB_AP_REP, the optional sub-key field MUST be used for IPCablecom. Its type and format MUST be the same as when it appears in the KRB_AP_REQ (see above).

The optional seq-number MUST be present, and MUST echo the value that was sent by the client in the KRB_AP_REQ. In this context, the seq-number field is used as a random nonce. The encrypted part of the KRB_AP_REP MUST be encrypted with the Kerberos etype value: des3-cbc-md5 as specified in 6.4.2.2.

6.5.2.3 Error handling

6.5.2.3.1 Error Reply

If the Application Server is able to successfully parse the AP Request and the ticket that is inside of it, but the AP Request is rejected, it MUST return an error message. This error message MUST be formatted as the concatenation of the following fields:

- Key Management Message ID – 1-byte value. Always set to 0x06.
- Protocol Version – 1-byte value. The high-order nibble is the major version number and the lower-order nibble is the minor version number. For IPCablecom the major version number MUST be 1 and the minor version number MUST be 0.
- Domain of Interpretation (DOI) – 1-byte value. Specifies the target protocol for which security parameters are established. See the table of DOI values in 6.5.2.
- KRB_ERROR – Kerberos error message as specified in IETF RFC 4120. It MUST include typed-data of REQ-SEQ to bind the error message to the sequence number from the authenticator in the AP-REQ message. The value encapsulated by the REQ-SEQ typed data MUST be the same as the value of the seq-number that was sent by the client in the KRB_AP_REQ. Also, the error message MUST include the optional e-cksum member, which is the keyed hash over the KRB_ERROR message. The checksum type MUST be rsa-md5-des3, as specified in 6.4.3.1.

³ The negative key type is used to indicate that it is application-specific and not defined in the Kerberos specification. When the Kerberos specification is updated to include this key type, the IPCablecom spec will be updated accordingly.

If the error is application-specific (not a Kerberos-related error), then the KRB_ERROR MUST include typed-data of type TD-APP-DEFINED-ERROR (value 106). The value of this typed-data is the following ASN.1 encoding (specified in IETF RFC 4120):

```
AppSpecificTypedData ::= SEQUENCE {
    oid          [0] OPTIONAL OBJECT IDENTIFIER,
                -- identifies the application
    data-value   [1] OCTET STRING
                -- application-specific data
}
```

Both the oid and the data-value fields inside AppSpecificTypedData are specified separately for each DOI.

- Upon receiving this error reply, the client MUST verify both the keyed checksum and the REQ-SEQ field, to make sure that they match the seq-number field from the authenticator in the AP Request.

If the Application Server is not able to successfully parse the AP Request and the ticket, it MUST drop the request and it MUST NOT return any response to the client. In case of a line error, the client will time out and re-send its AP Request. If the verification has failed, then the MTA MUST ignore this error message and continue waiting for the reply as if the error message was never received.

When a client receives an error message, in some cases this Recommendation calls for the client to take some recovery steps and then send a new AP Request message. When a client is responding to an error message, it is not a retry and MUST NOT be considered to be part of the client's back-off and retry procedure specified in 6.4.8. The client MUST reset its timers accordingly, to reflect that the AP Request in response to an error message is not a retry.

Although this Recommendation calls for an application server to return some specific error codes under certain error conditions, in the case when a server is repeatedly getting the same error from the same client IP address, it MAY at some point choose to stop sending back any further replies (errors or otherwise) to this client.

6.5.2.3.2 Clock skew error

When the Application Server clock and the client clock are off by more than the limit for a clock skew, an error code KRB_AP_ERR_SKEW MUST be returned. The value for the maximum clock skew allowed by the Application Server MUST NOT exceed 5 minutes. The optional client's time in the KRB_ERROR MUST be filled out, and the client MUST compute the difference (in seconds) between the two clocks based upon the client and server time contained in the KRB_ERROR message. The client SHOULD store this clock difference in non-volatile memory and MUST use it to adjust Kerberos timestamps in subsequent AP Request messages by adding the clock skew to its local clock value each time. The client MUST maintain a separate clock skew value for each realm and MAY share the same clock skew between the KDC and various application servers within that realm. The clock skew values are intended for uses only within the Kerberos protocol and SHOULD NOT otherwise affect the value of the local clock (since a clock skew is likely to vary from realm to realm).

In the case that an AP Request failed due to a clock skew error, a client MUST immediately retry after adjusting the Kerberos timestamp inside the AP Request message.

Additionally, the client MUST validate the time offset returned in the clock skew error, to make sure that it does not exceed a maximum allowable amount. This maximum time offset MUST not exceed 1 hour. This client check against a maximum time offset protects against an attack, where a rogue KDC attempts to fool an client into accepting an expired KDC certificate (later, during the next PKINIT exchange).

6.5.2.3.3 Handling ticket errors after a Wake Up

6.5.2.3.3.1 KRB_AP_ERR_BADKEYVER after a Wake Up

This clause addresses a scenario when an application server sends a Wake Up to a client and subsequently receives an AP Request that contains a ticket that is encrypted using an obsolete service key (results in the KRB_AP_ERR_BADKEYVER error code). This error normally requires the client to get another ticket and retry, but in this particular scenario the client has to retry in the middle of a key management transaction.

In this scenario, the application server MUST reply to the invalid AP Request with the KRB_ERROR message with the KRB_AP_ERR_BADKEYVER error code. Subsequent to the reply, the server MUST wait for another AP Request and MUST use the same time-out value that it would normally use when waiting for an AP Request. The client, upon getting back the above error code, MUST attempt to obtain a new ticket from the KDC (if the client has not done so already while waiting for server's reply) and if successful, MUST send another AP Request to the application server. If the client is unsuccessful in obtaining another ticket, it MUST not reply. If the server times out waiting for the second AP Request, it MUST proceed as if it timed out waiting for the original AP Request.

If the application server is able to validate the second AP Request, it MUST then proceed as specified in 6.5.3. If the second AP Request again results in the KRB_AP_ERR_BADKEYVER error, the server MUST abort key management with this client and not reply.

6.5.2.3.3.2 KRB_AP_ERR_SKEW after a Wake Up

An application server is not required to check for a clock skew in this case, but if it does generate the KRB_AP_ERR_SKEW, the same procedure MUST be followed as in 6.5.2.3.3.1, except that the client MUST retry after adjusting the timestamp (see 6.5.2.3.2) instead of getting a new ticket.

6.5.3 Kerberized IPsec

This clause specifies the Kerberized key management profile specific to IPsec ESP in transport mode. IPsec uses the term "Security Association" (SA) to refer to a set of security parameters. IPsec Security Associations are always unidirectional and they MUST always be established in pairs within IPCablecom.

An MTA MUST establish SAs with the IP address from where the corresponding Kerberized IPsec key management message (AP-REP or REKEY) has been received. Note that a CMS can notify an MTA that it is listening for NCS messages on a different port. Also, both the CMS and the MTA can send NCS messages from different ports, and the response must be sent to the port from which the message was sent. Kerberized Key Management does not allow for the negotiation of source or destination ports. Therefore SAs established to protect NCS signalling need to support multiple ports. One way to accomplish this is to establish two separate policies, outbound and inbound, in the IPsec Security Policy Database (see RFC 2367). Table 11 illustrates an example policy that would support changes in port numbers. Note that this table only illustrates inbound and outbound policies for NCS signalling between a specific MTA and a specific CMS. Table 11 is not a complete IPsec Security Policy Database. Other entries would be required to support communications over different protocols with the same host (e.g., Kerberized Key Management), communications with other hosts, or default policies for unknown hosts.

**Table 11 – Example IPsec Security Policy Database Entries
for NCS Signalling between MTA and CMS**

| Direction | Policy | Source IP | Source Port | Destination IP | Destination Port |
|---|-----------------|-------------------|--|-----------------------|--|
| Inbound – this applies to messages being received | Apply IPsec ESP | Remote IP address | Wildcard – any port | Local IP address | Bind to local port(s) that NCS messages will be sent from, and the provisioned NCS listening port. |
| Outbound – this applies to messages being sent | Apply IPsec ESP | Local IP address | Bind to local port(s) that messages will be sent from. | Remote IP address | Wildcard – any port |

The DOI value for IPsec MUST be set to 1.

The ASD (Application-Specific Data) field in the AP Request key management message MUST be the SPI (Security Parameter Index) for the client's inbound Security Associations. It is a 4-byte integer value, MSB first.

The ASD field in the AP Reply and Rekey key management messages MUST be the SPI for the server's inbound Security Associations. It is a 4-byte integer value, MSB first.

The sub-key for IPsec MUST be a 46-byte value defined as follows:

- If the AP-REQ does not include a sub-key, the 46-octet sub-key from AP-REP is taken as the subkey for IPsec.
- If the AP-REQ does include a sub-key but no AP-REP (in the case of Rekey) is sent, then the 46-octet AP-REQ sub-key is used as the sub-key for IPsec.
- Otherwise, both the AP-REQ and the AP-REP messages include 46-octet sub-keys, and their bit-by-bit XOR is the 46-byte sub-key for IPsec.

An MTA MUST NOT perform Kerberized Key Management or establish IPsec Security Associations with a CMS when the `pktcMtaDevCmsIpsecCtrl` flag for that CMS is set to FALSE in the `pktcMtaDevCmsTable`. Note that this flag may only be set in the MTA configuration file and cannot be updated using SNMPv3. In the case of an NCS Redirect or any other dynamic method for associating a new CMS with an MTA endpoint where there is not an entry in the `pktcMtaDevCmsTable` for the new CMS, the MTA MUST perform Kerberized Key Management and establish IPsec Security Associations with the new CMS.

The CMS MUST be capable of disabling its Kerberized Key Management interface. The CMS MUST NOT perform Kerberized Key Management or establish IPsec Security Associations when so configured.

6.5.3.1 Derivation of IPsec keys

After the Application Server sends out an AP Reply message, it is ready to derive a new set of IPsec keys. Similarly, after the client receives this AP Reply, it is ready to derive the same set of keys for IPsec. This clause specifies how the IPsec keys are derived from the Kerberos sub-key.

The size of the Kerberos sub-key MUST be 46 bytes – the same as with the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) pre-master secret.

The IPsec ESP keys MUST be derived in the following order:

- 1) Message authentication key for Client → Application Server messages;
- 2) Encryption key for Client → Application Server messages;
- 3) Message authentication key for Application Server → Client messages;

- 4) Encryption key for Application Server → Client messages.

For specific authentication and encryption algorithms that may be used by IPsec, refer to 6.1.2.

The derivation of the required keying material MUST be based on running a one-way pseudo-random function $F(S, \text{"IPsec Security Association"})$ recursively until the right number of bits has been generated. Here, S is the Kerberos sub-key and the ASCII string "IPsec Security Association" is taken without quotes and without a terminating null character. F is defined in 9.6.

6.5.3.2 Periodic re-establishment of IPsec Security Associations

An IPsec SA is defined with an expiration time T_{EXP} and a grace period GP_{IPsec} . The subclauses below specify how both the client and the Application Server handle the re-establishment of IPsec Security Associations (re-establish flag was TRUE in the AP Reply). When the re-establishment of IPsec SAs is required, there MUST always be at least one SA available for each direction and there MUST NOT be an interruption in the call signalling.

6.5.3.2.1 Periodic re-establishment of IPsec SAs at the client

If the re-establish flag is set, the client MUST attempt to establish a new set of IPsec SAs (one for each direction) starting at the time $T_{EXP} - GP_{IPsec}$. At this time, the client MUST send an AP Request as specified in 6.5. The destination IP Address of the AP Request message MUST be the destination IP Address of the outbound IPsec SA that is about to expire. After the client receives an AP Reply, it MUST perform the following steps:

- 1) Create new IPsec SAs, based on the negotiated ciphersuite, SPIs and on the established Kerberos subkey, from which the IPsec keys are derived as specified in 6.5. The expiration time for the outgoing SA MUST be set to T_{EXP} , while the expiration time for the incoming SA MUST be set to $T_{EXP} + GP_{IPsec}$.
- 2) From this point forward, the new SA MUST be used for sending messages to the Application Server. The old SA that the client used for sending signalling messages to the Application Server MAY be explicitly removed at this time, or it MAY be allowed to expire (using an IPsec timer) at the time T_{EXP} .
- 3) Continue accepting incoming signalling messages from the Application Server on both the old and the new incoming SAs, until the time $T_{EXP} + GP_{IPsec}$. After this time, the old incoming SA MUST expire. If a client receives a signalling message from the Application Server using a new incoming SA at an earlier time, it MAY at that time remove the old incoming SA.

If the client fails to get any reply from the server and has to retry one or more times with another AP Request, the re-establish flag MUST be set to FALSE in each retry. This implies that when CMS processes a retry, it will remove any existing outgoing IPsec SAs, including the ones that may have been created after the processing of the initial AP Request, and proceed as if it is processing the SAs on demand (see 6.5.3.5.1).

6.5.3.2.2 Periodic re-establishment of IPsec SAs at the Application Server

When an AP Request message is received with re-establish flag set, the Application Server MUST perform the applicable processing steps in 6.5.2. If the client is an MTA, the Application Server MUST also verify that the source IP address in the received datagram of the AP Request message is the same IP address as was used when the initial SA was established. The Application Server MUST ignore the AP Request if the IP addresses do not match.

In addition, the Application Server MUST perform the following steps, in the specified order immediately before an AP Reply is returned.

- 1) Create new IPsec SAs, based on the negotiated ciphersuite, SPIs and on the established Kerberos sub-key, from which the IPsec keys are derived as specified in 6.5.
- 2) Send back an AP Reply.
- 3) Continue sending signalling messages to the client using an old outgoing SA until the time T_{EXP} . During the same period, accept incoming messages from either the old or the new incoming SA.
- 4) At the time T_{EXP} both the old incoming and the old outgoing SAs MUST expire. At the time T_{EXP} , the Application Server MUST switch to the new SA for outgoing signalling messages to the client. If for some reason the new IPsec SAs were not established successfully, there would not be any IPsec SAs that are available after this time.

6.5.3.3 Expiration of IPsec SAs

An IPsec SA is defined with an expiration time T_{EXP} and a grace period GP_{IPsec} . This clause specifies how both the client and the Application Server MUST handle the expiration of IPsec Security Associations (re-establish flag was FALSE in the AP Reply).

At the Client:

- Outgoing SA expires at T_{EXP}
- Incoming SA expires at $T_{EXP} + GP_{IPsec}$

At the Application Server:

- Outgoing SA expires at T_{EXP}
- Incoming SA expires at $T_{EXP} + GP_{IPsec}$

Whenever an IPsec SA has been expired and a signalling message needs to be sent by either the client or the Application Server, the key management layer MUST be signalled to establish a new IPsec SA. It is established using the same procedures as the ones specified in 6.5.3.5.

6.5.3.4 Initial establishment of IPsec SAs

When a client is rebooted, it does not have any current IPsec SAs established with the Application Server, since IPsec SAs are not saved in non-volatile memory. In order to re-establish them, it MUST go through the recovery procedure that is described in 6.5.3.5.

6.5.3.5 On-demand establishment of IPsec SAs

This clause describes the recovery steps that MUST be taken in the case that an IPsec SA is somehow lost and needs to be re-established.

6.5.3.5.1 Client loses an outgoing IPsec SA

If a client attempts to send a signalling message to the Application Server without a valid IPsec SA, the IPsec layer in the client will realize the SA is missing and returns an error back to the signalling application⁴. In this case, the following recovery steps MUST be taken at the key management layer:

- 1) The client first makes sure that it has a valid Kerberos ticket for the Application Server. If not, it must first perform a PKINIT exchange as specified in 6.4.2.

⁴ In this case, there are no actual messages exchanged between the MTA and the application server (e.g., CMS).

- 2) Client sends a new AP Request to the Application Server and gets back an AP Reply, as specified in 6.5.2. After the receipt of an AP Reply, the client MUST:
 - create new IPsec SAs;
 - remove any old outgoing IPsec SAs;
 - be prepared to use both of the newly created IPsec SAs.
- 3) If the Kerberos ticket includes the optional caddr field and the caddr does not contain a matching source IP address for the AP Request datagram, the Application Server MUST ignore the request.
- 4) The Application Server also MUST NOT set the ACK-required flag in the AP Reply. Right after sending out an AP Reply, the Application Server MUST be prepared to both send and receive messages on the newly created SAs.
- 5) After receiving this AP Request (with Re-establish flag = FALSE), the Application Server MUST remove any existing outgoing IPsec SAs that it might already have for this client.

The key management application running on the client MUST send an explicit signal to the signalling application, when it completes the re-establishment of the IPsec SAs.

6.5.3.5.2 Client loses an incoming IPsec SA

When the client receives an IP packet from an Application Server on an unrecognized IPsec SA, the client MUST ignore this error and the packet MUST be dropped.

6.5.3.5.3 Application Server loses an outgoing IPsec SA

When an Application Server attempts to send a signalling message to the client, and the IPsec layer in the Application Server realizes a valid SA is missing, the IPsec layer MUST return an error back to the signalling application⁵. In this case, the following recovery steps MUST be taken at the key management layer:

- 1) Application Server sends a Wake Up message to the client.
- 2) The client makes sure that it has a valid Kerberos ticket for the Application Server. If not, it MUST first obtain it from the KDC.
- 3) Client sends a new AP Request to the Application Server, as specified in 6.5.2. If the Kerberos ticket includes the optional caddr field and the caddr does not contain a matching source IP address for the AP Request datagram, the Application Server MUST ignore the request.
- 4) For each AP Request, the client generates a nonce and puts it into the seq-number field. As specified in 6.5.2, the client will save this nonce for a period of time specified by the pktcMtaDevCmsSolicitedKeyTimeout MIB object, and wait for a matching AP Reply (this is not the same nonce as the Server-nonce received in the Wake Up). However, after this time-out, the client MUST NOT retry and MUST abort an attempt to establish a IPsec SA in response to a received Wake Up.

Once the client gets back a matching AP Reply, it will be in the format specified in 6.5.2. The ACK-required flag in the AP Reply MUST be set, to insure that the client replies with the SA Recovered message in the following step. If this client previously had any outgoing IPsec SAs with this Application Server IP address, they MUST be removed at this time. If the client previously had a timer set for automatic refresh of IPsec SAs with this Application Server IP address, that automatic refresh MUST be reset or disabled. The client MAY start using both of the newly created SAs. If the AP Reply had the Re-establish flag

⁵ In this case, there are no actual messages exchanged between the MTA and the CMS or other application server.

set, the client MUST be prepared to automatically re-establish new IPsec SAs, as specified in 6.5.3.2.

- 5) The Application Server can receive signalling messages from the client on the new incoming SA, but cannot yet start using an outgoing SA for sending messages to the client.
- 6) Immediately after the client establishes the new IPsec SAs, it MUST send a SA Recovered message to the Application Server.
- 7) Upon receipt of this message, the Application Server MUST immediately activate the new outgoing SA for sending signalling messages to the client.

The key management application running on the Application Server MUST send an explicit signal to the signalling application when it completes the re-establishment of the IPsec SAs.

6.5.3.5.4 Application Server loses an incoming IPsec SA

When the Application Server receives an IP packet from a client on an unrecognized IPsec SA, the Application Server MUST ignore this error and the packet MUST be dropped. In this case, any attempt at recovery (e.g., establishing a new SA) is prone to denial-of-service attacks.

6.5.3.6 IPsec-specific errors returned in KRB_ERROR

Inside AppSpecificTypedData the oid field MUST be set to:

```
enterprises (1.3.6.1.4.1) cableLabs (4491) clabProjects (2) clabProjPacketCable
(2) pktcSecurity (4) errorCodes (1) ipSec (1)
```

The data-value field MUST correspond to the following typed-data value:

```
pktcKrbIpsecError ::= SEQUENCE {
    e-code      [0]      INTEGER,
    e-text      [1]      GeneralString OPTIONAL,
    e-data      [2]      OCTET STRING OPTIONAL
}
```

The e-code field MUST correspond to one of the following error code values:

| | | |
|-------------------------|----|---|
| KRB_IPSEC_ERR_NO_POLICY | 1 | No IPsec policy defined for request |
| KRB_IPSEC_ERR_NO_CIPHER | 2 | No support for requested ciphersuites |
| KRB_IPSEC_NO_SA_AVAIL | 3 | No IPsec SA available (i.e., SAD is full) |
| KRB_IPSEC_ERROR_GENERIC | 16 | Generic KRB IPsec error |

The optional e-text field can be used for informational purposes (i.e., logging, network troubleshooting) and the optional e-data field is reserved for future use to transport any application data associated with a specific error.

6.5.4 Kerberized SNMPv3

This clause specifies the Kerberized key management profile specific to SNMPv3; see RFC 3414. In the case of SNMPv3, the security parameters are associated with the usmUserName (SNMPv3 user name), the agent's usmUserEngineID (SNMPv3 engine ID) and the manager's usmUserEngineID.

Multiple SNMP managers on different hosts but with the same user name are considered as unique Kerberos principals. Still, the SNMPv3 keys generated by any one of these SNMP managers MUST be shared across all the managers – as long as they apply to the same SNMPv3 user name and the same SNMPv3 engine ID (of the agent).

The security parameters consist of a single authentication key, a single privacy (encryption) key, SNMPv3 boot count and engine time. SNMPv3 privacy can be turned off by selecting a NULL encryption transform.

The DOI value for SNMPv3 MUST be set to 2.

The ASD field in the AP Request message MUST be set to the concatenation of the fields as shown in Table 12:

Table 12 – Required format for data in the AP Request

| Attribute | Length |
|-----------------------------|-----------------------------|
| Agent's snmpEngineID Length | 1 byte |
| Agent's snmpEngineID | Variable |
| Agent's snmpEngineBoots | 4 bytes, network byte order |
| Agent's snmpEngineTime | 4 bytes, network byte order |
| usmUserName Length | 1 byte |
| usmUserName | Variable |

The ASD field in the AP Reply message MUST be set to the concatenation of the fields as shown in Table 13:

Table 13 – Required format for data in the AP Reply

| Attribute | Length |
|-------------------------------|-----------------------------|
| Manager's snmpEngineId Length | 1 byte |
| Manager's snmpEngineId | Variable |
| Manager's snmpEngineBoots | 4 bytes, network byte order |
| Manager's snmpEngineTime | 4 bytes, network byte order |
| usmUserName Length | 1 byte |
| usmUserName | Variable |

For IPCablecom MTAs, the usmUserName contains in it the MTA MAC address (see ITU-T Rec. J.167). The manager MUST verify that this MAC address and the MTA FQDN specified in the MTA principal name match. The manager MUST also verify that any SNMP INFORM message containing a MAC address, from the MTA contains a correct MAC address – the same one that is in the usmUserName.

The usmUserName field inside the Application-specific data field in the AP Reply MUST be the same as the one in the preceding AP Request.

The Rekey message is not used for SNMPv3 key management.

The sub-key for SNMPv3 MUST be a 46-byte value.

6.5.4.1 Derivation of SNMPv3 keys

After the server sends out an AP Reply message, it is ready to derive a new set of SNMPv3 keys. Similarly, after the client receives this AP Reply, it is ready to derive the same set of keys for SNMPv3. This clause specifies how the SNMPv3 keys are derived from the Kerberos sub-key.

The size of the Kerberos sub-key MUST be 46 bytes.

The derived SNMPv3 Keys MUST be as follows, in the specified order:

- SNMPv3 authentication key;
- SNMPv3 privacy key.

For specific authentication and encryption algorithms that may be used by IPCablecom for SNMPv3, refer to 6.3.

The derivation of the required keying material MUST use a one-way pseudo-random function $F(S, \text{"SNMPv3 Keys"})$ recursively until the right number of bits has been generated. Here, S is the subkey and the string "SNMPv3 Keys" is taken without quotes and without a terminating null character. F is defined in 9.6.

6.5.4.2 Periodic re-establishment of SNMPv3 keys

Periodic re-establishment of SNMPv3 keys, where the next set of keys is created before the old one expired, is currently not supported by IPCablecom. The re-establish flag in the AP Reply key management message MUST be set to FALSE.

6.5.4.3 Expiration of SNMPv3 keys

Expiration of SNMPv3 keys is currently not supported by IPCablecom. The values of the Security Parameters Lifetime and Grace Period fields in the AP Reply MUST be set to 0.

6.5.4.4 Initial establishment of SNMPv3 keys

When a client is rebooted, it may not have any saved SNMPv3 keys established with the SNMP Manager. In order to re-establish them, it goes through the recovery procedure that is described in 6.5.4.5.1.

6.5.4.5 Error recovery

This clause describes the recovery steps that must be taken in the case that SNMPv3 keys are somehow lost and need to be re-established.

6.5.4.5.1 SNMP agent wishes to send with missing SNMPv3 keys

In the case of SNMP, an SNMP agent is not responsible for re-establishing SNMPv3 keys because it does not send unsolicited requests to the Provisioning Server after the initial provisioning is done. Still, an SNMP agent could attempt to re-establish SNMPv3 keys after it gets an SNMPv3 authentication error back from the SNMP manager. If the SNMP agent determines that it has incorrect SNMPv3 keys, it MUST perform the following steps before it is able to send out an SNMP message:

- 1) The agent first makes sure that it has a valid Kerberos ticket for the Application Server. If not, it must first obtain it as specified in 6.5.2.
- 2) The agent sends a new AP Request to the manager and gets back an AP Reply, as specified in 6.5.2. After the receipt of the AP Reply the agent is prepared to use the newly created SNMPv3 keys. In this scenario, the SNMP manager MUST NOT set an ACK-required flag in the AP Reply. Right after sending out an AP Reply, the manager is prepared to both send and receive messages with the new SNMPv3 keys. After receiving this AP Request (with Re-establish flag = FALSE), the manager MUST remove its previous set of SNMPv3 keys that it might already have for this agent (and for this SNMPv3 user name).

It is possible that the SNMP manager already initiated key management (with a Wake Up) but instead receives an unsolicited AP Request from the agent (with server-nonce = 0). This unlikely scenario might occur if the manager and the agent decide to initiate key management at about the same time. In this case, the SNMP manager MUST ignore the unsolicited AP Request message and continue waiting for the one that is in response to a Wake Up.

6.5.4.5.2 SNMP agent receives with missing SNMPv3 keys

If the SNMP agent receives a request from a SNMP manager and is unable to find SNMPv3 keys for the specified USM User Name, the agent MUST process the SNMP message according to RFC 3414 and RFC 3412.

6.5.4.5.3 SNMP manager wishes to send with missing SNMPv3 keys

SNMP manager attempts to send a message to the agent and does not find the desired user's SNMPv3 keys (or considers the existing SNMPv3 keys invalid or compromised). In this case, the following recovery steps MUST be taken at the key management layer:

- 1) The manager sends a Wake Up message to the agent.
- 2) The agent makes sure that it has a valid Kerberos ticket for the manager. If not, it MUST first obtain it from the KDC.
- 3) The agent sends a new AP Request to the manager, as specified in 6.5.2. For each AP Request, the agent generates a nonce and puts it into the seq-number field. As specified in 6.5.3.5.3, the agent will save this nonce for a period of time specified by the `pktcMtaDevProvSolicitedKeyTimeout` MIB object and wait for a matching AP Reply (this is not the same nonce as the server-nonce received in the Wake Up). However, after this time-out, the agent MUST NOT retry and MUST abort an attempt to establish SNMPv3 keys in response to a received Wake Up.

Once the agent gets back a matching AP Reply, it will be in the format specified in 6.5.2. The ACK-required flag in the AP Reply MUST be set, to ensure that the agent replies with the SA Recovered message in the following step. If this agent previously had SNMPv3 keys for the specified SNMPv3 user, they MUST be removed at this time.

- 4) After the receipt and validation of the AP Reply, the agent sends an SA Recovered message to the manager. At this time the agent will be ready to use the new SNMPv3 keys and will enable SNMPv3 security.
- 5) Upon receipt of the SA Recovered message, the manager will immediately activate the new set of SNMPv3 keys and will enable SNMPv3 security.

It is possible that the SNMP agent already initiated key management (with an unsolicited AP Request) but instead receives a Wake Up from the manager. This unlikely scenario might occur if the manager and the agent decide to initiate key management at about the same time. In this case, the SNMP agent MUST abort waiting for the reply to the unsolicited AP Request message and instead generate a new AP Request in response to the Wake Up.

If an SNMP agent receives a second Wake Up message from a different SNMP manager (FQDN or IP address) before the first key management session has been completed, the SNMP agent MUST ignore the second Wake Up message.

6.5.4.6 SNMPv3-specific errors returned in KRB_ERROR

Inside `AppSpecificTypedData`, the `oid` field MUST be set to:

```
enterprises (1.3.6.1.4.1) cableLabs (4491) clabProjects (2) clabProjPacketCable
(2) pktcSecurity (4) errorCodes (1) snmpv3 (2)
```

The data-value field MUST correspond to the following typed-data value:

```
pktcKrbSnmpv3Error ::= SEQUENCE {
    e-code    [0]    INTEGER,
    e-text    [1]    GeneralString OPTIONAL,
    e-data    [2]    OCTET STRING OPTIONAL
}
```

The e-code field MUST correspond to one of the following error code values:

| | | |
|---------------------------------------|----|---------------------------------------|
| <code>KRB_SNMPV3_ERR_USER_NAME</code> | 1 | Unrecognized SNMPv3 user name |
| <code>KRB_SNMPV3_ERR_NO_CIPHER</code> | 2 | No support for requested ciphersuites |
| <code>KRB_SNMPV3_ERR_ENGINE_ID</code> | 3 | Invalid SNMPv3 Engine ID specified |
| <code>KRB_SNMPV3_ERROR_GENERIC</code> | 16 | Generic KRB SNMPv3 error |

The optional e-text field can be used for informational purposes (i.e., logging, network troubleshooting) and the optional e-data field is reserved for future use to transport any application data associated with a specific error.

6.6 End-to-end security for RTP

RTP security is currently fully specified in 7.6.2.1. Key management for RTP requires that both the (encryption) Transform ID and the Authentication Algorithm are specified, analogous to the IPsec key management. This clause lists in Tables 14 and 15, respectively, the Transform IDs and Authentication Algorithms that are available for RTP security.

Table 14 – RTP Packet Transform Identifiers

| Transform ID | Value | Key size (in bits) | MUST support | Description |
|-----------------|---------|--------------------|--------------|---|
| RTP_ENCR_NULL | 0x50 | NA | Yes | Encryption turned off |
| RTP_AES | 0x51 | 128 | Yes | AES-128 in CBC mode with 128-bit block size |
| RTP_XDESX_CBC | 0x53 | 192 | No | DESX-XEX-CBC |
| RTP_DES_CBC_PAD | 0x54 | 128 | No | DES-CBC-PAD |
| RTP_3DES_CBC | 0x56 | 128 | No | 3DES-EDE-CBC |
| Reserved | 0x57-59 | – | – | |

The RTP_AES and RTP_ENCR_NULL Transform IDs MUST be supported. AES-128 0 MUST be used in CBC mode with a 128-bit block size and an Initialization Vector (IV) generated in accordance with 7.6.2.1.2.2.2. AES-128 requires 10 rounds of cryptographic operations (see FIPS PUB 197).

Table 15 – RTP IP/Cablecom Authentication Algorithms

| Authentication Algorithm | Value | Key size (in bits) | MUST support | Description |
|--------------------------|-------|------------------------------|--------------|---------------------------|
| AUTH_NULL | 0x60 | 0 | Yes | Authentication turned off |
| Reserved | 0x61 | – | – | |
| RTP_MMH_2 | 0x62 | Variable (see 7.6.2.1.2.1.1) | Yes | 2-byte MMH MAC |
| Reserved | 0x63 | – | – | |
| RTP_MMH_4 | 0x64 | Variable (see 7.6.2.1.2.1.1) | Yes | 4-byte MMH MAC |
| Reserved | 0x65 | – | – | |

The Authentication Algorithms AUTH_NULL, RTP_MMH_2 and RTP_MMH_4 MUST be supported.

6.7 End-to-end security for RTCP

RTCP security is currently fully specified in 7.6.2.2. Key management for RTCP requires that both the (encryption) Transform ID and the Authentication Algorithm be specified. This clause lists in Tables 16 and 17, respectively, the Transform IDs and Authentication Algorithms that are available for RTCP security.

Table 16 – RTCP Packet Transform Identifiers

| Transform ID | Value | Key size (in bits) | MUST support | Description |
|----------------|---------|--------------------|--------------|---|
| RTCP_ENCR_NULL | 0x70 | – | Yes | Encryption turned off. |
| AES-CBC | 0x71 | 128 | Yes | AES-128 in CBC mode with 128-bit block size |
| XDESX-CBC | 0x72 | 192 | No | DESX-XEX-CBC |
| DES-CBC-PAD | 0x73 | 128 | No | DES-CBC-PAD |
| 3DES-CBC | 0x74 | 128 | No | 3DES-EDE-CBC |
| Reserved | 0x75-7f | – | – | |

The AES-CBC and RTCP_ENCR_NULL Transform IDs MUST be supported. AES-128 0 MUST be used in CBC mode with a 128-bit block size and a randomly generated Initialization Vector (IV). AES-128 requires 10 rounds of cryptographic operations (see FIPS PUB 197).

Table 17 – RTCP Authentication Algorithms

| Transform ID | Value | Key size (in bits) | MUST support | Description |
|----------------|---------|--------------------|--------------|--|
| RTCP_AUTH_NULL | 0x80 | N/A | Yes | Authentication turned off. |
| HMAC-SHA1-96 | 0x81 | 160 | Yes | First 12 bytes of the HMAC-SHA1 per RFC 2404 |
| HMAC-MD5-96 | 0x82 | 128 | No | First 12 bytes of the HMAC-MD5 per RFC 2403 |
| Reserved | 0x83-8f | – | – | |

The HMAC-SHA1-96 and RTCP_AUTH_NULL authentication algorithms MUST be supported.

