ETSI TS 101 909-11 V1.2.1 (2002-07)

Technical Specification

Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 11: Security

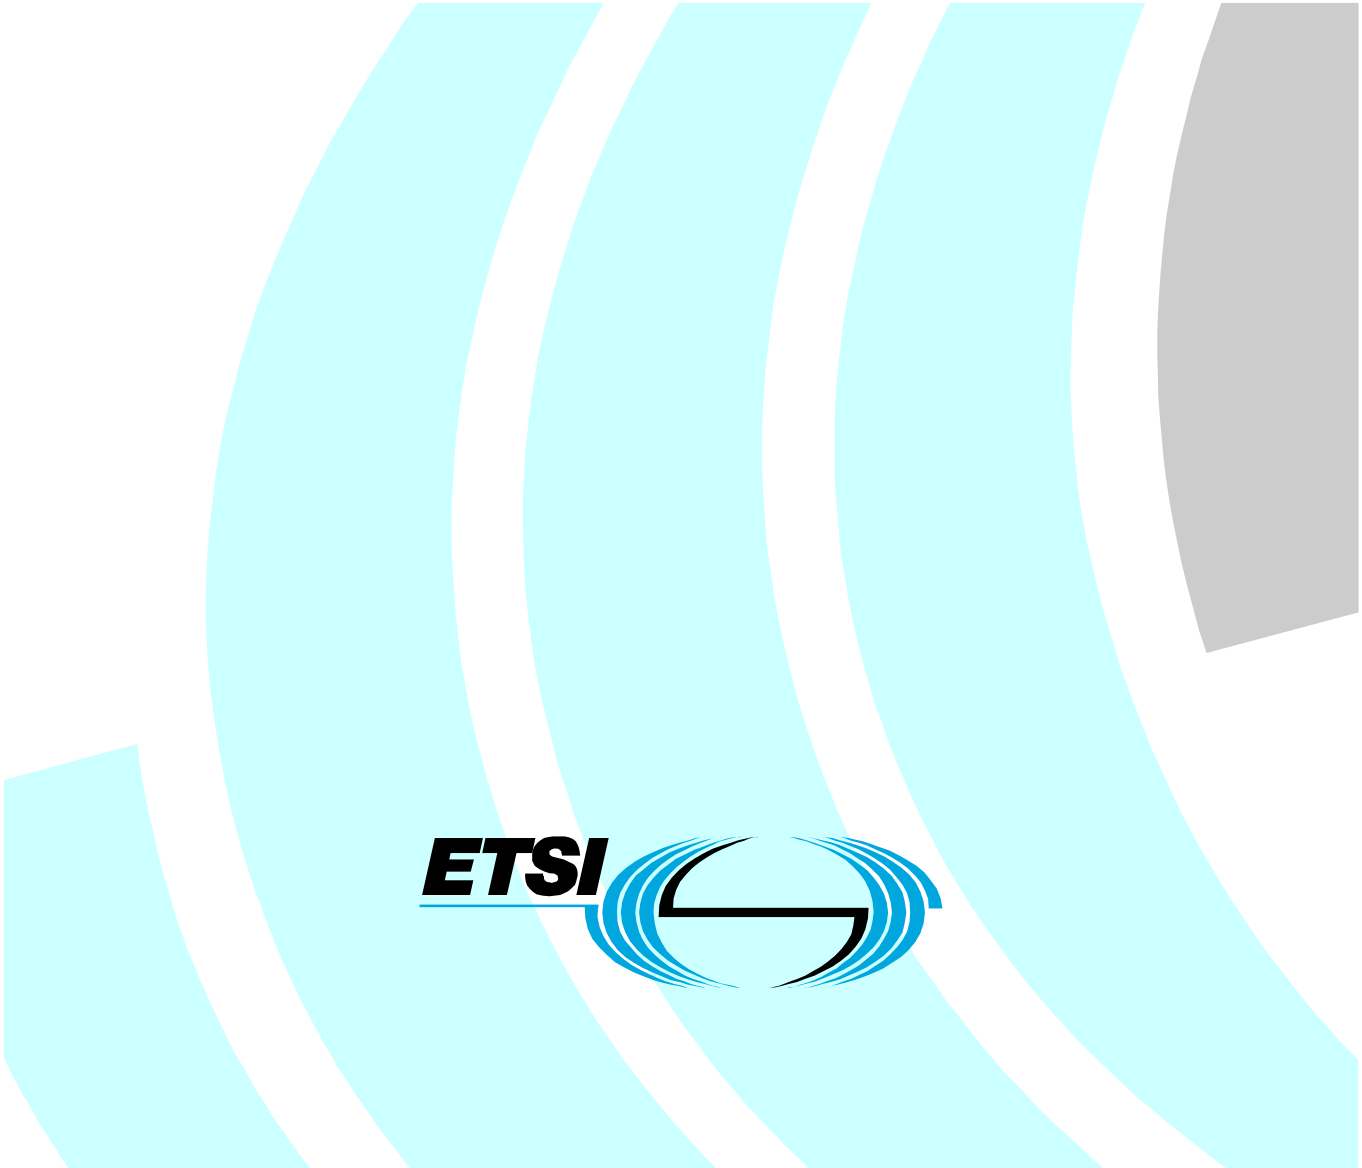

Reference
RTS/AT-020028-11

Keywords
access, broadband, cable, IP, multimedia, PSTN, security

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Access and Terminals (AT).

The present document is part 11 of a multi-part deliverable. Full details of the entire series can be found in part 1 [2].

Introduction

Security objectives definition

General security objectives

In general, security objectives can be divided into five main categories:

- **Confidentiality:** The avoidance of the disclosure of information without the permission of its owner.
- **Integrity:** The property that data has not been altered or destroyed in an unauthorized manner.
- **Accountability:** The principle whereby individuals are held responsible for the effect of any of their actions that might lead to a violation.
- **Availability:** The property of being accessible and usable upon demand by an authorized entity.
- **Non-repudiation:** A property by which one of the entities or parties in a communication cannot deny having participated in the whole or part of the communication.

Therefore, threat analysis, risk assessment and the proposed countermeasures of any technology or service will be based on these objectives.

Users' objectives

The objectives of users are not uniform. An enterprise does not always require the same as a private person. The following list gives examples of possible objectives, which may have implications on security:

- availability and correct functionality of service subscription (including reachability, availability and correct functionality);
- correct and verifiable billing;
- data integrity;
- data confidentiality/privacy;
- capability to use a service anonymously; and
- location confidentiality.

NOTE: This last objective may be modified for the provision of some location dependent services.

Service and network providers' objectives

The following list gives examples of objectives that may have implications on security:

- availability and correct functionality of network procedures;
- availability and correct functionality of service, network and element management;
- correct and verifiable billing and accounting, above all no possibility of fraud;
- non-repudiation for all network procedures and for all management activities; and
- preservation of reputation (above all preservation of users' and investors' trust).

Manufacturers' objectives

The following list gives examples of objectives that may have implications on security:

- fulfilling market objectives; and
- preservation of reputation.

Forms of attack

Eavesdropping

Eavesdropping is a threat against confidentiality and is performed by intercepting the physical (or logical) link between the sender and the receiver. The decision to intercept a line will essentially depend on whether the information to be obtained will be worth the technical (financial) expenditure and the risk of being detected. The answer to this question is largely determined by the attacker's means and interests.

In most cases, eavesdropping is used to obtain data (e.g. such as user identification and authentication data) to be able to perform more serious threats at another point in time.

Masquerade

A perpetrator can use masquerading to feign a false identity. For instance, the perpetrator will obtain a false identity by spying out the user ID and password, by manipulating the originator field of a message, by manipulating the I/O address within the network, or simply by using another person's telephone or computer.

A user who has been deceived as regards the identity of his communication partner can in that case easily be induced to disclose sensitive information. Many e-mail viruses work on this principle to propagate by using a forged identity and accessing a local address book to send attacks to known contacts.

A perpetrator can also use masquerading to try to tap an existing connection without having to authenticate himself, as this step has already been taken by the original participants in the communication (see also *Eavesdropping*).

Replay

A message, or part of a message, is repeated to produce unauthorized effect. For example, a valid message containing authentication information may be replayed by another entity in order to authenticate itself (as something that it is not).

Modification of information

In this case, data is corrupted or rendered useless through deliberate manipulation. The consequences of this are the rejection of authorized accesses to network resources.

Attackers may be interested in modifying either the information required during the registration or the call set up phase. Reasons for this might be to use a service for which the attacked user has to pay.

Generally modification of information may be a starting point for denial of service or masquerade and fraud attacks.

Unauthorized access

Access to network entities must be restricted and conformant to the security policy in place. If attackers get unauthorized access to any of the network entities this could generally lead to various other attacks like denial of service, eavesdropping or masquerade. Likewise it is possible that unauthorized access is also a consequence of the other threats mentioned above.

Stalking

Using information to infer the whereabouts of a principal especially for malicious or illegal purposes.

Denial of Service attacks

Denial of Service attacks (DoS), in particular Distributed Denial of Service (DDoS) attacks, strike at the physical networks used to host services with the goal of consuming all of the target's network capacity and other resources including processes, CPU time, disk space, inodes, ports and directories.

DoS attacks will be aimed at preventing users from using services or accessing devices that are normally available to them.

Countermeasures

Countermeasures have to be taken to contain the risk to an acceptable level. Most of these are simple in themselves.

Identification and authentication

In order to gain access to services users shall have to register an identity (private identity) and this identity shall be authenticated using methods based on secret key. Such methods are considered strong but do not authenticate a human user directly but only the device being used to gain service, however it may be possible to combine user input with a secret key method to more completely authenticate the user.

The mechanism shall be mutual and use as a time variant parameter a random number within a (mutual) challenge-response protocol.

Post authentication the private identity may be replaced with a temporary identity in like manner to the TMSI/IMSI relationship in GSM networks.

Encryption

Encryption in UCI shall be used for the transfer of data between service elements to maintain confidentiality. The signalling paths between service elements will generate the key used for encryption from the foregoing authentication exchange and an agreed method of getting time variance.

Integrity analysis

In order to protect information flows from being modified a means shall be provided to detect manipulation (e.g. cryptographic digests or checksums).

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 Architecture Framework document, TS 101 909-2 [3] defines the overall IPCablecom architecture, as well as the system elements, interfaces, and functional requirements for the entire IPCablecom network.

1.1 Goals

The present document describes the security relationships between the elements on the IPCablecom network. The general goals of the IPCablecom network security document 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 inter-operate 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 Assumptions

The following assumptions are made relative to the current scope of the IPCablecom Security document:

- Embedded Media Terminal Adaptors (MTAs) are within the current scope. Standalone MTAs will be addressed in later phases and security issues for standalone MTAs are thus for future study.
- NCS is the only call signalling method, on the access network, addressed in the present document.
- This version of the IPCablecom Security document specifies security for a single administrative domain and the communications between domains.
- Security for chained Radius servers is not currently in the scope.
- The present document 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 the present document.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication and/or edition number or version number) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

[1] ITU-T Recommendation J.112: "Transmission systems for interactive cable television services".

[2] ETSI TS 101 909-1: "Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 1: General".

[3] ETSI TS 101 909-2: "Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 2: Architectural framework for the delivery of time critical services over cable Television networks using cable modems".

[4] ETSI TS 101 909-3: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 3: Audio Codec Requirements for the Provision of Bi-Directional Audio Service over Cable Television Networks using Cable Modems".

[5] ETSI TS 101 909-4: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 4: Network Call Signalling Protocol".

[6] ETSI TS 101 909-5: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 5: Dynamic Quality of Service for the Provision of Real Time Services over Cable Television Networks using Cable Modems".

[7] ETSI TS 101 909-10: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 10: Event Message Requirements for the Provision of Real Time Services over Cable Television Networks using Cable Modems".

[8] ETSI TS 101 909-12: "Digital Broadband Cable Access to the Public Telecommunications Network; Part 12: Internet Signalling Transport Protocol".

[9] ETSI TS 101 909-7: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 7: Management Information Base (MIB) Framework".

[10] ETSI TS 101 909-6: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 6: Media Terminal Adapter (MTA) device provisioning".

[11] ETSI TS 101 909-8: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 8: Media Terminal Adapter (MTA) Management Information Base (MIB)".

[12] ETSI TS 101 909-9: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 9: Network Call Signalling (NCS) MIB Requirements".

[13] ETSI TS 101 909-13: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 13: Trunking Gateway Control Protocol".

[14] ITU-T Recommendation X.509 (1997 E): "Information Technology - Open Systems Interconnection - The Directory: Authentication Framework", June 1997".

[15] ITU-T Recommendation X.690: "Information technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", December, 1997".

[16] ETSI TS 101 909-19: "IPCablecom Audio Server Protocol Specification".

[17] ETSI TS 101 909-17: "Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 17: Inter-domain Quality of Service".

[18] IETF RFC 1889: "RTP: A Transport Protocol for Real-Time Applications", Proposed Standard, January, 1996.

[19] IETF RFC 2104: "HMAC: Keyed-Hashing for Message Authentication".

[20] IETF RFC 2139: "RADIUS: Accounting", April 1997.

[21] IETF RFC 2246: "The TLS Protocol Version 1.0", Proposed Standard, January 1999.

[22] IETF RFC 2327: "SDP: Session Description Protocol", Proposed Standard, April 1998.

[23] IETF RFC 2367: "PF-KEY: Key Management API, Version 2", July 1998.

[24] IETF RFC 2401: "Security Architecture for the Internet Protocol", Proposed Standard, November 1998.

[25] IETF RFC 2406: "IP Encapsulating Security Payload (ESP)", Proposed Standard, November 1998.

[26] IETF RFC 2407: "The Internet IP Security Domain of Interpretation for ISAKMP", Proposed Standard, November 1998.

[27] IETF RFC 2409: "The Internet Key Exchange (IKE)", Proposed Standard, November 1998.

[28] IETF RFC 2451: "The ESP CBC-Mode Cipher Algorithms", Proposed Standard, November 1998.

[29] IETF RFC 2459: "Internet X.509 Public Key Infrastructure Certificate and CRL Profile", Proposed Standard, January 1999.

[30] IETF RFC 2574: "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", Proposed Standard, April 1999.

[31] IETF RFC 2630: "Cryptographic Message Syntax", Proposed Standard, June 1999.

[32] IETF RFC 2747: "RSVP: Cryptographic Authentication", Proposed Standard, January 2000.

[33] IETF RFC 2437: "PKCS#1: RSA Cryptography Specification Version 2.0", Informational Standard, October 1998.

[34] IETF RFC 2403: "The Use of HMAC-MD5-96 within ESP and AH", Proposed Standard, November 1998.

[35] IETF RFC 2404: "The Use of HMAC-SHA-1-96 within ESP and AH", Proposed Standard, November 1998.

[36] FIPS 81: "DES Modes of Operation", December 1980.

[37] FIPS 180-1: "Secure Hash Algorithm (SHS)", April 1995.

[38] FIPS 197: "Advanced Encryption Standard (AES)", November 2001.

[39] ETSI ES 200 800 Digital Video Broadcasting (DVB); DVB interaction channel for Cable TV distribution systems (CATV).

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

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

access node: as used in the present document, a layer two termination device that terminates the network end of the CM connection

NOTE: It is technology specific. In ITU-T Recommendation J.112 [1], annex A, it is called the INA while in annex B it is the CMTS.

Application-Specific Data (ASD): application-specific field in the IPSec header that along with the destination IP address provides a unique number for each SA

audio server: server which plays informational announcements in IPCablecom network

NOTE: 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.

authentication: process of verifying the claimed identity of an entity to another entity

Authentication Header (AH): IPSec security protocol that provides message integrity for complete IP packets, including the IP header

authenticity: 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

authorization: act of giving access to a service or device if one has the permission to have the access

Baseline Privacy Interface Plus (BPI+): security portion of the J.112 standard that runs on the MAC layer

Call Agent (CA): part of the CMS that maintains the communication state, and controls the line side of the communication

Call Management Server (CMS): server which controls the audio connections, also called a Call Agent in MGCP/SGCP terminology

NOTE: This is one example of an Application Server.

Certification Authority (CA): trusted organization that accepts certificate applications from entities, authenticates applications, issues certificates and maintains status information about certificates

cipher: algorithm that transforms data between plaintext and ciphertext

Cipher-Block Chaining mode (CBC): 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

ciphersuite: set which must contain both an encryption algorithm and a message authentication algorithm (e.g. a MAC or an HMAC)

NOTE: In general, it may also contain a key management algorithm, which does not apply in the context of IPCablecom.

confidentiality: way to ensure that information is not disclosed to any one other then the intended parties

NOTE: Information is encrypted to provide confidentiality. Also known as privacy.

cryptanalysis: process of recovering the plaintext of a message or the encryption key without access to the key

DiffServ Code Point (DSCP): field in every IP packet that identifies the DiffServ Per Hop Behaviour

NOTE 1: 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.

NOTE 2: See annex C.

downstream: direction from the head-end toward the subscriber location

encryption: method used to translate information in plaintext into ciphertext

endpoint: Terminal, Gateway or MCU

event message: message capturing a single portion of a connection

gateway: devices bridging between the IPCablecom Voice Communication world and the PSTN

EXAMPLE: 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.

Hashed Message Authentication Code (HMAC): message authentication algorithm, based on either SHA-1 or MD5 hash and defined in IETF RFC 2104

header: protocol control information located at the beginning of a protocol data unit

integrity: way to ensure that information is not modified except by those who are authorized to do so

Internet Key Exchange (IKE): key management mechanism used to negotiate and derive keys for SAs in IPSec

IKE-: notation defined to refer to the use of IKE with pre-shared keys for authentication

kerberos: secret-key network authentication protocol that uses a choice of cryptographic algorithms for encryption and a centralized key database for authentication

key: mathematical value input into the selected cryptographic algorithm

key exchange: swapping of public keys between entities to be used to encrypt communication between the entities

key-management: process of distributing shared symmetric keys needed to run a security protocol

Media Access Control (MAC): sublayer of the Data Link Layer, which normally runs directly over the physical layer

Message Authentication Code (MAC): fixed-length data item that is sent together with a message to ensure integrity, also known as a MIC

non-repudiation: ability to prevent a sender from denying later that he or she sent a message or performed an action

Operations Systems Support (OSS): back-office software used for configuration, performance, fault, accounting, and security management

privacy: way to ensure that information is not disclosed to any one other then the intended parties

NOTE: Information is usually encrypted to provide confidentiality. Also known as confidentiality.

private key: key used in public key cryptography that belongs to an individual entity and must be kept secret

proxy: 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

PKCROSS: key which 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)

public key: key used in public key cryptography that belongs to an individual entity and is distributed publicly

NOTE: Other entities use this key to encrypt data to be sent to the owner of the key.

public key certificate: binding between an entity's public key and one or more attributes relating to its identity, also known as a digital certificate

public key cryptography: procedure that uses a pair of keys, a public key and a private key for encryption and decryption, also known as an asymmetric algorithm

NOTE: 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.

Public Key Infrastructure (PKI): process for issuing public key certificates, which includes standards, Certification Authorities, communication between authorities and protocols for managing certification processes

RC4: variable key length stream cipher offered in the ciphersuite, used to encrypt the media traffic in IPCablecom

Record Keeping Server (RKS): device which collects and correlates the various Event Messages

root private key: private signing key of the highest-level Certification Authority

NOTE: It is normally used to sign public key certificates for lower-level Certification Authorities or other entities.

Session Initiation Protocol (SIP): application-layer control (signalling) protocol for creating, modifying, and terminating sessions with one or more participants

Signalling Gateway (SG): signalling agent that receives/sends SCN native signalling at the edge of the IP network

NOTE: In particular the SS7 SG function translates variants ISUP and TCAP in an SS7-Internet Gateway to a common version of ISUP and TCAP.

Signalling System number 7 (SS7): architecture and set of protocols for performing out-of-band call signalling with a telephone network

Transaction Capabilities Application Protocol (TCAP): protocol within the SS7 stack that is used for performing remote database transactions with a Signalling Control Point

X.509 certificate: public key certificate specification developed as part of the ITU-T Recommendation X.500 standards directory

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AES	Advanced Encryption Standard
AH	Authentication Header
AN	Access Network
ASD	Application-Specific Data
BPI+	Baseline Privacy Interface Plus
CA	Call Agent
CA	Certification Authority
CBC	Cipher-Block Chaining mode
CM	Cable Modem
CMS	Call Management Server
CMS	Cryptographic Message Syntax
CVC	Code Verification Certificate
DES	Data Encryption Standard
DF	Delivery Function

NOTE: E.g. for lawful intercept.

DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Server
DSCP	DiffServ Code Point
DOCSIS	Data-Over-Cable Service Interface Specification

DQoS	Dynamic Quality of Service
DTMF	Dual-Tone Multi Frequency (tones)
EBP	Exterior Border Proxies
ESP	IPSec Encapsulating Security Payload
FQDN	Fully Qualified Domain Name

NOTE: Refer to IETF RFC 821 for details.

GC	Gate Controller
HMAC	Hashed Message Authentication Code
IKE	Internet Key Exchange
IPSec	Internet Protocol Security
ISTP	Internet Signalling Transport Protocol
IV	Initialization Vector
IVR	Interactive Voice Response system
KDC	Key Distribution Center
LNP	Local Number Portability
MAC	Message Authentication Code
MAC	Media Access Control
MD5	Message Digest 5
MG	Media Gateway
MGC	Media Gateway Controller
MGCP	Media Gateway Control Protocol
MIB	Management Information Base
MMH	Multilinear Modular Hash
MPC	Media Player Controller
MSB	Most Significant Bit
MTA	Media Terminal Adapter
NCS	Network Call Signalling
OID	Object IDentification
OSS	Operations Systems Support
PKCS	Public Key Cryptography Standards
PKI	Public Key Infrastructure
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RKS	Record Keeping Server
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
SIP	Session Initiation Protocol
SIP+	Session Initiation Protocol Plus

NOTE: An extension to SIP.

SNMP	Simple Network Management Protocol
SS7	Signalling System number 7
TCAP	Transaction Capabilities Application Protocol
TD	Timeout for Disconnect
TFTP	Trivial File Transfer Protocol
TGS	Ticket Granting Server

NOTE: It is a sub-system of the KDC used to grant Kerberos tickets.

UDP	User Datagram Protocol

4 Conventions

If the present document is implemented, the key words "MUST" and "SHALL" as well as "REQUIRED" are to be interpreted as indicating a mandatory aspect of the present document.

The key words indicating a certain level of significance of a particular requirements that are used throughout the present document are summarized in the table below:

"MUST"	This word or the adjective "REQUIRED" means that the item is an absolute requirement of the present document.
"MUST NOT"	This phrase means that the item is an absolute prohibition of the present document.
"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 event useful, but the full implications should be understood and the case carefully weighted 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, TS 101 909-2 [3].

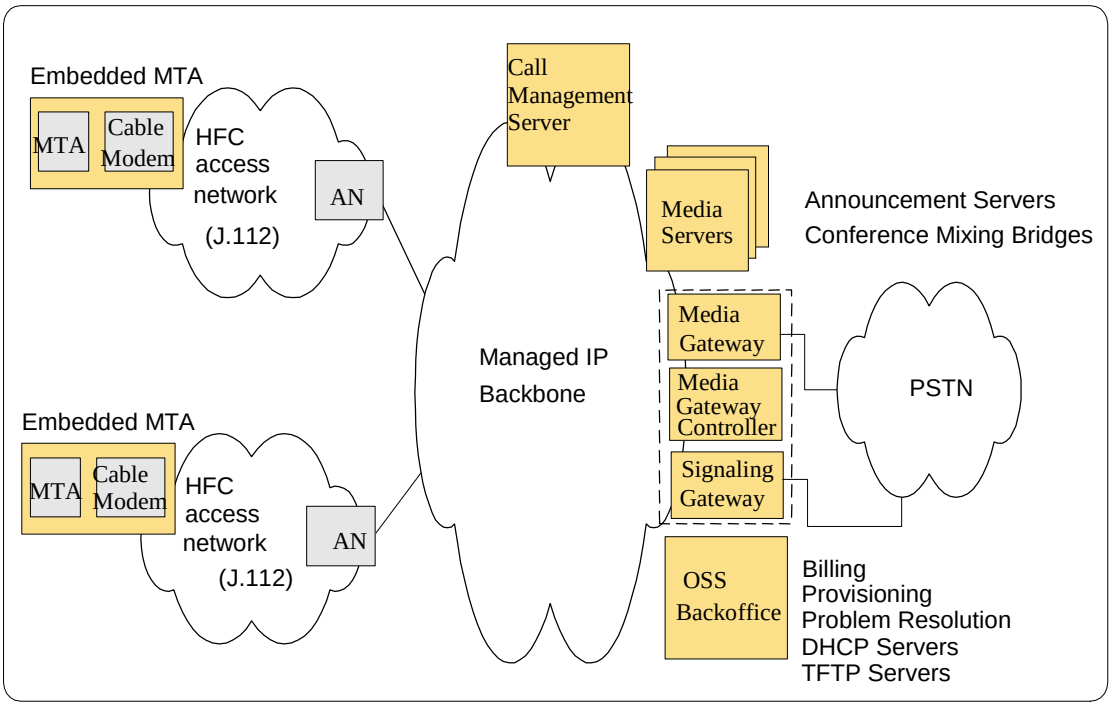

Figure 1: IPCablecom Single Zone Architecture

5.1.1 HFC network

In figure 1, the Access Network between the MTAs and the AN is an HFC network, which employs J.112 physical layer and MAC layer protocols. J.112 enhanced security (see clause 6.8) and QoS protocols are enabled over this link.

There is signalling between the MTA and the AN for the purpose of Dynamic QoS, TS 101 909-5 [6]. This includes both custom UDP messages defined specifically for IPCablecom DQoS, as well as standard RSVP messages, which run directly over the IP layer in the protocol stack. IPCablecom also provides guaranteed QoS for each voice communication between domains with Interdomain QoS, TS 101 909-17 [17].

Since J.112 key management and privacy runs between the Cable Modem and the AN, it does not help in authenticating the identity of the MTA. Therefore, all protocols that run between the MTA and the AN have additional security requirements that cannot be met by the methods in J.112. Also, since the MTA communicates with other devices on the IPCablecom network beyond the AN, these communications must be protected as well.

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;
- Third Party interfaces for Delivery Function (DF).

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

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). The present document specifies the steps in MTA provisioning, but provides detailed specifications only for the security parameters. Refer to TS 101 909-6 [10] 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, TS 101 909-5 [6].

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 COPS protocol to download QoS policy into the AN. After that, the QoS reservation is established via layer 2 signalling or J.112 QoS between the MTA and the AN on both sides of the connection. QoS reservations are also forwarded to the IP Backbone between the ANs, but the specifications of the Backbone reservations are currently for future study in IPCablecom. Therefore, the corresponding security specifications are also for future study.

5.1.3.3 Interdomain Quality of Service

IPCablecom provides guaranteed QoS for each voice communication between domains with Interdomain QoS, TS 101 909-17 [17]. DiffServ allows IP traffic to be marked with different DiffServ Code Points (DSCP) to obtain different queuing treatment on routers. Different queuing treatments in each router are called per-hop behaviour (PHB), which is a mechanism for enforcing QoS for different flows in the IP Backbone.

5.1.3.4 Billing system interfaces

The CMS, AN 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. The present document defines security requirements and specifications for the communication with RKS.

5.1.3.5 Call Signalling

The call signalling architecture defined within IPCablecom is Network Based Call Signalling (NCS) per TS 101 909-4 [5]. The CMS is used to control call setup, 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.

5.1.3.6 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 TS 101 909-13 [13] and TS 101 909-12 [8].

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. The present document 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 the present document. After a voice communication enters the PSTN, the security requirements as well as specifications are the responsibility of the PSTN.

5.1.3.7 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 clause 6.8 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-end.

NOTE 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.

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 message 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.

NOTE 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 TS 101 909-3 [4].

5.1.3.8 Audio Server services

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

5.1.3.8.1 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.

5.1.3.8.2 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.

5.1.3.9 Lawful Intercept

The event interface between the CMS and the DF provides descriptions of calls, which may be perform wiretapping (also known as user monitoring, or handover). This information includes the media stream encryption key and the corresponding encryption algorithm. Lawful intercept is the subject of TS 101 909-20.

5.2 Threats

The diagram below contains the interfaces that were analyzed for security threats.

To quote TS 101 909-2 [3], v1.2.1, clause 5, "Note that specific product implementations may combine functional components as needed. Not all components are required to be present in an IPCablecom Network."

There are additional interfaces identified in IPCablecom but for which protocols are not specified. In those cases, the corresponding security protocols are also not specified, and those interfaces are not listed in figure 2.

As well, the interfaces for which security is not required in IPCablecom are not listed.

IPCablecom Security Interfaces

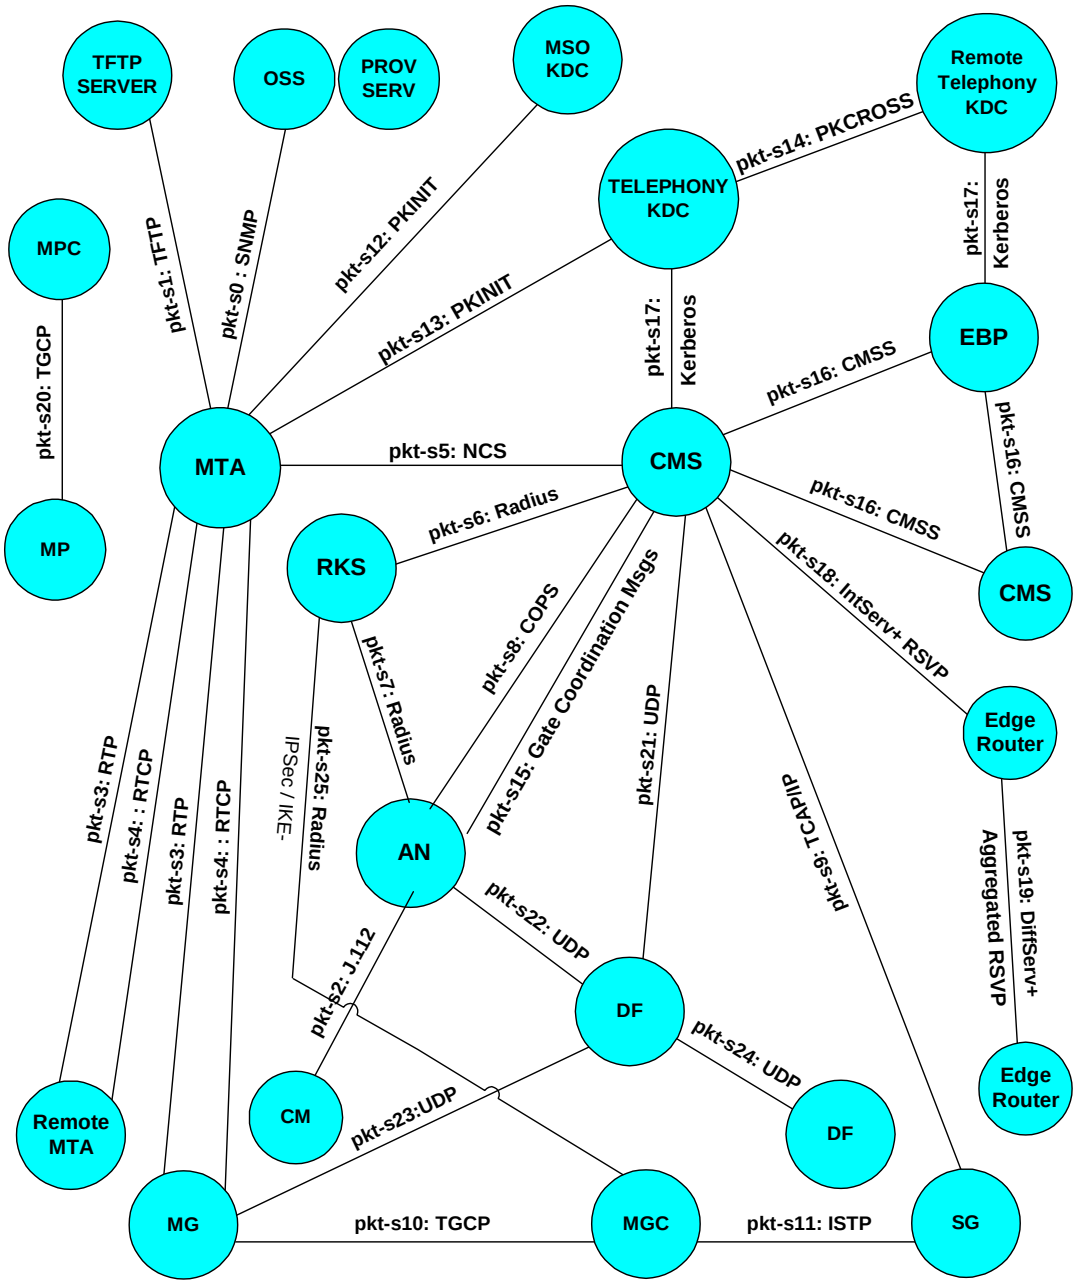

Figure 2: IPCablecom Secured Interfaces

Below is a summary of 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 clause 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.