6.8 BPI+

All E-MTAs and S-MTAs MUST use J.112-compliant cable modems and MUST implement BPI+ (ITU-T Rec. J.125). Baseline Privacy Plus (BPI+) provides security services to the J.112 data link layer traffic flows running across the cable access network, i.e., between CM and CMTS. These services are message confidentiality and access control. The BPI+ security services operating in conjunction with J.112 provide cable modem users with data privacy across the cable network and protect cable operators from theft of service.

The protected J.112 MAC data communication services fall into three categories:

- best-effort, high-speed, IP data services;
- QoS (e.g., constant bit rate) data services; and
- IP multicast group services.

When employing BPI+, the CMTS protects against unauthorized access to these data transport services by:

- 1) enforcing encryption of the associated traffic flows across the cable network; and
- 2) authenticating the J.112 MAC management messages that CMs use to establish QoS service flows.

BPI+ employs a client/server key management protocol in which the CMTS (the server) controls distribution of keying material to client CMs. The key management protocol ensures that only authorized CMs receive the encryption and authentication keys needed to access the protected services.

Baseline Privacy Plus has two component protocols:

- An encapsulation protocol for encrypting packet data across the cable network. This protocol defines:
 - 1) the frame format for carrying encrypted packet data within J.112 MAC frames;
 - 2) a set of supported *cryptographic suites*, i.e., pairings of data encryption and authentication algorithms; and
 - 3) the rules for applying those algorithms to a J.112 MAC frame's packet data.
- A key management protocol (Baseline Privacy Key Management, or "BPKM") provides the secure distribution of keying data from CMTS to CMs. Through this key management protocol, CM and CMTS synchronize keying data; in addition, the CMTS uses the protocol to enforce conditional access to network services.

Baseline Privacy Plus does not provide any security services beyond the J.112 cable access network. The majority of IPCablecom's signalling and media traffic flows, however, take paths that traverse the managed IP "backhaul" networks, which lie behind CMTSs. Since J.112 and IPCablecom service providers typically will not guarantee the security of their managed IP backhaul networks, the IPCablecom security architecture defines end-to-end security mechanisms for all these flows. End-to-end security is provided at the Network layer through IPsec, or, in the case of Client media flows, at the application/transport layer through RTP application layer security. Thus, IPCablecom does not rely on BPI+ to provide security services to its component protocol interfaces.

6.9 TLS

6.9.1 Overview

The TLS protocol (RFC 2246) provides privacy and data integrity over a reliable transport layer protocol such as TCP. The protocol is composed of two layers: the TLS Record Protocol and the TLS Handshake Protocol. The TLS Record Protocol is used to securely encapsulate upper layer protocols, while the TLS Handshake Protocol provides the key management functionality required to establish TLS sessions.

In IPCablecom, TLS is used to secure SIP-based signalling between SIP endpoints such as the CMS and EBPs.

6.9.2 IPCablecom profile for TLS with SIP

Unless specified within this Recommendation, IPCablecom SIP interfaces requiring TLS MUST be compliant with the TLS specification (RFC 2246) and any requirements specified in RFC 3261 relating to its usage in SIP.

TLS supports the negotiation and use of compression methods. However, since these methods are not specified within TLS, compression MUST NOT be used in IPCablecom.

6.9.2.1 TLS ciphersuites

In TLS, the ciphersuite includes the authenticated key agreement (AKE) method used in the TLS handshake, as well as encryption and authentication ciphers used to secure the record layer. Ciphersuites are negotiated with the TLS client presenting a list of supported ciphersuites in the Client Hello message, and the server responding with the selected ciphersuite in the Server Hello message.

Table 18 describes the TLS ciphersuites defined in RFC 2246 and RFC 3268 supported by IPCablecom:

Table 18 – TLS ciphersuites

| TLS Ciphersuite | Support | AKE method | Encryption | Auth. |
|-----------------------------------|----------------|--|-------------------|--------------|
| TLS_RSA_WITH_AES_128_CBC_SHA | MUST | RSA | AES-128 CBC | SHA |
| TLS_DHE_RSA_WITH_AES_128_CBC_SHA | MUST | Ephemeral Diffie-Hellman with RSA signatures | AES-128 CBC | SHA |
| TLS_RSA_WITH_3DES_EDE_CBC_SHA | SHOULD | RSA | 3DES CBC | SHA |
| TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA | SHOULD | Ephemeral Diffie-Hellman with RSA signatures | 3DES CBC | SHA |

6.9.2.2 IPCablecom TLS Certificates

TLS is a client-server based protocol with optional client authentication. However, in IPCablecom, mutual authentication using RSA-based certificates **MUST** be used. The TLS server **MUST** send a Certificate Request to the client. Both the TLS client and server certificates **MUST** conform to the IPCablecom Server Certificate as specified in 8.2.3.4.3.

IPCablecom Server Certificates include a server identifier (based on FQDN or IP address) embedded within the CN of the Subject Name field. Before accepting or continuing with a TLS connection, the TLS server or client **MUST** validate the remote server identifier to ensure it matches the IP address used for the TCP/TLS connection, in addition to any other local policy (i.e., provisioned list of allowed remote TLS endpoints based on FQDN or IP address).

In addition to the CableLabs Service Provider Root certificate, a TLS implementation **MAY** support a list of trusted CAs (Certificate Authorities) to facilitate inter-working between IPCablecom domains (i.e., between Server Providers).

6.9.2.3 Connection persistence and re-use

Since TCP connection and TLS session establishment (which relies on TCP) can be quite costly both in terms of performance and network latency, they are not suited for on-demand SIP signalling. As such, TLS sessions **SHOULD** be kept persistent as much as possible and SIP connection re-use **SHOULD** be supported.

6.9.2.4 Session caching

In TLS, it is possible to resume a previous session if it has been cached on both the TLS client and server. Resuming sessions drastically speeds up the session establishment, as fewer messages are exchanged and authentication is based on symmetric key cryptography.

In IPCablecom, TLS session caching **SHOULD** be supported. A TLS client initiating a TLS session **MUST** attempt to resume a cached session if it has retained a session for the remote server. The duration for which a TLS client or server must retain a cached session is a local policy and implementation specific.

7 Security profile

The IPCablecom architecture defines over half a dozen networked components and the protocol interfaces between them. These networked components include the media terminal adapter (MTA), call management server (CMS), signalling gateway (SG), media gateway (MG) and a variety of OSS systems (DHCP, TFTP and DNS servers, network management systems, provisioning servers, etc.). IPCablecom security addresses the security requirements of each constituent protocol interface by:

- identifying the threat model specific to each constituent protocol interface;
- identifying the security services (authentication, authorization, confidentiality, integrity, non-repudiation) required to address the identified threats;
- for each constituent protocol interface, specifying the particular security mechanism providing the required security services.

Clause 5.2 summarizes the threat models applicable to IPCablecom's protocol interfaces. In this clause, we identify the security service requirements of each protocol interface and security mechanisms providing those services.

The security mechanisms include both the security protocol (e.g., IPsec, RTP-layer security, SNMPv3 security) and the supporting key management protocol (e.g., IKE, PKINIT/Kerberos).

The security analysis in 5.3.3 is organized by functional categories. For each functional category, we identify the constituent protocol interfaces, the security services required by each interface, and the particular security mechanism employed to deliver those security services. Each per-protocol security description includes the detailed information sufficient to ensure interoperability. This includes cryptographic algorithms and cryptographic parameters (e.g., key lengths).

As a convenient reference, each functional category's security analysis includes a summary security profile matrix of the following form (Media security profile matrix shown) in Table 19:

Table 19 – RTP – RTCP Security Profile Matrix

| | RTP (MTA-MTA, MTA-PSTN GW) | RTCP (MTA-MTA, MTA-MG, MG-MG) |
|---------------------|---|--|
| Authentication | Optional (indirect) | Optional (indirect) |
| Access control | Optional | Optional |
| Integrity | Optional | Yes |
| Confidentiality | Yes | Yes |
| Non-repudiation | No | No |
| Security mechanisms | Application Layer Security via RTP IPCablecom Security Profile keys distributed over secured MTA-CMS links AES-128 encryption algorithm Optional 2-byte or 4-byte MAC based on MMH algorithm IPCablecom supports ciphersuite negotiation. | Application Layer Security via RTCP IPCablecom Security Profile keys distributed over secured MTA-CMS links RTCP ciphersuites are negotiated separately from the RTP ciphersuites and include both encryption and message authentication algorithms. Keys are derived from the end-end secret using the same mechanism as used for RTP encryption. |

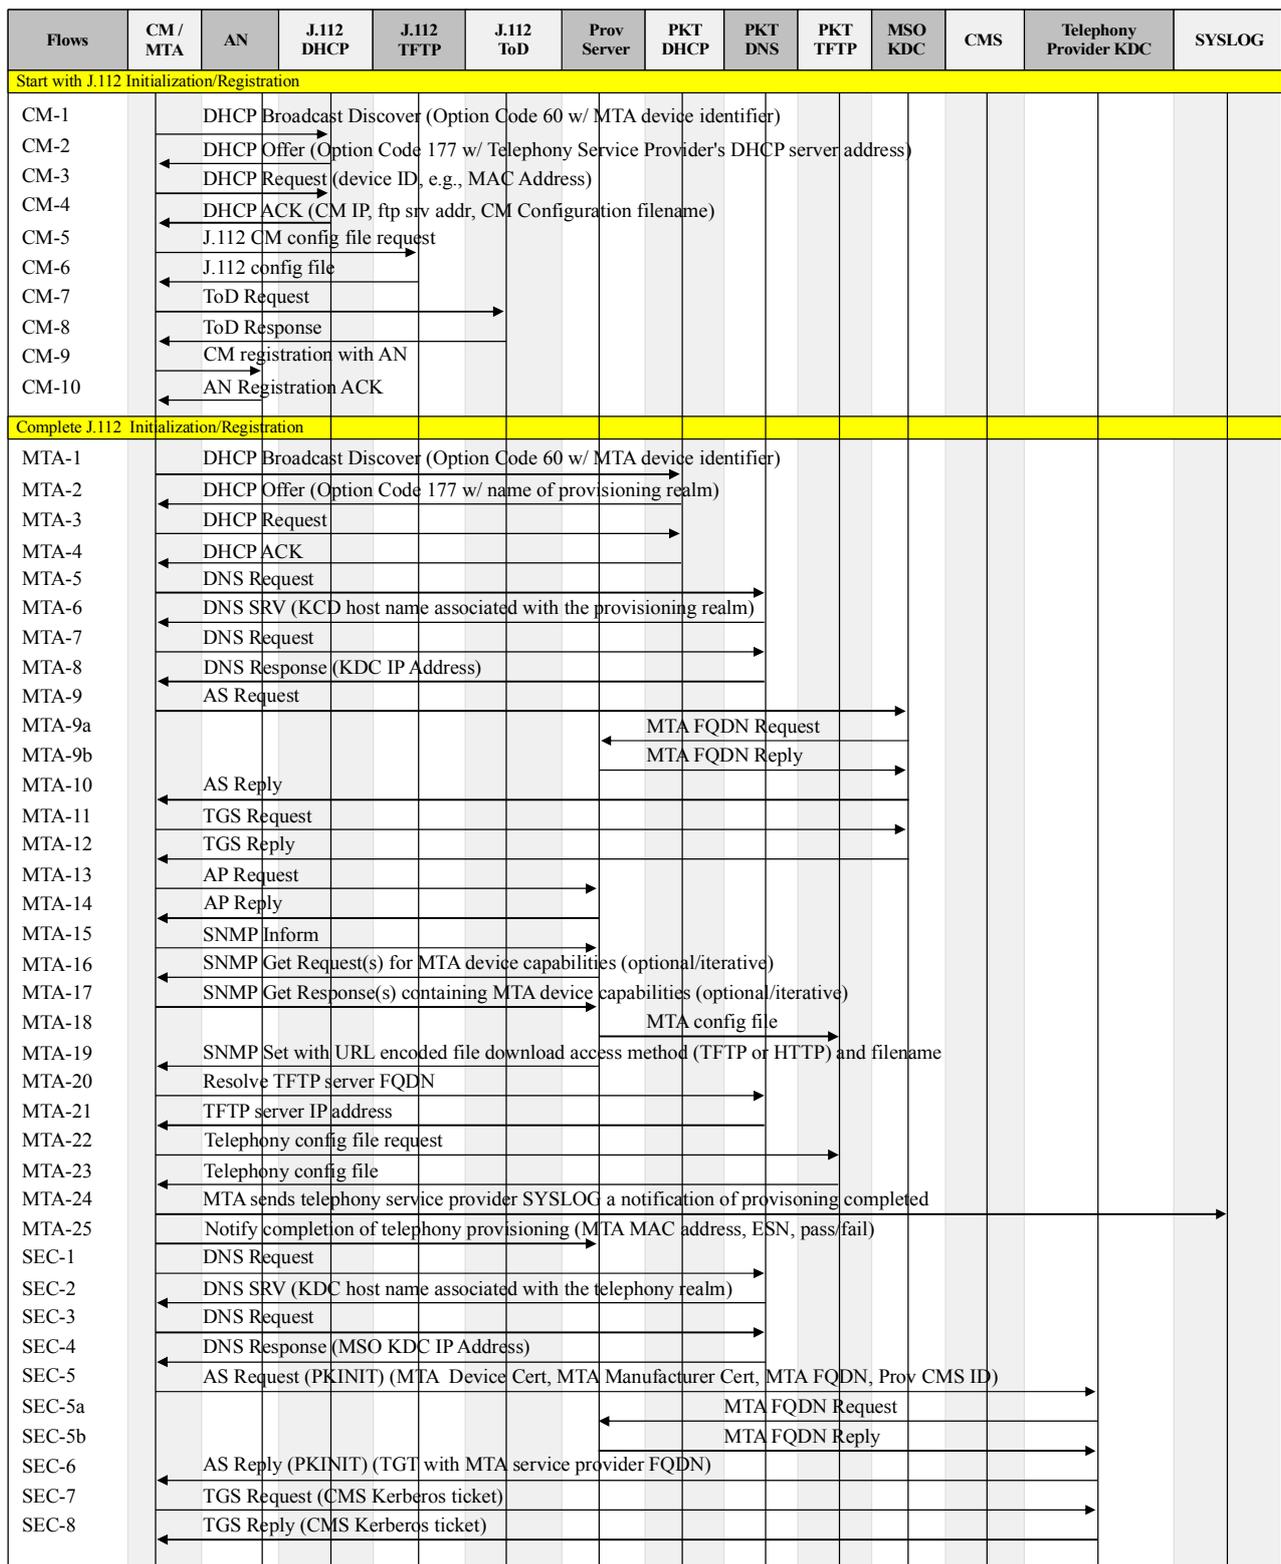
Each matrix column corresponds to a particular protocol interface. All but the last row corresponds to a particular security service; the cell contents in these rows indicate whether the protocol interface requires the corresponding security service. The final row summarizes the security mechanisms selected to provide the required services.

Note that the protocol interface column headings not only identify the protocol, but also indicate the network components the protocols run between.

7.1 Device and service provisioning

Device provisioning is the process by which an MTA is configured to support voice communications service. The MTA provisioning process is specified in ITU-T Rec. J.167.

Figure 10 illustrates only the flows involved with the Secure provisioning processes. The provisioning Recommendation lays out in detail these Secure Provisioning flows along with two non-secure MTA provisioning flows called Basic and Hybrid. The Secure Provisioning flows involving security mechanisms are described in this clause. Refer to ITU-T Rec. J.167 for the non-secure flows.



J.170_F10

Figure 10 – IPCablecom provisioning flows

As part of the provisioning process, the MTA performs Kerberos key management (AS Request/AS Reply and AP Request/AP Reply, and optional TGS Request/TGS Reply).

Table 20 describes the execution of the Kerberos key management step during MTA provisioning:

Table 20 – Kerberos key management during MTA provisioning

| Flow step | Security requirement | Lifetime | Step bypass permitted |
|---|--|---|--|
| <i>MTA-9/MTA-10</i> AS Request/AS Response (see 6.4.1) | TGT ticket if using TGS Request; Provisioning Server Ticket if otherwise | Max. 7 days | This step MUST NOT be performed if the MTA already possesses a valid ticket for the Provisioning Server. |
| <i>MTA-9a/MTA-9b –</i> MTA FQDN Request/MTA FQDN Reply (see 6.4.7) | MTA FQDN Request and Reply are protected using Kerberos tickets | | These steps will not occur if MTA-9 is skipped. Otherwise, this step cannot be bypassed. |
| <i>MTA-11/MTA 12</i> TGS Request/TGS Response (see 6.4.4) | Applies when a TGT is used. Obtains a Provisioning Server Ticket. | Lifetime is set to expire no later than the expiration time of the TGT ticket. | This step MUST NOT be performed if the MTA already possesses a valid ticket for the Provisioning Server. |
| <i>MTA-13/MTA-14</i> AP Request/AP Reply (see 6.5.2 and 6.5.4) | Initial SNMPv3 authentication and privacy keys for the MTA. The user name for the MTA is specified as "MTA-Prov- xx:xx:xx:xx:xx:xx", where xx:xx:xx:xx:xx:xx represents the MAC address of the MTA. AP Req/AP Rep messages do not specify the SNMPv3 key expiration time in the protocol, but the SNMP Manager may still set up expiration time locally; after the keys expire, the manager can send a Wake Up message to create a new set of SNMPv3 keys. | Expiration is not supported by IPCablecom. | None. New SNMPv3 keys and User Ids are created each time the MTA is re-initialized. It is assumed that SNMPv3 keys and User IDs are not saved in NVRAM. Also note that this step is used for Engine ID determination and SNMPv3 time synchronization – the two sides exchange initial values for SNMPv3 boots and engine time parameters. |

An MTA MUST get a new ticket before performing Kerberized Key Management with a particular Application Server if the ticket(s) it currently possesses is not valid. A ticket would no longer be valid if the KDC REALM or Application Server FQDN changes, if the MTA's IP address has changed, or if the current time, adjusted by the time offset for that REALM or Application Server, does not fall within the ticket validity period.

The PKINIT_{GP} for the Provisioning Server's realm is specified in the MTA MIB inside the realm table. When the MTA implementation requests a TGT in an AS Request and when the MTA needs to obtain tickets for one or more CMSs in the same realm as the Provisioning Server, the PKINIT_{GP} value specified in the MIB MUST be used to refresh the TGT. In all other cases, the AS Request for the TGT in the Provisioning Server's realm or for the Provisioning Server's ticket directly MAY be issued on demand.

The TGS Grace Period is not specified for the key management between the MTA and the Provisioning Server. The TGS Request for the Provisioning Server's ticket MAY be issued on demand.

7.1.1 Device provisioning

Device provisioning occurs when an MTA device is inserted into the network. A provisioned MTA device that is not yet associated with a billing record MAY have minimal voice communications service available.

Device provisioning involves the MTA making itself visible to the network, obtaining its IP configuration and downloading its configuration data.

The IPCablecom architecture supports three provisioning flows:

- Basic Flow;
- Hybrid Flow;
- Secure Flow.

The Basic and Hybrid Flows are completely insecure flows (i.e., there are no mechanisms in the flows that would prevent a user from provisioning their own MTA). The Basic and Hybrid Flows also do not provide a means to secure the SNMP management interface on the MTA. Service providers that choose to deploy MTAs with one of these insecure flows must accept that there are security risks. For example, a Denial-of-Service attack could be mounted by sending SNMP TRAPs and INFORMs to the operator's management system. The management system would have to process them, even though they are unauthenticated. Unfortunately, the inclusion of these insecure flows also poses security risks for Service Providers that choose to deploy MTAs with the Secure Flow.

MTAs that support the insecure flows may be provisioned by a user, even if the service provider is using the Secure Flows. Unauthorized provisioning of an MTA allows a user to provide their own configuration file. The MTA could then be used to communicate normally with a CMS. Alternatively, un-authorized provisioning of an MTA could be used to bypass service provider controls on secure software download (in the case of the S-MTA) and provide a software image that has some perceived value (such as a security vulnerability).

With respect to the Secure Flow, support for SNMPv2c coexistence for network management operations also introduces vulnerabilities to service providers that use the Secure Flow (unauthenticated TRAPs and INFORMs could be sent). The best way to address these vulnerabilities is to disable SNMPv2c coexistence.

Therefore it is recommended, as always, that service providers use multiple layers of security to ensure that their CMSs and back-office systems are protected against rogue MTAs.

7.1.1.1 Security services

7.1.1.1.1 MTA-DHCP server

Authentication and Message Integrity is desirable on this interface, in order to prevent denial-of-service attacks that cause an MTA to be improperly configured. Securing DHCP is considered an operational issue to be evaluated by each network operator. It is possible to use access control through the local DHCP relay inside the local loop. IPsec can be used for security between the DHCP relay and the DHCP server.

7.1.1.1.2 MTA-SNMP manager

This clause applies to all SNMPv3 messages between the MTA and an SNMPv3 Manager. Within the IPCablecom architecture, the Provisioning Server includes the SNMPv3 Manager function, although SNMPv3 traffic occurs both during and after the provisioning phase.

Authentication: the identity of the MTA that is sending configuration parameters and faults to the SNMP manager must be authenticated, to prevent denial of service attacks. For example, the Provisioning Server may be tricked into continuously creating bogus configuration files or into creating a configuration file based on incorrect MTA capabilities that in effect disable that MTA.

Also, during the provisioning sequence the MTA is told (via an SNMP Set) the parameters needed to find, authenticate and decrypt its configuration file. If this SNMP Set were forged, it would disrupt the MTA provisioning sequence.

Message Integrity: required to prevent denial-of-service attacks at the OSS and at the MTA – see the above description of the denial-of-service attacks under authentication.

Confidentiality: may be used to protect sensitive MTA configuration data. IPCablecom currently does not specify any such sensitive MTA parameters and so confidentiality is optional.

Access Control: write access to the MTA configuration parameters must be allowed only to the authorized OSS users, to prevent denial of service/misconfiguration attacks. Read access can be enforced in conjunction with confidentiality, which is optional (see above on confidentiality).

Note that DHCP is used to configure the MTA with the Kerberos realm name, which points it to a particular KDC. DHCP also configures the MTA with the location of the Provisioning Server. Since IPCablecom currently does not specify DHCP security, by faking DHCP responses it is possible to point MTAs to a wrong Provisioning Server and to a wrong KDC that permits security establishment with that Provisioning Server. (The MTA would only authenticate that wrong KDC if the CableLabs Service Provider Root CA signed the KDC certificate.) So, it is possible to bypass access control, but the attack has to be orchestrated by another MSO that had also been certified by IPCablecom.

7.1.1.1.3 MTA-Provisioning Server, via TFTP Server

Authentication: required to prevent denial-of-service attacks, that cause an MTA to be improperly configured.

Message integrity: required to prevent denial-of-service attacks that cause an MTA to be either improperly configured or configured with old configuration data that was replayed.

Confidentiality: optional; it is up to the Provisioning Server to decide whether or not to encrypt the file.

Access control: not required at the TFTP Server. If needed, MTA configuration file is encrypted with the Provisioning Server-MTA shared key.

Non-repudiation: is not required.

7.1.1.2 Cryptographic mechanisms

7.1.1.2.1 Call flows MTA-15, 16, 17: MTA-SNMP Manager: SNMP Inform/Get Requests/Responses

All SNMP traffic between the MTA and the SNMP Manager in both directions is protected with SNMPv3 security per RFC 3414 during the Secure Provisioning process. IPCablecom requires that SNMPv3 message authentication is always turned on with privacy being optional (see 6.3). The only SNMPv3 encryption algorithm is currently DES-CBC. This is the limitation of the SNMPv3 IETF standard, although stronger encryption algorithms are desirable. See 6.3. for the list of SNMPv3 cryptographic algorithms supported by IPCablecom.

7.1.1.2.2 Call Flow MTA-18: Provisioning Server-TFTP Server: Create MTA Configuration file

This clause describes the MTA Config file creation in the Secure Provisioning Flow. In this flow, the Provisioning Server builds a MTA device configuration file. This file **MUST** contain the following configuration info for each endpoint (port) in the MTA:

- CMS name (FQDN format);
- Kerberos realm for this CMS;
- Telephony Service Provider Organization Name;
- PKINIT Grace Period.

This file **MUST** be authenticated and **MAY** be encrypted. If the configuration file is encrypted, then the SNMPv3 privacy **MUST** be used in order to transport the configuration file encryption key securely. Once the Provisioning Server builds the configuration file, it will perform the following steps:

- 1) The provisioning Server decides to encrypt the file, it creates a configuration file encryption key and encrypts the file with this key. The encryption algorithm **MUST** be the same as the one that is used for SNMPv3 privacy. It then stores the key and the cipher. The file **MUST** be encrypted using the following procedure:
 - a) prepend the file contents with a random byte sequence, called a confounder. The size of the confounder **MUST** be the same as the block size for the encryption algorithm. In the case of DES it is 8 bytes;
 - b) append random padding to the result in a). The output of this step is of length that is a multiple of the block size for the encryption algorithm;
 - c) encrypt the result in b) using IV = 0. The output of this step is the encrypted configuration file.
- 2) It creates a SHA-1 hash of the configuration file and stores it. If the file was encrypted, the hash is taken over the encrypted file.
- 3) It sends the following items to the MTA in the SNMP SET in the flow MTA-19.
 - a) `pktcMtaDevConfigKey`, which is the configuration file encryption key MIB variable generated in step 1.
 - b) `pktcMtaDevConfigHash`, which is the SHA-1 of the configuration file MIB variable generated in step 2.
 - c) Name and location of the configuration file.

Steps 1 and 2 **MUST** occur only when a configuration file is created or an existing file is modified. If the `pktcMtaDevConfigKey` is set, then the MTA **MUST** use this key to decrypt the configuration file. Otherwise, MTA **MUST** assume that the file is not encrypted. SNMPv3 provides authentication when the `pktcMtaDevConfigHash` is set and therefore the configuration file is authenticated indirectly via SNMPv3.

In the event that SNMPv3 privacy is selected during the key management phase, but is using a different algorithm than the one that was selected to encrypt the configuration file (or the configuration file was previously in the clear), the configuration file **MUST** be re-encrypted and the TFTP server directory **MUST** be updated with the new file. Similarly, if the Provisioning Server decides not to encrypt the file this time, after it was previously encrypted, the TFTP server directory **MUST** be updated with the new file.

MTA endpoints **MAY** also be configured for IP Telephony service while the MTA is operational. In that case the same information that is normally assigned to an endpoint in a configuration file **MUST** be assigned with SNMP Set commands.

7.1.1.2.3 Call flows MTA-19, 20 and 21: Establish TFTP server location

This set of call flows is used to establish the IP address of the TFTP server from where the MTA will retrieve its configuration file. Although flow MTA-19 is authenticated via SNMPv3, MTA-20 and 21 are not authenticated.

Flow MTA-21 allows for denial-of-service attacks, where the MTA is pointed to a wrong TFTP server (IP address). The MTA cannot be fooled in accepting the wrong configuration file since checking the hash of the file authenticates the file – this denial-of-service attack will result in failed MTA provisioning.

The denial-of-service threats, where responses to DNS queries are forged, are currently not addressed by IPCablecom. It is mainly because DNS security (DNSSEC) is not yet available as a commercial product and would cause significant operational difficulty in the conversion of the DNS databases.

7.1.1.2.4 Call flows MTA-22, 23: MTA-TFTP Server: TFTP Get/Get Response

The TFTP get request is not authenticated and thus anyone can request an MTA configuration file. This file does not contain any sensitive data and may be encrypted with the Provisioning Server-MTA shared key if the Provisioning Server chooses to. In this case no one except the MTA can make use of this file.

This flow is open for a denial-of-service attack, where the TFTP server is made busy with useless TFTP get requests. This denial-of-service attack is not addressed at this time.

The TFTP get response retrieves a configuration file from the TFTP server. The contents of the configuration file are listed in 7.1.1.2.2.

7.1.1.2.5 Security flows

For each CMS specified in the `pktcMtaDevCmsTable` table with `pktcMtaDevCmsIpsecCtrl` value set to TRUE and assigned to a provisioned MTA endpoint, the MTA MUST perform the following security flows in Table 21 after the provisioning process and prior to any NCS message exchange. For each CMS specified in `pktcMtaDevCmsTable` with `pktcMtaDevCmsIpsecCtrl` set to FALSE, the MTA MUST NOT perform the following flows and MUST send and receive NCS messages without IPsec (i.e., NCS packets are sent in the "clear").

Table 21 – Post-MTA provisioning security flows

| Security flow | Flow description | If step fails, proceed here |
|--|---|-----------------------------|
| Get Kerberos tickets associated with each CMS with which the MTA communicates. | | |
| SEC-1 | <i>DNS SRV Request</i> The MTA requests the Telephony KDC host name for the Kerberos realm. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS. | SEC-1 |
| SEC-2 | <i>DNS SRV Reply</i> Returns the Telephony KDC host name associated with the provisioning realm. If the KDC's IP Address is included in the Reply, proceed to SEC-5. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS. | SEC-1 |

Table 21 – Post-MTA provisioning security flows

| Security flow | Flow description | If step fails, proceed here |
|---------------|--|---|
| SEC-3 | <p><i>DNS Request</i></p> <p>The MTA now requests the IP Address of the Telephony KDC. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | SEC-1 |
| SEC-4 | <p><i>DNS Reply</i></p> <p>The DNS Server returns the IP Address of the Telephony KDC. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | SEC-1 |
| SEC-5 | <p><i>AS Request</i></p> <p>For each different CMS assigned to voice communications endpoints, the MTA requests a TGT or a Kerberos ticket for the CMS by sending a PKINIT REQUEST message to the KDC containing the MTA Device Certificate and the MTA FQDN. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | Report alarm. Abort establishment of signalling security. |
| SEC-5a | <p><i>MTA FQDN Request</i></p> <p>The KDC requests the MTA's FQDN from the Provisioning Server.</p> <p>This step will not occur if the MTA skips SEC-5.</p> | |
| SEC-5b | <p><i>MTA FQDN Reply</i></p> <p>The Provisioning Server replies to the KDC request with the MTA's FQDN.</p> <p>This step will not occur if the MTA skips SEC-5.</p> | |
| SEC-6 | <p><i>AS Reply</i></p> <p>The KDC sends the MTA a PKINIT REPLY message containing the requested Kerberos ticket. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | Proceed to SEC-5 or abort signalling security depending upon error conditions. |
| SEC-7 | <p><i>TGS Request</i></p> <p>In the case where the MTA obtained a TGT in SEC-6, it now obtains the Kerberos ticket for the TGS request message. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | Report alarm. Abort establishment of signalling security. |
| SEC-8 | <p><i>TGS Reply</i></p> <p>Response to TGS Request containing the requested CMS Kerberos ticket. This step MUST NOT be performed if the MTA already possesses a valid ticket for the CMS.</p> | Proceed here to SEC-7/SEC-5 or abort signalling security depending upon error conditions. |

Table 21 – Post-MTA provisioning security flows

| Security flow | Flow description | If step fails, proceed here |
|---------------|---|--|
| SEC-9 | <p><i>AP Request</i></p> <p>The MTA requests a pair of IPsec simplex Security Associations (inbound and outbound) with the assigned CMS by sending the assigned CMS an AP REQUEST message containing the CMS Kerberos ticket.</p> | <p>Report alarm. Abort establishment of signalling security.</p> |
| SEC-10 | <p><i>AP Reply</i></p> <p>The CMS establishes the Security Associations and then sends an AP REPLY message with the corresponding IPsec parameters. The MTA derives IPsec keys from the subkey in the AP Reply and establishes IPsec SAs.</p> | <p>Proceed here to SEC-9/SEC-7/SEC-5 or abort signalling security depending upon error conditions.</p> |

Several tables in the MTA MIB are used to control security flows SEC-1 through SEC-10 (see Table 21).

The CMS table (pktcMtaDevCmsTable) and the realm table (pktcMtaDevRealmTable) are used for managing the MTA security signalling. The realm table defines the domains for the CMSs. The CMS table defines the CMSs within the domains. An endpoint is associated with one CMS at any given time. The following restrictions MUST be adhered to:

- a) The realm table in the configuration file MUST at a minimum include an entry for the realm that is identified in DHCP option 122, suboption 6.
- b) There MUST be a realm table entry for each CMS table entry. Multiple CMS table entries MAY utilize the same realm table entry.
- c) Each MTA endpoint defined in the NCS endpoint table (pktcNcsEndPntConfigTable) MUST be configured with a CMS FQDN (pktcNcsEndPntConfigCallAdgentId) that is also present in the CMS table (pktcMtaDevCmsFqdn).
- d) All members of a CMS cluster defined by the same FQDN MUST use the same configuration for establishing security associations as defined in the pktcMtaDevCmsTable.
- e) If NCS signalling selects a CMS (with an N: parameter selection) that is not defined by an entry in the CMS table, the same realm and CMS parameters, with the exception of the CMS FQDN and pktcMtaDevCmsIpsecCtrl, are used as defined in the current CMS table entry. The pktcMtaDevCmsIpsecCtrl flag for the new CMS MUST be set to TRUE.

The use of the security-relevant MIB tables immediately following step MTA-25 is as follows:

- 1) The MTA finds a list of CMSs with which it needs to establish IPsec SAs. This list MUST include every CMS that is assigned to a configured endpoint, as specified by the NCS MIB table pktcNcsEndPointConfigTable. This list of CMSs MUST include only CMSs that are listed in the pktcMtaDevCmsTable.
- 2) For each CMS in the above list, the MTA MUST attempt to establish IPsec security associations as follows:
 - a) Find the corresponding CMS table entry.
 - b) If the MTA does not already possess a valid ticket for the specified CMS, use the pktcMtaDevCmsKerbRealmName parameter in the CMS table entry to index into the pktcMtaDevRealmTable. Then, using the parameters associated with that realm, perform steps SEC-1 through SEC-6 and optionally SEC-7 and SEC-8 in order to obtain the desired CMS ticket.