Threats introduced by lawful interception functionality

A lawful interception facility requires the system to support anonymous call legs to anonymous call parties. These call legs shall be unidirectional. In addition lawful interception requires the system to support event recording and transmission of event records to anonymous third parties. Such messages have to be ignored in transit devices and should be invisible to transit devices and to unauthorized users. Identities in the network subject to interception shall not be able to be determined by unauthorized users.

In order to support the interception functions in a device the normal operation of that device is extended. The extension of the functionality of the device complicates the protocol and may be considered as an attack on the protocol.

- Unauthorized invocation of interception;
- Unauthorized delivery of interception.

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; and
- Quality of Service.

5.2.1.1 Cloning of MTAs

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 setup. The customer name, address, 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 Cloning of other network elements

It is conceivable that the cryptographic identity of another network element, such as a AN 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 AN 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 tamperproof protection for the cryptographic keys. (Note that 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.

5.2.2.2 Off-line cryptanalysis

Bearer channel information may be recorded and then analyzed 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 were already described under the 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 setup 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 of 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 IPCablecom 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.

- **Impersonate 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 analyzed 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 (IPCablecom 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.

5.3 Security architecture

5.3.1 Overview of security interfaces

Figure 3 summarizes all of the IPCablecom security interfaces, including key management.

IPCablecom Security Interfaces with Key Management

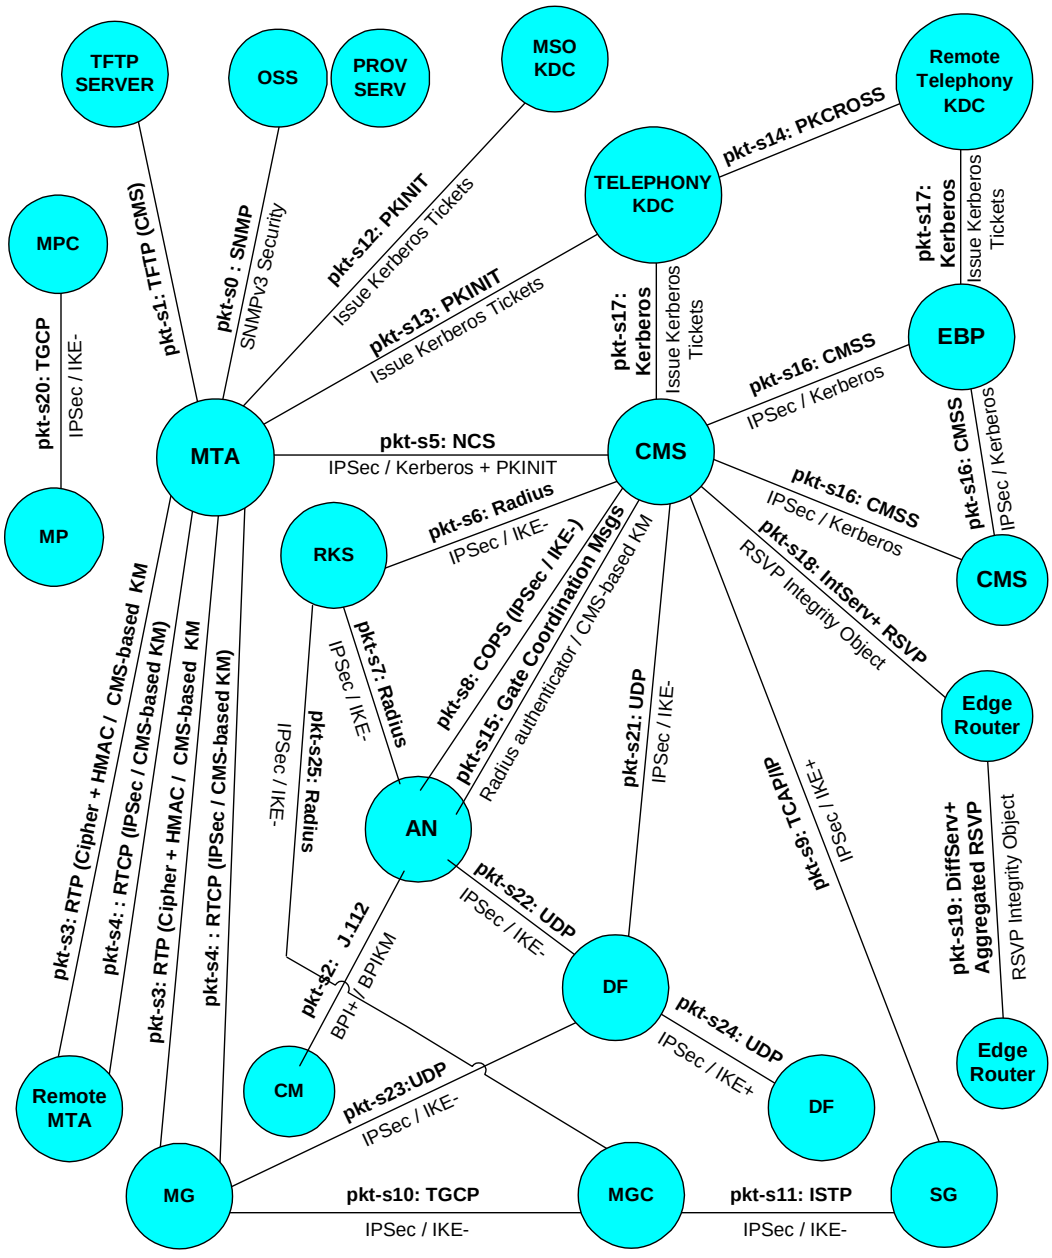

Figure 3: IPCablecom Security Interfaces with Key-Management

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 on this diagram.

The following abbreviations are used in figure 3:

IKE- IKE with pre-shared keys
IKE+ IKE with X.509 certificates

CMS-based KM Keys randomly generated and exchanged inside signalling messages.

Table 1 briefly describes each of the interfaces shown in the above diagram:

Table 1: IPCablecom Security interfaces table

Interface	Components	Description
pkt-s0	MTA - PS/OSS	SNMPv3: The initial SNMPv3 INFORM from the MTA to the Provisioning Server, followed by optional SNMP GET(s) by the SNMP Manager, is used to query MTA device capabilities. This occurs at the time where SNMPv3 keys may not be available, and security is provided with an RSA signature, formatted according to Cryptographic Message Syntax. Later, standard SNMPv3 security is enabled to the OSS.
pkt-s1	MTA - TFTP	TFTP: MTA Configuration file download. The MTA downloads a secure configuration file (with TFTP-get) that is signed by the TFTP server and sealed with the MTA public key, with a Cryptographic Message Syntax wrapper.
pkt-s2	CM - AN	J.112: 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: Radius billing events sent by the CMS to the RKS. Radius authentication keys are hardcoded to 0. IPSec is used for message integrity, as well as privacy. Key management is IKE-.
pkt-s7	RKS - AN	Radius: Radius events sent by the AN to the RKS. Radius authentication keys are hardcoded to 0. IPSec is used for message integrity, as well as privacy. Key management is IKE-.
pkt-s8	CMS - AN	COPS: COPS protocol between the GC and the AN, used to download QoS authorization to the AN. Security is provided with IPSec for message integrity, as well as privacy. Key management is IKE-.
pkt-s9	CMS - SG	TCAP/IP: CMS queries the PSTN Gateway for LNP (Local Number Portability) and other voice communications services. Security is provided with IPSec for message integrity as well as privacy. Key-Management is IKE-.
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-.
pkt-s11	MGC - SG	ISTP: IPCablecom interface to the PSTN Signalling Gateway. IPSec is used for both message integrity and privacy. Key management is IKE-.
pkt-s12	MTA - MSO KDC	PKINIT: An AS-REQ message is sent to the KDC as before, except public-key cryptography is used in the initial authentication step. The KDC verifies the certificate and issues a ticket granting ticket (TGT). The KDC authenticates the message using its public key signature.
pkt-s13	MTA - Telephony KDC	PKINIT: See pkt-s12 above.
pkt-s14	Telephony KDC - Remote Telephony KDC	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).
pkt-s15	CMS - AN	Gate Coordination messages for DQoS. Message integrity is provided with an application-layer (Radius) authenticator. Keys are distributed by local CMS over COPS.

Interface	Components	Description
pkt-s16	CMS - CMS CMS - EBP	SIP: IPSec is used for both message integrity and privacy. Key management is Kerberos.
pkt-s17	CMS - Telephony KDC EBP - Remote Telephony KDC	Kerberos: PKINIT requests (AS Request - AS Reply) for a TGT between Kerberos realms, both Intradomain and Interdomain. The TGT request generates a cross-realm TGT request using the TGS Request - TGS Reply and the PKCROSS mechanisms. The cross-realm TGS Reply is used to generate the AS Request - AS Reply needed to establish Security Associations between two domains.
pkt-s18	CMS - ER	IntServ + RSVP: Secured using RSVP Integrity Objects.
pkt-s19	ER - ER	Aggregated RSVP: Secure using RSVP Integrity Objects.
pkt-s20	MPC - MP	TGCP: IPSec is used for both message integrity and privacy. Key management is IKE-.
pkt-s21	DF - CMS	UDP: IPSec is used for both message integrity and privacy. Key management is IKE-.
pkt-s22	DF - AN	UDP: IPSec is used for both message integrity and privacy. Key management is IKE-.
pkt-s23	DF - MG	UDP: IPSec is used for both message integrity and privacy. Key management is IKE-.
pkt-s24	DF - DF	UDP: IPSec is used for both message integrity and privacy. Key management is IKE+.
pkt-s25	RKS-MGC	Radius: Radius events sent by the MGC to the RKS. Radius authentication keys are hardcoded to 0. IPSec is used for message integrity, as well as privacy. Key management is IKE-.

5.3.2 Security assumptions

5.3.2.1 AN downstream messages are trusted

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

Most messages secured in the present document 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 AN 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 AN. The CM is unable to cryptographically know that the network element it has connected to is the true AN for that network. However, even if a AN 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 AN.

5.3.2.2 Non-repudiation not supported

Non-repudiation, in the present document, 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 setup - to prove that a particular party had initiated that particular communication. However, due to very strict requirements on the setup 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 the present document 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

The present document does not attempt to address all or even most denial-of-service attacks. The cryptographic mechanisms defined in the present document 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.

The present document 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 setup 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 the present document.

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 the present document.

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 clause 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 clause 6.8 that is on the HFC J.112 links. The goal of this security document 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 AN 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; and
- 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; and
  - Identity fraud.
- An MTA running a bad code image that disrupts communications made by other MTAs or degrades network performance.

5.3.3.3 AN

The AN communicates both over the J.112 path and over the shared IP network. When the AN sends downstream messages over the J.112 path, it is assumed that a perpetrator cannot modify them or impersonate the AN. Implementing clause 6.8 over that path provides for privacy.

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

While the AN, 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 AN should be addressed by cryptographic authentication and authorization of the AN operators, as well as by physical and procedural security, which are all out of the scope of the IPCablecom documents.

A AN is responsible for:

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

A compromise of a AN may result in:

- Service theft by reporting invalid information to the RKS;
- Unauthorized levels of QoS;
- Loss of privacy, since the AN 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; and
- 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 clauses. 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 the present document), 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 documents.

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 AN;
- Distributing per communication keys for MTA-MTA signalling, bearer channel, and DQoS messages on the MTA-AN and AN-AN links; and
- 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 AN; and
  - 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; and
- 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; and
- 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; and
  - 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; and
- Reporting statistics to the RKS.

A compromise of the MG may result in:

- Service theft by reporting invalid information to the RKS; and
- 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; and
- 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 the present document:

- ASN.1 encoded messages and objects MUST conform to the Distinguished Encoding Rules per 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).

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 IETF RFC 2401 [24].

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 IETF RFC 2406 [25], 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 IETF RFC 2401 [24].

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 *IETF 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).

The following table describes the IPSec Transform Identifiers (all of which are specified in IETF RFC 2451 [28]) supported by IPCablecom.

Table 2: IPSEC ESP transform identifiers

Transform ID	Value (Hex)	Key size (in bits)	MUST support	Description
ESP_3DES	3	192	yes	3-DES in CBC mode
ESP_RC5	4	128	no	RC5 in CBC mode
ESP_IDEA	5	128	no	IDEA in CBC mode
ESP_CAST	6	128	no	CAST in CBC mode
ESP_BLOWFISH	7	128	no	BLOWFISH in CBC mode
ESP_AES	8	128	no	AES - RIJNDAEL in CBC mode
ESP_NULL	11	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 IETF RFC 2451 [28].

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.

The security specification for each IPCablecom interface lists whether or not encryption (confidentiality) is required. On each interface that requires confidentiality, the ESP_NULL transform MUST NOT be used.

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 IETF RFC 2406 [25]. 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 following IPSec authentication algorithms (all of which are specified in IETF RFC 2451 [28]).

Table 3: IPSEC authentication algorithms

Authentication algorithm	Value (Hex)	Key size (in bits)	MUST support	Description
HMAC-MD5	1	128	yes (also required by IETF RFC 2407)	MD5 HMAC
HMAC-SHA	2	160	yes	SHA-1 HMAC

The security specification for each IPCablecom interface lists whether or not message integrity is required. The HMAC-MD5 and HMAC-SHA 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 Parameter must be created as specified in IETF RFC 2406 [25].

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 clause 6.2), on other interfaces it is Kerberos/PKINIT (see clause 6.4.3.1), and in some cases IPSec keys are distributed over protected signalling interfaces (see clause 6.5.6).