- c) Perform IPsec key management according to flows SEC-9 and SEC-10. This step MAY occur at any time after step b above, but it must occur before any signalling messages are exchanged with that CMS.

The CMS table entry contains various timing parameters used in steps SEC-9 and SEC-10. In the case of time-outs or other errors, the MTA may retry using the timing parameters specified in the CMS table entry.

The above steps MUST also apply when an additional MTA endpoint is activated (see ITU-T Rec. J.167) or when an endpoint is configured (via SNMP sets) for a new CMS in the NCS MIB (see ITU-T Rec. J.166).

- 3) Any time before an MTA endpoint sends an signalling message to a particular CMS, it MUST ensure that the respective security association is present. If the MTA is unable to establish IPsec SAs with a CMS that is associated with a configured endpoint (by the NCS MIB), it MUST set the NCS MIB variable `pktcNcsEndPntStatusError` to `noSecurityAssociation` (2).

After the initial establishment of the IPsec security associations for CMSs, the MTA MIB is utilized in subsequent key management as follows:

When the MTA receives a Wake Up message, it MUST respond with an AP Request when the corresponding CMS FQDN is found in the `pktcMtaDevCmsTable` and MUST NOT respond otherwise.

Note that establishment of IPsec security associations due to a Wake Up does not result in any call signalling traffic between the MTA and the CMS.

7.1.1.2.5.1 Call flows SEC-5, 6: Get a Kerberos ticket for the CMS

The MTA uses PKINIT protocol to get a Kerberos ticket for the specified CMS (see 6.4.3). After the KDC receives a ticket request, it retrieves the MTA FQDN from the provisioning server so that it can verify the request before replying with a ticket. The Telephony KDC issues the Kerberos ticket for a group of one or more CMSs uniquely identified with the pair (Kerberos Realm, CMS Principal Name).

In the event that different MTA ports are configured for a different group of CMSs, the MTA MUST obtain multiple Kerberos tickets by repeating these call flows for each CMS. Note that there is no requirement that the MTA obtain all the tickets from a single KDC.

7.1.1.2.5.2 Call flows SEC-7, 8, 9: Establish IPsec SAs with the CMS

The MTA uses the Kerberos ticket to establish a pair of simplex IPsec Security Associations with the given CMS. In the event that different MTA ports are configured with different CMS (FQDN) names, multiple pairs of SAs will be established (one set for each CMS).

When a single Kerberos ticket is issued for clustered Call Agents, it is used to establish more than one pair of IPsec SAs.

A CMS FQDN MAY translate into a list of multiple IP addresses, as would be the case with the NCS clustered Call Agents. In those cases, the MTA MUST initiate Kerberized key management with one of the IP addresses returned by the DNS Server. The MTA MAY also establish SAs with the additional CMS IP addresses.

Additional IPsec SAs with the other IP addresses MAY be established later, as needed (e.g., the current CMS IP address does not respond).

7.1.1.3 Key management

7.1.1.3.1 MTA-SNMP manager

Key management for the MTA-Provisioning SNMPv3 user **MUST** use the Kerberized key management protocol as it is specified in 6.5.4. The MTA and the Provisioning Server **MUST** support this key management protocol. Additional SNMPv3 users **MAY** be created with the standard SNMPv3 cloning method per RFC 3414 or with the same Kerberized key management protocol.

In order to perform Kerberized key management, the MTA must first locate the KDC. It retrieves the provisioning realm name from DHCP and then uses a DNS SRV record lookup to find the KDC FQDN(s) based on the realm name (see 6.4.5.1). When there is more than one KDC (DNS SRV record) found, DNS assigns a priority (and possibly a weighting) to each one. The MTA will choose a KDC based on the DNS priority and weight labelling and will go through the list until it finds a KDC that is able to respond.

7.1.1.3.2 MTA-TFTP server

The optional encryption key for the MTA configuration file is passed to the MTA with an SNMP Set command (by the Provisioning Server) shown in the provisioning flow MTA-19. SNMPv3 security is utilized to provide message integrity and privacy. In the event that SNMPv3 privacy is not enabled, the MTA configuration file **MUST NOT** be encrypted and the file encryption key **MUST NOT** be passed to the MTA.

The encryption algorithm used to encrypt the file **MUST** be the same as the one used for SNMPv3 privacy. The same file encryption key **MAY** be re-used on the same configuration file while the MTA configuration file contents are unchanged. However, if the MTA configuration file changes or if a different encryption algorithm is selected for SNMPv3 privacy, the Provisioning Server **MUST** generate a new encryption key, **MUST** re-encrypt the configuration file and **MUST** update the TFTP server with the re-encrypted file.

7.1.1.4 MTA embedded keys

The MTA device **MUST** be manufactured with a public/private RSA key pair and an X.509 device certificate that **MUST** be different from the BPI+ device certificate.

7.1.1.5 Summary security profile matrix – Device provisioning

The following matrix in Table 22 applies only to the Secure Provisioning Flow and SNMPv3.

Table 22 – Security profile matrix – MTA device provisioning

| | SNMP | TFTP (MTA-TFTP server) |
|---------------------|---|--|
| Authentication | Yes | Yes: authentication of source of configuration data. |
| Access control | Yes: write access to MTA configuration is limited to authorized SNMP users. Read access can also be limited to the valid users when confidentiality is enabled. | Yes: write access to the TFTP server must be limited to the Provisioning Server but is out of scope for IPCablecom. Read access can be optionally indirectly enabled when the MTA configuration file is encrypted. |
| Integrity | Yes | Yes |
| Confidentiality | Optional | Optional (of MTA configuration information during the TFTP-get). |
| Non-repudiation | No | No |
| Security mechanisms | SNMPv3 authentication and privacy. Kerberized key management protocol defined by IPCablecom. | Hash of the MTA configuration file is sent to the MTA over SNMPv3, providing file authentication. When the file is encrypted, the key is also sent to the MTA over SNMPv3 (with SNMPv3 encryption turned on). |

7.1.2 Subscriber enrolment

The subscriber enrolment process establishes a permanent customer billing account that uniquely identifies the MTA to the CMS via the endpoint ID, which contains the MTA's FQDN. The billing account is also used to identify the services subscribed to by the customer for the MTA.

Subscriber enrolment MAY occur in-band or out-of-band. The actual specification of the subscriber enrolment process is out of scope for IPCablecom and may be different for each Service Provider. The device provisioning procedure described in the previous clause allows the MTA to establish IPsec Security Associations with one or more Call Agents, regardless of whether or not the corresponding subscriber had been enrolled.

As a result, when subscriber enrolment is performed in-band, a communication to a Customer Service Record (CSR) (or to an automated subscriber enrolment system) is protected using the same security mechanisms that are used to secure all other voice communication.

During each communication set-up (protected with IPsec ESP), the CMS MUST check the identity of an MTA against its authorization database to validate which voice communications services are permitted. If that MTA does not yet correspond to an enrolled subscriber, it will be restricted to permitting a customer to contact the service provider to establish service ("customer enrolment"). Some additional services, such as communications with emergency response organizations (e.g., 911), may also be permitted in this case. Since in-band customer enrolment is based on standard security provided for call signalling and media streams, no further details are provided in this clause. Refer to 7.6 and 6.6 on media streams.

7.2 Quality of Service (QoS) signalling

7.2.1 Dynamic Quality of Service (DQoS)

7.2.1.1 Reference architecture for embedded MTAs

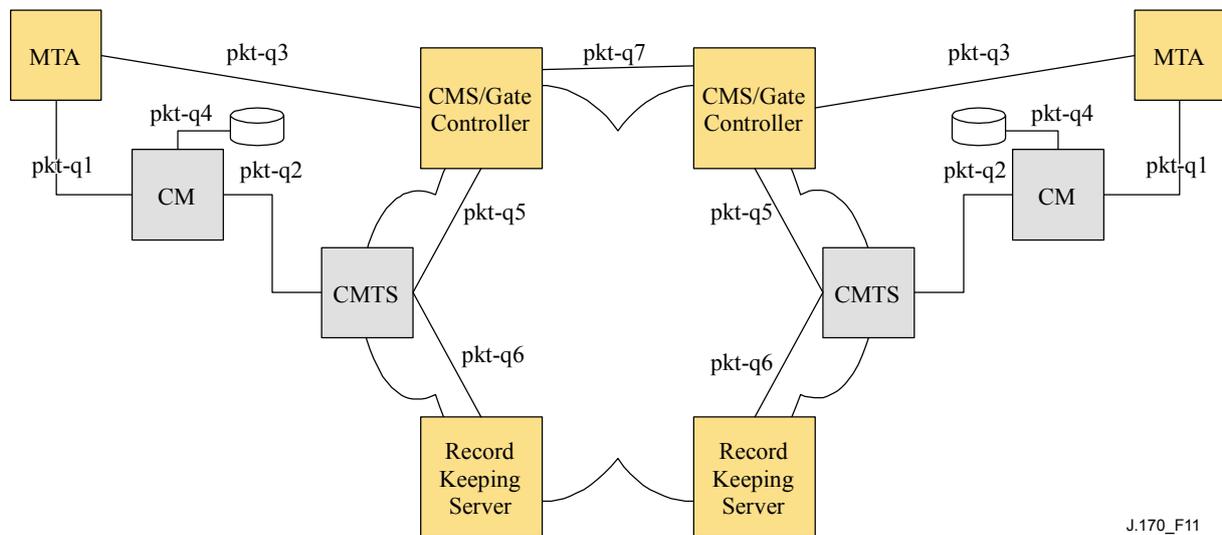


Figure 11 – QoS Signalling interfaces in the IPCablecom network

7.2.1.2 Security services

7.2.1.2.1 CM-CMTS J.112 QoS messages

Refer to ITU-T Rec. J.112.

7.2.1.2.2 Gate Controller – CMTS COPS messages

Authentication, access control and message integrity: required to prevent QoS theft and denial-of-service attacks.

Confidentiality: required to keep customer information private.

7.2.1.3 Cryptographic mechanisms

7.2.1.3.1 CM-CMTS J.112 QoS messages

The J.112 QoS messages are specified in the radio-frequency interface (RFI) specification in ITU-T Rec. J.112.

7.2.1.3.1.1 QoS service flow

A service flow is a J.112 MAC-layer transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CM or to downstream packets transmitted by the CMTS. A service flow is characterized by a set of QoS Parameters such as latency, jitter, and throughput assurances. In order to standardize operation between the CM and CMTS, these attributes include details of how the CM requests mini-slots and the expected behavior of the CMTS upstream scheduler.

ITU-T J.112 defines a Classifier, which consists of some packet matching criteria (IP source address, for example), a classifier priority, and a reference to a service flow. If a packet matches the specified packet matching criteria, it is then delivered on the referenced service flow.

Downstream Classifiers are applied by the CMTS to packets it is transmitting, and Upstream Classifiers are applied at the CM and may be applied at the CMTS to police the classification of upstream packets.

The network can be vulnerable to IP packet attacks; i.e., attacks stemming from an attacker using another MTA's IP source address and flooding the network with the packets intended for another MTA's destination address. A CMTS controlling downstream service flows will limit an MTA's downstream bandwidth according to QoS allocations. If the CMTS is flooded from the backbone network with extra packets intended for one of its MTAs, packets for that MTA may be dropped to limit the downstream packet rate to its QoS allocation. The influx of the attacker's packets may result in the dropping of good packets intended for the destination MTA.

To thwart this type of network attack, access to the backbone network should be controlled at the entry point. This can be accomplished using a variety of QoS classifiers, but is most effective when the packet source is verified by its source IP address. This will limit the ability of a rogue source to flood the network with unauthorized IP packets.

To address J.112 CMTS accesses to the network, the CMTS SHOULD apply upstream classifiers to police upstream packets from its network, including the verification of the source IP address.

For more information regarding the use of packet classifiers, refer to the J.112 CMTS-CMS Gate Coordination Messages (over UDP).

7.2.1.3.2 Gate Controller – CMTS COPS messages

To download a QoS policy for a particular communications connection, the Gate Controller function in the CMS MUST send COPS messages to the CMTS. These COPS messages MUST be both authenticated and encrypted with IPsec ESP. Refer to 6.1 on the details of how IPsec ESP is used within IPCablecom and for the list of available ciphersuites.

7.2.1.4 Key management

7.2.1.4.1 Gate Controller – CMTS COPS messages

Key management for this COPS interface is either IKE or Kerberos. Implementations MUST support IKE with pre-shared keys. Implementations MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys. For more information on the IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

When the Gate Controller detects a failure of all COPS connections associated with a particular outgoing IPsec SA, it MUST delete all associated SAs (IKE and IPsec SAs if IKE is used as the Key management protocol or only IPsec SAs if Kerberos is used as the Key management protocol).

Subsequently, every N times ($1 \leq N \leq 10$) that the Gate Controller tries to recover the connection, the SAs MUST be removed.

7.2.1.4.2 Security profile matrix summary

Table 23 – Security profile matrix – DQoS

| | COPS (CMTS-CMS) |
|---------------------|---|
| Authentication | Yes |
| Access control | Yes |
| Integrity | Yes |
| Confidentiality | Yes |
| Non-repudiation | No |
| Security mechanisms | IPsec with encryption and message integrity. IKE or Kerberos |

7.3 Billing system interfaces

7.3.1 Security services

7.3.1.1 CMS-RKS interface

Authentication, access control and message integrity: required to prevent service theft and denial-of-service attacks. Want to ensure that the billing events reported to the RKS are not falsified.

Confidentiality: required to protect subscriber information and communication patterns.

7.3.1.2 CMTS-RKS interface

Authentication, access control and message integrity: required to prevent service theft and denial-of-service attacks. Want to ensure that the billing events reported to the RKS are not falsified.

Confidentiality: required to protect subscriber information and communication patterns. Also, effective QoS information and network performance is kept secret from competitors.

7.3.1.3 MGC-RKS interface

Authentication, access control and message integrity: required to prevent service theft and denial-of-service attacks. Want to ensure that the billing events reported to the RKS are not falsified.

Confidentiality: required to protect subscriber information and communication patterns.

7.3.2 Cryptographic mechanisms

Both message integrity and privacy **MUST** be provided by IPsec ESP, using any of the ciphersuites that are listed in 6.1.2.

RADIUS itself defines MD5-based keyed MAC for message integrity at the application layer. And, there does not appear to be a way to turn off this additional integrity check at the application layer. For IPCablecom, the key for this RADIUS MAC **MUST** always be hardcoded to the value of 16 ASCII 0s. This, in effect, turns the RADIUS keyed MAC into an MD5 hash that can be used to protect against transmission errors but does not provide message integrity. No key management is needed for RADIUS MACs.

Billing event messages contain an 8-octet binary Element ID of the CMS, CMTS or the MGC. The RKS **MUST** verify each billing event by ensuring that the specified Element ID correctly corresponds to the IP address. This check is done via a lookup into a map of IP addresses to Element IDs. Refer to 7.3.3 on how this map is maintained. A combined element (such as a combined CMS/MGC) **MAY** use the same IP address and Security Association to convey Event Messages from both elements. Additionally, both elements may use the same Element ID. Refer to 7.3.3.1 for information on how to maintain a map of multiple elements and Element IDs.

7.3.2.1 RADIUS server chaining

RADIUS servers may be chained. This means that when the local RADIUS server that is directly talking to the CMS or CMTS client is not able to process a message, it forwards it to the next server in the chain.

IPCablecom specifies security mechanisms only on the links to the local RADIUS server. IPCablecom also requires authentication, access control, message integrity and privacy on the interfaces between the chained RADIUS servers, but the corresponding specifications are outside of the scope of IPCablecom.

Key management (in the following clause) applies to the local RADIUS server/RKS only.

7.3.3 Key management

7.3.3.1 CMS-RKS interface

The CMS and the RKS MUST negotiate a shared secret (CMS-RKS Secret) using IKE or Kerberos with symmetric keys (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys). For more information on the IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

The key management protocol MUST run asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired CMS-RKS Secret.

An RKS MUST maintain a mapping between an IP address and an Element ID for each host with which it has IPsec Security Associations. How this mapping is created depends on the IPsec key management protocol:

- 1) IKE with Pre-Shared Keys. One way to implement this mapping is to provide a local database of which Element ID(s) are associated with the source IP address.
- 2) IKE with Certificates. As specified in 8.2.3.4.3, a certificate of a server that sends billing event messages to an RKS contains its Element ID(s) in the CN attribute of the distinguished name. During IKE phase 1, the RKS MUST save a mapping between the IP address and its Element ID(s) that is contained in the certificate.
- 3) Kerberized Key Management. As specified in 6.4.5.5, a principal name of each server that reports billing event messages to the RKS includes its Element ID(s). After an RKS receives and validates an AP Request message, it MUST save a mapping between the IP address and its Element ID(s) that is contained in the principal name.

When an event message arrives at the RKS, the RKS MUST retrieve a source IP address based on the Element ID, using the mapping established during key management. The RKS MUST ensure that this address is the same as the source IP address in the IP packet header.

Later, when a billing event arrives at the RKS, it MUST be able to query the database of IPsec Security Associations and retrieve a source IP address, based on the Element ID. The RKS MUST ensure that it is the same as the source IP address in the IP packet header.

7.3.3.2 CMTS-RKS interface

CMTS and RKS MUST negotiate a shared secret (CMTS-RKS Secret) using IKE or Kerberos (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys). For more information on the IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

The key management protocol MUST be running asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired CMTS-RKS Secret.

An RKS maintains a mapping between an IP address and an Element ID for each host with which it has IPsec Security Associations, as specified in 7.3.3.1. This includes the CMTS.

When a billing event arrives at the RKS, it MUST retrieve a source IP address, based on the Element ID using the mapping established during key management. The RKS MUST ensure that it is the same as the source IP address in the IP packet header.

7.3.3.3 MGC-RKS interface

MGC and RKS MUST negotiate a shared secret (MGC-RKS Secret) using IKE or Kerberos (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using pre-shared keys). For more information on the

IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

The key management protocol MUST be running asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired MGC-RKS Secret.

An RKS maintains a mapping between an IP address and an Element ID for each host with which it has IPsec Security Associations, as specified in 7.3.3.1. This includes the MGC.

When an event arrives at the RKS, it MUST retrieve a source IP address, based on the Element ID based on the mapping established during key management. The RKS MUST ensure that it is the same as the source IP address in the IP packet header.

7.3.4 Billing system summary security profile matrix

Table 24 – Security profile matrix – RADIUS

| | RADIUS Accounting (CMS-RADIUS Server/RKS) | RADIUS Accounting (CMTS-RADIUS Server/RKS) | RADIUS Accounting (MGC-RADIUS Server/RKS) |
|---------------------|--|--|--|
| Authentication | Yes | Yes | Yes |
| Access control | Yes | Yes | Yes |
| Integrity | Yes | Yes | Yes |
| Confidentiality | Yes | Yes | Yes |
| Non-repudiation | No | No | No |
| Security mechanisms | IPsec ESP with encryption and message integrity enabled. Key management using IKE or Kerberos | IPsec ESP with encryption and message integrity enabled. Key management using IKE or Kerberos | IPsec ESP with encryption and message integrity enabled. Key management using IKE or Kerberos |

7.4 Call signalling

7.4.1 Network Call Signalling (NCS)

7.4.1.1 Reference architecture

Figure 12 shows the network components and the various interfaces to be discussed in this clause.

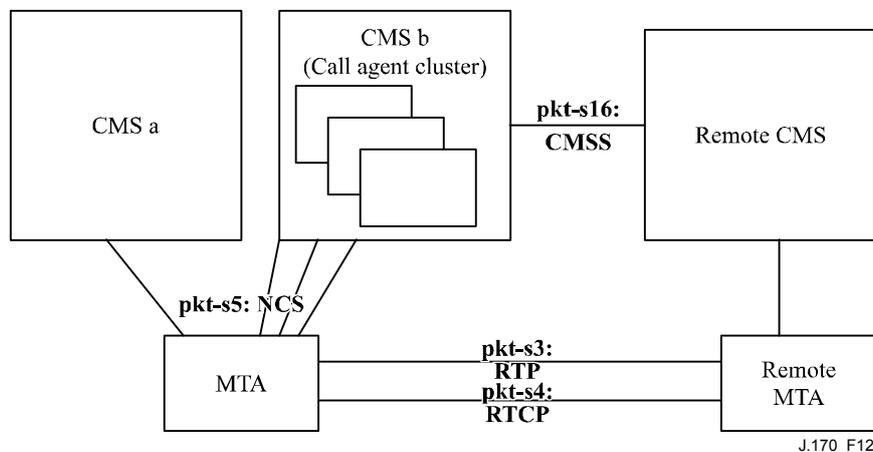


Figure 12 – NCS reference architecture

Figure 12 shows a CMS containing a cluster of Call Agents, which are identifiable by one CMS FQDN. It also shows, even though this is not a likely scenario in early deployments, that different CMSs could potentially manage different endpoints in a single MTA.

The security aspects of interfaces pkt-s3 and pkt-s4 (RTP bearer channel and RTCP) are described in 6.6. The protocol interface pkt-s16 (CMS to CMS) is SIP, with IPCablecom extensions, as specified in ITU-T Rec. J.178.

When a call is made between two endpoints in different zones, the call signalling has to traverse the path between two different CMSs. The signalling protocol between CMSs is SIP with IPCablecom specific extensions. See ITU-T Rec. J.178 for more details. Initially, the initiating CMS may not have a direct signalling path to a terminating CMS. The call routing table of the initiating CMS may point it to an intermediate SIP proxy. That SIP proxy, in turn, may point to another SIP proxy. In general, we make no assumptions about the number of SIP proxies in the signalling path between the CMSs. Once the two CMSs have discovered each other's location, they have the option to continue SIP signalling directly between each other. The SIP proxies that route traffic between domains are called Exterior Border Proxies (EBPs). EBPs enforce access control on all signalling messages routed between domains. They also provide application level security on sensitive information contained within SIP messages. While not depicted in Figure 12, CMSS may also be used between a CMS and an MGC.

As SIP proxies and CMSs may be in different PacketCable domains (and consequently different trust domains), there must be a signalling path and trust relationship between two domains, before any direct SIP signalling can take place. IPCablecom Server certificates are used for TLS mutual authentication in CMSS and provide the trust infrastructure for SIP signalling. A CMS or EBP may be configured to trust only specific Service Provider CA certificates and/or FQDNs (i.e., access list) of external CMSs and EBPs. Generally, trust between different Service Provider domains is provided by the EBPs.

7.4.1.2 Security services

The same set of requirements applies to both CMS-MTA and CMS-CMS signalling interfaces.

Authentication: Signalling messages should be authenticated in order to prevent a third party masquerading as either an authorized MTA, CMS, MGC or SIP Proxy.

Confidentiality: NCS messages carry dialled numbers and other customer information, which must not be disclosed to a third party. Thus, confidentiality of signalling messages should be required. The signalling messages carry media stream keying material that must be kept private on each signalling hop, and should also be kept private end-to-end between the initiating and target CMSs, to avoid exposure at SIP signalling proxies. There is no standard well-supported mechanism to support end-to-end privacy of keying material, however, so only hop-by-hop confidentiality is supported in IPCablecom.

Message integrity: This should be assured in order to prevent tampering with signalling messages – e.g., changing the dialled numbers.

Access control: Services enabled by the NCS signalling should be made available only to authorized users – thus, access control is required at the CMS.

7.4.1.3 Cryptographic mechanisms

IPsec ESP MUST be used to secure the NCS signalling between the CMS and MTA. IPsec keys MUST be derived using the mechanism described in 6.5.3.1.

TLS MUST be used to secure the SIP signalling (CMSS) between CMSs and between CMSs and SIP proxies (EBPs).

The first SIP signalling round trip between the initiating and target CMSs may transit through any number of intermediate SIP signalling proxies. Since TLS is applied separately on each signalling hop, the contents of the SIP signalling message is decrypted and re-encrypted at each SIP signalling proxy. The full contents of the SIP signalling message, including media stream keying material, are available in the clear at each intermediate SIP signalling proxy.

7.4.1.3.1 MTA-CMS interface

Each signalling message coming from the MTA and containing the MTA domain name (included in the NCS endpoint ID field) must be authenticated by the CMS. This domain name is an application-level NCS identifier that will be used by the Call Agent to associate the communication with a paying subscriber. In order to perform this authentication, the CMS MUST maintain an IP address to FQDN map for each MTA IP address that has a current SA. This map MUST be built during the key management process described in the following clause and does not need to reside in permanent storage.

7.4.1.3.2 CMS-CMS, CMS-MGC, CMS-SIP Proxy and SIP Proxy-SIP Proxy interfaces

When a CMS, MGC or a SIP Proxy receives a SIP signalling message, it SHOULD map the source IP address to the identity (FQDN) of the CMS or SIP Proxy and to the local policy associated with that FQDN. This lookup would utilize an IP address FQDN map for all MGCs and SIP Proxies that have current TLS sessions with this host. This map is built during key management described in the following clause and does not need to reside in permanent storage.

7.4.1.4 Key management

7.4.1.4.1 MTA-CMS key management

The MTA MUST use Kerberos with PKINIT to obtain a CMS service ticket (see 6.4.3). The MTA SHOULD first obtain a TGT (Ticket Granting Ticket) via the AS Request/AS Reply exchange with the KDC (authenticated with PKINIT). In the case that the MTA obtained a TGT, it performs a TGS Request/TGS Reply exchange to obtain the CMS service ticket (see 6.4.4).

After the MTA has obtained a CMS ticket, it MUST execute a Kerberized key management protocol (that utilizes the CMS ticket) with the CMS to create SAs for the pkt-s10 interface. This Kerberized key management protocol is specified in 6.5. Clause 6.5 also describes the mechanism to be deployed to handle timed-out IPsec keys and Kerberos tickets. The mechanism for transparently handling key switch-over from one key lifetime to another key lifetime is also defined.

The key distribution and time-out mechanism is not linked to any specific NCS message. Rather, the MTA will obtain the Kerberos ticket from the KDC when started and will refresh it based on the time-out parameter. Similarly, the MTA will obtain the sub-key (and thus IPsec ESP keys) based on the IPsec time-out parameters. In addition, when the IPsec ESP keys are timed out and the MTA needs to transmit data to the CMS, it will perform key management with the CMS and obtain the new keys. It is also possible for the IPsec SAs to expire at the CMS while it has data to send to the MTA. In this case, clause 6.5.3.5.3 describes the technique for the CMS to initiate key management and establish new security associations.

7.4.1.4.1.1 Call Agent clustering

At the time that the CMS receives a Kerberos ticket for establishing an IPsec SA, it MUST extract the MTA FQDN from the MTA principal name in the ticket and map it to the IP address. This map is later used to authenticate the MTA endpoint ID in the NCS signalling messages.

In the case a CMS, or an application server, is constructed as a cluster of Call Agents with different IP addresses, all Call Agents should share the same service key for decrypting a Kerberos ticket. Thus, the MTA will need to execute a single PKINIT Request/Reply sequence with the KDC and

multiple AP Request/Reply sequences for each Call Agent in the cluster. The Kerberos messages are specified in 6.4.4.

Optimized key management is specified for the case when, in the middle of a communication, a clustered Call Agent sends a message to an MTA from a new IP address, where it does not yet have an IPsec SA with that MTA (see 6.5.2.1).

In this optimized approach, the CMS sends a Rekey message instead of the Wake Up. This Rekey message is authenticated with a SHA-1 HMAC, using a Server Authentication Key, derived from a session key used to encrypt the last AP Reply sent from the same CMS (or another CMS with the same Kerberos Principal Name).

Additionally, the Rekey message includes IPsec parameters, to avoid the need for the AP Reply message. The MTA responds with a different version of the AP Request that includes the MTA-CMS Secret, normally sent by the CMS in the AP Reply. As a result, after the MTA responds with the AP Request, a new IPsec SA can be established with no further messages. The total price for establishing a new SA with this optimized approach is a single round-trip time. This is illustrated in Figure 13.

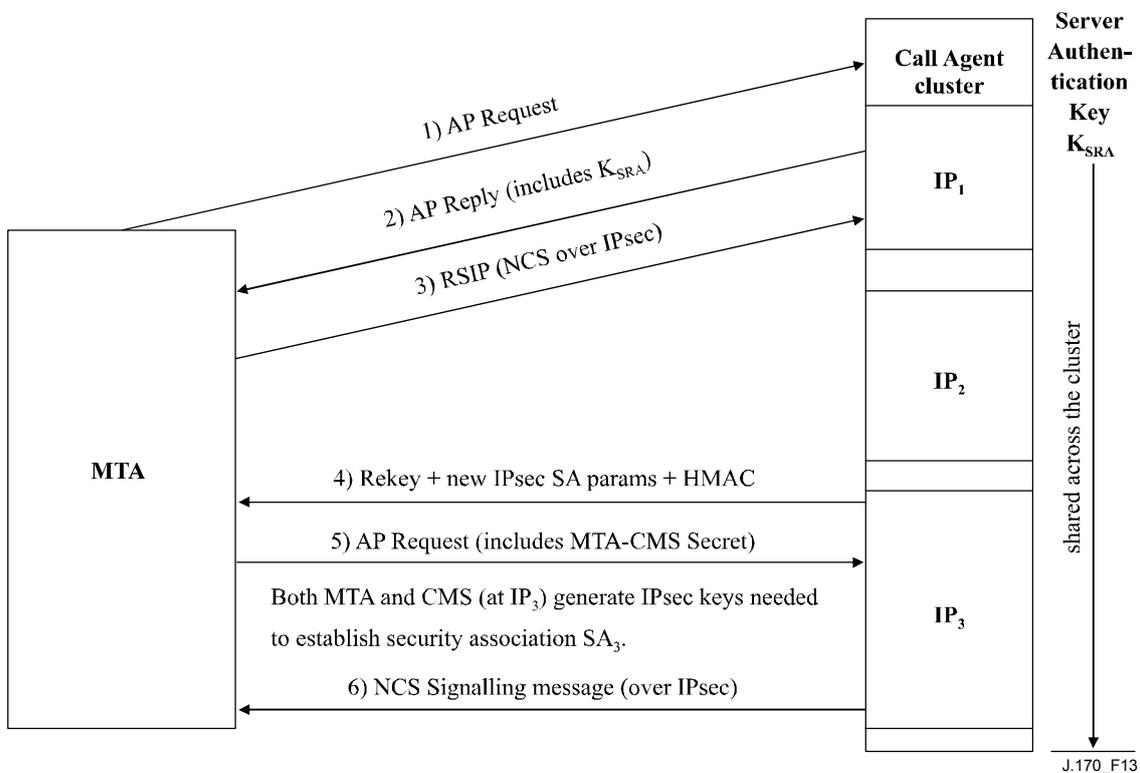


Figure 13 – Key management for NCS clusters

In Figure 13, an NCS clustered Call Agent suddenly decides to send an NCS message from a new IP address that did not previously have any SA established with that MTA.

The first security association SA₁ with CMS at IP₁ was established with a basic AP Request/AP Reply exchange. HMAC key K_{SRA} for authenticating Rekey message from the CMS was derived from the session key used to encrypt the AP Reply (as shown in steps 1 through 3 in Figure 13).

When a new SA₃ needs to be established between the MTA and CMS at IP₃, the key management is as follows in steps 4 and 5 in Figure 13:

- 4) The CMS at IP₃ sends a Rekey message, similar in functionality to the Wake Up message, but with a significantly different content. It contains:
 - IPsec parameters (also found in the AP Reply): SPI, selected ciphersuite, SA lifetime, grace period, and re-establish flag. The purpose of adding these IPsec parameters to Rekey is to eliminate the need for the subsequent AP Reply message.
 - SHA-1 HMAC using K_{SRA}.
- 5) AP Request that includes the MTA-CMS secret, normally sent in the AP Reply message. This is a legal Kerberos mode, where the key is contained in the AP Request and AP Reply is not used at all.

For more details, refer to 6.5.3.

7.4.1.4.1.2 MTA controlled by multiple CMSs

In the case a single MTA is controlled by multiple CMSs and each CMS is associated with a different Kerberos realm, the MTA will need to execute multiple PKINIT Request/Reply exchanges with the KDC, one for each realm, optionally followed by a TGS Request/Reply exchange. Then, an MTA would execute multiple AP Request/Reply exchanges in order to create the security association with the individual CMSs.

7.4.1.4.1.3 Transferring from one CMS to another via NCS signalling

When control of an MTA endpoint is transferred from one CMS to another via NCS signalling, the following steps are taken:

- 1) The new CMS might not have been included in the CMS table. In that case, the corresponding table entry **MUST** be locally created. Refer to 7.1.1.2.5 for instructions on how to create the new CMS table entry.
- 2) If the MTA does not already have IPsec SAs established with this CMS (e.g., via an earlier Wake Up), it **MUST** attempt to establish them at this time.
- 3) If the MTA now possesses valid IPsec security associations with the new CMS, the NCS signalling software is notified and the security association can be utilized. Further signalling traffic for this affected endpoint related to the prior CMS security association **MUST NOT** be sent.

7.4.1.4.2 CMS-CMS, CMS-MGC, CMS-SIP Proxy, SIP Proxy-SIP Proxy key management

When a CMS MGC, or a SIP Proxy has data to send to another CMS, MGC, or SIP Proxy and does not already have a TLS session with that host, it **MUST** first establish a TLS session with the other CMS, MGC, or SIP Proxy (see 6.9).

A CMS or a SIP Proxy **SHOULD** create TLS sessions ahead of time (before they are needed) whenever possible and maintain persistent connections.

7.4.1.4.2.1 Example of inter-domain Call Set-up with TLS sessions

Figure 14 depicts a typical SIP signalling flow for an inter-domain call set-up in which EBPs are used (refer to ITU-T Rec. J.178 for further details on CMSS call flows). It illustrates several points in the End-End call set-up where different TLS sessions and TCP connections may be required and also emphasizes the importance of connection persistence and re-use to minimize TCP connection and TLS session establishment during the call set-up (i.e., in order to minimize performance impacts and call set-up delays). It should be noted that a peering relationship between two SIP User Agents (i.e., CMSs and/or EBPs) often results in two TCP connections, one for SIP transactions initiated in each direction. This is due to the fact most TCP connections are initiated using

ephemeral source ports, and SIP transactions are initiated by sending SIP requests to a User Agent's well-known SIP port. As for securing each TCP connection with TLS, TLS clients typically cache TLS sessions based on specific remote IP address and port pairs; therefore, it is unlikely that TLS session caching using a common TLS master key can be used for both of the TLS sessions.

The inter-domain signalling flow begins with CMS "A" sending an INVITE to EBP "A" (CMS "A" is initiating a SIP INVITE transaction and will signal the well-known SIP port on EBP "A"). A TLS session is required and may need to be established for the TCP connection if one does not already exist (which can be re-used) for this new transaction. Similarly, a TLS session is required for each hop in this INVITE transaction, and may require a TLS session to be established between EBP "A" and EBP "B", and also between EBP "B" and CMS "B".

The 183 (Session Progress) response from CMS "B" is routed back through the EBPs to CMS "A" using the previously established TLS sessions. Once CMS "A" receives this 183 response, it sends a PRACK (Provisional ACK) directly to CMS "B". This PRACK is a SIP request which initiates a new transaction, and requires that a TLS session be established with CMS "B"'s well-known SIP port. CMS "B" sends a 200 OK response back to CMS "A" (using the same TLS session and TCP connection) as the final response to this PRACK transaction. Upon receiving a response to its initial INVITE, CMS "A" will also send an UPDATE request to CMS "B" over the previously established TLS session, to indicate that resource reservation has been completed. CMS "B" will respond with a 200 OK, completing this UPDATE transaction.

Upon receiving the UPDATE, CMS "B" reserves any necessary resource and sends back a 180 (Ringing) provisional response to CMS "A" over the previously established TLS session. This provisional response will also initiate PRACK/200 OK transaction between the two CMSs, over the same TLS session.

Once the terminating end answers the call, CMS "B" sends the 200 OK final response to the INVITE. However, this response is sent back via the EBPs using the same TLS sessions used for the INVITE. CMS "A" will acknowledge receipt of this 200 OK response by sending an ACK (a SIP request) directly to CMS "B". This ACK is sent using the previously established TLS session used for the first PRACK.

Once the call is established, the example illustrates the case where the terminating end goes on hook. CMS "B" initiates a BYE/200 OK transaction by sending a BYE (SIP request) to CMS "A". As this BYE request is sent to the well-known SIP signalling port of CMS "A", it is very likely that CMS "B" will need to use a different TCP connection and TLS session than the ones used for sending SIP requests from CMS "A" to CMS "B" (assuming CMS "A" is using an ephemeral port for its TCP connection to CMS "B").

As can be seen from this example, two TCP connections with two distinct TLS sessions may be required between two CMSs. It is important to support persistent and re-usable connections and TLS session caching in order to minimize impacts on CMS performance and call latency.

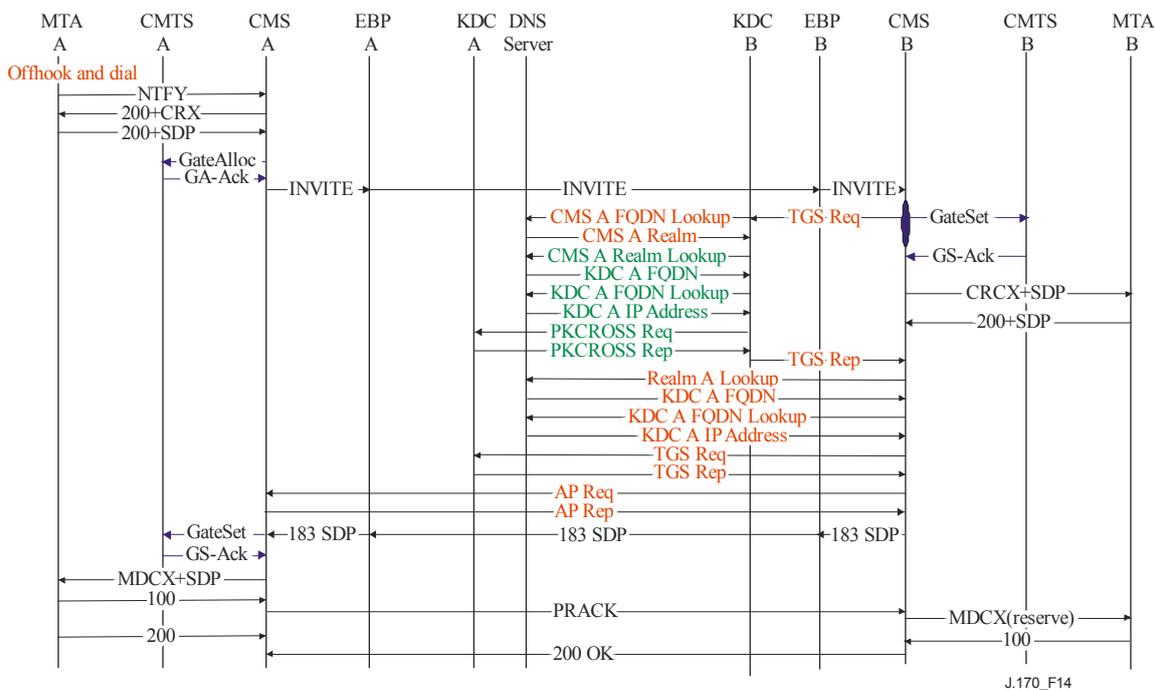


Figure 14 – CMS – CMS Signalling Flow with Security

Figure 14 illustrates the beginning of an inter-domain call set-up. In this example, the signalling between the two CMSs is routed through two EBPs (Exterior Border Proxies). After one round-trip between the CMSs (after the 183 SDP message is received by CMS "A"), the rest of the CMS-CMS signalling is done directly, without the involvement of the intermediate EBPs.

Figure 14 assumes that there are no prior SAs between the two CMSs and that CMS "B" does not possess a service ticket for CMS "A". It shows the key management flows necessary to establish the necessary SAs for this call.

CMS "B" sends a TGS Request to its local KDC (KDC "B") to get a ticket for CMS "A". This TGS Request is sent at the same time as the Gate-Set message to the local CMTS. The key management flows continue in parallel with the subsequent DQoS flows (Gate-Set ACK) and also in parallel with some NCS signalling flows (CRCX from CMS "B" to MTA "B" and the 200+SDP response from MTA "B"). After the 200+SDP response, no more parallelism is possible.

The next signalling message that CMS "B" sends out is the 183 SDP to EBP "B" that has to wait until CMS "B" obtains a ticket for CMS "A", in order to encrypt the media stream keying material inside SDP. To ensure synchronization, CMS "B" also waits until it establishes IPsec SAs with CMS "B", before sending the 183 SDP. This guarantees that when CMS "A" later sends a signalling message (PRACK) directly to CMS "B", IPsec SAs will already be established.

Since CMS "A" is in a different realm (and not in KDC "B's" database), the TGS Request from CMS "B" causes KDC "B" to perform a DNS lookup to retrieve CMS "A's" realm name. After the DNS lookup is complete, KDC "B" will attempt to locate a ticket for KDC "A" in its local ticket cache. If it finds that ticket, it will immediately return to CMS "B" the cross-realm TGT needed to authenticate to KDC "A". If (in a rare circumstance) KDC "B" does not currently have a ticket for KDC "A", it will first have to perform DNS lookups to locate it and then perform a PKCROSS exchange with KDC "A" to obtain the KDC ticket.

Once CMS "B" finally receives a cross-realm TGT for KDC "A", it has to send another TGS Request for KDC "A" and then finally obtain the service ticket for CMS "A". In order for CMS "B" to contact KDC "A", it will first have to perform two more DNS queries (one to get the KDC "A's" FQDN and another to get its IP address).

After CMS "B" had obtained a ticket for CMS "A", it will initiate Kerberized IPsec key management with CMS "A" in order to set up SAs. After IPsec SAs are established, CMS "B" continues with signalling by sending 183 SDP SIP message to CMS "A", via the external border proxy EBP "B". The keying material inside 183 SDP is encrypted as specified in 7.4.1.4.2.1. Further SIP signalling between the two CMSes is sent directly, without the participation of the border proxies and using the just established IPsec SAs.

7.4.2 Call Signalling Security Profile Matrix

Table 25 – Security profile matrix – Network call signalling

| | MTA-CMS | CMS-CMS | CMS-SIP Proxy/ SIP Proxy-SIP Proxy |
|---------------------|---|---|---|
| Authentication | Optional | Yes | Yes |
| Access control | Optional | Yes | Yes |
| Integrity | Optional | Yes | Yes |
| Confidentiality | Optional | Yes | Yes |
| Non-repudiation | No | No | No |
| Security mechanisms | IPsec ESP with encryption and message integrity enabled Authentication via Kerberos with PKINIT Kerberized key management defined by IPCablecom Security may be disabled through the provisioning process. | TLS with encryption and message integrity Authentication via X.509 certificates (or symmetric keys when TLS session caching is used) | TLS with encryption and message integrity Authentication via X.509 certificates (or symmetric keys when TLS session caching is used) |

7.5 PSTN Gateway interface

7.5.1 Reference architecture

An IPCablecom PSTN Gateway consists of three functional components:

- a Media Gateway Controller (MGC) which may or may not be part of the CMS;
- a Media Gateway (MG); and
- a Signalling Gateway (SG).

These components are described in detail in ITU-T Rec. J.171.x.

7.5.1.1 Media Gateway Controller

The Media Gateway Controller (MGC) is the PSTN gateway's overall controller. The MGC receives and mediates call-signalling information between the IPCablecom and the PSTN domains (from the SG), and it maintains and controls the overall state for all communications.

7.5.1.2 Media Gateway

Media Gateways (MG) provide the bearer connectivity between the PSTN and the IPCablecom network.

7.5.1.3 Signalling Gateway

IPCablecom provides support for SS7 signalling gateways. The SG contains the SG to MGC interface. Refer to ITU-T Rec. J.171.x for more detail on signalling gateways.

The SS7 Signalling Gateway performs the following security-related functions:

- Isolates the SS7 network from the IP network. Guards the SS7 network from threats such as information leakage, integrity violation, denial-of-service, and illegitimate use.
- Provides mechanisms for certain trusted entities ("TCAP Users") within the IPCablecom network, such as Call Agents, to query external PSTN databases via TCAP messages sent over the SS7 network.

7.5.2 Security services

7.5.2.1 MGC-MG interface

Authentication: Both the MG and the MGC must be authenticated, in order to prevent a third party masquerading as either an authorized MGC or MG.

Access control: MG resources should be made available only to authorized users – thus, access control is required at the MG.

Integrity: must be assured in order to prevent tampering with the TGCP signalling messages – e.g., changing the dialled numbers.

Confidentiality: TGCP signalling messages carry dialled numbers and other customer information, which must not be disclosed to a third party. Thus, confidentiality of the TGCP signalling messages is required.

7.5.3 Cryptographic mechanisms

7.5.3.1 MGC-MG interface

IPsec ESP MUST be used to both authenticate and encrypt the messages from MGC to MG and vice versa. Refer to 6.1 for details of how IPsec ESP is used within IPCablecom and for the list of available ciphersuites.

7.5.4 Key management

7.5.4.1 MGC-MG interface

Key management for MGC-MG interface is either IKE or Kerberos. Implementations MUST support IKE with pre-shared keys. Implementations MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys. For more information on the IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

The key management protocol ensures that there is always a valid, non-expired MGC-MG Secret.

7.5.5 MGC-MG summary security profile matrix

Table 26 – Security profile matrix – TCAP/IP and TGCP

| | TGCP (MG-MGC) |
|---------------------|--------------------------|
| Authentication | Yes |
| Access control | Yes |
| Integrity | Yes |
| Confidentiality | Yes |
| Non-repudiation | No |
| Security mechanisms | IPsec, IKE or Kerberos |

7.6 Media stream

This security Recommendation allows for end-to-end ciphersuite negotiation so that the communicating parties can choose their preferred encryption and authentication algorithms for the particular communication.

7.6.1 Security services

7.6.1.1 RTP

Authentication: End-to-end authentication cannot be required, because the initiating party may want to keep their identity private. Optional end-to-end exchanges for both authentication and additional key negotiation are possible but are outside of the scope for IPCablecom.

Encryption: The media stream between MTAs and/or MGs should be encrypted for privacy. Without encryption, the stream is vulnerable to eavesdropping at any point in the network.

Key Distribution via the CMS, a trusted third party, assures the MTA (or MG) that the communication was established through valid signalling procedures, and with a valid subscriber. All this guarantees confidentiality (but not authentication).

Message integrity: It is desirable to provide each packet of the media stream with a message authentication code (MAC). A MAC ensures the receiver that the packet came from the legitimate sender and that it has not been tampered with en route. A MAC defends against a variety of potential known attacks, such as replay, clogging, etc. It also may defend against as-yet-undiscovered attacks. Typically, a MAC consists of 8 or more octets appended to the message being protected. In some situations, where data bandwidth is limited, a MAC of this size is inappropriate. As a tradeoff between security and bandwidth utilization, a short MAC consisting of 2 or 4 octets is specified and selectable as an option to protect media stream packets. Use of the MAC during an end-to-end connection is optional; whether it is used or not is decided during the end-to-end ciphersuite negotiation (see 7.6.2.3.1).

Low complexity: Media stream security must be easy to implement. Of particular concern is a PSTN gateway, which may have to apply security to thousands of media streams simultaneously. The encryption and MAC algorithms used with the PSTN gateway must be of low complexity so that it is practical to implement them on such a scale.

7.6.1.2 RTCP

Authentication: See the above clause.

Encryption: Within IPCablecom, RTCP messages are not permitted to contain the identity of the RTCP termination endpoint. Snooping on RTCP messages, therefore, does not reveal any subscriber-specific information but may reveal network usage and reliability statistics. RTCP encryption is optional.

Message integrity: RTCP signalling messages (e.g., BYE) can be manipulated to cause denial-of-service attacks and alteration of reception statistics. To prevent these attacks, message integrity should be used for RTCP.

7.6.2 Cryptographic mechanisms

MTAs and MGs MUST have an ability to negotiate a particular encryption and authentication algorithm. If media security parameters are negotiated and RTP encryption is on (Transform ID is not RTP_ENCR_NULL), each media RTP packet MUST be encrypted for privacy. If RTP encryption is on, encryption MUST be applied to the RTP payload and MUST NOT be applied to the RTP header. Security MUST NOT be applied to RTP packets if the negotiated RTP ciphersuite is AUTH_NULL and RTP_ENCR_NULL.

Each RTP packet MAY include an optional message authentication code (MAC). The MAC algorithm can also be negotiated. The MAC computation MUST span the packet's unencrypted header and encrypted payload. The receiver MUST perform the same computation as the sender and it MUST discard the received packet if the value in the MAC field does not match the computed value.

Keys for the encryption and MAC calculation MUST be derived from the End-End secret, which is exchanged between sending and receiving MTAs as described in 7.6.2.3.1.

7.6.2.1 RTP messages

Figure 15 shows the format of an encoded RTP packet. IPCablecom MUST adhere to the RTP packet format as defined by RFC 1889 and RFC 1890 after being authenticated and decrypted (where the MAC bytes, if included, are stripped off as part of the authentication).

The packet's header consists of 12 or more octets, as described in RFC 1889. The only field of the header that is relevant to the encoding process is the timestamp field.

The RTP header has the format (RFC 1889) as shown in Figure 15.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|--|---|---|---|----|---|---|---|---|----|---|---|---|---|---|---|---|-----------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0 | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | | | | | |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 |
| V=2 | | P | X | CC | | | | M | PT | | | | | | | | Sequence Number | | | | | | | | | | | | | | |
| Timestamp | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Synchronization Source (SSRC) Identifier | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Contributing Source (CSRC) Identifier | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 15 – RTP packet header format

The first twelve octets are present in every RTP packet, while the list of CSRC identifiers is present only when inserted by a mixer. See Figure 16.

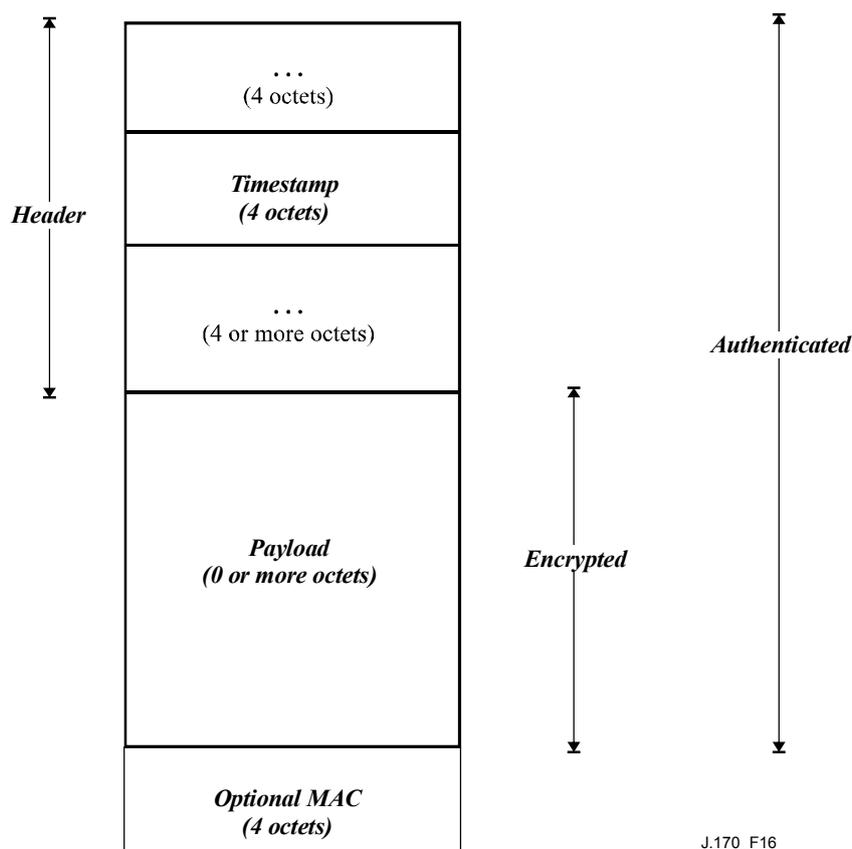


Figure 16 – Format of encoded RTP packet

In IP-Cablecom, an RTP packet will carry compressed audio from the sender's voice codec, or it will carry a message describing one or more events such as a DTMF tone, trunk or line signalling, etc. For simplicity, the former is referred to as a "voice packet" and the latter as an "event packet".

A voice packet's payload consists of compressed audio from the sender's voice codec. The length of the payload is variable and depends on the voice codec as well as the number of codec frames carried by the packet.

An event packet's payload consists of a message describing the relevant event or events. The format of the message is outside the scope of this Recommendation. The length of the payload is variable, but it will not exceed a known maximum value.

For either type of packet, the payload MUST be encrypted. If the optional MAC is selected, the MAC field is appended to the end of the packet after the payload.

Parameters representing RTP packet characteristics are defined as follows:

- N_c : the number of octets in one frame of compressed audio. Each codec has a well-defined value of N_c . In the case of a codec that encodes silence using short frames, N_c refers to the number of octets in a non-silent frame.
- N_u : the number of speech samples in one frame of uncompressed audio. The number of speech samples represented by a voice packet is an integral multiple of N_u .
- N_f : the frame number. The first frame of the sender's codec has a value of zero for N_f . Subsequent frames increment N_f by one. N_f increments regardless of whether a frame is actually transmitted or discarded as silent.

- M_f : the maximum number of frames per packet. M_f is determined by the codec's frame rate and by the sender's packetization rate. The packetization rate is specified during communications set-up. For NCS signalling, it is a parameter in the LocalConnectionOptions – see ITU-T Rec. J.162.

For example, suppose the speech sample rate is 8000 samples/s, the frame rate is 10 ms, the packetization rate is 30 ms, and the compressed audio rate is 16 000 bits/s. Then $N_c = 20$, $N_u = 80$, $M_f = 3$, and N_f counts the sequence 0, 1, 2.

- N_e : the maximum number of bytes that might be sent within the duration of one codec frame. It is assumed that an event packet can have a payload as large as that of a voice packet, but no longer. In the case of a block cipher, the cryptographic keys do not change after midstream codec changes. When a codec change does not require a corresponding key change, the value of N_e MUST be calculated as follows:

$$N_e = \text{MAX} \{ N_{cK} \} \text{ for } K = 1, \dots, N$$

where N_1, N_2, \dots, N_K are the different frame sizes for codecs that are supported by a particular endpoint.

Otherwise, $N_e = N_c$, where N_c is the frame size for the current codec.

- N_m : the number of MAC octets. This value is 0, if the optional MAC is not selected; or 2 or 4, representing the MAC size if the optional MAC is selected.

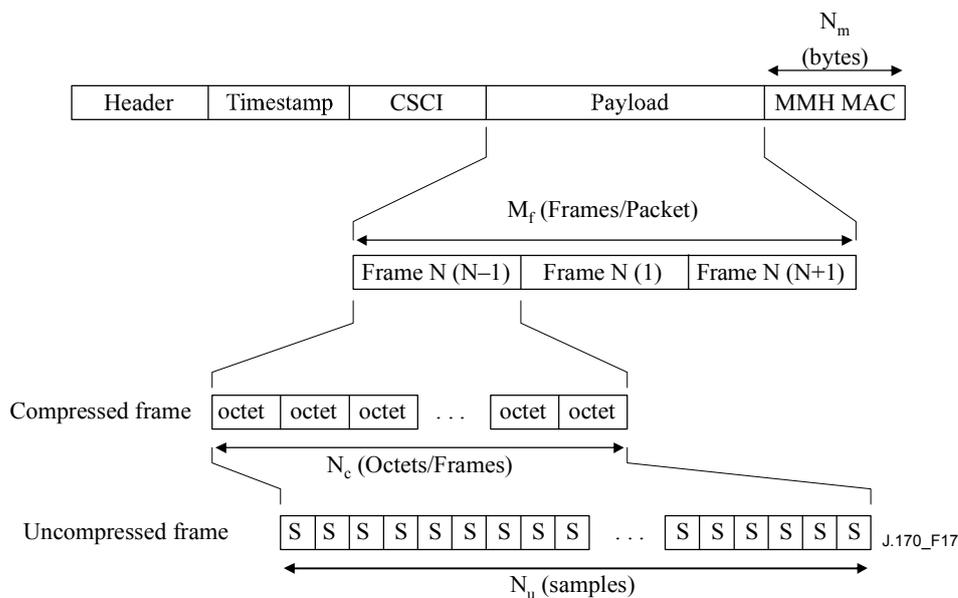


Figure 17 – RTP packet profile characteristics

7.6.2.1.1 RTP timestamp

According to RFC 1889, the timestamp field is a 32-bit value initially chosen at random. To IPCablecom, the timestamp MUST increment according to the codec sampling frequency. The timestamp in the RTP header MUST reflect the sampling instant of the first octet in each RTP packet presented as offset from the initial random timestamp value. The timestamp field MAY be used by the receiver to synchronize its decryption process to the encryption process of the sender.

Based on the definition of the timestamp and the packet parameters described in the previous clause, the timestamp MUST equate to the value: $((N_f \times N_u) + (\text{RTP initial timestamp})) \text{ modulo } 2^{32}$, where N_f is the frame number of the first frame included in the packet.

7.6.2.1.2 Packet encoding requirements

Prior to encoding the packets of an RTP stream, the sending MTA MUST derive the keys and parameters from the End-End Secret it shares with the receiving MTA, as specified in 7.6.2.3.3.

An MTA MUST derive two distinct sets of these quantities, one set for processing outgoing packets and another set for processing incoming packets.

7.6.2.1.2.1 Encryption and MMH MAC option

7.6.2.1.2.1.1 Deriving an MMH MAC Key

The MMH MAC Key size MUST be determined before generating the MMH MAC Key. The following algorithm specifies how to derive the MMH MAC Key when being used with block ciphers:

$$\text{MMH MAC key size} = (M_f \times N_e) + N_h + N_m - 2 + P$$

where:

M_f is the maximum number of frames per packet

N_e is maximum number of octets in one frame of compressed audio

N_h is the maximum number of octets in the RTP header, as defined in 7.6.2.1

N_m is the number of octets in the MAC

Therefore, $(M_f \times N_e) + N_h$ represents the maximum size of an RTP packet, and $N_m - 2$ represents the additional two octets that are added to the key size when a four-octet MMH MAC is used. (The key size is the same as the maximum RTP packet size when a two-octet MMH MAC is used.) P is 0 or 1, as needed to make the MMH MAC key size an even number so that it is a multiple of the word size (2 bytes) used in the MMH MAC algorithm.

The number of octets in the RTP header ranges from 12 to 72, inclusive, depending on the number of CSRC identifiers that are included (see RFC 1889). An implementation MUST choose N_h at least as large as required to accommodate the maximum number of CSRC identifiers that may occur during a session. An implementation MUST set N_h to 72 if the maximum number of CSRC identifiers is otherwise unknown.

Since the key derivation procedure generates the MMH MAC key last (see 7.6.2.3.3.1), it is not necessary to generate a complete MMH MAC key at the start of the RTP session. Implementations MAY generate less than the full MMH MAC key and generate the rest later, as needed. For example, instead of using a value of N_e that reflects all possible codecs supported by an endpoint, an implementation might initially derive an MMH key of size $(M_f \times N_c) + N_h + N_m - 2 + P$, where N_c is the frame size for the currently selected codec. Later, after a codec change that results in a larger value of N_c , additional bytes for the MMH key may be generated.

7.6.2.1.2.1.2 RTP timestamp wrap-around

Let us say that the initial RTP timestamp value is T_0 . A timestamp wrap-around occurs when:

- an RTP packet with sequence number i has a timestamp value $2^{32} - \xi_1$ for $0 < \xi_1 \leq \Delta T_{\text{MAX}}$, where ΔT_{MAX} is the maximum difference between two consecutive RTP timestamps;
- an RTP packet with a sequence number $i+1$ has a timestamp value ξ_2 for $0 \leq \xi_2 < \Delta T_{\text{MAX}}$.

The wrap-around point is between the RTP packets i and $i+1$.

Each endpoint MUST keep a count N_{WRAP} of RTP timestamp wrap-arounds, with a range from 0 to $2^{16} - 1$ and initialized to zero at the start of the connection. N_{WRAP} MUST be incremented by the sender right after the wrap-around point. N_{WRAP} MUST also be incremented by the receiver before it decrypts any RTP packets after the wrap-around point.

7.6.2.1.2.2 Block cipher encryption of RTP packets

The AES block cipher must be supported for encryption of RTP packets. The following subclauses specify how to support any block cipher, including AES.

7.6.2.1.2.2.1 Block termination

If an implementation supports block ciphers, residual block termination (RBT) MUST be used to terminate streams that end with less than a full block of data to encrypt (see 9.3).

7.6.2.1.2.2.2 Initialization Vector

An Initialization Vector (IV) is required when using a block cipher in CBC mode to encrypt RTP packet payloads. The size of an IV is the same as the block size for the particular block cipher. For example, the IV size for DESX and 3-DES is 64 bits, while for AES-CBC it is 128 bits. In order to calculate the IV, each endpoint MUST keep track of N_{WRAP} – the count of timestamp wrap-arounds during this RTP session; see 7.6.2.1.2.1.2. The IV MUST be calculated new for each RTP packet as specified below:

- 1) Take the first N bits of the header, where $N = \min(\text{cipher block size, RTP header size})$.
- 2) In the result of the previous step, replace the first 16 bits of the header with the 16-bit value of N_{WRAP} , MSB first.
- 3) Pad the result of previous step with 0s on the right, so that the resulting bit string is equal in size to the cipher block size.
- 4) XOR the result of the previous step with the RTP Initialization Key (defined in 7.6.2.3.3.1). The size of the RTP Initialization Key is the same as the cipher block size.
- 5) Encrypt the result of the previous step using the same block cipher that is used to encrypt RTP packets, but in ECB mode. The result of this step is the Initialization Vector for this RTP packet.

7.6.2.1.2.2.3 MMH-MAC pad derivation when using a block cipher

The MMH-MAC algorithm requires a one-time pad for each RTP packet. The MMH-MAC Pad MUST be derived by performing the MMH Function on the Block Cipher's IV. For a 2-byte MMH-MAC, use the MMH Function described in 9.7.1.1; for a 4-byte MMH-MAC, use the MMH Function described in 9.7.1.2.

The IV value is calculated according to 7.6.2.1.2.2.2 for block ciphers that require an IV. Even if the block cipher does not require an IV, one MUST be derived according to 7.6.2.1.2.2.2 and used as the basis of the MMH-MAC pad derivation.

A key is also required by the MMH digest function in order to calculate the pad. The MMH MAC key derived in 7.6.2.3.3.1 MUST be truncated according to 9.7.2.2 and MUST then be used as the key to the MMH digest. Accordingly, the MMH MAC key is truncated to:

$$\langle \text{size of IV} \rangle + N_m - 2$$

Where $\langle \text{size of IV} \rangle$ is 16 bytes for AES, N_m is the size of the MMH MAC in bytes, as defined in 7.6.2.1, and $N_m - 2$ represents the additional two octets that are added to the key size when a four-octet MMH MAC is used. (The truncated key size is the same as the IV size when a two-octet MMH MAC is used.)

7.6.2.1.3 Packet decoding requirements

Prior to decoding the packets of an RTP stream, the receiving MTA MUST derive the keys and parameters from the End-End Secret it shares with the sending MTA, as specified in 7.6.2.3.3.

The derived quantities MUST match the corresponding quantities at the sending MTA.

7.6.2.1.3.1 Timestamp tolerance check

Before processing a received packet, the receiver SHOULD perform a sanity check on the timestamp value in the RTP header, consisting of the items 1 and 2 below:

- 1) Beginning with the RTP timestamp in the first packet received from a sender, the receiver calculates an expected value for the timestamp of the sender's next RTP packet based on timestamps received in the sender's previous packets for the session.
- 2) The next packet is rejected without being processed if its timestamp value is outside a reasonable tolerance of the expected value. (Timestamps from rejected packets are not to be used to predict future packets). The tolerance value is defined to be:
 - a) sufficiently tight to ensure that an invalid timestamp value cannot derail the receiver's state so much that it cannot quickly recover to decrypting valid packets;
 - b) able to account for known differences in the expected and received timestamp values, such as might occur at call startup, codec switch-over and due to sender/receiver clock drift.

If the timestamp value in the RTP headers from a sender never comes back within the acceptable range, the receiver discontinues the session.

At the receipt of each packet, the receiver adjusts its time relationship with the sender within the acceptable tolerance range of estimated values.

7.6.2.1.3.2 Packet authentication

If authentication is used on an RTP packet stream, verification of the MAC MUST be the first step in the packet decoding process. When the timestamp tolerance check is performed, the MAC MAY be verified on packets with valid RTP timestamps immediately after the check is completed.

If the MAC does not verify, the packet MUST be rejected.

7.6.2.2 RTCP messages

7.6.2.2.1 RTCP format

RFC 1889 defines the packet format of RTCP messages as in Figure 18.

| v=2 | p | count | pkt type | Length |
|------|---|-------|----------|--------|
| SSRC | | | | |
| | | | | |
| | | | | |

Figure 18 – RTCP packet format

The RTCP packet type could be SR (sender reports), RR (receiver reports), SDES (source description), BYE (leaving conference), and APP (application-specific function). The length varies depending on the message type, but generally around 40 bytes.

7.6.2.2.2 RTCP encryption

The RTCP messages MUST always be encrypted in their entirety when the negotiated encryption algorithm is a block cipher in CBC mode. RTCP messages MUST NOT be encrypted when the negotiated encryption algorithm is RTCP_ENCR_NULL. However, the encoded RTCP messages MUST still be formatted according to 7.6.2.2.2 when RTCP_ENCR_NULL is selected in conjunction with a non-NULL authentication algorithm (e.g., HMAC-SHA1-96 or HMAC-MD5-96). Security MUST NOT be applied to RTCP packets if the negotiated RTCP ciphersuite is RTCP_AUTH_NULL and RTCP_ENCR_NULL. After the message is encrypted, an

additional header and MAC (Message Authentication Code) are added. The result packet has the format in Figure 19.

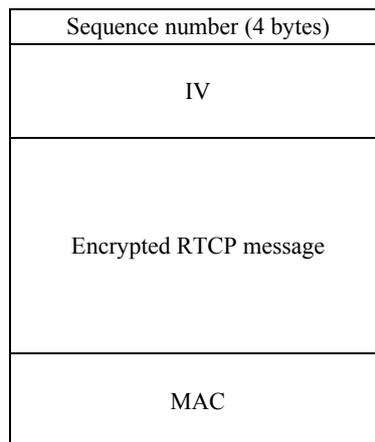


Figure 19 – RTCP encrypted packet format

The first 4 bytes MUST be the sequence number, MSB first. The initial sequence number for each direction of traffic MUST be 0. Afterwards, the sequence number for each direction MUST be incremented by 1. Generally, one RTCP message is sent every 5 seconds for each channel. Thus, 32 bits for the sequence number field would be big enough for any connections without wrapping around.