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 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 of 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-AN and CMS-RKS interfaces.

The PF_KEY interface (see IETF RFC 2367 [23]) 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 IETF 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 (fewer than 100 000);
- 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 setup. The only exception would be some unexpected error, where Security Parameter 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 IETF RFC 2409 [27]).

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 of the key can be 16-20 bytes. If the shared secret is derived from a passphrase that is a string of random alpha-numerics (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 passphrases (or any language, for that matter) 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 Parameters 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 IETF RFC 2409 [27].

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 clause 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 of clause 6.1.2.1.

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 IETF RFC 2409 [27] The use of Diffie-Hellman groups within IPCablecom is identical to that specified in IETF RFC 2409 [27].

6.3 SNMPv3

All SNMP-based network management within IPCablecom MUST run over SNMPv3 with security specified by IETF RFC 2574 [30]. 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 clause 6.5.7). In addition, SNMPv3 key management techniques specified in IETF RFC 2574 [30] 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 MUST be supported.

Table 4: SNMPv3 transform identifiers

Transform ID	Value (Hex)	Key size (in bits)	MUST be supported	Description
SNMPv3_DES	0x21	64	yes	DES in CBC mode
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 IETF RFC 2574 [30]. 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 following SNMPv3 authentication algorithms are supported (both of which are specified in IETF RFC 2574 [30]).

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 IETF RFC 2574)	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 Definitions

Application Server: In this clause, the generic term Application Server means any Kerberized application server (e.g. CMS, Provisioning Server).

PKINIT: Kerberos Public Key INITial authentication extension.

6.4.2 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 session key of a Kerberos ticket. On this interface, Kerberos (see annex B) is utilized with the PKINIT public key extension (also see annex C).

Kerberized IPSec consists of three distinct phases:

1) A client SHOULD obtain a TGT (Ticket Granting Ticket) from the KDC (Key Distribution Center). 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 3 phases of Kerberos-based key management for IPSec:

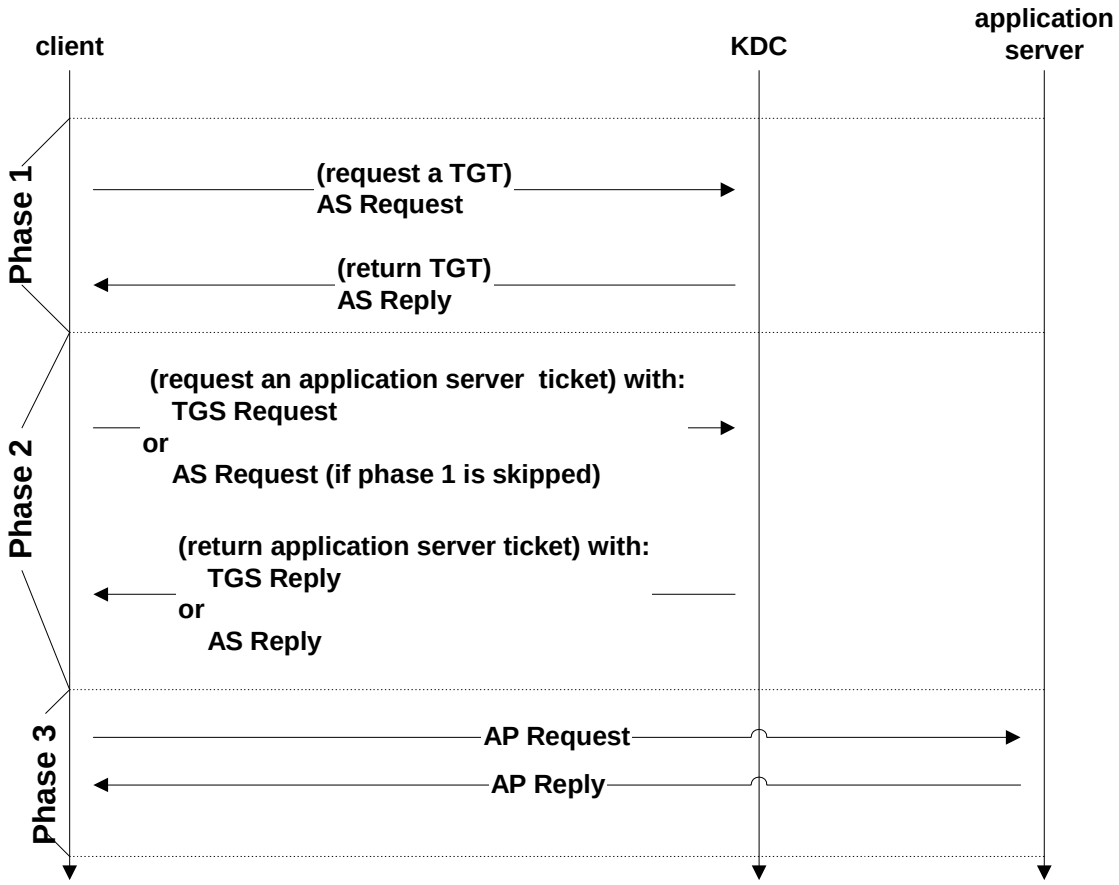

Figure 4: Kerberos-Based Key Management for IPSEC

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 clause 6.4.3.2). 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 clause 6.4.3). 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 PacketCable.

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 cluster) 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.

6.4.3 PKINIT exchange

The diagram below illustrates how a client MAY use PKINIT to either obtain a TGT (phase 1) or a Kerberos ticket for an Application Server (phase 2).

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 annex B.

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 clause 6.4.4.

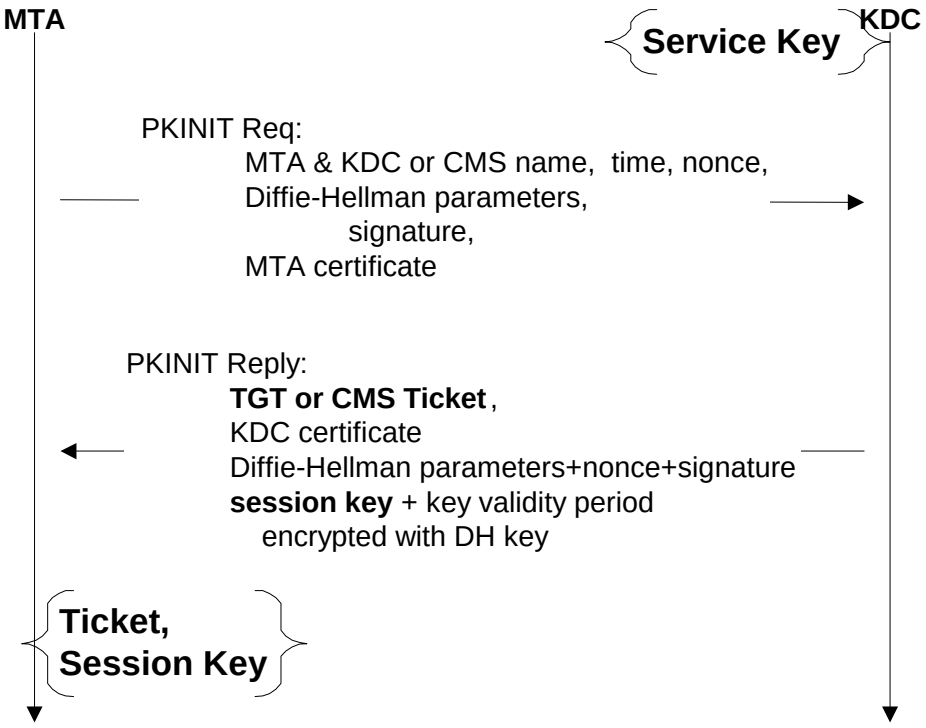

Figure 5: PKINIT Exchange

The above diagram 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 annex B). Its format is based on the MTA's X.500 name in the certificate, as specified in annex C;
- KDC or Application Server (Kerberos principal) name - found in the KDC-REQ-BODY Kerberos structure (see annex B). For the format used in IPCablecom, see clause 6.4.6.3;

- Time - found in the PKAuthenticator structure, specified by PKINIT (see annex C);
- Nonce - found in the PKAuthenticator structure, specified by PKINIT (see annex C). 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 (see annex C) and their use in IPCablecom is specified in clause 6.4.3.1.1;
- PKINIT Reply;
- TGT or Application Server Ticket - found in the KDC-REP Kerberos structure (see annex B);
- KDC Certificate, Diffie-Hellman parameters, signature - these are all specified by PKINIT (see annex C) and their use in IPCablecom is specified in clause 6.4.3.1.2;
- Nonce - found in the KdcDHKeyInfo structure, specified by PKINIT (see annex C). 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 annex B). This nonce must be the same as the one found in the KDC-REQ-BODY of the PKINIT Request; and
- Session key, key validity period - found in the EncKDCRepPart Kerberos structure (see annex B).

In this diagram, 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 is either a CMS or a KDC (in which case the ticket is the TGT).

6.4.3.1 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 ($Ticket_{EXP}$) and on the PKINIT Grace Period ($PKINIT_{GP}$) The MTA MUST initiate the PKINIT exchange at the time: $Ticket_{EXP}$ - $PKINIT_{GP}$. On the interfaces where $PKINIT_{GP}$ 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 $Ticket_{EXP}$ 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 AP Request/AP 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.

Once an MTA receives an AS Reply (with the PKINIT Reply in it), 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 (which is usually the case with existing Kerberos implementations). Thus, the MTA will be able to re-use the same Kerberos ticket after a reboot, avoiding the need to perform PKINIT again, with the associated overhead of public key operations.

Since an MTA is not required to save the ticket, the MTAs that do not follow the above document should not adversely affect the performance of call signalling. Therefore, 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 the CMS 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 min within that period (7 days). The PKINIT Grace Period $PKINIT_{GP}$ MUST be at least 15 min.

6.4.3.1.1 PKINIT request

The PKINIT request message (PA-PK-AS-REQ) in annex C 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, see annex C, 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 {
eContentType ContentType,
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.)

The PKAuthenticator data type is specified by PKINIT as follows:

```
PKAuthenticator ::= SEQUENCE {
    cusec           [0] INTEGER,
                    -- for replay prevention as in IETF RFC 1510
    ctime           [1] KerberosTime,
                    -- for replay prevention as in IETF RFC 1510
    nonce           [2] INTEGER,
    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.

- 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. For details, refer to annex C. IPCablecom requirements for the Diffie-Hellman parameters (i.e. prime and generator) MUST follow the IKE specification in IETF RFC 2409 [27].

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; and

3) check the validity of the checksum in the PKAuthenticator.

6.4.3.1.2 PKINIT reply

The PKINIT reply message (PA-PK-AS-REP) in annex C 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 annex C. 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 clause 6.4.3.1.1):
  - 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 {
 nonce [0] INTEGER,
 subjectPublicKey [2] BIT STRING
}
```

Where the nonce must be the same nonce that was passed in by the client in the PKINIT Request and subjectPublicKey is the Diffie-Hellman public value generated by the KDC. The Diffie-Hellman-derived key is used to directly encrypt part of the AS Reply.

- 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 Certificate MUST also be present. The Service Provider Certificate MUST also be present in this field. 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.

- 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; and
3) check the validity of the signature in the SignerInfo field.

6.4.3.1.2.1 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 annex C. Any error code that is defined in annex C for PKINIT MAY be returned. See also annex H.

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 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.

- Annex A defines event messages for each of the Kerberos error codes. An application server, such as a KDC or a CMS, would generate these event messages, (see TS 101 909-10 [7]).

6.4.3.1.2.1.1 Clock skew error

When the KDC clock and the client clock are off by more than the limit for a clock skew (usually 5 minutes), an error code KRB_AP_ERR_SKEW MUST be returned. 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.

6.4.3.2 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. None of the Kerberos ticket flags are currently supported within IPCablecom.

In the AS Reply, key-expiration, starttime, renew-till and caddr 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 MUST be encrypted with the encryption type set to des3-cbc-md5. The following data MUST be concatenated and processed in the following sequence before being encrypted with 3-DES CBC, IV=0:

- an 8-byte random byte sequence, called a confounder;
- an MD5 checksum, calculated as specified in annex B;
- AS Reply part that is to be encrypted; and
- 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.

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; and
2) check that the ticket has not expired.

6.4.3.3 Profile for Kerberos tickets

In Kerberos tickets, caddr, authorization-data, starttime and renew-till optional fields MUST NOT be present. None of the Kerberos ticket flags are currently supported within IPCablecom. 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.

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 the above clause 6.4.3.2.

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; and
2) check that the ticket has not expired.

6.4.4 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 preauthenticator, it MUST instead use a different preauthenticator, of type PA-ENC-TS-ENC. This preauthenticator 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 preauthenticator MUST be encrypted with the client key using the encryption type des3-cbc-md5, as described in clause 6.4.3.2. 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 preauthenticator is used to prevent replays. The KDC checks for replays upon the receipt of this preauthenticator; 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 preauthenticator differs from the current KDC time by more than pktcKdcToMtaMaxClockSkew then KDC MUST reply with a clock skew error message and the MTA MUST respond to this error message as specified in clause 6.4.3.1.2.1.1.

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 preauthenticator match any recently-seen such tuples, the KRB_AP_ERR_REPEAT error MUST be returned. The KDC MUST remember any such preauthenticator presented within pktcKdcToMtaMaxClockSkew, so that a replay attempt is guaranteed to fail.

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

Symmetric-key AS Request/AS Reply exchange is illustrated in figure 6:

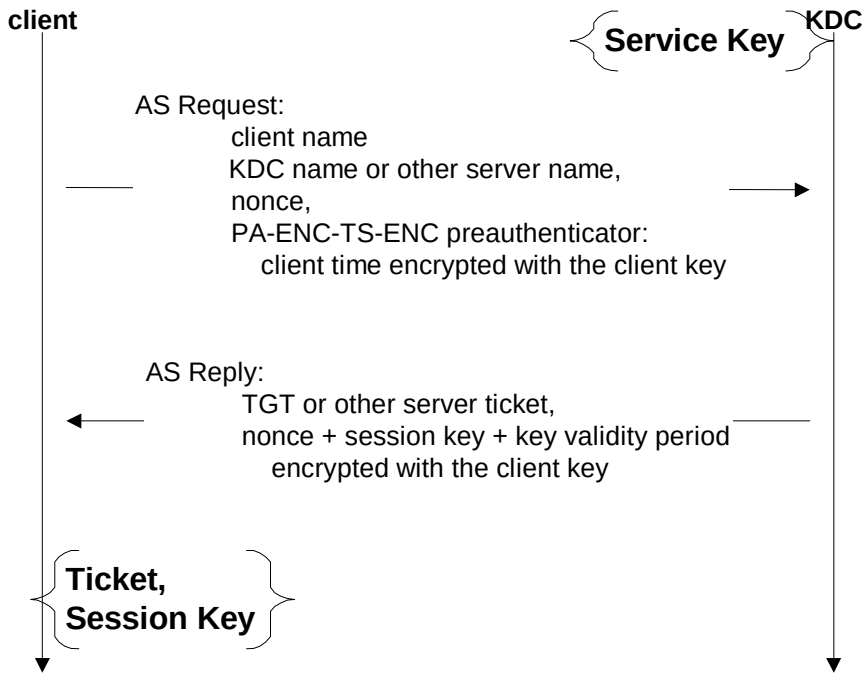

Figure 6: Symmetric-Key AS Request/AS Reply Exchange

6.4.4.1 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 preauthentication (see clause 6.4.3.2) with the exception of the type of the preauthenticator 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 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 key management would take only a single roundtrip (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 annex B. 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; and
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 re-use 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.

6.4.5 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 it is specified in annex B.

A TGS Request includes a KRB_AP_REQ data structure (the same structure used in an AP Request: see clause 6.4.5.1). 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:

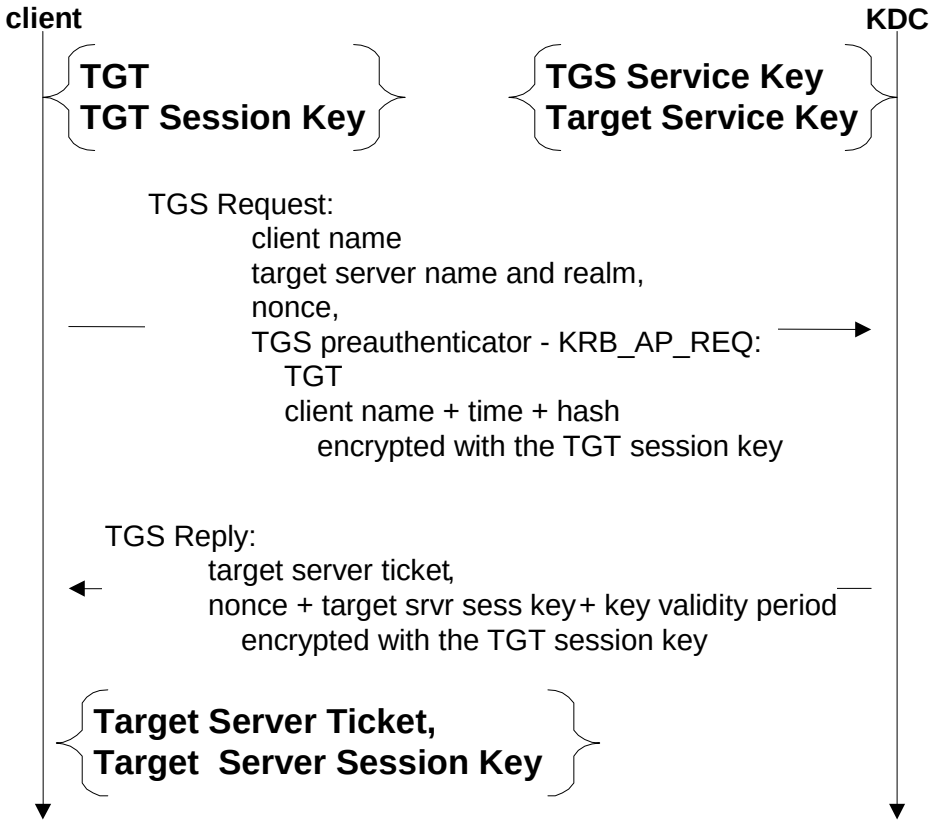

Figure 7: Kerberos TGS Request/TGS Reply Exchange

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

- TGS Request;
- Client (principal) name, target server (principal) name and realm, nonce - found in the KDC-REQ-BODY Kerberos structure (see annex B);
- TGS preauthenticator - found in the KDC-REQ Kerberos structure, inside the padata field (see annex B). The preauthenticator type in this case is PA-TGS-REQ;
- KRB_AP_REQ - the value of the preauthenticator 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 annex B); and
- Target server session key, nonce, key validity period - found in the EncKDCRepPart Kerberos structure (see annex B).

In general, the TGS Request/Reply exchange may be performed on-demand - whenever an Application Server ticket is needed to establish Security Parameters. However, there may be IPCablecom elements (e.g. MTAs) that are required to always possess a valid ticket for a particular Application Server (e.g. CMS) to improve efficiency in the load balancing scenarios. In those cases the client MUST initiate the TGS Request/Reply exchange at the time: $Ticket_{EXP}$ - TGS_{GP}. Here, $Ticket_{EXP}$ is the expiration time of the current Application Server ticket and TGS_{GP} is the TGS Grace Period.

Once a client receives a TGS 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 re-use the same Kerberos ticket after a reboot, avoiding the need to perform the TGS Request again.

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.

6.4.5.1 TGS request profile

The optional padata element in the KDC-REQ data structure MUST consist of exactly one element - a preauthenticator of type PA-TGS-REQ. The value of this preauthenticator 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: des3-cbc-md5.

The optional fields enc-authorization-data, additional-tickets and rtime in the KDC-REQ-BODY MUST NOT be present in the TGS Request. All other optional fields in the TGS Request MAY be present for IPCablecom. None of the Kerberos ticket flags are currently supported within IPCablecom.

Upon receipt of a TGS Request, the KDC MUST:

1) check the validity of the TGT; and

2) check the validity of the checksum in the authenticator.

6.4.5.2 TGS reply profile

In the AS Reply, key-expiration, starttime, renew-till and caddr optional fields MUST NOT be present. The encrypted part of the AS Reply MUST be encrypted with the encryption type set to des3-cbc-md5, using the same procedure as described in clause 6.4.3.2.

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.

6.4.5.3 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 annex B. 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 clause 6.4.4.1.

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.

6.4.6 Kerberos server locations and naming conventions

6.4.6.1 Kerberos realms

A realm name MAY use the same syntax as a domain name. For a full specification of Kerberos realms, refer to annex B.

6.4.6.2 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 annex E. For example, let us say that realm ASDF.COM contains two KDCs: kdc1.asdf.com and kdc2.asdf.com. The DNS SRV records in this case are:

_kerberos._udp.ASDF.COM. IN SRV 0 0 88 kdc1.asdf.com.

_kerberos._udp.ASDF.COM. IN SRV 1 0 88 kdc2.asdf.com.

After the above DNS SRV records are retrieved, the client will try kdc1.asdf.com first, based on its priority. (Priority for kdc1.asdf.com is 0, while priority for kdc2.asdf.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).

When in an ASCII representation, the principal identifier is as follows:

KDCQuery@<realm>

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

6.4.6.3 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 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.

For the full syntax of Kerberos principal names, refer to annex B.

6.4.6.4 Provisioning server

When an IPCablecom MTA Provisioning Server is acting in the role of an SNMP manager, its 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.

When in an 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.

The PKINITGP is not specified for the key management between the MTA and the Provisioning Server. When the MTA implementation requests a TGT in an AS Request and when the Provisioning Server realm is the same as the realm for one of the MTA signalling endpoints, the PKINITGP value of that signalling endpoint 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 KDC Grace Period is also not specified for the key management between the MTA and the Provisioning Server. The KDC Request for the Provisioning Server's ticket MAY be issued on-demand.

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) "MTA-FQDN-Map" (without quotes);

2) the FQDN of the Provisioning Server.

When in an ASCII representation, the principal identifier MUST be as follows:

MTA-FQDN-Map/<Prov Server FQDN>@<realm>

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

Names of other Kerberized services

All Kerberized services within IPCablecom 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 last component of the service principal name is always the FQDN of the corresponding host.

6.4.7 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>@<realm>

where <realm> is the Kerberos realm of the MTA and <MTA FQDN> is the MTA FQDN.

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

6.4.8 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 and consists of the following messages that the 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.

6.4.8.1 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 -> 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 in the format specified below.

The MTA FQDN Request MUST be formatted as follows:

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. 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 clause 6.4.3.2.

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 clause 6.4.4.1.

The data that is wrapped inside KRB-SAFE MUST be formatted as follows:

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	1 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 annex B, also described in clause 6.5.3. 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_ MTAMAP _ERR_ PUBKEY _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.

6.4.8.2 MTA FQDN reply

The MTA FQDN reply MUST be formatted as follows:

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 clause 6.4.3.2.

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 clause 6.4.4.1.

The data that is wrapped inside KRB-SAFE MUST be formatted as follows:

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	1 for this version
MTA FQDN	variable	MTA FQDN

After the KDC receives this reply message, it MUST validate the integrity of both the KRB-AP-REP and KRB-SAFE objects (see annex B) 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).

6.4.8.3 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 annex B. 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. The encryption method for des3-cbc-md5 is specified in clause 0.

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 clause 6.5.5.2.

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 annex B 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) kerberosApplication (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.