The IV (Initialization Vector) MUST immediately follow the sequence number. The IV MUST be randomly generated by the sender for each RTCP message and the IV size MUST be the same as the block size for the selected block cipher. The Initialization Vector (IV) MUST NOT be included when RTCP_ENCR_NULL is used.

The original cleartext RTCP message encrypted in its entirety MUST immediately follow the IV. The MAC (Message Authentication Code) computed over the concatenation of the sequence number, the IV, and the encrypted message MUST follow the encrypted RTCP message. The size of the MAC is algorithm-dependent.

7.6.2.2.3 Sequence numbers

The receiver of RTCP messages SHOULD keep a sliding window of the RTCP sequence numbers. The size of the sliding window W_{RTCP} depends on the reliability of the UDP transport and is locally configured at each endpoint. W_{RTCP} SHOULD be 32 or 64. The sliding window is most efficiently implemented with a bit mask and bit shift operations.

When the receiver is first ready to receive RTCP packets, the first sequence number in this window MUST be 0 and the last MUST be $W_{RTCP} - 1$. All sequence numbers within this window MUST be accepted the first time but MUST be rejected when they are repeated. All sequence numbers that are smaller than the "left" edge of the window MUST be rejected.

When an authenticated RTCP packet with a sequence number that is larger than the "right" edge of the window is received, that sequence number is accepted and the "right" edge of the window is replaced with this sequence number. The "left" edge of the window is updated in order to maintain the same window size.

When for a window $(S_{right} - W_{RTCP} + 1, S_{right})$, sequence number S_{new} is received and $S_{new} > S_{right}$, then the new window becomes:

$$(S_{new} - W_{RTCP} + 1, S_{new})$$

7.6.2.2.4 Block termination

Residual block termination (RBT) MUST be used to terminate RTCP messages that end with less than a full block of data to encrypt (see 9.3).

7.6.2.2.5 RTCP message encoding

Each RTCP message MUST be encoded using the following procedure:

- 1) A random IV is generated.
- 2) The entire RTCP message is encrypted with the selected block cipher and the just generated IV.
- 3) The current sequence number, the IV, and the encrypted RTCP message are concatenated in that order.
- 4) The MAC is computed (using the selected MAC algorithm) over the result in step 3) and appended to the message.

7.6.2.2.6 RTCP message decoding

Each RTCP message MUST be decoded using the following procedure:

- 1) Regenerate the MAC code and compare with the received value. If the two do not match, the message is dropped.
- 2) The sequence number is verified based on the sliding window approach specified in 7.6.2.2.3. If the sequence number is rejected, the message is dropped. The sliding window is also updated as specified in 7.6.2.2.3.
- 3) The RTCP message is decrypted with the shared encryption key and with the IV that is specified in the message header.

7.6.2.3 Key management

The key management specified here for end-to-end communication is identical in the cases of the MTA-to-PSTN and MTA-to-MTA communications. In the case of the MTA-to-PSTN communications, one of the MTAs is replaced by a MG (Media Gateway).

The descriptions below refer to MTA-to-MTA communications only for simplicity. In this context, an MTA actually means a communication endpoint, which can be an MTA or a MG. In the case that the endpoint is a MG, it is controlled by an MGC instead of a CMS.

During call set-up MTA_0 (the initiating MTA) and MTA_1 (the terminating MTA) exchange randomly generate keying material, carried inside the call signalling messages. Call signalling messages are themselves protected by IPsec ESP or TLS at each hop. This keying material is then used to generate the AES-CBC keys used to protect both RTP and RTCP messages between the two MTAs.

MTA_0 generates two randomly generated values: End-End Secret₀ (46-bytes) and Pad₁ (46-bytes).

MTA_1 generates two randomly generated values: End-End Secret₁ (46-bytes) and Pad₀ (46-bytes).

MTA_0 uses End-End Secret₁ and Pad₁ to derive encryption and authentication keys to be applied to its outbound traffic and used by MTA_1 to decrypt and authenticate it.

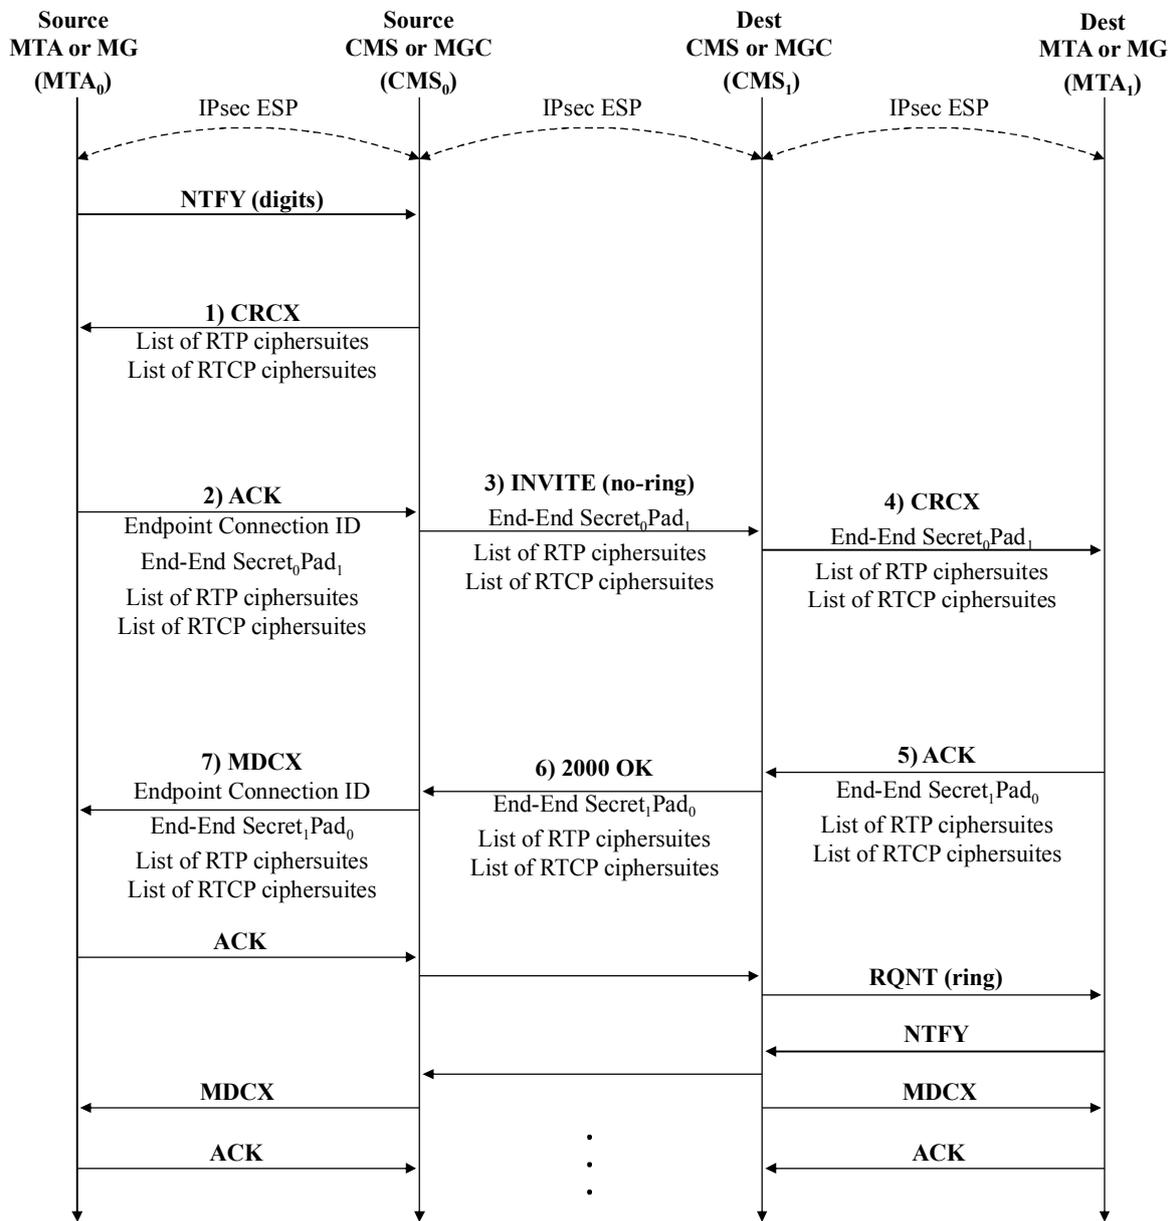
MTA_1 uses End-End Secret₀ and Pad₀ to derive encryption and authentication keys to be applied to its outbound traffic, and used by MTA_0 to decrypt and authenticate it. As a result, both MTA_0 and MTA_1 contribute randomly generated bytes to all of the keying material for both RTP and RTCP traffic.

The distribution of the end-to-end keying material is specific to the call signalling from ITU-T Rec. J.162 and is described in the following subclauses.

7.6.2.3.1 Key management over NCS

Figure 20 shows the actual NCS messages that are used to carry out the distribution of end-to-end keys. Each NCS message that is involved in the end-to-end key management is labelled with a number of the corresponding key management interface.

The name of each NCS message is in bold. Below the NCS message name is the information needed in the NCS message, in order to perform end-to-end key distribution. Messages between the CMSs are labelled as SIP+ messages.



J.170_F20

AKS Acknowledge
 CRCX Create Connection
 MDCX Modify Connection
 NTFY Notify
 RQNT Request for Notification

Figure 20 – End-End secret distribution over NCS

Figure 20 shows that before the start of this scenario, both the source and destination MTAs had already established an IPsec ESP session with their local CMS. It is also assumed that CMS-CMS signalling is secure.

This allows the End-End Secrets to be distributed securely, with privacy, integrity and anti-replay mechanisms already in place. The CMSs have access to this keying material but are trusted by the MTAs.

7.6.2.3.1.1 NULL ciphersuite combinations and ordering

RTP_ENCR_NULL MUST only be used in conjunction with AUTH_NULL. RTP packets, with authentication but no encryption, are not allowed.

RTCP_AUTH_NULL MUST only be used in conjunction with RTCP_ENCR_NULL. RTCP messages with encryption and without authentication are not allowed.

Both RTP and RTCP security must be enabled or disabled together. The following five combinations MUST NOT be generated.

- RTP NULL encryption and RTP non-NULL authentication;
- RTCP non-NULL encryption & RTCP NULL authentication;
- RTP non-NULL encryption and RTCP NULL authentication;
- RTP NULL encryption and RTCP non-NULL authentication;
- RTP NULL encryption and RTCP non-NULL encryption.

If the MTA receives LocalConnectionOptions parameter that meet the above combinations, the MTA MUST return the error code 524 (Internal inconsistency in LocalConnectionOptions). Otherwise, if the MTA receives RemoteConnectionDescriptor parameter that meet the above combinations, then the MTA MUST return the error code 505 (Unsupported RemoteConnectionDescriptor).

For both RTP and RTCP ciphersuite lists exchanged during ciphersuite negotiation, the combination of NULL encryption and NULL authentication algorithms MUST always be included last. For example, the list of RTP ciphersuites "60/50;62/51;64/51" is not allowed, while the list of RTP ciphersuites "62/51;64/51;60/50", or "60/50" is allowed. If the list of ciphersuites in LocalConnectionOptions includes the NULL authentication and NULL encryption combination (60/50 for RTP, and 80/70 for RTCP), but this combination is not the last in the list, the MTA MUST return error code 524 (Internal inconsistency in LocalConnectionOptions). Otherwise, if this combination is not last in a RemoteConnectionDescriptor, error code 505 (Unsupported RemoteConnectionDescriptor) MUST be returned.

7.6.2.3.1.2 Ciphersuite negotiation for MTAs

This Recommendation only defines security for RTP/RTCP media streams; therefore, ciphersuite negotiation applies only to RTP/RTCP media streams. Use of security for any other type of media streams is not specified.

An MTA MUST perform RTP and RTCP ciphersuite negotiation when processing any of the following:

- a CreateConnection command;
- a ModifyConnection command with a RemoteConnectionDescriptor parameter;
- a ModifyConnection command where the LocalConnectionOptions parameter includes ciphersuite fields.

An MTA MUST NOT perform ciphersuite negotiation in any other case. The steps involved in ciphersuite negotiation are the following:

- 1) An approved list of ciphersuites is formed by taking the intersection of the internal list of ciphersuites and ciphersuites allowed by the LocalConnectionOptions parameter, subject to the constraints specified in 7.6.2.3.1.1. The internal list of ciphersuites contains the ciphersuites that the MTA supports and which this Recommendation requires. If the LocalConnectionOptions parameter was not included, or if the ciphersuite fields were not provided in the LocalConnectionOptions parameter, the approved list of ciphersuites contains the previously agreed upon approved list, or if no such list exists, the internal list of ciphersuites.
- 2) If the approved list of ciphersuites is empty, an error response MUST be generated, error code 532 (Unsupported value(s) in LocalConnectionOptions).
- 3) Otherwise, a negotiated list of ciphersuites is formed by taking the intersection of the approved list of ciphersuites and ciphersuites allowed by the RemoteConnectionDescriptor parameter (if present), subject to the constraints specified in 7.6.2.3.1.1. If a RemoteConnectionDescriptor was not provided, the negotiated list of ciphersuites thus contains the approved list of ciphersuites. If a RemoteConnectionDescriptor parameter is provided without fields containing the RTP and RTCP ciphersuite lists, then the RTP AUTH_NULL/RTP_ENCR_NULL and RTCP_AUTH_NULL/RTCP_ENCR_NULL ciphersuites are assumed for the remote endpoints, and the regular ciphersuite negotiation process continues (i.e., the negotiated list of ciphersuites is formed by taking the intersection of the approved list of ciphersuites and the RTP AUTH_NULL/RTP_ENCR_NULL and RTCP_AUTH_NULL/RTCP_ENCR_NULL ciphersuites).
- 4) If the negotiated list of ciphersuites is empty, a ciphersuite negotiation failure has occurred and an error response MUST be generated. If a RemoteConnectionDescriptor parameter was provided, two different error codes can be returned:
 - a) If the endpoint does not support any of the ciphersuites allowed by the RemoteConnectionDescriptor, error code 505 (Unsupported RemoteConnectionDescriptor) MUST be used.
 - b) If the endpoint does support at least one of the ciphersuites, but the negotiated list of ciphersuites ended up being empty, error code 506 (Unable to satisfy both LocalConnectionOptions and RemoteConnectionDescriptor) MUST be used.
- 5) Otherwise, ciphersuite negotiation has succeeded, and the negotiated list of ciphersuites is returned in the LocalConnectionDescriptor parameter. Note that both LocalConnectionOptions and the RemoteConnectionDescriptor parameters can contain a list of ciphersuites that MUST be ordered by preference provided by the CMS in the RemoteConnectionDescriptor parameter. When both are supplied, the MTA SHOULD adhere to the preferences provided by the CMS in the RemoteConnectionDescriptor parameter, and otherwise, the MTA SHOULD adhere to the preferences provided in the LocalConnectionOptions parameter. If the MTA receives a RemoteConnectionDescriptor parameter with AUTH_NULL/RTP_ENCR_NULL for RTP or RTCP_AUTH_NULL/RTCP_ENCR_NULL for RTCP that is not last in the list, it MUST return the error code 505 (Unsupported RemoteConnectionDescriptor).

The following requirements apply during ciphersuite negotiation:

- A CMS MUST be capable of sending the allowable lists of ciphersuites for RTP and/or RTCP in the LocalConnectionOptions parameter of a CreateConnection command (CRCX) or a ModifyConnection command (MDCX) in the order of preference specified by the operator subject to the constraints specified in 7.6.2.3.1.1.

- Whenever possible, a MTA SHOULD select the first supported ciphersuite for RTP and the first supported ciphersuite for RTCP in the RemoteConnectionDescriptor parameter. This allows the MTA to immediately start sending RTP and RTCP packets to the other MTA. An MTA MAY instead select alternate ciphersuites specified by the other MTA.
- When returning a LocalConnectionDescriptor and the negotiated list of RTP and RTCP ciphersuites is NULL, an MTA MUST NOT include an End-End Secret or Pad.
- When returning a LocalConnectionDescriptor and the negotiated list of RTP and RTCP ciphersuites contains at least one non-NULL selection each, an MTA MUST include an End-End Secret (for incoming RTP and RTCP packets) and MAY include a Pad value (for outgoing RTP and RTCP packets). The following rules apply:
 - 1) The MTA MUST generate a new End-End Secret when responding to a CreateConnection command.
 - 2) The MTA MUST generate a new End-End Secret when responding to a ModifyConnection command if the remote connection address (e.g., IP Address) or the remote transport address (e.g., port) are not identical to what was previously assigned.
 - 3) The MTA MUST use the existing End-End Secret when responding to a ModifyConnection command where there was no previous RemoteConnectionDescriptor provided.
 - 4) The MTA MUST generate a new Pad when responding to a CreateConnection command without a RemoteConnectionDescriptor.
 - 5) The MTA MUST generate a new Pad when generating a new End-End Secret in response to a ModifyConnection command without a RemoteConnectionDescriptor.
 - 6) If not otherwise required, the MTA MAY generate a new Pad when generating a new End-End Secret.
 - 7) The MTA MUST NOT generate a new Pad when not generating a new End-End Secret.
- If, in response to a CreateConnection command, the list of ciphersuites selected for RTP contains at least one non-NULL encryption or authentication algorithm, before sending the response message, an MTA MUST:
 - 1) Establish inbound RTP security based on the preferred (first) RTP ciphersuite, its End-End Secret (which it generated), and a Pad value (if included in the RemoteConnectionDescriptor), as described in 7.6.2.3.3.1.
 - 2) If a RemoteConnectionDescriptor was included and it contains media security attributes, establish outbound RTP security based on the selected RTP ciphersuite, End-End Secret (generated by the other MTA), and a Pad value (which it may have generated) as described in 7.6.2.3.3.1.
 - 3) If connection mode allows, be ready to receive RTP packets, which may arrive any time after the Response message is sent.
- If, in response to a CreateConnection command, the list of ciphersuites for RTCP contains at least one non-NULL encryption algorithm, before sending the response message, an MTA MUST:
 - 1) Establish inbound RTCP security based on the preferred (first) RTCP ciphersuite, its End-End Secret (which it generated), and a Pad value (if included in the RemoteConnectionDescriptor), as described in 7.6.2.3.3.1.
 - 2) If a RemoteConnectionDescriptor was included and it contained media security attributes, establish outbound RTCP security based on the selected RTCP ciphersuite, End-End Secret (generated by the far-end MTA), and a Pad value (which it may have generated) as described in 7.6.2.3.3.1.

3) Be ready to receive RTCP packets, which may arrive any time after the Response message is sent.

- If, in response to a ModifyConnection command that includes a RemoteConnectionDescriptor, and negotiated lists of ciphersuites for RTP and RTCP contain at least one non-NULL encryption or authentication algorithm each, before sending the response message, an MTA MUST:

- 1) If a Pad was included in the RemoteConnectionDescriptor and it is different from a Pad that may have previously been received, remove any existing inbound RTP keys and generate new ones, based on the keys that are generated from both the End-End Secret (generated locally) and the Pad (generated by the other MTA). The MTA MUST re-initialize the RTP timestamp if new keys are generated. The ciphersuites used for these inbound keys are taken from the RemoteConnectionDescriptor parameter just received from the CMS.
- 2) If a Pad was included in the RemoteConnectionDescriptor and it is different from a Pad that may have previously been received, remove any existing inbound RTCP keys and generate new ones, based on the keys that are generated from both the End-End Secret (generated locally) and the Pad (generated by the other MTA). The MTA MUST re-initialize RTCP sequence numbers if new keys are generated. The ciphersuites used for these inbound keys are taken from the RemoteConnectionDescriptor parameter just received from CMS.
- 3) If the RemoteConnectionDescriptor parameter was received without a Pad, check if the first RTP ciphersuite field in the RemoteConnectionDescriptor parameter differs from the one that the MTA originally selected. Also, check to see if a Pad had been previously received. If the ciphersuites differ, or if a Pad had been previously received, perform the following steps:
 - a) Remove any existing inbound RTP key.
 - b) If the new RTP ciphersuite is non-NULL, generate new inbound RTP keys and RTP timestamp from the same End-End Secret (generated locally) as the last time, as specified in 7.6.2.3.3.1.
- 4) If the RemoteConnectionDescriptor parameter was received without a Pad, check if the first RTCP ciphersuite field in the RemoteConnectionDescriptor parameter differs from the one that the MTA originally selected. Also, check to see if a Pad had been previously received. If the ciphersuites differ, or if a Pad had been previously received, perform the following steps:
 - a) Remove any existing inbound RTCP key.
 - b) If the new RTCP ciphersuite is non-NULL, generate new inbound RTCP keys from the same End-End Secret (generated locally) as the last time, as specified in 7.6.2.3.3.1, and reset the RTCP sequence number to 0.
- 5) If the End-End Secret included in the RemoteConnectionDescriptor has changed or the negotiated RTP ciphersuite has changed, perform the following steps:
 - a) Remove any existing outbound RTP keys.
 - b) If the new list of RTP ciphersuites is non-NULL, generate new outbound RTP keys, based on the End-End Secret (generated by the other MTA) and the Pad (generated locally), and generate a new RTP timestamp.

- 6) If the End-End Secret included in the RemoteConnectionDescriptor has changed or the negotiated RTCP ciphersuite has changed, perform the following steps:
 - a) Remove any existing outbound RTCP keys
 - b) If the new list of RTCP ciphersuites is non-NULL, generate new outbound RTCP keys, based on the End-End Secret (generated by the other MTA) and the Pad (generated locally), and reset the RTCP sequence number to 0.
 - 7) Be ready to send RTCP messages to and receive RTCP messages from the remote MTA. If connection mode allows, be ready to send and receive RTP messages with the remote MTA. If the list of ciphersuites for RTP was sent within a ModifyConnection command, the CMS MAY send an inactive directive to the MTA in the same command. The MTA should be returned to active status only when the new ciphersuite negotiation is complete.
- If, in response to a ModifyConnection command that does not include a RemoteConnectionDescriptor, and negotiated lists of ciphersuites for RTP and RTCP contain at least one non-NULL encryption or authentication algorithm each, before sending the response message, an MTA MUST:
 - 1) If the first RTP ciphersuites field in the negotiated list differs from the one that the MTA previously selected, then perform the following steps:
 - a) Remove any existing inbound RTP keys.
 - b) Generate new inbound RTP keys from the previous End-End Secret (locally generated) and Pad (generated by the other MTA), and generate a new RTP timestamp.
 - 2) If the first RTCP ciphersuites field in the negotiated list differs from the one that the MTA previously selected, then perform the following steps:
 - a) Remove any existing inbound RTCP keys.
 - b) Generate new inbound RTCP keys from the previous End-End Secret (locally generated) and Pad (generated by the other MTA), and reset the RTCP sequence number to 0.
 - 3) Be ready to send RTCP messages to and receive RTCP messages from the remote MTA. If connection mode allows, be ready to send and receive RTP messages with the remote MTA. If the list of ciphersuites for RTP was sent within a ModifyConnection command, the CMS MAY send an inactive directive to the MTA in the same command. The MTA should be returned to active status only when the new ciphersuite negotiation is complete.
 - If an MTA receives a ModifyConnection command, and the resulting intersection of ciphersuites results in NULL encryption and authentication algorithms for RTP and RTCP, then the MTA MUST remove any existing RTP and RTCP keys and do not perform security on the RTP and RTCP packets.
 - If an MTA returns a LocalConnectionDescriptor parameter, it MUST return the latest negotiated list of ciphersuites.

The following message flow is informative. Each of the numbered flows in Figure 20 is described below:

- 1) CMS₀ → MTA₀
 CMS₀ may send the allowable lists of ciphersuites for the new communication to MTA₀ in the CreateConnection (CRCX) command, inside the LocalConnectionOptions parameter, if the CMS has been configured to do so. The ciphersuites are provided in the order of preference specified by the operator subject to the constraints specified in 7.6.2.3.1.1. There

can be two lists of ciphersuites, one list for RTP security and one for RTCP security. Each of these two lists may be included to specify the list of allowable ciphersuites, however ciphersuite negotiation will take place for both RTP and RTCP irrespective of whether the lists are included or not.

If RTP and/or RTCP ciphersuites are included but do not adhere to the rules provided in 7.6.2.3.1.1, the MTA returns an error, e.g., 524 (Internal inconsistency in LocalConnectionOptions).

2) MTA₀ → CMS₀

MTA₀ performs ciphersuite negotiation according to the ciphersuite negotiation procedure described above, and returns a non-empty list of RTP ciphersuites in the response message. This list contains the list of MTA₀'s list of allowed ciphersuites in the order of preference specified by CMS₀ if the LocalConnectionOptions ciphersuites parameter(s) is included in step 1, as specified above. If RTP or RTCP ciphersuite negotiation fails, MTA₀ returns an error code as specified above.

If the lists of negotiated ciphersuites for RTP and RTCP contain at least one non-NULL combination each, MTA₀ generates the End-End Secret₀ and Pad₁ value and returns them along with the ciphersuites in the LocalConnectionDescriptor parameter. For further details on the NCS message syntax, refer to ITU-T Rec. J.162. Note that the NULL authentication and NULL encryption combinations will be at the end of each ciphersuite list.

The response message also includes the ConnectionId and the EndpointId for MTA₀ as described in ITU-T Rec. J.162. The pair (ConnectionId EndpointId) uniquely identifies this connection, where the EndpointId is an NCS identifier for MTA₀.

If the list of ciphersuites for RTP contains at least one non-NULL encryption or authentication algorithm, before sending the response message, MTA₀ must:

- a) establish inbound RTP security based on its preferred (first) RTP ciphersuite and End-End Secret₀, as described in 7.6.2.3.3.1;
- b) if connection mode allows, be ready to receive RTP packets, which may arrive any time after this message is sent by the MTA₀. If the list of ciphersuites for RTP was sent within a ModifyConnection command, the CMS may send an inactive directive to the MTA in the same command. The MTA should be returned to active status only when the new ciphersuite negotiation is complete.

If the list of ciphersuites for RTCP contains at least one non-NULL encryption algorithm, before sending the response message, MTA₀ must:

- a) establish inbound RTCP security based on its preferred (first) RTCP ciphersuite and End-End Secret₀, as described in 7.6;
- b) be ready to receive RTCP packets, which may arrive any time after this message is sent by MTA₀.

If MTA₁ decides to use an alternate ciphersuite listed by MTA₀, MTA₀ will later have to update its RTP and RTCP keys. If MTA₁ decides to send MTA₀ packets before ciphersuite negotiation had completed, processing on those packets at MTA₀ will fail (since it assumed a different ciphersuite). If media stream security is disabled (AUTH_NULL/RTP_ENCR_NULL ciphersuite list for RTP and RTCP_AUTH_NULL/RTCP_ENCR_NULL for RTCP), MTA₀ will later have to discard its keys and send and receive RTP and RTCP packets without any security.

3) CMS₀ → CMS₁

CMS₀ must send End-End Secret₀ (if included), Pad₁ (if included) and the list of RTP and RTCP ciphersuites to CMS₁ (local to MTA₁) as selected by MTA₀. CMS₁ will later forward this information to MTA₁. Note that End-End Secret₀ and Pad₁ will not be included if the

RTP and RTCP ciphersuites lists both contain only the NULL authentication and NULL encryption combination.

4) CMS₁ → MTA₁

CMS₁ sends a CreateConnection to MTA₁. CMS₁ may provide lists of approved RTP and RTCP ciphersuites, if the CMS has been configured to do so. The ciphersuites are provided in the order of preference specified by the operator subject to the constraints specified in 7.6.2.3.3.1. The RemoteConnectionDescriptor must be included in this CRCX command. It must contain End-End Secret₀ (if sent in step 3) and Pad₁ (if sent in step 3) received from MTA₀ (via CMS₀). It must also contain the ciphersuites preferred by MTA₀.

5) MTA₁ → CMS₁

MTA₁ has received a CRCX message that contains both LocalConnectionOptions and RemoteConnectionDescriptor parameters and must follow the ciphersuite negotiation procedure described above to negotiate RTP and RTCP ciphersuites. This list will consist of MTA₁'s allowed ciphersuites in the order of preference specified by CMS₁ if the LocalConnectionOptions ciphersuites parameter is included in step 4. If RTP and RTCP ciphersuite negotiation succeeds and there is at least one RTP ciphersuite and at least one RTCP ciphersuite, then MTA₁ returns the negotiated list of ciphersuites in the subsequent response message, in the LocalConnectionDescriptor parameter, in the form of SDP attributes. Note that if media stream security is being disabled, the NULL authentication and NULL encryption combination will be the only entry in both the RTP and RTCP ciphersuites lists. If RTP or RTCP ciphersuite negotiation fails, MTA₁ must return an error code as specified above.

In the event that MTA₁ receives SDP in the RemoteConnectionDescriptor parameter without ciphersuites media attributes, MTA₁ assumes that the lists of RTP and RTCP ciphersuites supported by the remote endpoint is RTP AUTH_NULL/RTP_ENCR_NULL and RTCP_AUTH_NULL/RTCP_ENCR_NULL.

If the RTP and RTCP ciphersuites provided do not adhere to the rules provided in 7.6.2.3.1.1, the MTA returns an error, e.g., 524 (Internal inconsistency in LocalConnectionOptions).

Whenever possible, MTA₁ SHOULD select the first supported ciphersuite for RTP and the first supported ciphersuite for RTCP in the RemoteConnectionDescriptor parameter. This allows MTA₁ to immediately start sending RTP and RTCP packets to MTA₀. MTA₁ MAY instead select alternate ciphersuites specified by MTA₀.

MTA₁ returns a response message, which includes lists of the selected ciphersuites inside the LocalConnectionDescriptor parameter, in the form of SDP attributes. The first ciphersuite in each list (one for RTP and one for RTCP) must be the one that was selected by MTA₁. Additional ciphersuites in each list are alternatives in a prioritized order. If at any time MTA₀ wants to switch to one of the alternatives that were selected by MTA₁, it would have to go through a new key negotiation. The response message must also include the ConnectionId (generated by MTA₁) as specified in ITU-T Rec. J.162. Thus, both End-End Secret₀ and End-End Secret₁ are now associated with a pair (EndpointId, ConnectionId).

If the lists of ciphersuites for RTP and RTCP contain at least one non-NULL selection each, then MTA₁ must generate the End-End Secret₁ for the incoming RTP and RTCP packets, and return it along with the ciphersuite list in the LocalConnectionDescriptor. If the lists of ciphersuites for RTP and RTCP contain at least one non-NULL selection each, MTA₁ should also generate Pad₀ and return it in the same LocalConnectionDescriptor parameter.

Although the option of not generating Pad₀ is provided in order to better support early media flows from MTA₁, it results in MTA₁ using a send key that is completely dependent

on a random value generated by MTA₀. In other words, privacy of the media stream generated by MTA₁ in this case depends on the strength of MTA₀'s random number generator.

If the list of ciphersuites for RTP contains at least one non-NULL encryption or authentication algorithm, before sending the response message, MTA₁ must:

- a) establish inbound RTP security based on its selected RTP ciphersuite, End-End Secret₁ and Pad₁, as described in 7.6.2.3.3.1;
- b) establish outbound RTP security based on its selected RTP ciphersuite and End-End Secret₀, as described in 7.6.2.3.3.1. If Pad₀ was generated by MTA₁, the outbound RTP security will also be based on Pad₀;
- c) if connection mode allows, be ready to receive RTP packets, which may arrive from MTA₀ any time after this message is sent.

If the list of ciphersuites for RTCP contains at least one non-NULL encryption or authentication algorithm, before sending the response message, MTA₁ must:

- a) establish inbound RTCP security based on its selected RTCP ciphersuite, End-End Secret₁ and Pad₁ as described in 7.6.2.3.3.1;
- b) establish outbound RTCP security, based on its selected RTCP ciphersuite and End-End Secret₀, as described in 7.6.2.3.3.1. If Pad₀ was generated by MTA₁, the outbound RTCP security will also be based on Pad₀;
- c) be ready to receive RTCP messages, which may arrive from MTA₀ any time after this message is sent by.

Any time after sending this response message to the CMS₁, MTA₁ may begin sending RTP and RTCP packets to MTA₀. However, in the case that MTA₁ generated Pad₀ or selected a different ciphersuite from the one preferred by MTA₀, MTA₀ will not be able to decrypt packets from MTA₁, until MTA₀ has received MTA₁'s SDP.

6) CMS₁ → CMS₀

CMS₁ must forward the End-End Secret₁ (if included), Pad₀ (if included) and the selected ciphersuites sent from MTA₁ to CMS₀. Note that End-End Secret₀ and Pad₁ will not be included if the RTP and RTCP ciphersuites lists both contain only the NULL authentication and NULL encryption algorithm combination.

7) CMS₀ → MTA₀

CMS₀ may send to MTA₀ in the ModifyConnection (MDCX) command, inside the LocalConnectionOptions parameter, the lists of allowed RTP and RTCP ciphersuites. These ciphersuites should be what CMS₀ policy allows. (The reason that CMS₀ is not required to send the lists of ciphersuites is because it might have already sent them to MTA₀ in a CreateConnection (CRCX) command. CMS₀ would send the ciphersuites again for consistency.)

In the event that MTA₀ receives SDP in the RemoteConnectionDescriptor parameter without fields containing ciphersuites media attributes, MTA₀ assumes that the RTP and RTCP ciphersuite lists supported by the remote endpoint are AUTH_NULL/RTP_ENCR_NULL for RTP and RTCP_AUTH_NULL/RTCP_ENCR_NULL for RTCP.