6.4.8.4 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 annex B, 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
```

6.4.9 Server key management time out procedure

The MTA MUST perform a configurable backoff and resend procedure for any KDC or application server requests that have not been acknowledged by the server. The handling of this procedure and any subsequent fault MUST be consistent with the procedure described in TS 101 909-4.

If the MTA has reached the maximum number of retries with a particular KDC or Application Server IP address failing to get a reply, it MUST follow the failover procedure in clause 6.4.9 of the NCS, TS 101 909-4. This procedure allows the MTA to retry with the alternative server IP addresses whenever a single server FQDN is mapped to multiple IP addresses.

6.4.10 Service key versioning

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

If a CMS service key is changed because it has been compromised, the CMS 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 CMS MUST respond as specified in annex B with a KRB_AP_ERR_BADKEYVER error. The CMS 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 clause 6.5.5.1). 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 CMS ticket which is no longer valid and fetch a new one from its KDC.

6.4.11 Kerberos cross-realm operation

It is possible that Security Parameters will have to be established between two entities in two different administrative domains. Each IPCablecom administrative domain corresponds to a Kerberos realm.

Kerberos key management requires that one of these entities (acting as a client) obtain a service ticket from the remote KDC (in the server's realm). In the case that the client entity does not use PKINIT, it cannot authenticate directly to the remote KDC.

The client MUST first obtain a cross-realm TGT for the remote KDC from its local KDC. Then, the client MUST authenticate to the remote KDC with this cross-realm TGT in order to obtain a service ticket for the remote server. And, in order for the local KDC to issue a cross-realm TGT there needs to be some trust established between the two realms.

There are two ways within Kerberos to establish trust between two realms. The first way is to establish a pre-shared key between the two realms, called a cross-realm key. The second way is the use of public key authentication and automatic establishment of the cross-realm key via the PKCROSS protocol, see annex D. PKCROSS utilizes PKINIT for establishing the inter-realm key and associated inter-realm policy to be applied in issuing cross-realm service tickets between realms and domains in support of Intradomain and Interdomain CMS-to-CMS signalling (called CMSS for CMS-to-CMS Signalling). An IPCablecom KDC MUST support PKCROSS and MAY also support pre-shared cross-realm keys.

Clause 6.4.1 outlines the various phases of Kerberos-based key management. Whenever cross-realm authentication is involved, that introduces a new key management phase. This phase would run after phase 1 and before phase 2 - we will call it phase 1.5. In phase 1.5, a client in realm A requests a cross-realm TGT for realm B from its local KDC. If a cross-realm key does not already exist between realms A and B, the local KDC MUST automatically run PKCROSS to establish the cross-realm key. The local KDC (Realm A) MUST then return the requested cross-realm TGT to the client.

Phase 2 will proceed as before, except that the client in realm A authenticates itself to the remote KDC in realm B with the cross-realm TGT obtained in phase 1.5, rather than with the TGT obtained in phase 1. Phase 3 in this case is unchanged.

Kerberos cross-realm authentication is illustrated in figure 8:

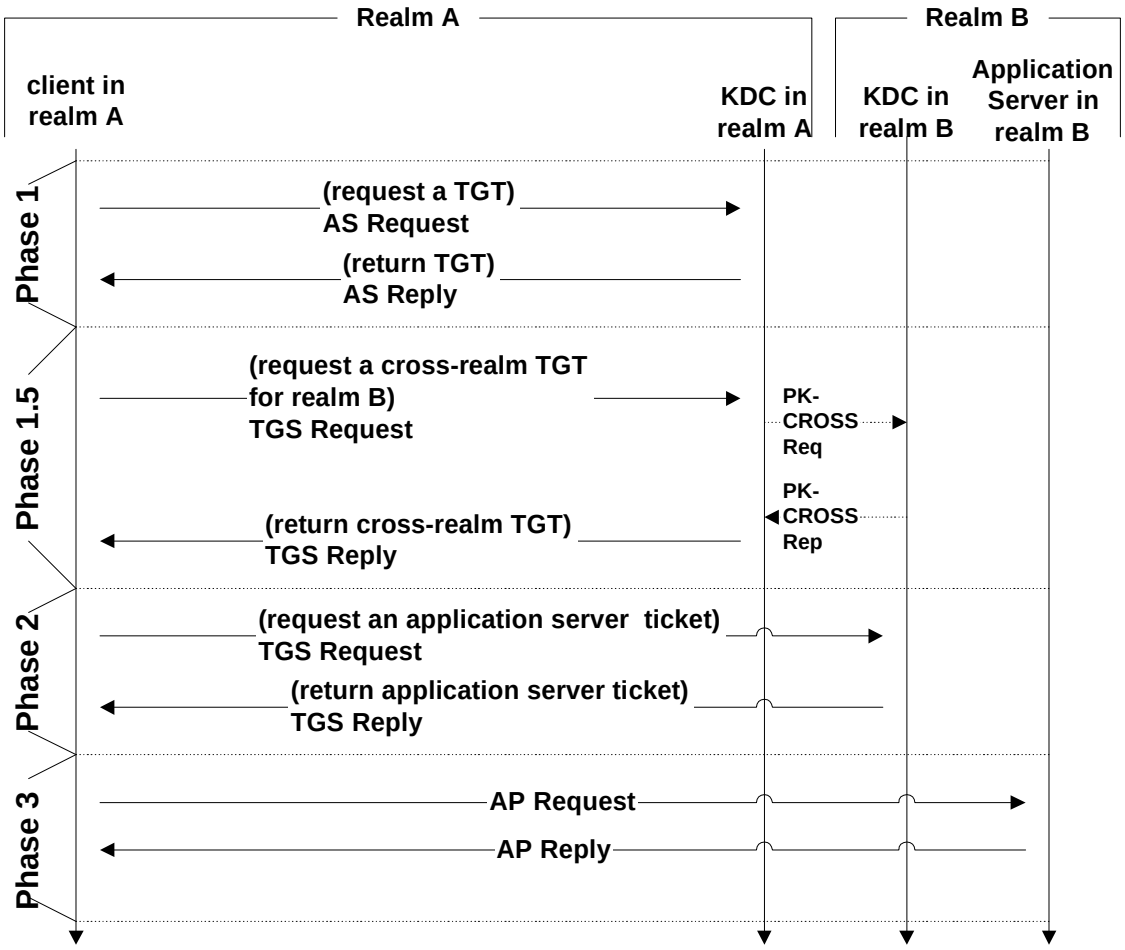

Figure 8: Kerberos Cross-Realm Operation

6.4.11.1 IPCablecom profile for cross-realm operation

There is a new exchange required by the cross-realm operation in phase 1.5, where the client obtains a cross-realm TGT. A client MUST request a cross-realm TGT with a TGS Request. The IPCablecom profile for a TGS Request/TGS Reply exchange is specified in clause 6.4.5. The only difference is that the cross-realm TGT is not encrypted with the normal TGS key of the local KDC - instead it is encrypted with a cross-realm key (see annex D).

In the case where a cross-realm key does not already exist, the TGS Request for a cross-realm TGT MUST trigger a PKCROSS exchange between the two KDCs, resulting in the automatic establishment of the cross-realm key. This is specified in the clause 6.4.11.3.

6.4.11.2 Referrals

In the above diagram, in phase 1.5 the client might not know the realm of the remote server and thus will be unable to generate a TGS Request for a cross-realm TGT.

For IPCablecom, when the client does not know the realm of the server, it MUST assume that it is in the local realm and send a TGS Request for the service ticket to the local KDC. In this case, the local KDC MUST attempt to determine the realm of the server with a query for a DNS TXT record. Below is an example of the DNS TXT records that map host names to a Kerberos realm:

_kerberos.asdf.com. IN TXT "ASDF.COM"

_kerberos.CMS1.asdf.com. IN TXT "FOO.ASDF.COM"

Let us suppose that in this case, the KDC gets a TGS Request for the service on the host CMS2.asdf.com. It would first query:

_kerberos.CMS2.asdf.com. IN TXT

Finding no match, it would then query:

_kerberos.asdf.com. IN TXT

And find an answer of ASDF.COM. This would be the realm that CMS2.asdf.com resides in.

If another TGS Request asks for the Kerberized service on the host CMS1.asdf.com, the KDC would query:

_kerberos.CMS1.asdf.com IN TXT

And find an answer of FOO.ASDF.COM.

Once the KDC has obtained the name of the remote realm using the above procedure, it MUST obtain a cross-realm ticket for that realm, which may require a PKCROSS exchange as described in clause 6.4.11.3.

In order to perform a PKCROSS exchange, the local KDC has to determine the FQDN of the remote KDC(s). If the local KDC does not already know the identity of the KDC(s) serving the remote realm, it MUST use the procedure described in clause 6.4.6.2 (in which case the KDC acts as a "Kerberos client").

PKCROSS exchange

Kerberos cross-realm authentication requires that administrators maintain separate keys for every realm for which a direct trust relationship is possible. Indirect, transitive trust is also possible, but it relies on trust of intermediate realms, and it is unnecessarily complex due to location of intermediate realms and establishment of transitive trust policies. For more information on Kerberos transitive trust issues, see annex B. Direct trust relationships require n(n-1) keys to be established and administered, which rapidly becomes an unwieldy administrative burden for maintaining keys and policies.

PKCROSS leverages a public key infrastructure (PKI) to establish trust between Kerberos realms, while it mitigates administrative issues of PKI by limiting the number of PKI endpoints to just Kerberos realms. In this way, Kerberos may be utilized for key management of large numbers of centrally administered principals, while PKI may be utilized for inter-realm key management of a small number of Kerberos realms. Thus, PKCROSS enables the dynamic establishment of cross-realm Kerberos keys. This exchange uses the PKINIT protocol and enables a remote Kerberos realm to issue a cross-realm key and policy information to another realm. This key and policy information is then returned to the remote realm in the form of a special cross-realm ticket.

Figure 8 above, depicts the flows for cross-realm authentication. The following description explains more of the PKCROSS exchange:

Table 10: PKCROSS exchange

KDCL	local KDC
KDCR	remote KDC
XTKT(L,R)	PKCROSS ticket that the remote KDC issues to the local KDC - contains the dynamic Cross-Realm key
TGT(C,R)	cross-realm TGT that the local KDC issues to the client for presentation to the remote KDC

- Phase 1.5:

KDCL issues a PKINIT request to KDCR with the PKCROSS flag (bit 9) set in the AS-REQ kdc-options field.

KDCR replies with XTKT(L,R) that is encrypted under KDCR's PKCROSS key. Note that, within the PKCROSS protocol, the PKCROSS key, as defined in the PKCROSS specification, is used in place of the TGS key. Also, KDCR MAY place policy information in ticket extensions (for example, this policy information may reflect service level agreements).

KDCL applies the policies dictated by KDCR in XTKT(L,R) and it issues TGT(C,R) to the client. This TGT(C,R) is encrypted under the key that resides in XTKT(L,R). XTKT(L,R) is added to TGT(C,R) as a ticket extension. (TicketExtensions is an optional field in a Kerberos Ticket).

- Phase 2:

When the client presents TGT(C,R) to KDCR, KDCR extracts XTKT(L,R) and is then able to decrypt TGT(C,R) and verify policy on the ticket, then it issues a service ticket to the client.

6.4.11.3 Determining the location of a remote KDC

The FQDN of a remote KDC is determined with a DNS SRV record lookup. This mechanism is identical to the mechanism used by a client to locate a KDC in the local realm.

6.5 Kerberized key management

6.5.1 Definitions

Security Parameters: A security relationship established according to the protocol for which Kerberos is being used for key management (e.g. IPSec, SNMPv3).

6.5.2 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 subkey from the AP Reply is 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 insure 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).

6.5.3 Kerberized key management messages

Figure 9 illustrates an AP Request/AP Reply exchange:

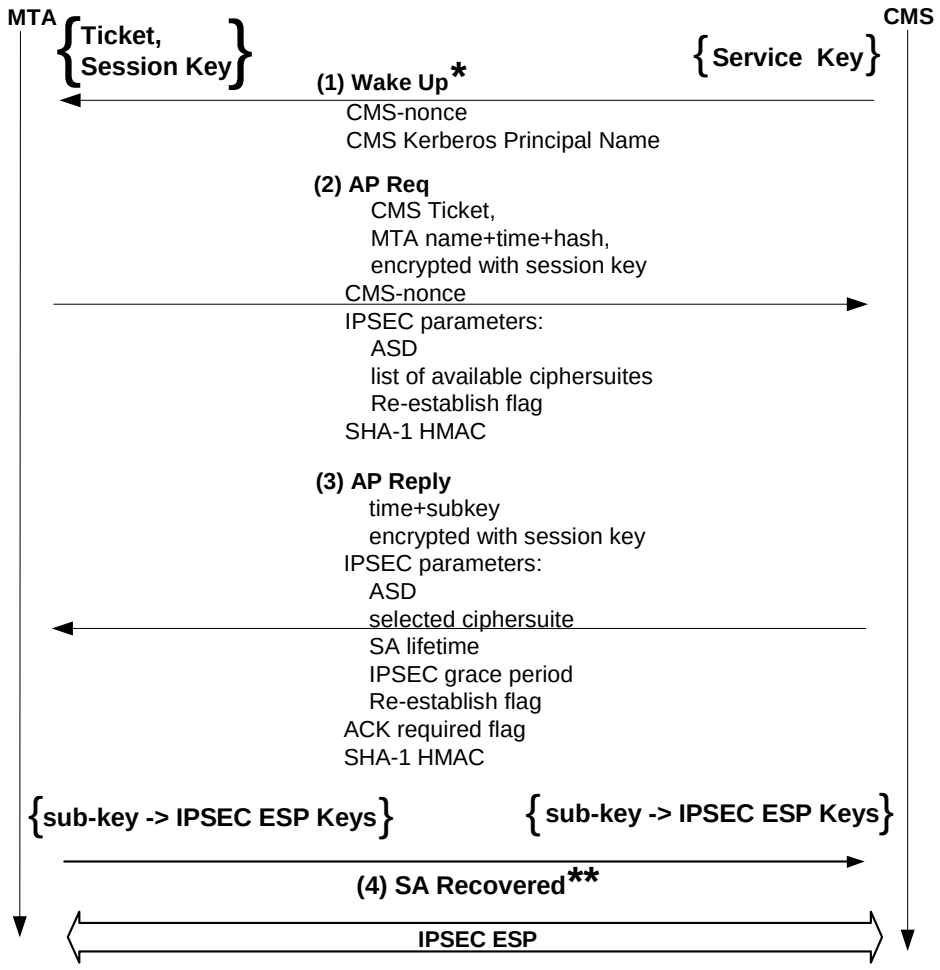

* This message is sent whenever key management is initiated by the CMS

** This message is optional, sent whenever the ACK-required flag is set in the preceeding AP Reply

Figure 9: 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 name, 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 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 0's.
- Server Kerberos principal name: a printable, null-terminated ASCII string, representing a fully qualified Kerberos Principal Name of the Application Server, as specified in annex B.

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 backoff procedure described in clause 6.4.9. When the "Timeout Procedure" has completed without success, the Application Server MUST discard this server-nonce, 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 that the Client receives a Wake Up message, it MUST respond with this AP Request.

In addition, the present document specifies the use of this message by the Client to periodically establish a new set of security parameters with the Application Server - see clause 6.5.6.4. 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 clause 6.5.6.5.

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 table 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 annex B.
- 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, next clause 6.5.4, then the value MUST be identical to that of the Server-nonce field in the Rekey message. Otherwise, the value MUST be all 0's.
- 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 Security Parameter 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 name 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 clause 6.4.9.

If the AP Request was generated in response to a message sent by the Application Server (Wake Up or Rekey), then the client MUST keep the nonce and server Kerberos principal name for pktcClntSolicitedKeyMgmtTimeout (for the MTA it is the pktcMtaDevSolicitedKeyMgmtTimeout MIB variable for IPSec key management). After the timeout has been exceeded or when the "Timeout Procedure" has completed without success, the client MUST discard this (nonce, server Kerberos principal name) 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 clause 6.4.9.

In the case that the Server-nonce is 0 (not filled in), the Application Server MUST verify that this AP Request is not a replay using the procedure specified in the Kerberos standard (see annex B):

- If the timestamp in the AP Request differs from the current Application Server time by more than pktcSrvrToMtaMaxClockSkew then Application Server MUST reply with an error message specified in clause 6.5.5.2.

- 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 is returned. The Application Server MUST remember any authenticator presented within pktcSrvrToMtaMaxClockSkew, so that a replay attempt is guaranteed to fail.

- If the Application Server loses track of any authenticator presented within pktcSrvrToMtaMaxClockSkew, 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 (see annex B). 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 subkey (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 table in clause 6.5.3.

- 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 annex B.

- 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.

- 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 as specified in clause 6.5.6.2.1. When the value is FALSE (0), the old set of security parameters MUST be allowed to expire as specified in clause 6.5.6.3.

- 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 and 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 pktcSrvrKeyMgmtTimeout3 in order to validate a Security Parameter Recovered response from the Client. After pktcSrvrKeyMgmtTimeout3 the Application Server MUST discard this value, after which it will no longer accept a matching Security Parameter Recovered.

(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. Seetable in clause 6.5.3.

- 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 subkey 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.

6.5.4 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 roundtrip. This is illustrated in figure 10:

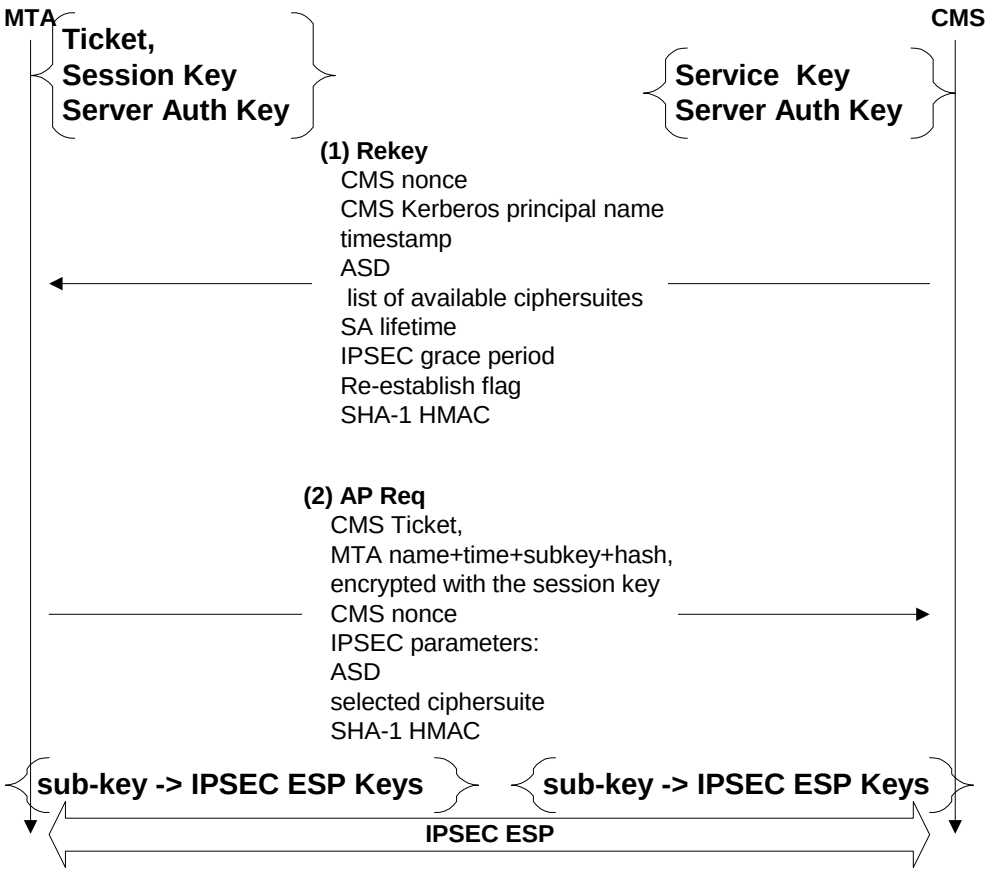

Figure 10: Rekey Message to Establish a Security Parameter

The messages listed in this diagram are defined as follows:

(1) Rekey - sent by the Application Server to establish a new Security Parameter. 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 table in clause 6.5.3.
- 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 0's.
- Server Kerberos Principal Name - a printable, null-terminated ASCII string, representing a fully qualified Kerberos Principal Name of the Application Server, as specified in annex B. 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 or SNMPv3).

- 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 as specified in clause 6.5.6.3. When the value is FALSE (0), the old set of security parameters MUST be allowed to expire as specified in clause 6.5.6.2.2.

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.

NOTE: 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 an SA Recovered message), which is not used in this case for performance reasons - to avoid an extra upstream 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 clause 6.4.9. When the "Timeout Procedure" has completed without success, the Application Server MUST discard the server-nonce, 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 annex B:

- If $|T_{CMS} - (T_{MTA} + Skew)| > pktcSrvrToMtaMaxClockSkew$ 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 that was recorded the last time this MTA sent an AP Request to this Application Server and received a clock skew error. PktcSrvrToMtaMaxClockSkew is currently in the MTA MIB (see TS 101 909-8 [11]) as the variable pktcMtaDevCmsMaxClockSkew.
- If the Server-nonce, principal name and timestamp fields match any recently seen (within the pktcSrvrToMtaMaxClockSkew) 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 subkey (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 clause 6.4.4. 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 clause 6.5.

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

6.5.5 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 are not supported.

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 - random value generated by the Client

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

- subkey - used to generate security parameters for the target protocol. The subkey type MUST be set to -1. The actual subkey length is dependent on the target protocol.

NOTE: 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 PacketCable spec will be updated accordingly.

All other optional fields within the encrypted authenticator are not supported within IPCablecom.

The authenticator itself MUST be encrypted using 3-DES CBC with the Kerberos etype value: des3-cbc-md5 as it is specified in clause 6.4.3.2.

IPCablecom does not use combined etypes that specify an encryption algorithm and a checksum, since a keyed HMAC (outside of the KRB_AP_REP) is used instead.

In the encrypted part of the KRB_AP_REP, the optional subkey 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.

6.5.5.1 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 table in clause 6.5.3.

- KRB-ERROR - Kerberos error message as specified in annex B. 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. 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 des3-cbc-md5, as it is specified in clause 6.4.4.1.

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 annex B):

```
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 it matches 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 it's 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.

6.5.5.2 Clock skew error

When the Application Server clock and the client clock are off by more than the limit for a clock skew (usually 5 minutes), an error code KRB_AP_ERR_SKEW MUST be returned. 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).

6.5.6 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 Association are always uni-directional and they MUST always be established in pairs within IPCablecom.

The DOI value for IPSec MUST be set to 1.

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

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

The subkey for IPSec MUST be a 46-byte value

6.5.6.1 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 subkey.

The size of the Kerberos subkey MUST be 46 bytes (the same as with the SSL or 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; and

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

For specific authentication and encryption algorithms that may be used by IPCablecom for IPSec, refer to clause 6.1.

The derivation of the required keying material MUST be based on running a one-way pseudo-random function F(S, "IPSec Security Associations") recursively until the right number of bits has been generated. Here, S is the Kerberos subkey and the ASCII string "IPSec Security Association" is taken without quotes and without a terminating null character. F is defined in clause 9.7.

6.5.6.2 Periodic re-establishment of IPSec security associations

An IPSec SA is defined with an expiration time TEXP and a grace period GPIPSEC. The clauses 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.