In the event that CMS₀ received SDP from from MTA₁, the RemoteConnectionDescriptor parameter must be included in this ModifyConnection command. If present, it must contain the RTP and RTCP ciphersuites (and alternatives) selected by MTA₁. If ciphersuites are included in LocalConnectionOptions parameter or a RemoteConnectionDescriptor parameter is included with the ModifyConnection command, MTA₀ must perform ciphersuite negotiation as described above.

If the RemoteConnectionDescriptor is not sent in this MDCX command, MTA₀ will still be able to receive RTP and RTCP messages but will be unable to send anything to MTA₁.

After receiving this message, MTA₀ must:

- a) if Pad₀ was received, remove its inbound RTP keys and replace them with new ones, based on the keys that are generated from both End-End Secret₀ and Pad₀. MTA₀ MUST re-initialize the RTP timestamp for the new keys. The ciphersuites used for these inbound keys are taken from the RemoteConnectionDescriptor just received from CMS₀;
- b) if Pad₀ was received, remove its inbound RTCP keys and replace them with new ones, based on the keys that are generated from both End-End Secret₀ and Pad₀. Re-initialize RTCP sequence numbers for the new keys. The ciphersuites used for these inbound keys are taken from the RemoteConnectionDescriptor just received from CMS₀;
- c) if the RemoteConnectionDescriptor was received without Pad₀, check if the first RTP ciphersuite in the RemoteConnectionDescriptor differs from the one that MTA₀ selected in step b. If they differ, perform the following steps:
 - Remove the inbound RTP key.
 - If the new RTP ciphersuite is non NULL, generate new inbound RTP keys and RTP timestamp from the same End-End Secret₀ as the last time, as specified in 7.6.2.3.3.1;
- d) if the RemoteConnectionDescriptor parameter was received without Pad₀, check if the first RTCP ciphersuite field in the RemoteConnectionDescriptor parameter differs from the one that MTA₀ selected in step b. If they differ, perform the following steps:
 - Remove the inbound RTCP key.
 - If the new RTCP ciphersuite is non NULL, generate a new key based on the key generated from the same End-End Secret₀ as the last time, but for the new authentication and/or encryption algorithms;
- e) if the RemoteConnectionDescriptor parameter was received, establish outbound RTP keys, based on End-End Secret₁ and Pad₁;
- f) if the RemoteConnectionDescriptor parameter was received, establish outbound RTCP keys, based on End-End Secret₁ and Pad₁;
- g) be ready to send and receive RTCP messages with MTA₁. If connection mode allows, be ready to send and receive RTP messages with MTA₁.

For the full syntax of the NCS messages, please refer to the NCS signalling Recommendation (ITU-T Rec. J.162).

7.6.2.3.2 Ciphersuite format

Each ciphersuite for both RTP security and RTCP security MUST be represented as follows:

| | |
|--|---|
| Authentication Algorithm (1 byte) – represented by 2 ASCII hex characters (using characters 0-9, A-F) | Encryption Transform ID (1 byte) – represented by 2 ASCII hex characters (using characters 0-9, A-F) |
|--|---|

For the list of available transforms and their values, refer to 6.6 for RTP security, and to 6.7 for RTCP security. For the exact syntax of how the Authentication Algorithm and the Encryption Transform ID are included in the signalling messages, refer to ITU-T Rec. J.162 for NCS.

7.6.2.3.3 Derivation of end-to-end keys

7.6.2.3.3.1 Initial key derivation

The End-End Secrets MUST be 46 bytes long. The Pad parameters MUST be 46 bytes long.

Keys are independently derived by each MTA from either just the End-End Secret or from the End-End Secret and Pad concatenated together. The Pad may or may not be available – see the call flow details specified in 7.6.2.3.1.

The keys derived from one End-End Secret (and possibly a Pad) MUST be used to secure RTP and RTCP messages directed to only one of the MTAs. There is a separate End-End Secret and a separate Pad value for each direction, negotiated through NCS signalling. The keys MUST be derived as follows, in the specified order:

- 1) RTP (media stream security). Derive a set of the following keys with the derivation function $F(S, \text{"End-End RTP Security Association"})$. Here, S is concatenation of the following binary values, each in MSB-first order:

- a) End-End Secret;
- b) Pad (optional, if it was negotiated through signalling).

The string "End-End RTP Security Association" is taken without quotes and without a terminating null character. Function F (specified in 9.6) is used to recursively generate enough random bytes to produce all of the keys and other parameters that are specified below, in the listed order:

- a) RTP privacy key.
- b) RTP Initial Timestamp (integer value, 4 octets, Big-Endian byte order).
- c) RTP Initialization Key (required when using a block cipher to encrypt the RTP payload). The length MUST be the same as the selected cipher's block size. This value is used to derive the IV according to 7.6.2.1.2.2. The resulting IV is used for the block cipher in CBC mode (if applicable) and for the random pad used to calculate the MMH MAC.
- d) RTP packet MAC key (if MAC option is selected). The requirements for the MMH MAC key can be found in 7.6.2.1.2.1.1

- 2) RTCP security. Derive a set of the following keys in the specified order with the derivation function $F(S, \text{"End-End RTP Control Protocol Security Association"})$. Here, S is concatenation of the following binary values:

- a) End-End Secret.
- b) Pad (optional, if it was negotiated through signalling).

Function F (specified in 9.6) is used to recursively generate enough random bytes to produce all of the keys that are specified below, in the listed order:

- a) RTCP authentication key.
- b) RTCP encryption key.

7.6.2.4 RTP-RTCP summary security profile matrix

Table 27 – Security profile matrix – RTP and RTCP

| | RTP (MTA-MTA, MTA-MG) | RTCP (MTA-MTA, MTA-MG, MG-MG) |
|--|--|---|
| Authentication | Optional (indirect) (Note) | Optional (indirect) |
| Access control | Optional | Optional |
| Integrity | Optional | Optional |
| Confidentiality | Optional | Optional |
| Non-repudiation | No | No |
| Security mechanisms | <p>Application Layer Security via RTP IPCablecom Security Profile End-to-End Secret distributed over secured MTA-CMS links. Final keys derived from this secret. AES-128 in CBC mode encryption algorithm Optional 2-byte or 4-byte MAC based on MMH algorithm. RTP encryption and authentication can be optionally turned off with the selection of NULL encryption and NULL authentication algorithms. RTP security and RTCP security are disabled together.</p> <p>IPCablecom requires support for ciphersuite negotiation.</p> | <p>RTCP messages are secured by RTCP application layer security mechanisms specified in the profile. RTCP ciphersuites are negotiated separately from the RTP ciphersuites and include both encryption and message authentication algorithms. RTCP encryption can be optionally turned off with the selection of a null encryption algorithm. Both RTCP encryption and authentication can be optionally turned off with the selection of NULL encryption and NULL authentication algorithms. RTCP security and RTP security are disabled together. Keys are derived from the End-End Secret using the same mechanism as used for RTP encryption.</p> |
| NOTE – MTAs do not authenticate directly. Authentication refers to the authentication of identity. | | |

7.7 Audio Server services

7.7.1 Reference architecture

Figure 21 shows the network components and the various interfaces to be discussed in this clause; see ITU-T Rec. J.175.

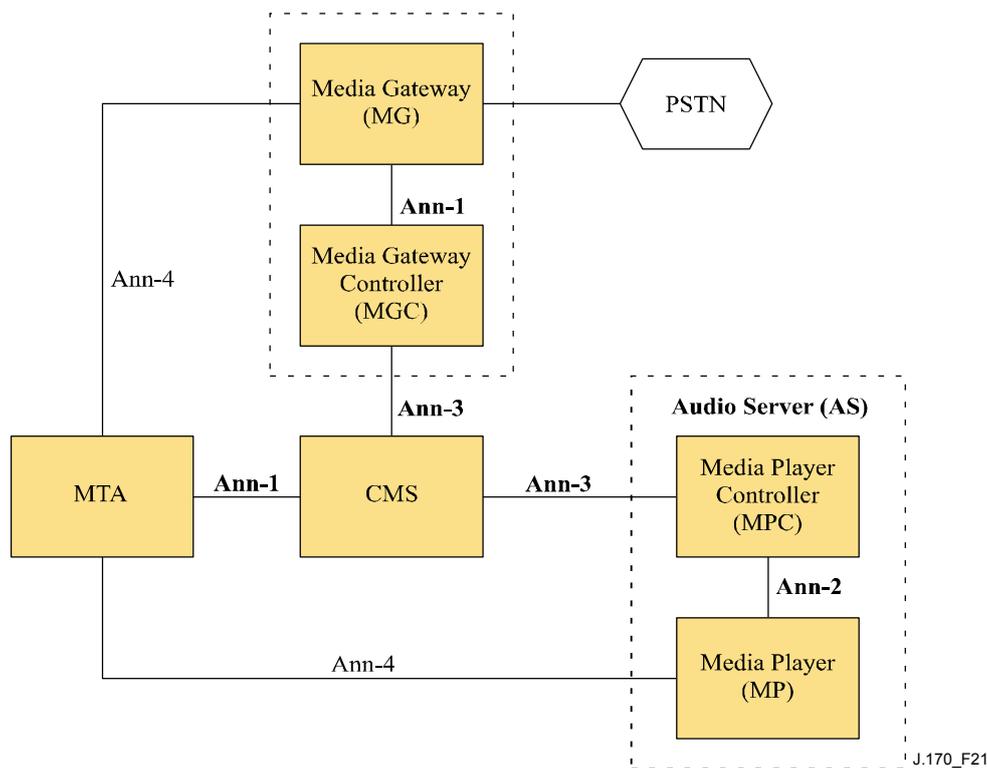


Figure 21 – Audio Server components and interfaces

Figure 21 shows a network-based Media Player (MP). It has an optional Audio Server Protocol (ASP) interface (Ann-2) to the Media Player Controller (MPC), in the case that MPC and MP are not integrated into a single physical entity. Security on this interface is specified in this clause.

There is also an NCS signalling interface (Ann-1) between the MTA and CMS and between the Media Gateway Controller (MGC) and the Media Gateway (MG). Refer to 7.4.1 for NCS signalling security. There is also a signalling interface (Ann-3) between the CMS and the MPC and the CMS and the MGC. This interface is proprietary for IPCablecom, and thus the corresponding security interface is not specified (although this clause lists recommended security services for Ann-3).

Finally, there is a media stream (RTP and RTCP) interface (Ann-4) between the MTA and the MP. This is a standard media stream interface, for which security is defined in 6.6.

The Audio Server architecture also allows local playout of announcements at the MTA. In those cases, an announcement is initiated with NCS signalling between the MTA and the CMS (interface Ann-1). No other interfaces are needed for MTA-based announcement services.

7.7.2 Security services

7.7.2.1 MTA-CMS NCS signalling (Ann-1)

Refer to the security services in NCS signalling in 7.4.1.2.

7.7.2.2 MPC-MP Signalling (Ann-2)

Authentication: All signalling messages must be authenticated, in order to prevent a third party masquerading as either an authorized MPC or MP. A rogue MPC could configure the MP to play obscene or inappropriate messages. A rogue MP could likewise play obscene or inappropriate messages that the MPC did not intend it to play. If MP is unable to authenticate to the MPC, the MPC should not pass it the key for media packets, preventing unauthorized announcement playout.

Confidentiality: If a snooper is able to monitor ASP signalling messages on this interface, he or she might determine which services are used by a particular subscriber or which destinations a

subscriber is communicating to. This information could then be sold for marketing purposes or simply used to spy on other subscribers. Thus, confidentiality is required on this interface.

Message integrity: must be assured in order to prevent tampering with signalling messages. This could lead to payout of obscene or inappropriate messages – see authentication above.

Access control: An MPC should keep a list of valid Media Players and which announcements each supports. Along with authentication, this ensures that wrong announcements are not played out.

7.7.2.3 MTA-MP (Ann-4)

Security services on this media packet interface are listed in 7.6.1.

7.7.3 Cryptographic mechanisms

7.7.3.1 MTA-CMS NCS signalling (Ann-1)

Refer to the cryptographic mechanisms in NCS signalling in 7.4.1.3.

7.7.3.2 MPC-MP signalling (Ann-2)

IPsec ESP MUST be used to both authenticate and encrypt the messages from MPC to MP and vice versa. Refer to 6.1 for details of how IPsec ESP is used within IPCablecom and for the list of available ciphersuites.

7.7.3.3 MTA-MP (Ann-4)

Cryptographic mechanisms on this media packet interface are specified in 7.6.2.

7.7.4 Key management

7.7.4.1 MTA-CMS NCS signalling (Ann-1)

Refer to the key management in NCS signalling in 7.4.1.

7.7.4.2 MPC-MP signalling (Ann-2)

The MPC and the MP negotiate a shared secret (MPC-MP Secret) using IKE or Kerberos (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys). For more information on the use of IKE, refer to 6.2.2. For more information on the use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5.

The key management protocol MUST be running asynchronous to the signalling messages and will guarantee that there is always a valid, non-expired MPC-MP Secret. This shared secret MUST be unique to this particular MPC and MP.

7.7.4.3 MTA-MP (Ann-4)

Key management on the media packet interface is specified in 7.6.2.3. This case is very similar to the key management for the MTA-MG media interface. The flow of signalling messages and the syntax of carrying keys and ciphersuites MUST be the same, except that here MG is replaced with the MP and MGC (which delivers the key to MG) is replaced with MPC (which delivers the key to MP).

7.7.5 MPC-MP summary security profile matrix

The CMS to MPC protocol is not defined in IPCablecom and thus is outside the scope of this Recommendation. The corresponding column in the matrix in Table 28 provides only the security requirements on that interface. Security specifications on that interface will be added in future revisions of this Recommendation.

Table 28 – Security profile matrix – Audio Server services

| | Ann-1: NCS (MTA-CMS) & (MG-MGC) | Ann-2: ASP (MPC- MP) | Ann-3: Unspecified (CMS-MPC) & (CMS-MGC) interface security requirements (Note) | Ann-4: RTP (MTA-MP) | Ann-4: RTCP (MTA-MP) |
|--|---|---|--|---|--|
| Authentication | Yes | Yes | Yes | Yes (indirect) | Yes (indirect) |
| Access control | Yes | Yes | Yes | Optional | Optional |
| Integrity | Yes | Yes | Yes | Optional | Yes |
| Confidentiality | Yes | Yes | Yes | Yes | Yes |
| Non-repudiation | No | No | No | No | No |
| Security mechanisms | IPsec ESP in transport mode Encryption and message integrity both enabled Kerberos with PKINIT key management for MTA-CMS interface IKE or Kerberos for MG-MGC interface | IPsec IKE or Kerberos | | Application Layer Security via RTP Packet Cable Security Profile keys distributed over secured MTA-CMS and MP-MPC links AES-128 encryption algorithm Optional 2-byte or 4-byte MAC based on MMH algorithm. | RTCP messages are secured by RTCP application layer security mechanisms specified in the profile. Keys are derived from the End-End Secret using the same mechanism as used for RTP encryption. |
| NOTE – Although (CMS-MPC) is a proprietary interface, these are security requirements for the CMS-MPC interface. | | | | | |

7.8 Electronic surveillance interfaces

7.8.1 Reference architecture

The IPCablecom system for electronic surveillance consists of elements and interfaces, as shown in Figure 22.

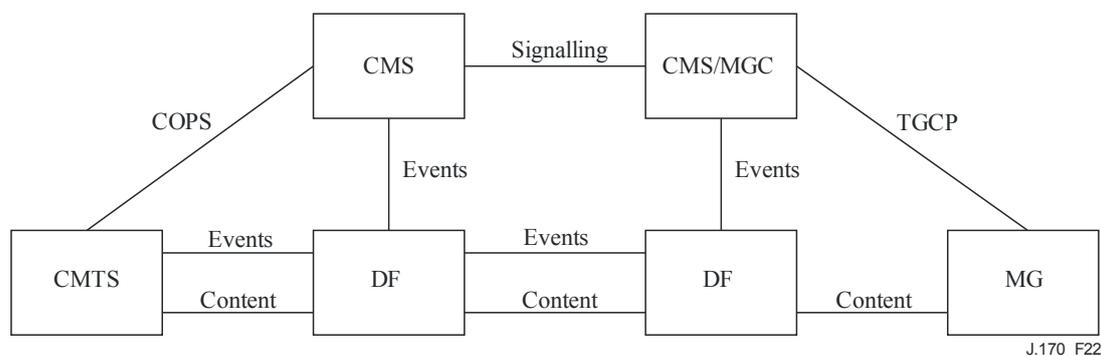


Figure 22 – Electronic surveillance security interfaces

The DF (Delivery Function) in Figure 22 is responsible for redirecting duplicated media packets to law enforcement, for the purpose of wiretapping.

The event interface between the CMS or the MGC and the DF provides descriptions of calls, which is necessary to perform wiretapping. That information includes the media stream encryption key and the corresponding encryption algorithm. This event interface uses RADIUS and is similar to the CMS-RKS interface.

The COPS interface between the CMS and the CMTS is used to signal the CMTS to start/stop duplicating media packets to the DF for a particular call. This is the same COPS interface that is used for (DQoS) Gate Authorization messages. For the corresponding security services, refer to 7.2.1.2.2, 7.2.1.3.2 and 7.2.1.4.1.

The TGCP signalling interface between the MGC and MG is used to signal the MG to start/stop duplicating media packets to the DF for a particular call. This is the same TGCP signalling interface that is used during call set-up on the PSTN Gateway side. For the corresponding security services, refer to 7.8.2.1, 7.8.3.1 and 7.8.4.1.

The event interface between the CMTS and DF is needed to tell the DF when the actual call begins and when it ends. In IPCablecom, the start and end of the actual call is signalled with RADIUS event messages generated by the CMTS.

The interface between the CMTS and DF for call content is where the CMTS encapsulates copies of the RTP media packets – including the original IP header – inside UDP and forwards them to the DF. Since the original media packets are already encrypted (and optionally authenticated), no additional security is defined on this interface. Similarly, there is no additional security applied to the call content interface between the MG and DF: the MG simply encapsulates copies of the encrypted RTP packets inside UDP and forwards them to the DF.

The event interface between the two DFs is used to forward call information in the case where a wiretapped call is forwarded to another location that is wiretapped using a different DF. This interface utilizes the RADIUS protocol – the same as all other event message interfaces.

The interface between the two DFs for call content is used to forward media packets (including the original IP header) in the case where a wiretapped call is forwarded to another location that is wiretapped using a different DF. Since the original media packets are already encrypted (and optionally authenticated), no additional security is defined on this interface.

7.8.2 Security services

7.8.2.1 Event interfaces CMS-DF, MGC-DF, CMTS-DF and DF-DF

Authentication, Access control and Message integrity: required to prevent service theft and denial-of-service attacks. There is a need to ensure that the DF (law enforcement) has the right parameters for wiretapping (prevent denial-of-service). Also, there is a need to authenticate the DF, to make sure that the copy of the media stream is directed to the right place (protect privacy).

Confidentiality: required to protect subscriber information and communication patterns.

7.8.2.2 Call content interfaces CMTS-DF, MG-DF and DF-DF

Authentication and Access control: already performed during the phase of key management for protection of event messages – see the above clause. In order to protect privacy, a party that is not properly authorized should not receive the call content decryption key.

Message integrity: optional for voice packets, since it is generally hard to make undetected changes to voice packets. No additional security is required here – an optional integrity check would be placed into the media packets by the source (MTA or MG).

Confidentiality: required to protect call content from unauthorized snooping.

However, no additional security is required in this case – the packets had been previously encrypted by the source (MTA or MG).

7.8.3 Cryptographic mechanisms

7.8.3.1 Interface between CMS and DF

This interface MUST be protected with IPsec ESP in transport mode, where each packet is both encrypted and authenticated – identical to the security for the CMS-RKS interface specified in 7.3.3.1.

As with the CMS-RKS interface, the MAC value normally used to authenticate RADIUS messages is not used (message integrity is provided with IPsec). The key for this RADIUS MAC MUST always be hardcoded to 16 ASCII 0s.

7.8.3.2 Interface between CMTS and DF for Event Messages

This interface MUST be protected with IPsec ESP in transport mode, where each packet is both encrypted and authenticated – identical to the security for the CMTS-RKS interface specified in 7.3.3.2.

As with the CMTS-RKS interface, the MAC value normally used to authenticate RADIUS messages is not used (message integrity is provided with IPsec). The key for this RADIUS MAC MUST always be hardcoded to 16 ASCII 0s.

7.8.3.3 Interface between DF and DF for event messages

This interface MUST be protected with IPsec ESP in transport mode, where each packet is both encrypted and authenticated – identical to the security for the CMS-RKS interface specified in 7.3.3.1.

As with the CMS-RKS interface, the MAC value normally used to authenticate RADIUS messages is not used (message integrity is provided with IPsec). The key for this RADIUS MAC MUST always be hardcoded to 16 ASCII 0s.

7.8.3.4 Interface between MGC and DF

This interface MUST be protected with IPsec ESP in transport mode, where each packet is both encrypted and authenticated – identical to the security for the MGC-RKS interface specified in 7.3.3.3.

As with the MGC-RKS interface, the MAC value normally used to authenticate RADIUS messages is not used (message integrity is provided by IPsec). The key for this RADIUS MAC MUST always be hardcoded to 16 ASCII 0s.

7.8.4 Key management

7.8.4.1 Interface between CMS and DF

CMS and DF MUST negotiate a pair of IPsec SAs (inbound and outbound) using IKE or Kerberos (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys).

7.8.4.2 Interface between CMTS and DF

The CMTS and the DF MUST negotiate a pair of SAs (inbound and outbound) using IKE or Kerberos (implementations MUST support IKE with pre-shared keys; they MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys).

The key management protocol will be running asynchronous to the event message generation and will guarantee that there is always a valid, non-expired pair of SAs.

7.8.4.3 Interface between DF and DF

The two DF hosts MUST negotiate a shared secret (DF-DF Secret) using IKE with certificates. The IPCablecom profile for IKE with certificates is specified in 6.2.2. IKE will be running asynchronous to the event message generation. In the case where an event message needs to be sent to a DF with which there is not a valid SA, the IPsec layer MUST automatically signal IKE to proceed with the key management exchanges and build a pair of IPsec SAs (inbound and outbound).

Not all interfaces between the same pair of DFs will require IPsec. For example, the call content interface does not run over IPsec. In order for the IPsec SAs to be established only for the DF-DF event message interface, each DF MUST allocate a set of UDP ports on which it will both send and receive DF-DF event messages. IPsec policy database for each DF MUST specify either an enumeration or a range of local UDP ports for which IPsec is enabled and which will be used exclusively for DF-DF event messages. If there are multiple calls that are simultaneously wiretapped and forwarded between the same pair of DFs (on different UDP ports) – they MUST all be protected with a single pair of IPsec SAs (inbound-outbound). Whenever a DF attempts to send on one of those UDP ports, it will either use an existing IPsec SA for a particular destination DF, or it will trigger IKE to establish a pair of SAs (inbound-outbound) for the specific target DF. When the CMS tells a DF to forward event messages to another DF, it specifies the destination DF with an IP address. This means that the DF identity that needs authentication during an IKE exchange is the IP address. An IKE certificate for a DF contains the IP address of that DF. This IP address in the certificate MUST be used by IKE to validate the DF's IP address – to prevent IP address spoofing attacks.

After a pair of DF-DF SAs has been idle for some period of time, a DF MAY decide to remove it. In this case, the DF MUST send an ISAKMP Delete message to the other DF – to notify the other side of the SA deletion. Upon receiving a Delete message, the other DF MUST also remove that pair of SAs.

It will still be possible (with very small probability) that a DF uses an IPsec SA to send an event message to another DF; but when the event message arrives, the target DF has already deleted the corresponding SA and has to drop the message. If there is still a problem after several timeouts and retries (e.g., ISAKMP Delete message was lost in transit), the sending DF MUST remove all of the corresponding IPsec SAs and re-run IKE to set up new SAs.

7.8.4.4 Interface between MGC and DF

MGC and the DF MUST negotiate a pair of IPsec SAs (inbound and outbound) using IKE with pre-shared keys.

IKE will be running asynchronous to the event message generation and will guarantee that there is always a valid, non-expired pair of SAs.

At the DF, MGC Element IDs MUST somehow be associated with the corresponding IP addresses. One possibility is to associate each pre-shared key directly with the Element ID. IKE negotiations will use an ISAKMP identity payload of type ID_KEY_ID to identify the pre-shared key. The value in that identity payload will be the Element ID used in event messages.

Later, when an event message arrives at the DF, it will be able to query the database of SAs and retrieve a source IP address, based on the Element ID. The DF will make sure that it is the same as the source IP address in the IP packet header. One way to query this database is through SNMP, using an IPsec MIB.

7.8.5 Electronic surveillance security profile matrix

Table 29 – Security profile matrix – Electronic surveillance

| | CMS-DF events, MGC-DF events & CMTS-DF events | DF-DF events | CMTS-DF content & MG-DF content | DF-DF content |
|---------------------|--|--|---|---|
| Authentication | Yes | Yes | Yes (indirect) | Yes (indirect) |
| Access control | Yes | Yes | Optional | Optional |
| Integrity | Yes | Yes | Optional | Optional |
| Confidentiality | Yes | Yes | Yes | Yes |
| Non-repudiation | No | No | No | No |
| Security mechanisms | IPsec with encryption and message integrity enabled Key management is IKE or Kerberos | IPsec with encryption and message integrity enabled Key management is IKE with certificates | RTP packets are already encrypted and authenticated by the source (MTA or MG) | RTP packets are already encrypted and authenticated by the source (MTA or MG) |

7.9 CMS provisioning

7.9.1 Reference Architecture

Provisioning is defined as the operations necessary to provide a specified service to a customer. IPCablecom service provisioning can be viewed as two distinct operations: MTA provisioning and CMS subscriber provisioning. CMS provisioning refers to the interface between the Provisioning Server and the CMS.

7.9.2 Security Services

Authentication: provisioning Server needs to be authenticated to prevent a third party from masquerading as a provisioning server to enable services for unauthorized MTAs. CMS needs to be authenticated to prevent someone from impersonating the CMS to receiving provisioning messages, thereby compromising privacy and deny service to provisioned MTAs.

Access Control: required along with authentication to prevent unauthorized access to provisioning data as well as denial-of-service.

Integrity: must be assured to disallow tampering with provisioning messages, in order to prevent a class of denial-of-service attacks.

Confidentiality: provisioning messages contains private customer information, thus confidentiality is required.

7.9.3 Cryptographic mechanisms

IPsec ESP MUST be used to both authenticate and encrypt the messages from CMS to Provisioning Server and vice versa. Refer to 6.1.2 for details of how IPsec ESP is used within IPCablecom and for the list of available ciphersuites.

7.9.4 Key management

Key management for the CMS-Provisioning Server interface is either IKE or Kerberos. Implementations MUST support IKE with pre-shared keys. Implementations MAY support IKE with X.509 certificates and they MAY support Kerberos using symmetric keys. For more information on the IPCablecom use of IKE, refer to 6.2.2. For more information on the IPCablecom use of Kerberos with symmetric keys, refer to 6.4.3 and 6.5

7.9.5 Provisioning Server – CMS summary security profile matrix

Table 30 – Security profile matrix – CMS provisioning

| | CMS-Provisioning Server |
|---------------------|--------------------------------|
| Authentication | Yes |
| Access control | Yes |
| Integrity | Yes |
| Confidentiality | Yes |
| Non-repudiation | No |
| Security mechanisms | IPsec IKE or Kerberos |

8 IPCablecom certificates

IPCablecom uses digital certificates, which comply with ITU-T Rec. X.509 and RFC 2459.

8.1 Generic structure

8.1.1 Version

The version of the certificates MUST be v3. All certificates MUST comply with RFC 2459 except where the non-compliance with the RFC is explicitly stated in this clause.

8.1.2 Public Key type

RSA Public Keys are used throughout the hierarchy. The subjectPublicKeyInfo.algorithm.algorithm Object Identifier (OID) used MUST be 1.2.840.113549.1.1.1 (rsaEncryption).

The public exponent for all RSA IPCablecom keys MUST be F4 – 65537.

8.1.3 Extensions

The following four extension MUST be used as specified in the subclauses below. Any other certificate extensions MAY also be included, but must be marked as non-critical.

8.1.3.1 subjectKeyIdentifier

The subjectKeyIdentifier extension included in all IPCablecom CA certificates as required by RFC 2459 (e.g., all certificates except the device and ancillary certificates) MUST include the keyIdentifier value composed of the 160-bit SHA-1 hash of the value of the BIT STRING subjectPublicKey (excluding the tag, length and number of unused bits from the ASN.1 encoding) (see RFC 2459).

8.1.3.2 authorityKeyIdentifier

The authorityKeyIdentifier extension MUST be included in all IPCablecom certificates with the exception of root certificates, and MUST include a keyIdentifier value that is identical to the subjectKeyIdentifier in the CA certificate.

8.1.3.3 keyUsage

The keyUsage extension MUST be used for all IPCablecom CA certificates and MUST be marked as critical with a value of keyCertSign and cRLSign. A KeyUsage extension MAY be included in end-entity certificates and SHOULD be marked as critical if included as specified in RFC 2459.

8.1.3.4 basicConstraints

The basicConstraints extension MUST be used for all IPCablecom CA certificates and MUST be marked as critical. The values for each certificate for basicConstraints MUST be marked as specified in each of the certificate description tables.

8.1.4 Signature algorithm

The signature mechanism used MUST be SHA-1 with RSA encryption. The specific OID is 1.2.840.113549.1.1.5.

8.1.5 SubjectName and IssuerName

If a string cannot be encoded as a PrintableString, it MUST be encoded as a UTF8String (tag [UNIVERSAL 12]).

When encoding an X.500 Name:

- 1) each RelativeDistinguishedName (RDN) MUST contain only a single element in the set of X.500 attributes;
- 2) the order of the RDNs in an X.500 name MUST be the same as the order in which they are presented in this Recommendation.

It should be noted that ITU-T Rec. X.509 and RFC 2459 define constraints (i.e., upper bounds) on the length of the attribute values. For example, the maximum length for common name (CN), organization name (O) and organizational unit (OU) name values is 64 characters. Where this Recommendation mandates the inclusion of a static string in one of these values (i.e., CN=<Company> System Operator CA), companies MUST ensure that the addition of their identifying information does not cause the total length of the value to exceed the upper bound. In the case where a company's name causes the length of the value to exceed the upper bound, the vendor MUST truncate or abbreviate their information to ensure the total length does not exceed the upper bound.

8.1.6 Certificate profile notation

The tables below use the following notation:

- Extension details are specified by [c:critical, n:non-critical; m:mandatory, o:optional].
- Optional subject naming attributes are surrounded by square brackets (e.g., [L = <city>]).
- Variable naming attribute values are surrounded by angle brackets. (e.g., CN = <Company Name> IPCablecom CA). Values not surrounded by angle brackets are static and cannot be modified.

8.2 Certificate trust hierarchy

There are two distinct certificate hierarchies used in IPCablecom: the MTA Device Certificate Hierarchy and the IPCablecom Telephony Certificate Hierarchy (see Figure 23).

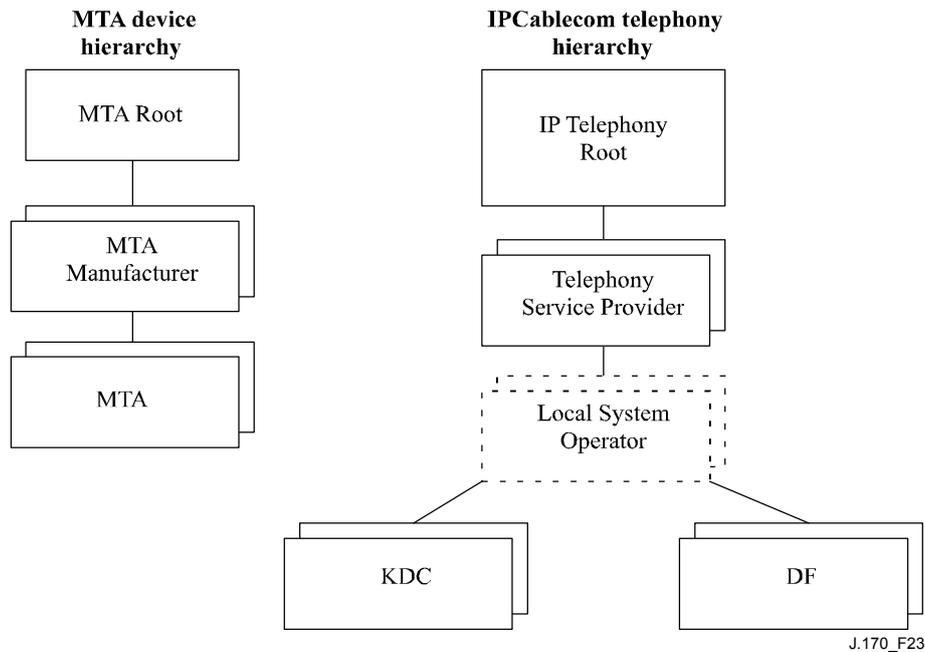


Figure 23 – IPCablecom certificate hierarchy

8.2.1 Certificate validation

Within IPCablecom, certificate validation, in general, involves validation of a whole chain of certificates. As an example, when the Provisioning Server validates an MTA Device Certificate, the actual chain of three certificates is validated:

MTA Root Certificate + MTA Manufacturer Certificate + MTA Device Certificate

The signature on the MTA Manufacturer Certificate is verified with the MTA Root Certificate and the signature on the MTA Device Certificate is verified with the MTA Manufacturer Certificate. The MTA Root Certificate is self-signed and is known in advance to the Provisioning Server. The public key present in the MTA Root Certificate is used to validate the signature on this same certificate.