6.5.6.2.1 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 TEXP - GPIPSEC. At this time, the Client MUST send an AP Request as specified in clause 6.5. After the Client receives an AP Reply, it MUST perform the following steps:

1) Create new IPSec SAs, based on the negotiated ciphersuite, ASDs and on the established Kerberos subkey, from which the IPSec keys are derived as specified in clause 6.4.4. 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.

6.5.6.2.2 Periodic re-establishment of IPSec SAs at the application server

When an AP Request message is received and right before an AP Reply is returned, the Application Server MUST perform the following steps, in the specified order.

1) Create new IPSec SAs, based on the negotiated ciphersuite, ASDs and on the established Kerberos subkey, from which the IPSec keys are derived as specified in clause 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.

6.5.6.3 Expiration of IPSec SAs

An IPSec SA is defined with an expiration time TEXP and a grace period GPIPSEC. 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 clause 6.5.6.5.

6.5.6.4 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 clause 6.5.

6.5.6.5 On-demand establishment of IPSec SAs

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

6.5.6.5.1 Client loses an outgoing IPSec SA

When a Client attempts to send a signalling message to the Application Server without a valid IPSec SA. At that time, the IPSec layer in the Client realizes the SA is missing and returns an error back to the signalling application.

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

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 clause 6.4.3.

2) Client sends a new AP Request to the Application Server and gets back an AP Reply, as specified in clause 6.5.4. After the receipt of the AP Reply the Client MUST be prepared to use both of the newly created IPSec SAs.

3) The Application Server MAY set an ACK-required flag in the AP Reply. In that case, right after sending out an AP Reply, the Application Server MUST be prepared to receive messages on the incoming SA but cannot yet start using an outgoing SA for sending messages to the Client. In this case, the IPSec SA setup continues with the following steps 6 and 7.

4) The Application Server also MAY NOT set the ACK-required flag in the AP Reply. In that case, right after sending out an AP Reply, the Application Server MUST be prepared to both send and receive messages on the newly created SAs. In this case, steps 3 and 4 below are skipped.

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.

6) Immediately after the Client establishes the new IPSec SAs, it sends a SA Recovered message to the Application Server.

7) Upon receipt of this message, the Application Server will immediately activate the new outgoing SA for sending signalling messages to the Client.

The key management application running on the Client MUST send an explicit signal when it completes the re-establishment of the IPSec SAs. The signalling application at the Client MUST retry sending an AP Request after some period of time when it has not received an explicit signal from the key management application running on the same Client, when it completes the establishment of IPSec SAs.

6.5.6.5.2 Client loses an incoming IPSec SA

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

6.5.6.5.3 Application server loses an outgoing IPSec SA

When a 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 application.

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

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 clause 6.5. For each AP Request, the Client generates a nonce and puts it into the seq-number field. As specified in clause 6.5, the Client will save this nonce for a short period of time 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 timeout, the Client MUST NOT retry and MUST abort an attempt to establish a IPSec SA in response to a received Wake Up.

4) Once the Client gets back a matching AP Reply, it will be in the format specified in clause 6.5. 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 set, the Client MUST be prepared to automatically re-establish new IPSec SAs, as specified in clause 6.5.6.3.

5) In the event that 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 when it completes the re-establishment of the IPSec SAs. The signalling application at the Application Server MUST retry a WAKE UP send after some period of time when it has not received an IPSec SA established signal.

6.5.6.5.4 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.

6.5.6.6 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) kerberosApplication (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.

6.5.7 Kerberized SNMPv3

This clause specifies the Kerberized key management profile specific to SNMPv3, see IETF RFC 2574 [30]. In the case of SNMPv3, the security parameters are associated with the usmUserName (SNMPv3 user name), agent's usmUserEngineID (SNMPv3 engine ID) and 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. Within IPCablecom, SNMPv3 authentication MUST always turned on. In addition, SNMPv3 privacy MAY also be used (it can be turned off by selecting a NULL encryption transform).

The DOI value for SNMPv3 MUST be set to 2.

The Application-specific data field in the AP Request and AP Reply key management messages MUST be set to the SNMPv3 Engine ID corresponding to the SNMP agent. The SNMP Manager MUST verify the SNMPv3 Engine ID contained in the AP Request based on the contents of the client principal name contained in the ticket. For IPCablecom MTAs, the manager MUST verify that the MTA FQDN specified in the principal name corresponds to the SNMPv3 Engine ID. The SNMPv3 Engine ID in the AP Reply MUST be the same as the one in the preceding AP Request.

The Application-specific data field in the AP Request and AP Reply key management messages MUST be set to the concatenation of the following:

This format MUST be used for the data in the AP Request message:

Table 11: 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
Agent's snmpEngineTime	4 bytes
usmUserName Length	1 byte
usmUserName	variable

This format MUST be used for the data in the AP Reply message:

Table 12: 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
Manager's snmpEngineTime	4 bytes
usmUserName Length	1 byte
usmUserName	variable

For IPCablecom MTAs, the usmUserName contains in it the MTA MAC address (see TS 101 909-6 [10]). 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 the SNMP INFORM message 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 subkey for SNMPv3 MUST be a 46-byte value.

6.5.7.1 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 subkey.

The size of the Kerberos subkey MUST be 46 bytes.

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

- SNMPv3 authentication key; and
- SNMPv3 privacy key.

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

The derivation of the required keying material MUST use a one-way pseudo-random function F(S, "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 clause 9.7.

6.5.7.2 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.

6.5.7.3 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.

6.5.7.4 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 clause 6.5.7.5.1.

6.5.7.5 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.

6.5.7.5.1 SNMP agent wishes to send with missing SNMPv3 keys

An SNMP agent is capable of initiating protocol exchanges with the manager, e.g. with the SNMP Trap and SNMP Inform messages. If the SNMP agent determines that it is missing SNMPv3 keys, it MUST perform the following steps before it is able to send out an SNMP message:

1) (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 clause 6.4.3.

2) (2) The agent sends a new AP Request to the manager and gets back an AP Reply, as specified in clause 6.4.4. 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.

6.5.7.5.2 SNMP agent receives with missing SNMPv3 keys

The SNMP agent receives a request from a manager and is unable to find SNMPv3 keys for the specified user. The agent MUST ignore this error and the message MUST be dropped. In this case, any attempt at recovery (e.g. establishing new SNMPv3 keys) is prone to denial-of-service attacks.

6.5.7.5.3 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) 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) (3) Agent sends a new AP Request to the manager, as specified in clause 6.5.3. For each AP Request, the agent generates a nonce and puts it into the seq-number field. As specified in clause 6.5.6.5.3, the agent will save this nonce for a short period of time 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 timeout, 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 clause 6.5.3. The ACK-required flag in the AP Reply MUST be set, to insure 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) (4) After the receipt and validation of the AP Reply, the agent sends 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) (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 for the same SNMPv3 user name before the first key management session has been completed, the SNMP agent MUST ignore the second Wake Up message.

6.5.7.6 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) kerberosApplication (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:

KRB_SNMPV3_ERR_USER_NAME	1	Unrecognized SNMPv3 user name
KRB_SNMPV3_ERR_NO_CIPHER	2	No support for requested ciphersuites
KRB_ SNMPV3_ERR_ENGINE_ID	3	Invalid SNMPv3 Engine ID Specified
KRB_ SNMPV3_ERROR_GENERIC	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 clause 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 the Transform IDs and Authentication Algorithms that are available for RTP security.

Table 13: RTP packet transform identifiers

Transform ID	Value	Key size (in bits)	MUST support	Description
reserved	0x50	-	-	
RTP_AES	0x51	128	yes	Rijndael with the parameters used in the AES submission
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
RTP_RC4	0x57	128	no	RC4 stream cipher
reserved	0x58-59	-	-	

The RTP_AES Transform ID MUST be supported.

Table 14: RTP IPCablecom 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 clause 6.6)	yes	2-byte MMH MAC
reserved	0x63	-	-	
RTP_MMH_4	0x64	variable (see clause 6.6)	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 clause 7.6.2.2. Key Management for RTCP requires that both the (encryption) Transform ID and the Authentication Algorithm be specified. This clause lists the Transform IDs and Authentication Algorithms that are available for RTCP security.

Table 15: RTCP Packet transform identifiers

Transform ID	Value	Key size (in bits)	MUST support	Description
reserved	0x70	-	-	-
AES-CBC	0x71	128	yes	Rijndael with the parameters used in the AES submission
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 Transform ID MUST be supported.

Table 16: RTCP authentication algorithms

Transform ID	Value	Key size (in bits)	MUST support	Description
reserved	0x80	-	-	-
HMAC-SHA1-96	0x81	160	Yes	First 12 bytes of the HMAC-SHA1 per IETF RFC 2404 [35]
HMAC-MD5-96	0x82	128	No	First 12 bytes of the HMAC-MD5 per IETF RFC 2403
reserved	0x83-8f	-	-	

The HMAC-SHA1-96 authentication algorithm MUST be supported.

6.8 Additional requirements for cable modems

6.8.1 Additional requirements for cable modems based on ITU-T Recommendation J.112 annex A

6.8.1.1 Requirements

All MTAs MUST use ITU-T Recommendation J.112, annex A, compliant CMs and MUST implement the security option described in ITU-T Recommendation J.112 annex A. The security option provides security services to the data link layer traffic streams running across the cable access network, i.e. between CM and INA. These services are message confidentiality and access control which provide CM users with data privacy across the cable network and protect cable operators from theft of service.

The INA MUST encrypt all packet cable traffic in the downstream direction.

The CM MUST encrypt all packet cable traffic in the upstream direction.

6.8.1.2 Security mechanisms provided

The protected ITU-T Recommendation J.112 annex A data communications 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.

Employing the ITU-T Recommendation J.112 [1], annex A, security option which meets the above requirements, the INA protects against unauthorized access to these data transport services by enforcing encryption of the associated traffic flows across the cable network. Key management is performed using an extended set of MAC messages. For unicast streams, the keys are derived using Diffie-Hellman between the INA and CM. For mulitcast streams, a client/server approach is adopted with the INA controlling the distribution of the keys to the CMs. During the key exchange for both unicast and multicast streams the CM is authenticated by the INA.

6.8.1.3 Packet data encryption

ITU-T Recommendation J.112 annex A encryption services are defined as a set of optional services within the MAC. Encryption can be selectively applied to the various payload data streams. For each secure stream two session keys can be used for encrypting and decrypting it of which only one of the keys is used to process any particular payload unit. Each key can be used for processing both upstream and downstream payload data.

Having two keys allows negotiation of a new key to take place while payload data is processed using the old one, and then an immediate switch-over can be performed once the new key is agreed upon, without interrupting payload traffic. The INA initiates the key exchanges, and can start using a session key for downstream traffic encryption once the key exchange is complete. For upstream traffic encryption, the NIU should use whichever key was used by the INA in the most recent payload unit.

A payload stream is identified by either of:

- A 24-bit (UNI) ATM virtual circuit VPI/VCI: this is used for ATM-based IB downstream, OOB downstream, and upstream payload data. The ATM circuit can be one-to-one, or one end-point of a multicast circuit.
- A 48-bit MAC-address: this is used for Multiprotocol Encapsulation downstream payload data. The MAC-address can be the physical address of the STB or a pseudo address used for MAC-address based multicasting.

For ATM-based payload streams, the unit of encryption is a single ATM cell. The 48-byte cell payload is encrypted using the security context implied by the 24-bit VPI/VCI of the cell header. For Multiprotocol Encapsulation payload streams, the unit of encryption is a single Multiprotocol Encapsulation clause. The datagram_data_bytes (between the MAC-address and the CRC/checksum) are encrypted using the security context implied by the 48-bit MAC address in the clause header.

Bits in the ATM header (GFC field) and in the MPE clause (scrambling field) are used to identify whether the payload is encrypted and the session key it was encrypted with.

The currently supported algorithms are 40 and 56-bit DES in CBC mode. Additional algorithms may be support in later revisions of ES 200 800 [39].

6.8.1.4 Key management

This is achieved by using Diffie-Hellman, which requires no up-front shared secret, or a simpler protocol based on a long-term shared secret between INA and CM called a cookie. The cookie is also used for authenticating the CM to the INA during the key exchanges. Three mechanisms exist for establishing a shared key between the INA and the CM and all three are initiated by the INA, these are:

Main Key Exchange: This uses Diffie-Hellman to derive a shared secret between the INA and CM, which is independent of the cookie value. It can also be used to update the cookie value held in the CM.

Quick Key Exchange: This uses the existing cookie value to derive a shared secret key.

Explicit Key Exchange: This is used by the INA to deliver a pre-determined session key to the CM. The session key is encrypted under a temporary key derived from the cookie value.

As stated earlier, two session keys can be in place for a given stream, therefore during the key exchange the particular session key being generated is identified by the INA. The above mechanisms are also used to update the keys for an active stream. See ES 200 800 [39] for further information about key management.

6.8.2 Additional requirements for cable modems based on ITU-T Recommendation J.112 annex B

All Clients MUST use J.112 compliant cable modems and MUST implement BPI+. 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 AN. 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 communications 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 AN 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 AN (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 AN to CMs. Through this key management protocol, CM and AN synchronize keying data; in addition, the AN 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 "back haul" networks, which lie behind ANs. Since J.112 and IPCablecom service providers typically will not guarantee the security of their managed IP back haul 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 Radius

Radius protocol requires an authenticator field for all messages, which provides message integrity. No other security services or key management are defined within the Radius standard IETF RFC 2139.

A 16-byte Authenticator field is calculated as follows:

- Request Authenticator: MD5 hash calculated over a stream of octets consisting of the Request Code + Identifier + Length + 16 zero octets + request attributes + shared secret (where + indicates concatenation).

- Response Authenticator: MD5 hash calculated over a stream of octets consisting of the Response Code + Identifier + Length + Request Authenticator field from the Accounting-Request packet being replied to + the response attributes if any + shared secret

The shared secrets for the Response and Request Authenticator fields do not have to be the same.

IPCablecom interfaces that utilize Radius require that the authentication algorithm (ciphersuite) be specified (see clause 6.1.2.2). Currently, only the standard Radius authentication mechanism (as described above) is supported and the ID for this authentication algorithm is 100 (decimal).

Upon receipt of a RADIUS message, the recipient MUST check the value of the authenticator.

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; and
- For each constituent protocol interface, specifying the particular security mechanism providing the required security services.

Clause 5.2 describes 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 per-protocol security analysis is organized by functional categories (see clause 5.2.1.5). 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):

Table 17: 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. Since a CMS can perform multiple functions, the security profile matrices indicate which of the CMS functional components is participating in the identified protocol interface.

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 TS 101 909-6 [10].

Figure 11 illustrates the flows involved with the provisioning processes. The provisioning document lays these flows out in detail. The flows involving security mechanisms are described in this clause of the document.