Usually, the first certificate in the chain is not explicitly included in the certificate chain that is sent over the wire. In the cases where the first certificate is explicitly included it MUST already be known to the verifying party ahead of time and MUST NOT contain any changes to the certificate with the possible exception of the certificate serial number, validity period and the value of the signature. If changes other than the certificate serial number, validity period and the value of the signature, exist in the IP Telephony Root certificate that was passed over the wire in comparison to the known IP Telephony Root certificate, the device making the comparison MUST fail the certificate verification.

The exact rules for certificate chain validation must fully comply with RFC 2459, where they are referred to as "Certificate Path Validation". In general, X.509 certificates support a liberal set of rules for determining if the issuer name of a certificate matches the subject name of another. The rules are such that two name fields may be declared to match even though a binary comparison of the two name fields does not indicate a match. RFC 2459 recommends that certificate authorities restrict the encoding of name fields so that an implementation can declare a match or mismatch using simple binary comparison. IPCablecom security follows this recommendation. Accordingly, the DER-encoded `tbsCertificate.issuer` field of an IPCablecom certificate MUST be an exact match to the DER-encoded `tbsCertificate.subject` field of its issuer certificate. An implementation MAY compare an issuer name to a subject name by performing a binary comparison of the DER-encoded `tbsCertificate.issuer` and `tbsCertificate.subject` fields.

The clauses below specify the required certificate chain, which must be used to verify each certificate that appears at the leaf node (i.e., at the bottom) in the IPCablecom certificate trust hierarchy illustrated in Figure 23.

Validity period nesting is not checked and intentionally not enforced. Thus, the validity period of a certificate need not fall within the validity period of the certificate that issued it.

8.2.2 MTA Device certificate hierarchy

The device certificate hierarchy is rooted in an IPCablecom MTA Root certificate, which is used as the issuing certificate of a set of manufacturer's certificates. The manufacturer's certificates are used to sign the individual device certificates.

The information contained in Table 31 contains the IPCablecom specific values for the required fields according to RFC 2459. These IPCablecom specific values MUST be followed according to Table 31, except that Validity Periods SHOULD be as given in the tables. If a required field is not specifically listed for IPCablecom, then the guidelines in RFC 2459 MUST be followed.

8.2.2.1 MTA Root Certificate

This certificate MUST be verified as part of a certificate chain containing the MTA Root Certificate, MTA Manufacturer Certificate and the MTA Device Certificate.

Table 31 – MTA Root Certificate

| | |
|-------------------|--|
| Subject name form | C=US, O=CableLabs, OU=PacketCable, CN=PacketCable Root Device Certificate Authority |
| Intended usage | This certificate is used to sign MTA Manufacturer Certificates and is used by the Provisioning Server. This certificate is not used by the MTAs and thus does not appear in the MTA MIB. |
| Signed by | Self-signed |
| Validity period | 20+ years. It is intended that the validity period be long enough so that this certificate is never re-issued. |
| Modulus length | 2048 |
| Extensions | KeyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier, basicConstraints[c,m](cA=true, pathLenConstraint=1) |

8.2.2.2 MTA Manufacturer Certificate

This certificate MUST be verified as part of a certificate chain containing the MTA Root Certificate, MTA Manufacturer Certificate (see Table 32) and the MTA Device Certificate.

The state/province, city and manufacturer's facility are optional attributes. A manufacturer may have more than one manufacturer's certificate, and there may exist one or more certificates per manufacturer. All certificates for the same manufacturer MUST be provided to each MTA either at manufacture time or during a field update. The MTA MUST select an appropriate certificate for its use by matching the issuer name in the MTA Device Certificate with the subject name in the MTA Manufacturer Certificate. If present, the authorityKeyIdentifier of the device certificate MUST be matched to the subjectKeyIdentifier of the manufacturer certificate as described in RFC 2459.

The <CompanyName> field that is present in O and CN MAY be different in the two instances.

Table 32 – MTA Manufacturer Certificate

| | |
|-------------------|--|
| Subject name form | C=<country>, O=<CompanyName>, [S=<state/province>], [L=<city>], OU=PacketCable, [OU=<Manufacturer's Facility>], CN=<CompanyName> PacketCable CA |
| Intended usage | This certificate is issued to each MTA manufacturer and can be provided to each MTA as part of the secure code download as specified by this Recommendation (either at manufacture time, or during a field update). This certificate appears as a read-only parameter in the MTA MIB. This certificate along with the MTA Device Certificate is used to authenticate the MTA device identity (MAC address) during provisioning. |
| Signed by | MTA Root Certificate CA |
| Validity period | 20 years |
| Modulus length | 2048 |
| Extensions | keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier, authorityKeyIdentifier basicConstraints[c,m](cA=true, pathLenConstraint=0) |

8.2.2.3 MTA Device Certificate

This certificate MUST be verified as part of a certificate chain containing the MTA Root Certificate, MTA Manufacturer Certificate and the MTA Device Certificate (see Table 33).

The state/province, city and manufacturer's facility are optional attributes. The Manufacturer's Facility OU field, (if present) MAY be different from the Manufacturer's Facility OU field (if present) in the MTA Manufacturer certificate.

The MAC address MUST be expressed as six pairs of hexadecimal digits separated by colons, e.g., "00:60:21:A5:0A:23". The alpha hex characters (A-F) MUST be expressed as uppercase letters.

The MTA device certificate should not be replaced or renewed.

Table 33 – MTA Device Certificate

| | |
|-------------------|---|
| Subject name form | C=<country>, O=<Company Name>, [S=<state/province>], [L=<city>], OU=PacketCable, [OU=<Product Name>], [OU=<Manufacturer's Facility>], CN=<MAC Address> |
| Intended usage | This certificate is issued by the MTA manufacturer and installed in the factory. The provisioning server cannot update this certificate. This certificate appears as a read-only parameter in the MTA MIB. This certificate is used to authenticate the MTA device identity (MAC address) during provisioning. |
| Signed by | MTA Manufacturer Certificate CA |
| Validity period | At least 20 years |
| Modulus length | 1024, 1536 or 2048 |
| Extensions | keyUsage[n,o](digitalSignature, keyEncipherment), authorityKeyIdentifier[n,m](keyIdentifier=<subjectKeyIdentifier value from CA certificate>) |

8.2.3 IPCablecom Telephony Certificate Hierarchy

The Service Provider Certificate Hierarchy is rooted in an IP Telephony Root certificate. That certificate is used as the issuing certificate of a set of service provider's certificates. The service provider's certificates are used to sign an optional local system certificate. If the local system

certificate exists, then that is used to sign the ancillary equipment certificates; otherwise, the ancillary certificates are signed by the Service Provider's CA.

The information contained in Table 34 contains the IPCablecom specific values for the required fields according to RFC 2459. These IPCablecom specific values **MUST** be followed according to Table 34 except that Validity Periods **SHOULD** be as given in the tables. If a required field is not specifically listed for IPCablecom, then the guidelines in RFC 2459 **MUST** be followed.

8.2.3.1 IP Telephony Root Certificate

Before any Kerberos key management can be performed, an MTA and a KDC need to perform mutual authentication using the PKINIT extension to the Kerberos protocol. An MTA authenticates a KDC after it receives a PKINIT Reply message containing a KDC certificate chain. In authenticating the KDC, the MTA verifies the KDC certificate chain, including KDC's Service Provider Certificate signed by the IP Telephony Root CA.

Table 34 – IP Telephony Root Certificate

| | |
|-------------------|--|
| Subject name form | C=US, O=CableLabs, OU=PacketCable, CN=PacketCable Root IP Telephony Certificate Authority |
| Intended usage | This certificate is used to sign Telephony Service Provider certificates. This certificate is installed into each MTA at the time of manufacture or with a secure code download as specified by this Recommendation and cannot be updated by the Provisioning Server. Neither this root certificate nor the corresponding public key appears in the MTA MIB. |
| Signed by | Self-signed |
| Validity period | 20+ years. It is intended that the validity period be long enough so that this certificate is never re-issued. |
| Modulus length | 2048 |
| Extensions | keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier, basicConstraints[c,m](cA=true, pathLenConstraint=2) |

8.2.3.2 Telephony Service Provider Certificate

This is the certificate (see Table 35) held by the telephony service provider, signed by the IP Telephony Root CA. It is verified as part of a certificate chain that includes the IP Telephony Root Certificate, Telephony Service Provider Certificate, optional Local System Certificate and an end-entity server certificate. The authenticating entities normally already possess the IP Telephony Root Certificate and it is not transmitted with the rest of the certificate chain.

The fact that a Telephony Service Provider CA Certificate is always explicitly included in the certificate chain allows a Service Provider the flexibility to change its certificate without requiring re-configuration of each entity that validates this certificate chain (e.g., MTA validating a PKINIT Reply). Each time the Service Provider CA Certificate changes, its signature **MUST** be verified with the IP Telephony Root Certificate. However, a new certificate for the same Service Provider **MUST** preserve the same value of the OrganizationName attribute in the SubjectName.

The <Company> field that is present in O and CN **MAY** be different in the two instances.

Table 35 – Telephony Service Provider Certificate

| | |
|-------------------|---|
| Subject name form | C=<country>, O=<Company>, OU=PacketCable, CN=<Company> PacketCable System Operator CA |
| Intended usage | <p>This certificate corresponds to a top-level Certification Authority within a domain of a single Service Provider. In order to make it easy to update this certificate, each network element is configured with the OrganizationName attribute of the Service Provider Certificate SubjectName. This is the only attribute in the certificate that must remain constant.</p> <p>In the case of an MTA, there is a read-write parameter in the MIB that identifies the OrganizationName attribute for each Kerberos realm (that may be shared among multiple MTA endpoints). The MTA does not accept Service Provider certificates that do not match this value of the OrganizationName attribute in the SubjectName.</p> <p>An MTA needs to perform the first PKINIT exchange with the MSO KDC right after a reboot, at which time its MIB tables are not yet configured. At that time, the MTA MUST accept any Service Provider OrganizationName attribute, but it MUST later check that the value added into the MIB for this realm is the same as the one in the initial PKINIT Reply.</p> |
| Signed by | Signed by IP Telephony Root Certificate |
| Validity period | 20 years |
| Modulus length | 2048 |
| Extensions | keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier[n,m], authorityKeyIdentifier[n,m](keyIdentifier=<subjectKeyIdentifier value from CA certificate>), basicConstraints[c,m](cA=true, pathLenConstraint=1) |

8.2.3.3 Local System CA Certificate

This is the certificate (see Table 36) held by the local system. The existence of this certificate is optional, as the Telephony Service Provider CA may be used to directly sign all network server end-entity certificates. A certificate chain with a Local System Certificate MUST consist of the IP Telephony Root CA Certificate, Service Provider CA Certificate, Local System CA Certificate and an end-entity certificate.

The <Company> field that is present in O and CN MAY be different in the two instances.

Table 36 – Local System Certificate

| | |
|-------------------|--|
| Subject name form | C=<Country>, O=<Company>, OU=PacketCable, OU=<Local System Name>, CN=<Company> PacketCable Local System CA |
| Intended usage | <p>Telephony Service Provider CA may delegate the issuance of certificates to a regional Certification Authority called Local System CA (with the corresponding Local System Certificate).</p> <p>Network servers are allowed to move freely between regional Certification Authorities of the same Service Provider. Therefore, the MTA MIB does not contain any information regarding a Local System Certificate (which might restrict an MTA to KDCs within a particular region).</p> |
| Signed by | Telephony Service Provider Certificate |
| Validity period | 20 years |
| Modulus length | 1024, 1536, 2048 |
| Extensions | keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier-[n,m], authorityKeyIdentifier-[n,m](keyIdentifier=<subjectKeyIdentifier value from CA certificate>), basicConstraints[c,m](cA=true, pathLenConstraint=0) |

8.2.3.4 Operational ancillary certificates

All of these are signed by either the Local System CA or by the Telephony Service Provider CA. Other ancillary certificates may be added to this Recommendation at a later time.

8.2.3.4.1 Key Distribution Centre Certificate

This certificate (see Table 37) MUST be verified as part of a certificate chain containing the IP Telephony Root Certificate, Service Provider Certificate and the Ancillary Device Certificates.

The PKINIT specification in IETF RFC 4556 requires the KDC certificate to include the subjectAltName v3 certificate extension, the value of which must be the Kerberos principal name of the KDC.

Table 37 – Key Distribution Centre Certificate

| | |
|-------------------|--|
| Subject name form | C=<Country>, O=<Company>, OU=PacketCable, OU=[<Local System Name>], OU= Key Distribution Center, CN=<DNS Name> |
| Intended usage | To authenticate the identity of the KDC server to the MTA during PKINIT exchanges. This certificate is passed to the MTA inside the PKINIT replies and is therefore not included in the MTA MIB and cannot be updated or queried by the Provisioning Server. |
| Signed by | Telephony Service Provider Certificate or Local System Certificate |
| Validity period | 20 years |
| Modulus length | 1024, 1536 or 2048 |
| Extensions | keyUsage[n,o](digitalSignature), authorityKeyIdentifier The keyUsage tag is optional. When it is used it SHOULD be marked as critical. subjectAltName[n,m](see PKINIT Spec) |

8.2.3.4.2 Delivery Function (DF) Certificate

This certificate (see Table 38) MUST be verified as part of a certificate chain containing the IP Telephony Root Certificate, Service Provider Certificate and the Ancillary Device Certificates.

This certificate is used to sign phase 1 IKE intra-domain exchanges between DFs (which are used in Electronic Surveillance). Although Local System Name is optional, it is REQUIRED when the Local System CA signs this certificate. The IP address MUST be specified in standard dotted-quad notation, e.g., 245.120.75.22.

Table 38 – DF Certificate

| | |
|-------------------|---|
| Subject name form | C=<Country>, O=<Company>, OU=[<Local System Name>], OU=IPCablecom Electronic Surveillance, CN=<IP address> |
| Intended usage | To authenticate IKE key management, used to establish IPsec Security Associations between pairs of DFs. These Security Associations are used when a subject that is being legally wiretapped forwards the call and event messages containing call info have to be forwarded to a new wiretap server (DF). |
| Signed by | Telephony Service Provider Certificate or Local System Certificate |
| Validity period | 20 years |
| Modulus length | 2048 |
| Extensions | KeyUsage[n,o](digitalSignature), authorityKeyIdentifier[n,m](keyIdentifier=<subjectKeyIdentifier value from CA certificate>) subjectAltName[n,m](dNSName=<DNSName>) extendedKeyUsage[n,m](iKEIntermediate) |

8.2.3.4.3 IPCablecom Server Certificates

These certificates MUST be verified as part of a certificate chain containing the Service Provider Root Certificate, Service Provider Certificate, Local System Operator Certificate (if used) and the Ancillary Device Certificates.

These certificates are used to identify various servers in the IPCablecom system. For example, they may be used to sign phase 1 IKE exchanges or to authenticate a PKINIT exchange. Although the Local System Name is optional, it is REQUIRED when the Local System CA signs this certificate. IP address values MUST be specified in standard dotted decimal notation: e.g., 245.120.75.22. DNS Name values MUST be specified as a fully qualified domain name (FQDN): e.g., device.packetcable.com.

Table 39 – IPCablecom Server Certificates

| IPCablecom Server Certificates | |
|---------------------------------------|--|
| Subject name form | <p>C=<Country> O=<Company> OU=PacketCable OU=[<Local System Name>] OU=<Sub-System Name>[&<Sub-System Name>] CN=<Server Identifier>]</p> <p>or</p> <p>CN=[<Element ID>][&<Element ID>]</p> <p>The CN will contain either a <Server Identifier> or one or more <Element ID>s. If the CN contains a <Server Identifier>, the value of <Server Identifier> MUST be the server's FQDN or its IP address, optionally followed by a colon (:) and an Element ID with no white space either before or after the colon.</p> <p><Element ID> is the identifier that appears in billing event messages; it MUST be included in a certificate of every server that is capable of generating event messages. This includes a CMS, CMTS and MGC. There MAY be multiple <Element ID> fields, each separated by the character "&". ITU-T Rec. J.164 defines the Element ID as a 5-octet right-justified, space-padded ASCII-encoded numerical string. When converting the Element ID for use in a certificate, any spaces MUST be converted to ASCII zeroes (0x30).</p> <p>For example, a CMTS that has the Element ID "311" will have a common name "00311".</p> <p>The value of <Sub-System Name> MUST be one of the following:</p> <ul style="list-style-type: none"> – for Border Proxy: bp – for Cable Modem Termination System: cmts – for Call Management Server: cms – for Media Gateway: mg – for Media Gateway Controller: mgc – for Media Player: mp – for Media Player Controller: mpc – for Provisioning Server: ps – for Record Keeping Server: rks – for Signalling Gateway: sg <p>Components that contain combined elements (such as a CMS with an integrated MGC) MUST indicate this in the Subject Name by including all Sub-System Names, joined with the character "&", in the OU field. In the case of combined elements, a single Element ID or multiple Element IDs may be used. If multiple Element IDs are used, all Element IDs MUST be included in the CN, and the order of these Element IDs MUST correspond to the order of the Sub-System Name fields in the OU. The following is an example OU and CN for a combined CMS and MGC. The CMS with Element ID "311" and a MGC with Element ID "312".</p> |

Table 39 – IPCablecom Server Certificates

| IPCablecom Server Certificates | |
|---------------------------------------|---|
| Subject name form | OU=cms&mgc CN=00311&00312 The following is an example OU and CN for a combined CMS and MGC. In this case, the CMS and MGC share a single Element ID of "311". OU=cms&mgc CN=00311&00311 |
| Intended usage | These certificates are used to identify various servers in the IPCablecom system. For example they may be used to sign phase 1 IKE exchanges or to authenticate a device in a PKINIT exchange. |
| Signed by | Telephony Service Provider Certificate or Local System Certificate |
| Validity period | Set by MSO policy |
| Modulus length | 2048 |
| Extensions | keyUsage[c,o](digitalSignature, keyEncipherment) authorityKeyIdentifier[n,m](keyIdentifier=<subjectKeyIdentifier value from CA cert>) subjectAltName[n,o](dNSName=<DNSName> iPAddress=<IP Address Name>) The keyUsage tag is optional. When it is used it MUST be marked as critical. The subjectAltName extension MUST be included for all servers that are capable of generating event messages. For all other servers, the subjectAltName extension MAY be included. If the subjectAltName extension is included, it MUST include the corresponding name value as specified in the CN field of the subject. |

8.2.3.4.4 TLS Certificates

These certificates MUST be verified as part of a certificate chain containing the Service Provider Root Certificate, Service Provider Certificate, Local System Operator Certificate (if used) and the Ancillary Device Certificates.

These certificates are used to authenticate TLS handshake exchanges (and encrypt when using RSA key exchange). Although the Local System Name is optional, it is REQUIRED when the Local System CA signs this certificate. DNS Name values MUST be specified as a fully qualified domain name (FQDN): e.g., device.packetcable.com.

Table 40 – TLS Certificates

| Server Certificates | |
|----------------------------|---|
| Subject name form | C=<Country> O=<Company> OU=[<Local System Name>] OU= IPCablecom CN=[<Server Identifier>] The value of <Server Identifier> MUST be the server's FQDN. Note that only a single FQDN can be included in the CN field. |
| Intended usage | These certificates are used to authenticate TLS handshake exchanges (and encrypt when using RSA key exchange). |
| Signed by | Telephony Service Provider Certificate or Local System Certificate |
| Validity period | Set by operator policy |
| Modulus length | 1024, 1536, 2048 |
| Extensions | KeyUsage[c,m](digitalSignature, keyEncipherment) extendedKeyUsage[n,m] (id-kp-serverAuth, id-kp-clientAuth) authorityKeyIdentifier[n,m](keyIdentifier=<subjectKeyIdentifier value from CA cert>) |

8.2.4 Certificate revocation

For future study.

9 Cryptographic algorithms

This clause describes the cryptographic algorithms used in the IPCablecom security Recommendation. When a particular algorithm is used, the algorithm MUST follow the corresponding specification.

9.1 AES

AES-128 is a 128-bit block cipher that MUST be implemented according to the Advanced Encryption Standard (AES) proposal specified in AES – The Rijndael Block Cipher. AES-128 is used in CBC mode with a 128-bit block size in IPCablecom. AES-128 requires 10 rounds of cryptographic operations in encryption or decryption. The Initialization Vector for CBC mode is specified for each use of AES in IPCablecom.

In 1997, the National Institute of Standards and Technology (NIST) initiated a process to select a symmetric-key encryption algorithm to be used to protect sensitive (unclassified) Federal information in furtherance of NIST's statutory responsibilities. In 1998, NIST announced the acceptance of fifteen candidate algorithms and requested the assistance of the cryptographic research community in analysing the candidates. This analysis included an initial examination of the security and efficiency characteristics for each algorithm. NIST reviewed the results of this preliminary research and selected MARS, RC6(TM), Rijndael, Serpent and Twofish as finalists. Having reviewed further public analysis of the finalists, NIST has decided to propose Rijndael as the Advanced Encryption Standard.

9.2 DES

The Data Encryption Standard (DES) is specified in FIPS PUB 81. For Media Stream encryption, IPCablecom does not require error checking on the DES key, and the full 64 bits of key provided to the DES algorithm will be generated according to 7.6.2.3.3.1.

9.2.1 XDESX

An option for the encryption of RTP packets is DESX-XEX. XDESX, or DESX, has been proven as a viable method for overcoming the weaknesses in DES while not greatly adding to the implementation complexity. The strength of DESX against key search attacks is presented in FIPS PUB 81. The CBC mode of DESX-XEX is shown in Figure 24 below, where DESX-XEX is executed within the block called "block cipher". Inside the block, DESX-XEX is performed as shown in Figure 26 below using a 192-bit key. K1 is the first 8 bytes of the key, and K2 represents the second 8 bytes of the key; and K3 the third 8 bytes of the key.

9.2.2 DES-CBC-PAD

This variant of DES is also based on the analysis of DESX presented in *How to protect DES Against Exhaustive Key Search*. When using DESX in CBC mode, an optimized architecture is possible. It can be described in terms of the DES-CBC configuration plus the application of a random pad on the final DES-CBC output blocks. This configuration uses 128 bits of keying material, where 64 bits are applied to the DES block according to FIPS PUB 81, and an additional 64 bits of keying material are applied as the random pad on the final DES-CBC output blocks.

In this case, the same IV used to initialize the CBC mode is used as keying material for the random pad. Each block of DES-CBC encrypted output is XORed with the 64-bit Initialization Vector that was used to start the CBC operation. If a short block results from using residual block termination (see 9.3), the left-most-bits of the IV are used in the final XOR padding operation. This mode of DES-CBC is shown in Figure 25 below, where DES is executed in the block called "block cipher". A 64-bit key value is used.

9.2.3 3DES-EDE

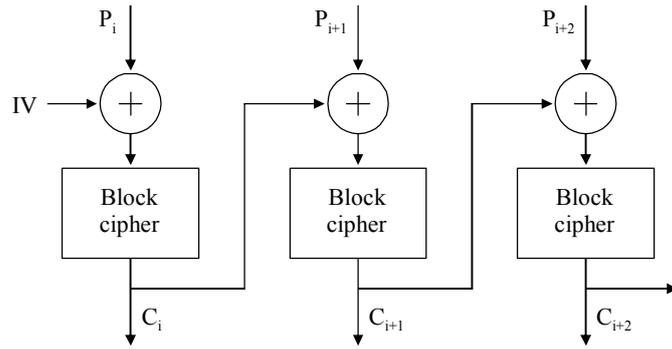
Another option for the encryption of RTP packets for IPCablecom is 3DES-EDE-CBC. The CBC mode of 3DES-EDE is shown in Figure 24 below, where 3DES-EDE is executed within the block called "block cipher." Inside the block, 3DES-EDE is performed as shown in Figure 27 below using a 128-bit key. K1 is the first 8 bytes of the key, and K2 represents the second 8 bytes of key; and K3 = K1.

9.3 Block termination

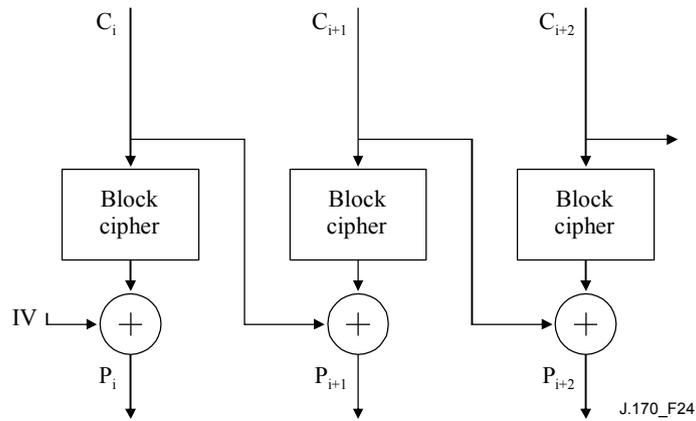
If block ciphers are supported, a short block (n bits $<$ block size, depending on the cipher algorithms) MUST be terminated by residual block termination as shown in Figure 28. Residual block termination (RBT) is executed as follows:

Given a final block having n bits, where n is less than the block size, the n bits are padded up to a block by appending (block size $- n$) bits of arbitrary value to the right of the n -bits. The resulting block is encrypted using B-bit CFB mode, with the next-to-last ciphertext block serving as the initialization vector for the CFB operation (see B. Schneier's *Applied Cryptography*). Here, B stands for the cipher-specific block size. The leftmost n bits of the resulting ciphertext are used as the short cipher block. In the special case where the complete payload is less than the cipher block size, the procedure is the same as for a short final block, with the provided initialization vector serving as the initialization vector for the operation. Residual block termination is illustrated in Figure 28 for both encryption and decryption operations.

CBC encryption architecture



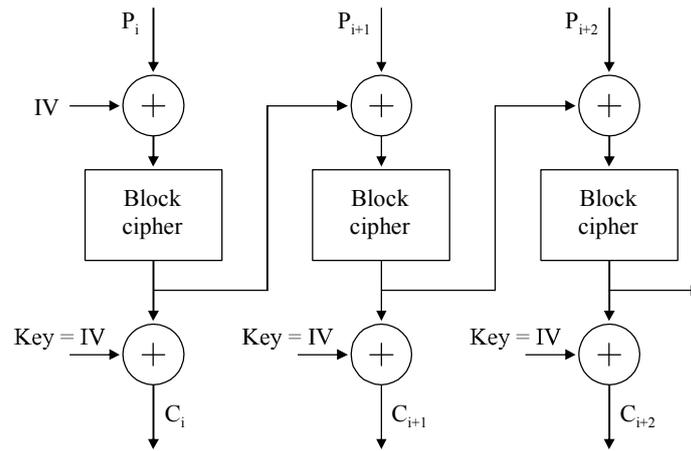
CBC decryption architecture



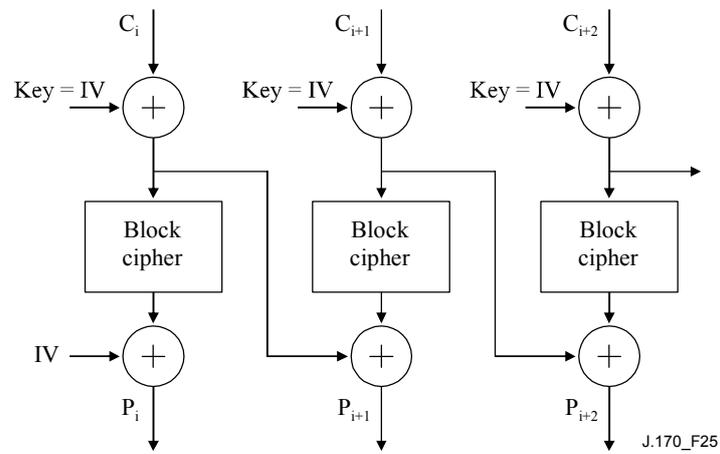
J.170_F24

Figure 24 – CBC Mode

CBC-PAD encryption architecture



CBC-PAD decryption architecture



J.170_F25

Figure 25 – CBC Pad Mode

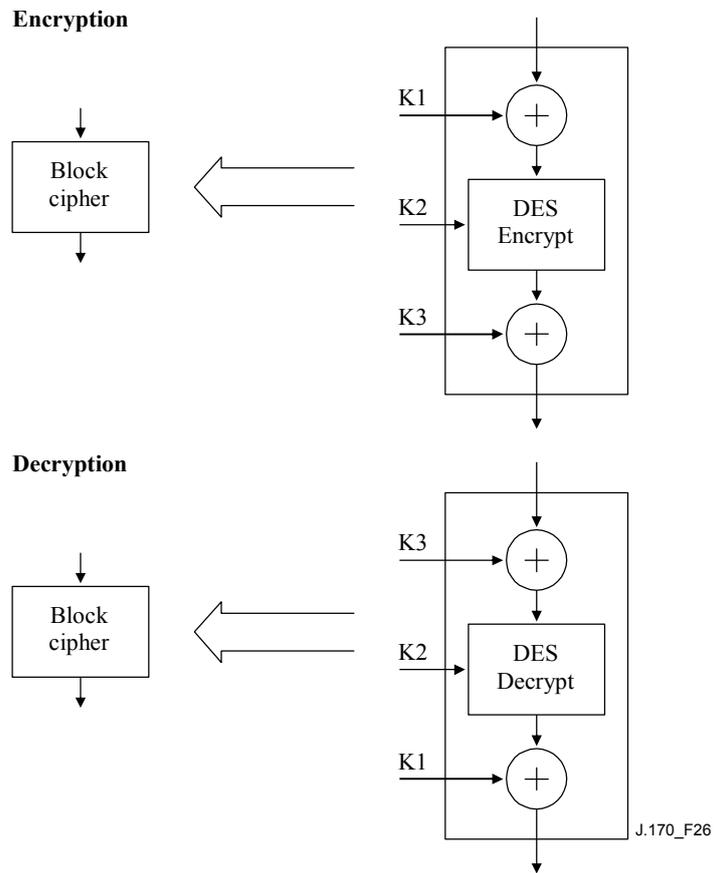


Figure 26 – DESX-XEX as Block Cipher

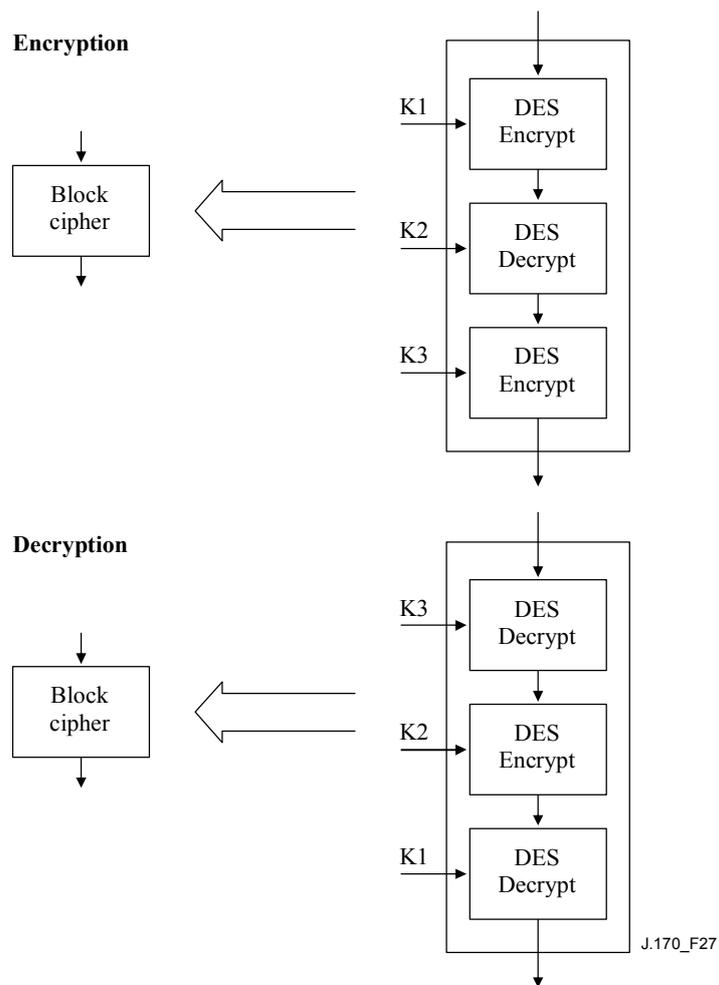
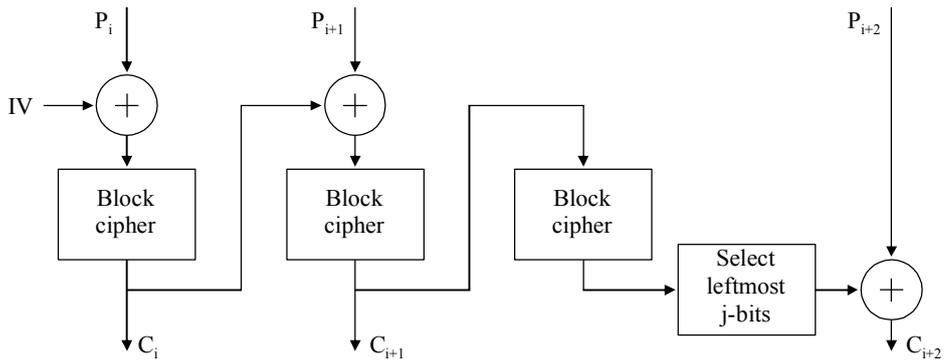


Figure 27 – 3DES-EDE as Block Cipher

Encryption
CBC w/ residual block termination



Decryption
CBC w/ residual block termination

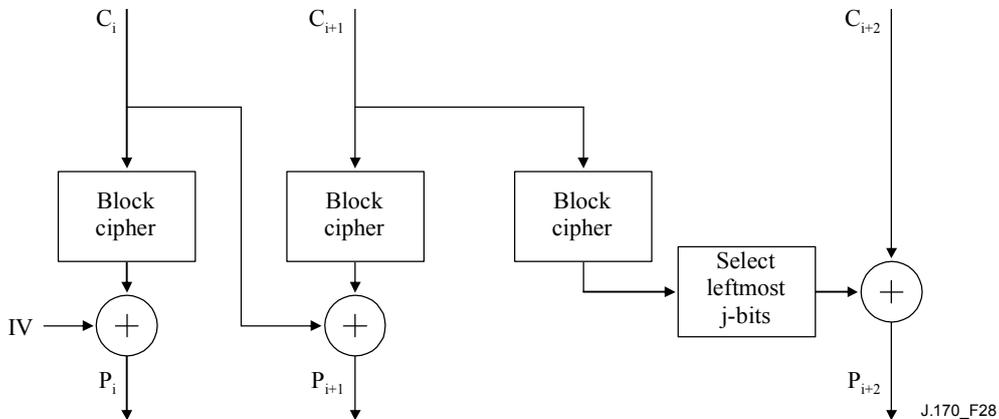


Figure 28 – CBC with Residual Block Termination

9.4 RSA signature

All public key signatures for IPCablecom MUST be generated and verified using the RSA signature algorithm described in RFC 2437. The format for all IPCablecom RSA signatures MUST be compliant with the Cryptographic Message Syntax of RFC 2630.

9.5 HMAC-SHA-1

The keyed hash employed by the HMAC-Digest Attribute MUST use the HMAC message authentication method per RFC 2104 with the SHA-1 hash algorithm per FIPS PUB 180-1.

9.6 Key derivation

Key derivation clauses in this Recommendation refer to a function $F(S, \text{seed})$, where S is a shared secret from which keying material is derived, and seed is a constant string of bytes. Below is the specification of $F(S, \text{seed})$, borrowed from TLS (RFC 2246):

$$F(S, \text{seed}) = \text{HMAC_SHA-1}(S, A(1) + \text{seed}) + \\
\text{HMAC_SHA-1}(S, A(2) + \text{seed}) + \\
\text{HMAC_SHA-1}(S, A(3) + \text{seed}) + \dots$$

where $+$ indicates concatenation.

$A()$ is defined as: $A(0) = \text{seed}$
 $A(i) = \text{HMAC_SHA-1}(S, A(i - 1))$

$F(S, \text{seed})$ is iterated as many times as is necessary to produce required quantity of data. Unused bytes at the end of the last iteration will be discarded.

9.7 The MMH-MAC

In this clause the MMH Function and the MMH Message Authentication Code (MAC) are described. The MMH-MAC is the message authentication code option for the media flows. As discussed in 7.6.2, the MMH-MAC is computed over the RTP header and the payload is generated by the codec. The MMH Function will be described next, followed by a description of the MMH-MAC.

9.7.1 The MMH Function

The Multilinear Modular Hash (MMH) Function described below is a variant of the MMH Function described in *MMH: Software Message Authentication in Gbit/second Rates*. Some of the computations described below use signed arithmetic whereas the computations in *MMH: Software Message Authentication in Gbit/second Rates* use unsigned arithmetic. The signed arithmetic variant described here was selected for its computational efficiency when implemented on DSPs. All of the properties shown for the MMH function in *MMH: Software Message Authentication in Gbit/second Rates* continue to hold for the signed variant.

The MMH Function has three parameters: the word size, the number of words of input, and the number of words of output. $\text{MMH}[\omega, s, t]$ specifies the hash function with word size ω , s input words, and t output words. For IP-Cablecom the word size is fixed to 16 bits: $\omega = 16$. The number of output words will be either 1 or 2: $t \in \{1, 2\}$. The MMH Hash Function will first be described for $t = 1$, i.e., one output word.

9.7.1.1 MMH[16,s,1]

For the remainder of this clause, $\text{MMH}[16, s, 1]$ is denoted by H . In addition to s words of input, H also takes as input a key of s words. When H is used in computing the MMH-MAC, the key is randomly generated and remains fixed for several inputs as described in 9.7.2. The key is denoted by k and the i th word of the key by k_i : $k = k_1, k_2, \dots, k_s$. Likewise the input message is denoted by m and the i th word of the input message by m_i : $m = m_1, m_2, \dots, m_s$.

To describe H , the following definitions are needed. For any even positive integer n , S_n is defined to be the following set of n integers: $\{-n/2, \dots, 0, \dots, (n/2) - 1\}$. For example, $S_{2^{16}} = \{-2^{15}, \dots, 0, \dots, 2^{15} - 1\}$ is the set of signed 16-bit integers. For any integer z , $z \text{ smod } n$ is the unique element ω of S_n such that $z \equiv \omega \pmod{n}$. For example, if z is a 32-bit signed integer in 32-bit two's complement representation, then $z \text{ smod } 2^{16}$ can be computed by taking the 16 least significant bits of z and interpreting those bits in 16-bit two's complement representation.

For any positive integer q , Z_q denotes the following set of q integers: $\{0, 1, \dots, q - 1\}$.

As described above, H takes as input a key of s words. Each of the s words is interpreted as a 16-bit signed integer, i.e., an element of $S_{2^{16}}$. H also takes as input a message of s words. Each of the s words is interpreted as a 16-bit signed integer, i.e., an element of $S_{2^{16}}$. The output of H is an unsigned 16-bit integer, i.e., an element of $Z_{2^{16}}$. Alternatively, the range of H is $S_{2^{16}}^S \times S_{2^{16}}^S$ and the domain is $Z_{2^{16}}$.

H is defined by a series of steps. For $k, m \in S_{2^{16}}^S$:

- 1) define H_1 as $H_1(k, m) = \sum_{i=1}^S k_i \cdot m_i \text{ smod } 2^{32}$;
- 2) define H_2 as $H_2(k, m) = H_1(k, m) \text{ mod } p$ where p is the prime number $p = 2^{16} + 1$;

3) define H as $H(k,m) = H_2(k,m) \bmod 2^{16}$.

Equivalently:

$$H(k,m) = \left(\left(\left(\sum_{i=1}^S k_i \cdot m_i \right) \bmod 2^{32} \right) \bmod p \right) \bmod 2^{16}$$

Each step is discussed in detail below.

Step 1) $H_1(k,m)$ is the inner product of two vectors each of s 16-bit signed integers. The result of the inner product is taken $\bmod 2^{32}$ to yield an element of $S_{2^{32}}$ ⁶. That is, if the inner product is in twos complement representation of 32 or more bits, the 32 least significant bits are retained and the resulting integer is interpreted in 32-bit twos complement representation.

Step 2) This step consists of taking an element x of $S_{2^{32}}$ and reducing it $\bmod p$ to yield an element of Z_p . If x is represented in 32-bit twos complement notation, then this reduction can be accomplished very simply as follows. Let a be the unsigned integer given by the 16 most significant bits of x . Let b be the unsigned integer given by the 16 least significant bits of x . There are two cases depending upon whether x is negative.

Case i) If x is non-negative then $x = a2^{16} + b$ where $a \in \{0, \dots, 2^{15} - 1\}$ and $b \in \{0, \dots, 2^{16} - 1\}$.

From the modular equation:

$$a2^{16} + b \equiv a2^{16} + b - a(2^{16} + 1) \pmod{(2^{16} + 1)}$$

it follows that $x \equiv b - a \pmod{p}$. The quantity $b - a$ is in the range $\{-2^{15} + 1, \dots, 2^{16} - 1\}$. Therefore if $b - a$ is non-negative then $x \bmod p = b - a$. If $b - a$ is negative then $x \bmod p = b - a + p$.

Case ii) If x is negative then $x = a2^{16} + b - 2^{32}$ where $a \in \{2^{15}, \dots, 2^{16} - 1\}$ and $b \in \{0, \dots, 2^{16} - 1\}$. From the modular equation:

$$a2^{16} + b - 2^{32} \equiv b + a2^{16} - a(2^{16} + 1) - 2^{32} + 2^{16}(2^{16} + 1) \pmod{(2^{16} + 1)}$$

it follows that $x \equiv (b - a + 2^{16}) \pmod{p}$. The quantity $b - a + 2^{16}$ is in the range $\{2^{15} + 1, \dots, 2^{17} - 1\}$. Therefore, if $b - a < p$ then $x \bmod p = b - a$. If $b - a \geq p$ then $x \bmod p = b - a - p$.

Step 3) This step takes an element of Z_p and reduces it $\bmod 2^{16}$. This is equivalent to taking the 16 least significant bits.

9.7.1.2 MMH[16,s,2]

This clause describes the MMH Function with an output length of two words, which in this case is 32 bits. For convenience, let $H' = \text{MMH}[16,s,2]$. H' takes a key of $s + 1$ words. Let $k = k_1, \dots, k_{s+1}$. Furthermore, define $k^{(1)}$ to be the s words of k starting with k_1 , i.e., $k^{(1)} = k_1, \dots, k_s$. Define $k^{(2)}$ to be the s words of k , starting with k_2 , i.e., $k^{(2)} = k_2, \dots, k_{s+1}$. For any $k \in S_{2^{16}}^{s+1}$ and any $m \in S_{2^{16}}^s$, $H'(k,m)$ is computed by first computing $H(k^{(1)},m)$ and then $H(k^{(2)},m)$ and concatenating the results. That is, $H'(k,m) = H(k^{(1)},m) \circ H(k^{(2)},m)$.

⁶ The entire sum need not be computed before performing the $\bmod 2^{32}$ operation. The $\bmod 2^{32}$ operation can be computed on partial sums since $(x + y) \bmod 2^{32} = (x \bmod 2^{32} + y \bmod 2^{32}) \bmod 2^{32}$.

9.7.2 The MMH-MAC

This clause describes the MMH-MAC. The MMH-MAC has three parameters: the word size, the number of words of input, and the number of words of output. MMH-MAC[ω, s, t] specifies the message authentication code with word size ω , s input words and t output words. For IPCablecom the word size is fixed to 16 bits: $\omega = 16$. The number of output words will be either 1 or 2: $t \in \{1, 2\}$.

For convenience, let $M = \text{MMH-MAC}[16, s, t]$. When using M , a sender and receiver share a key k of $s + t - 1$ words. In addition, they share a sequence of key streams of t words each, one one-time pad for each message sent. Let $r^{(i)}$ be the key stream used for the i th message sent and received. For the i th message, $m^{(i)}$, the message authentication code is computed as:

$$M(k, r^{(i)}, m^{(i)}) = H(k, m^{(i)}) + r^{(i)}$$

Here $H = \text{MMH}[16, s, t]$, $r^{(i)}$ is in $Z_{2^{16}}$ and addition is mod 2^{16} .

9.7.2.1 MMH-MAC when using a block cipher

When calculating the MMH-MAC when encryption is performed by one of the available block ciphers, the block cipher is used to calculate the t words of $r^{(i)}$ key stream (pad) as defined in 7.6.2.1.2.2.3.

9.7.2.2 Handling variable-size data

In order to handle data of all possible sizes up to a maximum value, the following rules MUST be followed for computing an MMH function:

- If the data is not a multiple of the word size, pad the data up to a multiple of the word size (16 bits) with zero-bytes. In other words, if the length of message m is not a multiple of word size w , but rather of length b octets, $b = n \times w + r$ with $n \geq 0$ and $0 < r < w$, then pad message m at the end with $w - r$ zero-bytes before passing it as the input to M .
- If the key is larger than what is needed for a particular message, truncate the key. In other words, if a message m is not of length s words, but rather of length $v < s$ words, then truncate the value of the key k to $v + t - 1$ words before it is used to calculate the MMH hash. (For MMH hash with 1 word output, $t = 1$ and k is truncated to v words. For 2-word output, $t = 2$ and k is truncated to $v + 1$ words.)

9.8 Random number generation

Good random number generation is vital to most cryptographic mechanisms. Implementations SHOULD do their best to produce true-random seeds; they should also use cryptographically strong pseudo-random number generation algorithms. RFC 1750 gives some suggestions; other possibilities include use of a per-MTA secret installed at manufacture time and used in the random number generation process.

10 Physical security

10.1 Protection for MTA key storage

An MTA MUST maintain in permanent write-once memory an RSA key pair. An MTA SHOULD deter unauthorized physical access to this keying material.

The level of physical protection of keying material required by this Recommendation for an MTA is specified in terms of the security levels defined in FIPS PUBS 140-2, *Security Requirements for Cryptographic Modules*. An MTA SHOULD, at a minimum, meet FIPS PUBS 140-2 Security Level 1 requirements.

This Recommendation's minimal physical security requirements for an MTA will not, in normal practice, jeopardize a customer's data privacy. Assuming the subscriber controls the access to the MTA with the same diligence they would protect a cellular phone, physical attacks on that MTA to extract keying data are likely to be detected by the subscriber.

An MTA's weak physical security requirements, however, could undermine the cryptographic protocol's ability to meet its main security objective: to provide a service operator with strong protection from theft of high value network.

This Recommendation specifies requirements to protect against unauthorized access to these network services by enforcing an end-to-end message integrity and encryption of signalling flows across the network and by employing an authenticated key management protocol. If an attacker is able to legitimately subscribe to a set of services and also gain physical access to an MTA containing keying material, then in the absence of strong physical protection of this information, the attacker can extract keying material from the MTA, and redistribute the keys to other users running modified illegitimate MTAs, effectively allowing theft of network services.

There are two distinct variations of "active attacks" involving the extraction and redistribution of cryptographic keys. These include the following:

- 1) An "RSA active clone" would actively participate in IPCablecom key exchanges. An attacker must have some means by which to remove the cryptographic keys that enable services from the clone master, and install these keys into a clone MTA. An active clone would work in conjunction with an active clone master to passively obtain the clone master's keying material and then actively impersonate the clone master. A single active clone may have numerous active clone master identities from which to select to obtain access to network services. This attack allows, for example, the theft of non-local voice communications.
- 2) An DH active clone would also actively participate in the IPCablecom key exchanges and like the RSA active clone, would require an attacker to extract the cryptographic keys that enable the service from the clone master and install these keys into a clone MTA. However, unlike the RSA active clone, the DH active clone must obtain the clone master's random number through alternate means or perform the key exchange and risk detection. Like an RSA active clone, an DH active clone may have numerous clone master identities from which to select to obtain access to the network services.
- 3) An "active black box" MTA, holding another MTA's session or IPsec keys, would use the keys to obtain access to network-based services or traffic flows similar to the RSA active clone. Since both session keys and IPsec keys change frequently, such clones have to be periodically updated with the new keying material, using some out-of-band means.

An active RSA clone, for example, could operate on a cable access network within whatever geographic region the cloned parent MTA was authorized to operate in. Depending upon the degree to which a service operator's subscriber authorization system restricted the location from which the MTA could operate, the clone's scope of operation could extend well beyond a single J.112 MAC domain.

An active clone attack may be detectable by implementing the appropriate network controls in the system infrastructure. Depending on the access fraud detection methods that are in place, a service operator has a good probability of detecting a clone's operation should it attempt to operate within the network. The service operator could then take defensive measures against the detected clone. For example, in the case of an active RSA clone, it could block the device's future network access by including the device certificate on the certificate hot list. Also, the service operator's subscriber authorization system could limit the geographic region over which a subscriber, identified by its cryptographic credentials, could operate. Additionally, the edge router functionality in the CMTS

could limit any access based upon the IP address. These methods would limit the region over which an active RSA clone could operate and reduce the financial incentive for such an attack.

The architectural guidelines for IPCablecom security are determined by balancing the revenues that could be lost due to the classes of active attacks against the cost of the methods to prevent the attack. At the extreme side of preventive methods available to thwart attacks, both physical security equivalent to FIPS PUB 140-1 Security Level 3 and network-based fraud detection methods could be used to limit the access fraud that allows theft of network-based services. The network-based intrusion detection of active attacks allows operators to consider operational defenses as an alternative to increased physical security. If the revenues threatened by the active attacks increase significantly to the point where additional protective mechanisms are necessary, the long-term costs of operational defenses would need to be compared with the costs of migrating to MTAs with stronger physical security. The inclusion of physical security should be an implementation- and product-differentiation-specific decision.

Although the scope of the current IPCablecom Recommendations do not specifically define requirements for MTAs to support any requirements other than voice communications, the goal of the IPCablecom effort is to provide for the eventual inclusion of integrated services. Part of these integrated services may include the "multicast" of high-value content or extremely secure multicast corporate videoconference sessions.

Two additional attacks enabling a compromise of these types of services are defined:

- 1) An "RSA passive clone" passively monitors the parent MTA's key exchanges and, having a copy of the parent MTA's RSA private key, is able to obtain the same traffic keying material the parent MTA has access to. The clone then uses the keying material to decrypt downstream traffic flows it receives across the shared medium. This attack is limited in that it only allows snooping, but if the traffic were of high value, the attack could facilitate the theft of high-value multicast traffic.
- 2) A "Passive black box" MTA, holding another MTA's short-term (relative to the RSA key) keys, uses the keying material to gain access to encrypted traffic flows similar to the RSA passive clone.

The passive attacks, unlike the active attacks, are not detectable using network-based intrusion detection techniques since these units never make themselves known to the network while performing the attack. However, this type of service theft has unlimited scale since the passive clones and black boxes, even though they operate on different cable access networks (sometimes referred to as the same J.112 MAC domain) as the parent MTA from whom the keys were extracted, gain access to the protected data the parent MTA is currently receiving since the encryption of the data most likely occurred at the source. (These are general IP multicast services, not to be confused with the specific J.112/BPI+ multicast implementation, where passive clones would be restricted to a single downstream CMTS segment.) The snooping of the point-to-point data is limited to the J.112 MAC domain of the parent MTA. Passive attacks may be prevented by ensuring that the cryptographic keys that are used to enable the services cannot be tampered with in any manner.

In setting goals and guidelines for the IPCablecom security architecture, an assessment has to be made of the value of the services and content that can be stolen or monitored by key extraction and redistribution to passive MTAs. The cost of the solution should not be greater than the lost revenue due to theft of the service or subscribers terminating the service due to lack of privacy. However at this time, there is no clear cost that can be attributed to either the lost revenue from high-value multicast services or the loss of subscribers due to privacy issues unique to this type of network. Therefore, it was concluded that passive key extraction and redistribution attacks would pose an indeterminate financial risk to service operators, and that the cost of protection (i.e., incorporation of stronger physical security into the MTA) should be balanced against the value of the risk. As with the active attacks, the decision to include additional functionality to implement physical

security in the MTA should be left as an implementation and product differentiation issue and not be mandated as a requirement of this Recommendation.

10.2 MTA key encapsulation

As stated in the previous clause, FIPS PUB 140-2 Security Level 1 specifies very little actual physical security and that an MTA MUST deter unauthorized "physical" access to its keying material. This restricted access also includes any ability to directly read the keying material using any of the MTA interfaces.

One of the (many) requirements of FIPS PUB 140-2 Security Level 3 is that "the entry or output of plaintext Critical Security Parameters (CSPs) be performed using ports that are physically separated from other ports or interfaces that are logically separated using a trusted path from other interfaces. Plaintext CSPs may be entered into or output from the cryptographic module in encrypted form (in which case they may travel through enclosing or intervening systems)". As also mentioned in the previous clause, this Recommendation is not requiring compliance with any of the FIPS PUB 140-2 Security Level 3 requirements.

However, it is strongly recommended that any persistent keying material SHOULD be encapsulated such that there is no way to extract the keying material from the MTA using any of the MTA interfaces (either required in the IPCablecom specifications or proprietary provided by the vendor) without modifications to the MTA.

In particular, an MTA subscriber may also be connected to the Internet via a cable modem (which may be embedded in the same MTA). In that case, hackers may potentially exploit any weakness in the configuration of the subscriber's local network and steal MTA's secret and private keys over the network. If instead, the MTA subscriber is connected to a company Intranet, the same threat still exists, although from a smaller group of people.

11 Secure software upgrade

IPCablecom includes only Embedded MTAs. E-MTAs are embedded with J.112/J.122 cable modems (including J.125 BPI+). Embedded MTAs MUST have their software upgraded by the cable modem according to the J.112/J.122 and J.125 requirements.

Annex A

Oakley groups

PKINIT states that DH parameters SHOULD be taken from the first or second Oakley groups as defined in RFC 2409. Additionally, this Recommendation requires that DH groups are used exactly as defined in RFC 2409.

RFC 2409 defines several so-called "Oakley groups." Only the first two are relevant to this Recommendation. RFC 2409 requires implementations to support the first group, and recommends that they support the second. This annex is included because RFC 2409 does not give values of q (the $p-1$ factor) for the groups, and these are necessary in order to encode the dhpublicnumber type used in the subjectPublicKeyInfo data structure in PKINIT.

The first two Oakley groups are defined as follows:

First Oakley Group:

Prime (p):

```
FFFFFFFF FFFFFFFFF C90FDAA2 2168C234 C4C6628B 80DC1CD1
29024E08 8A67CC74 020BBEA6 3B139B22 514A0879 8E3404DD
EF9519B3 CD3A431B 302B0A6D F25F1437 4FE1356D 6D51C245
E485B576 625E7EC6 F44C42E9 A63A3620 FFFFFFFFF FFFFFFFFF
```

Generator (g or b):

2.

Factor (q):

```
7FFFFFFFF FFFFFFFFF E487ED51 10B4611A 62633145 C06E0E68
94812704 4533E63A 0105DF53 1D89CD91 28A5043C C71A026E
F7CA8CD9 E69D218D 98158536 F92F8A1B A7F09AB6 B6A8E122
F242DABB 312F3F63 7A262174 D31D1B10 7FFFFFFFF FFFFFFFFF
```

Second Oakley Group:

Prime (p):

```
FFFFFFFF FFFFFFFFF C90FDAA2 2168C234 C4C6628B 80DC1CD1
29024E08 8A67CC74 020BBEA6 3B139B22 514A0879 8E3404DD
EF9519B3 CD3A431B 302B0A6D F25F1437 4FE1356D 6D51C245
E485B576 625E7EC6 F44C42E9 A637ED6B 0BFF5CB6 F406B7ED
EE386BFB 5A899FA5 AE9F2411 7C4B1FE6 49286651 ECE65381
FFFFFFFF FFFFFFFFF
```

Generator (g or b):

2.

Factor (q):

```
7FFFFFFFF FFFFFFFFF E487ED51 10B4611A 62633145 C06E0E68
94812704 4533E63A 0105DF53 1D89CD91 28A5043C C71A026E
F7CA8CD9 E69D218D 98158536 F92F8A1B A7F09AB6 B6A8E122
F242DABB 312F3F63 7A262174 D31BF6B5 85FFAE5B 7A035BF6
F71C35FD AD44CFD2 D74F9208 BE258FF3 24943328 F67329C0
FFFFFFFF FFFFFFFFF
```

Annex B

Kerberos Network Authentication Service

See IETF RFC 4120 (2005).

Annex C

PKINIT specification

See IETF RFC 4556 (2006).

Appendix I

IPCablecom administration guidelines and best practices

This appendix describes various administration guidelines and best practices recommended by IPCablecom. These are included to help facilitate network administration and/or strengthen overall security in the IPCablecom network.

I.1 Routine CMS service key refresh

IPCablecom recommends that the CMS service keys be routinely changed (refreshed) at least once every 90 days in order to reduce the risk of key compromises. The refresh period should be a provisioned parameter that can be used in one of the following ways:

- In the case of manual key changes, an administrator is prompted or reminded to manually change a CMS service key.
- In the case of autonomous key changes (using Kerberos Set/Change Password), it will define the refresh period.

Note that in the case of autonomous key refreshes, whereby administrative overhead and scalability are not an issue, it may be desirable to use a refresh period that is less than 90 days (but at least the maximum ticket lifetime). This may further reduce the risk of key compromise.

Appendix II

Example of MMH algorithm implementation

This appendix gives an example implementation of the MMH MAC algorithm. There may be other implementations that have advantages over this example in particular operating environments. This example is for informational purposes only and is meant to clarify the Recommendation.

The example implementation uses the term "MMH16" for the case where the MAC length is 2 octets and "MMH32" for the case where the length is 4 octets.

A main program is included for exercising the example implementation. The output produced by the program is included.

```
/*
Demo of IPCablecom MMH16 and MMH32 MAC algorithms.

This program has been tested using Microsoft C/C++ Version 5.0.
It is believed to port easily to other compilers, but this has
not been tested. When porting, be sure to pick the definitions
for int16, int32, uint16, and uint32 carefully.
*/
#include <stdio.h>
/*
Define signed and unsigned integers having 16 and 32 bits.
This is machine/compiler dependent, so pick carefully.
*/
typedef short int16;
typedef unsigned short uint16;
typedef int int32;
typedef unsigned int uint32;
/*
Define this symbol to see intermediate values.
Comment it out for clean display.
*/
#define VERBOSE
int32 reduceModF4(int32 x) {
    /*
    Routine to reduce an int32 value modulo F4, where F4 = 0x10001.
    Result is in range [0, 0x10000].
    */
    int32 xHi, xLo;
    /* Range of x is [0x80000000, 0x7fffffff]. */
    /*
    If x is negative, add a multiple of F4 to make it non-negative.
    This loop executes no more than two times.
    */
    while (x < 0) x += 0x7fff7fff;
    /* Range of x is [0, 0x7fffffff]. */
    /* Subtract high 16 bits of x from low 16 bits. */
    xHi = x >> 16;
    xLo = x & 0xffff;
    x = xLo - xHi;
    /* Range of x is [0xffff8001, 0x0000ffff]. */
    /* If x is negative, add F4. */
    if (x < 0) x += 0x10001;
    /* Range of x is [0, 0x10000]. */
    return x;
}
uint16 mmh16(
    unsigned char *message,
    unsigned char *key,
```

```

unsigned char *pad,
int msgLen) {
/*
Compute and return the MMH16 MAC of the message using the indicated key and
pad.
The length of the message is msgLen bytes; msgLen must be even.
The length of the key must be at least msgLen bytes.
The length of the pad is two bytes. The pad must be freshly picked from a
secure random source.
*/
int16 x, y;
uint16 u, v;
int32 sum;
int i;
sum = 0;
for (i=0; i<msgLen; i+=2) {
/* Build a 16-bit factor from the next two message bytes. */
x = *message++;
x <<= 8;
x |= *message++;
/* Build a 16-bit factor from the next two key bytes. */
y = *key++;
y <<= 8;
y |= *key++;
/* Accumulate product of the factors into 32-bit sum */
sum += (int32)x * (int32)y;
#ifdef VERBOSE
printf(" x %04x y %04x sum %08x\n", x & 0xffff, y & 0xffff, sum);
#endif
}
/* Reduce sum modulo F4 and truncate to 16 bits. */
u = (uint16) reduceModF4(sum);
#ifdef VERBOSE
printf(" sum mod F4, truncated to 16 bits: %04x\n", u & 0xffff);
#endif
/* Build the pad variable from the two pad bytes */
v = *pad++;
v <<= 8;
v |= *pad;
#ifdef VERBOSE
printf(" pad variable: %04x\n", v & 0xffff);
#endif
/* Accumulate pad variable, truncate to 16 bits */
u = (uint16)(u + v);
#ifdef VERBOSE
printf(" mmh16 value: %04x\n", u & 0xffff);
#endif
return u;
}
uint32 mmh32(
unsigned char *message,
unsigned char *key,
unsigned char *pad,
int msgLen) {
/*
Compute and return the MMH32 MAC of the message using the indicated key and
pad.
The length of the message is msgLen bytes; msgLen must be even.
The length of the key must be at least (msgLen + 2) bytes.
The length of the pad is four bytes. The pad must be freshly picked from a
secure random source.
*/
uint16 x, y;
uint32 sum;

```

```

    x = mmh16(message, key, pad, msgLen);
    y = mmh16(message, key+2, pad+2, msgLen);
    sum = x;
    sum <<= 16;
    sum |= y;
    return sum;
}
void show(char *name, unsigned char *src, int nbytes)
{
    /*
    Routine to display a byte array, in normal or reverse order
    */
    int i;
    enum { BYTES_PER_LINE = 16 };
    if (name) printf("%s", name);
    for (i=0; i<nbytes; i++) {
        if ((i % BYTES_PER_LINE) == 0) printf("\n");
        printf("%02x ", src[i]);
    }
    printf("\n");
}
int main()
{
    uint16 mac16;
    uint32 mac32;
    unsigned char message[] = {
        0x4e, 0x6f, 0x77, 0x20, 0x69, 0x73, 0x20, 0x74, 0x68,
        0x65, 0x20, 0x74, 0x69, 0x6d, 0x65, 0x2e,
    };
    unsigned char key[] = {
        0x35, 0x2c, 0xcf, 0x84, 0x95, 0xef, 0xd7, 0xdf, 0xb8,
        0xf5, 0x74, 0x05, 0x95, 0xeb, 0x98, 0xd6, 0xeb, 0x98,
    };
    unsigned char pad16[] = {
        0xae, 0x07,
    };
    unsigned char pad32[] = {
        0xbd, 0xe1, 0x89, 0x7b,
    };
    unsigned char macBuf[4];
    printf("Example of MMH16 computation\n");
    show("message", message, sizeof(message));
    show("key", key, sizeof(message));
    show("pad", pad16, 2);
    mac16 = mmh16(message, key, pad16, sizeof(message));
    macBuf[1] = (unsigned char)mac16; mac16 >>= 8;
    macBuf[0] = (unsigned char)mac16;
    show("MMH16 MAC", macBuf, 2);
    printf("\n");
    printf("Example of MMH32 computation\n");
    show("message", message, sizeof(message));
    show("key", key, sizeof(message)+2);
    show("pad", pad32, 4);
    mac32 = mmh32(message, key, pad32, sizeof(message));
    macBuf[3] = (unsigned char)mac32; mac32 >>= 8;
    macBuf[2] = (unsigned char)mac32; mac32 >>= 8;
    macBuf[1] = (unsigned char)mac32; mac32 >>= 8;
    macBuf[0] = (unsigned char)mac32;
    show("MMH32 MAC", macBuf, 4);
    printf("\n");
    return 0;
}

```

Here is the output produced by the program:

Example of MMH16 computation

```
message
4e 6f 77 20 69 73 20 74 68 65 20 74 69 6d 65 2e
key
35 2c cf 84 95 ef d7 df b8 f5 74 05 95 eb 98 d6
pad
ae 07
  x 4e6f y 352c sum 104a7614
  x 7720 y cf84 sum f9bac294
  x 6973 y 95ef sum ce0a23f1
  x 2074 y d7df sum c8f3d4fd
  x 6865 y b8f5 sum abfb55a6
  x 2074 y 7405 sum bab087ea
  x 696d y 95eb sum 8f00bff9
  x 652e y 98d6 sum 663aa46d
sum mod F4, truncated to 16 bits: 3e33
pad variable: ae07
mmh16 value: ec3a
```

MMH16 MAC

```
ec 3a
```

Example of MMH32 computation

```
message
4e 6f 77 20 69 73 20 74 68 65 20 74 69 6d 65 2e
key
35 2c cf 84 95 ef d7 df b8 f5 74 05 95 eb 98 d6
eb 98
pad
bd e1 89 7b
  x 4e6f y 352c sum 104a7614
  x 7720 y cf84 sum f9bac294
  x 6973 y 95ef sum ce0a23f1
  x 2074 y d7df sum c8f3d4fd
  x 6865 y b8f5 sum abfb55a6
  x 2074 y 7405 sum bab087ea
  x 696d y 95eb sum 8f00bff9
  x 652e y 98d6 sum 663aa46d
sum mod F4, truncated to 16 bits: 3e33
pad variable: bde1
mmh16 value: fc14
  x 4e6f y cf84 sum f125323c
  x 7720 y 95ef sum bfca091c
  x 6973 y d7df sum af427949
  x 2074 y b8f5 sum a640e84d
  x 6865 y 7405 sum d590b646
  x 2074 y 95eb sum c81e04c2
  x 696d y 98d6 sum 9da1dde0
  x 652e y eb98 sum 95912b30
sum mod F4, truncated to 16 bits: 959f
pad variable: 897b
mmh16 value: 1f1a
```

MMH32 MAC

```
fc 14 1f 1a
```


SERIES OF ITU-T RECOMMENDATIONS

| | |
|-----------------|--|
| Series A | Organization of the work of ITU-T |
| Series D | General tariff principles |
| Series E | Overall network operation, telephone service, service operation and human factors |
| Series F | Non-telephone telecommunication services |
| Series G | Transmission systems and media, digital systems and networks |
| Series H | Audiovisual and multimedia systems |
| Series I | Integrated services digital network |
| Series J | Cable networks and transmission of television, sound programme and other multimedia signals |
| Series K | Protection against interference |
| Series L | Construction, installation and protection of cables and other elements of outside plant |
| Series M | Telecommunication management, including TMN and network maintenance |
| Series N | Maintenance: international sound programme and television transmission circuits |
| Series O | Specifications of measuring equipment |
| Series P | Terminals and subjective and objective assessment methods |
| Series Q | Switching and signalling |
| Series R | Telegraph transmission |
| Series S | Telegraph services terminal equipment |
| Series T | Terminals for telematic services |
| Series U | Telegraph switching |
| Series V | Data communication over the telephone network |
| Series X | Data networks, open system communications and security |
| Series Y | Global information infrastructure, Internet protocol aspects and next-generation networks |
| Series Z | Languages and general software aspects for telecommunication systems |