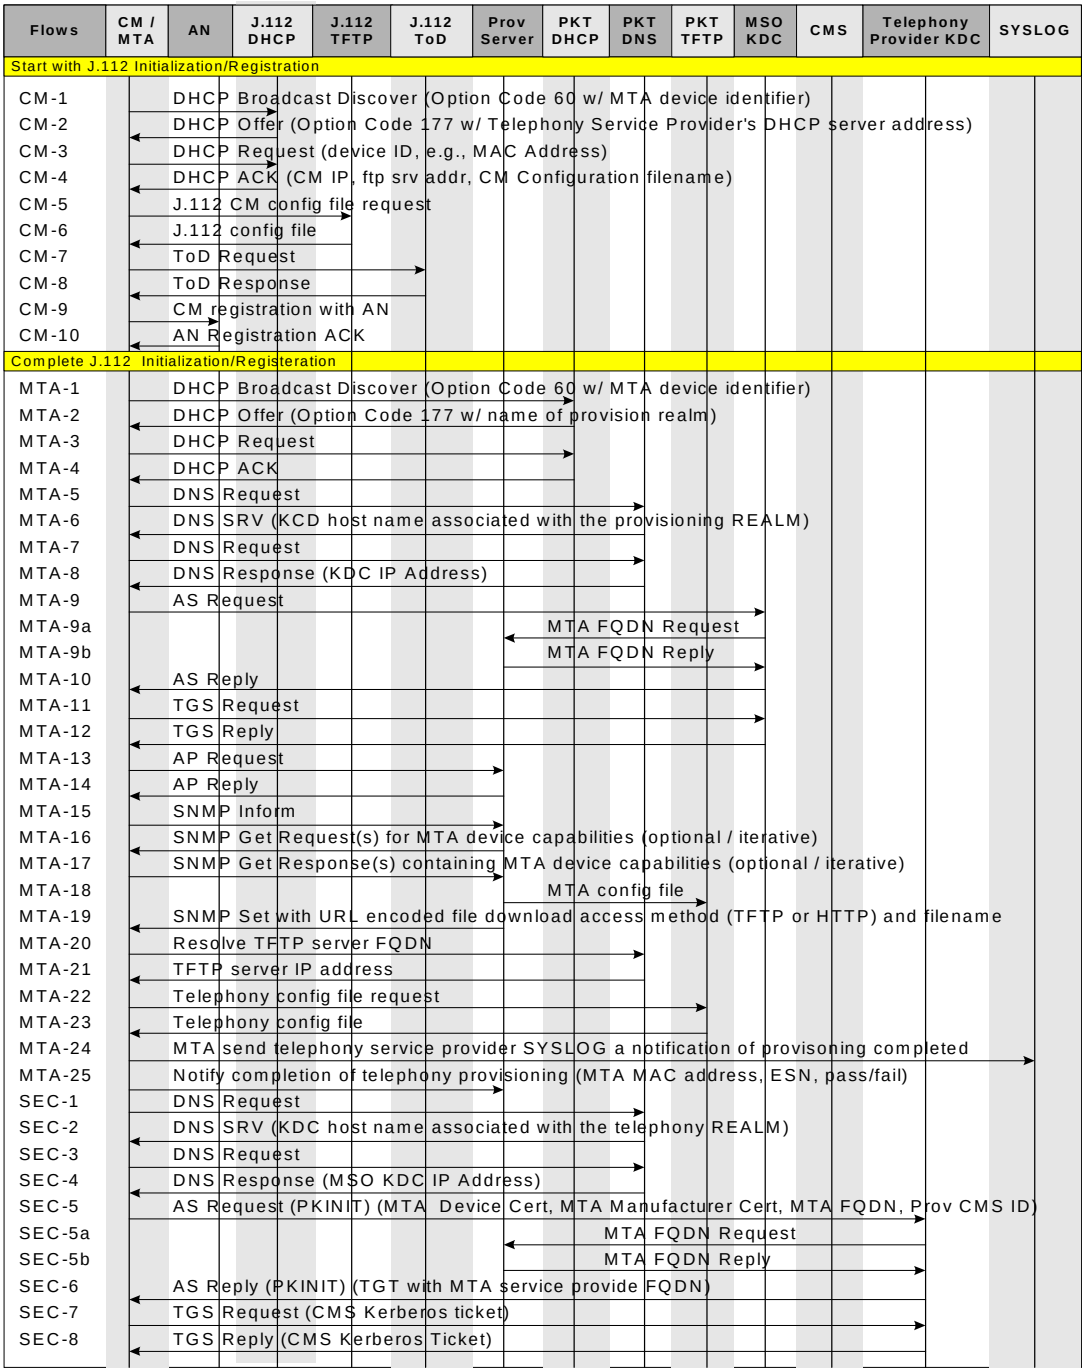

Figure 11: IPCablecom provisioning flows

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

Table 18: Kerberos key management during MTA provisioning

Flow Step	Security Requirement	Life Time	Step bypass permitted
MTA-9/MTA-10 -- AS Request/AS Response (see clause 6.4.3.1)	TGT ticket if using TGS Request, Provisioning Server Ticket if otherwise		If current ticket is stored in NVRAM and has not expired
MTA-11/MTA 12 -- TGS Request/TGS Response (see clause 6.4.5)	Applies when a TGT is used. Obtains a Provisioning Server Ticket	Life time less than life time of TGT ticket	If TGT is not used or if current Provisioning Server Ticket is stored in NVRAM and has not expired
MTA-13/MTA-14 -- AP Request/AP Response (see clause 6.5.3 and clause 6.5.7)	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 reinitialized. 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.

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.

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, where an MTA is 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

All SNMP traffic between the MTA and the SNMP Manager in both directions is protected with SNMPv3 security per IETF RFC 2574 [30]. IPCablecom requires that SNMPv3 message authentication is always turned on with privacy being optional. 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 IETF RFC 2574 [30] for the list of SNMPv3 cryptographic algorithms supported by IPCablecom.

7.1.1.1.3 MTA-provisioning server, via TFTP server

Authentication: the identity of the OSS that generated the MTA configuration file is authenticated with comparing the hash of the configuration file that was generated by the Provisioning Server and transported to the MTA securely via SNMPv3 against the hash of the configuration file downloaded from the TFTP server. This is required to prevent denial-of-service attacks, where an MTA is improperly configured.

The identity of the MTA requesting the file is not authenticated. Authentication of the MTA is not required; if there is a need to keep the file contents private, it may be encrypted with the Provisioning Server-MTA shared key and no one else will be able to use it.

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

Confidentiality: is 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 Flow MTA-15: 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 IETF RFC 2574 [30]. IPCablecom requires that SNMPv3 message authentication is always turned on with privacy being optional. 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 clause 6.5.7 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 Config File

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; and
- 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 do the following steps:

1) If 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 (1). 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 (2) 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 configuration file format is described in clause 7.1.1.2.2 above.

7.1.1.2.5 Security flows

The following security flows MUST be performed immediately following the provisioning process. These flows MUST be performed for every entry in the table PktcMtaDefCmsTable.

Table 19: Post-MTA provisioning security flows

Sec Flow	Flow description	If Step Fails, Proceed Here
Get Kerberos tickets associated with each CMS with which the MTA communicates.		
SEC-1	1 DNS Request The MTA requests the Telephony KDC host name for the Kerberos realm.	SEC-1
SEC-2	DNS SRV Returns the Telephony KDC host name associated with the provisioning REALM.	SEC-1
SEC-3	DNS The MTA now requests the IP Address of the Telephony KDC.	SEC-1
SEC-4	DNS The DNS Server returns the IP Address of the Telephony KDC.	SEC-1
SEC-5	AS Request 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.	Report alarm. Abort establishment of signalling security.
SEC-6	AS Reply The KDC sends the MTA a PKINIT REPLY message containing the requested Kerberos Ticket.	If Step Fails, Proceed Here to SEC-5 or abort signalling security depending upon error conditions. The MTA should report the recommended alarm SEC-EV-1 or SEC-EV-2 as defined in annex A
SEC-7	TGS Request In the case where the MTA obtained a TGT in SEC-6, it now obtains the Kerberos ticket for the TGS request message.	Report alarm. Abort establishment of signalling security.
SEC-8	TGS Reply Response to TGS Request containing the requested CMS Kerberos Ticket.	If Step Fails, Proceed Here to SEC-7/SEC-5 or abort signalling security depending upon error conditions. The MTA should report the recommended alarm SEC-EV-2 as defined in annex A
SEC-9	AP Request 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.	Report alarm. Abort establishment of signalling security.
SEC-10	AP Reply 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.	If Step Fails, Proceed Here to SEC-9/SEC-7/SEC-5 or abort signalling security depending upon error conditions. The MTA should report the recommended alarm SEC-EV-3 as defined in annex A
SEC-11	The MTA responds with an SA Recovered message that lets the CMS know, the MTA is now ready to receive on its incoming IPSec Security Association. This message is sent when requested by the flag in the AP Reply. This flag should not be used in the initial provisioning flows.	Report alarm. Abort establishment of signalling security.

Several tables in the MTA MIB control security flows SEC-1 through SEC-11 listed in table 19.

The CMS table (pktcMtaDevCmsTable) and the realm table (pktcMtaDevRealmTable) are used for managing the MTA security signalling security. 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 177, suboption 6.

b) There MUST be a realm table entry for each CMS table entry. Multiple table CMS 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 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 are used as defined in the current CMS table entry. If different parameters are desired, the CMS table entry MUST be pre-staged.

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 ticket for the specified CMS, use the pktcMtaDevCmsKerbRealmName parameter in the CMS table entry to index into 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 TS 101 909-6 [10]) or when an endpoint is configured (via SNMP sets) for a new CMS in the NCS MIB (see TS 101 909-9 [12]).

3) Any time before an MTA endpoint sends an RSIP 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 indicate to the NCS signalling software on the MTA that the security association is presently not valid.

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 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 or other application server (see clause 6.4.3). The Telephony KDC issues the Kerberos Ticket for a group of one or more CMSs or other application servers, 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 or other application servers, the MTA MUST obtain multiple Kerberos Tickets by repeating these call flows for each ticket. It is also possible, that the MTA is configured to request Kerberos Tickets from different Telephony KDC servers, depending on the CMS or other application server group.

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 or other application server. In the event that different MTA ports are configured with different CMS or other application server (FQDN) names, multiple sets of SAs will be established (one set for each CMS).

Since a Kerberos Ticket is issued for a group of CMSs or other application servers, it is possible that a single Kerberos Ticket is used to establish more than one set of IPSec SAs.

In IPCablecom, a CMS or other application server 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 initially establish SAs with one of the IP addresses returned by the DNS Server. The MTA MAY also establish SAs with the additional CMS or other application server 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 clause 6.5.7. 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 IETF RFC 2574 [30] 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 clause 6.4.6.1). When there is more than one KDC (DNS SRV record) found, DNS assigns a priority to each one. The MTA will choose a KDC based on the DNS priority labeling 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

Table 20: 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 enrollment

The subscriber enrollment process establishes a permanent customer billing account that uniquely identifies the MTA to the CMS via the MTA's MAC address. The billing account is also used to identify the services subscribed to by the customer for the MTA.

Subscriber enrollment MAY occur in-band or out-of-band. The actual specification of the subscriber enrollment 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 enrollment is performed in-band, a communication to a CSR (or to an automated subscriber enrollment system) is protected using the same security mechanisms that are used to secure all other voice communication.

During each communication setup (protected with IPSec ESP), the CMS or other application server 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 enrollment"). Some additional services, such as communications with emergency response organizations (e.g. 911), may also be permitted in this case. Since in-band customer enrollment is based on standard security provided for call signalling and media streams, no further details are provided in this clause. Refer to clause 7.4 and to clause 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

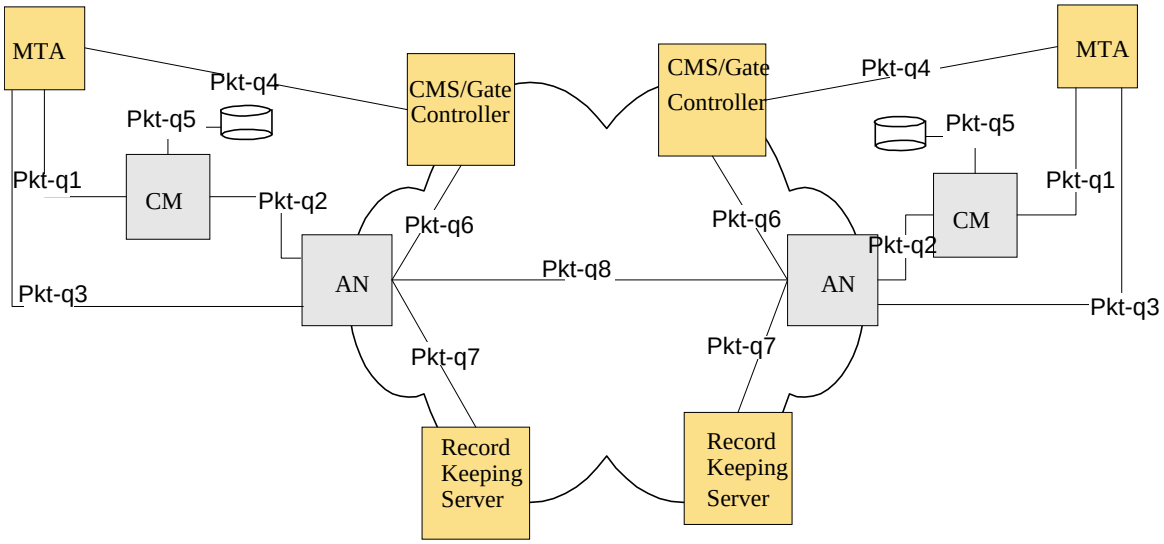

Figure 12: QoS Signalling Interfaces in IPCablecom Network

7.2.1.2 Security services

7.2.1.2.1 CM-AN J.112 QoS messages

Refer to the ITU-T Recommendation J.112 [1].

7.2.1.2.2 *AN*-CMS gate coordination Messages (over UDP)

The AN-CMS Gate Coordination messages are required in order to prevent some theft of service scenarios, described in the DQoS document, TS 101 909-5 [6].

Authentication: required to prevent theft of service and denial-of-service attacks. Direct authentication is not possible due to the large number of associations and performance requirements (for post-dial and post-pickup delays). Each end is indirectly authenticated, via a trusted third party - the CMS.

Message Integrity: required on this interface. Without it, the same service theft scenarios are still possible, along with the denial-of-service attacks.

7.2.1.2.3 Gate controller - AN 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-AN J.112 1.1 QoS Messages

The J.112 QoS messages are specified in the J.112 RFI Recommendation.

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 AN. 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 AN, these attributes include details of how the CM requests mini-slots and the expected behaviour of the AN upstream scheduler.

A Classifier is a set of matching criteria applied to each packet entering the cable network. A classifier 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 AN to packets it is transmitting, and Upstream Classifiers are applied at the CM and may be applied at the AN 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 MTAs IP source address and flooding the network with the packets intended for another MTAs destination address. A AN controlling downstream service flows will limit an MTAs downstream bandwidth according to QoS allocations. If the AN 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 from flooding the network with unauthorized IP packets.

To address J.112 AN accesses to the network, the AN 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 AN-CMS Gate Coordination Messages (over UDP)

These are DQoS handshake messages between the two sides of the communication, to ensure that the QoS resources had been reserved on both sides. DQoS handshake messages MUST be formatted as Radius messages, which MUST include an authenticator field. Security for this interface is provided solely by the Radius-specific authenticator, based on an MD5 hash, as defined by IETF RFC 2139 [20]. This provides message integrity, but not privacy. The key for the Radius authenticator MUST be exactly 16 bytes long, and there MUST be a separate key used for each direction.

7.2.1.3.2 Gate Controller - AN 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 AN. These COPS messages MUST be both authenticated and encrypted with IPSec ESP. Refer to clause 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 AN-CMS gate coordination messages (over UDP)

The keys for this interface MUST be securely distributed by the local CMS over the existing IPSec links. The key MUST be included in the Gate Authorization message (COPS) sent from the local CMS to the AN. For the on-net-to-on-net communications, the keys for Gate Coordination messages on each side of the communication are different. There is only one message that is involved in the key management for this interface: the Gate-Set Authorization message sent to the AN by the local CMS, which MUST include the following parameters:

- Transaction ID (associates request an response together);
- Gate ID;
- Authentication algorithm (1 byte) - only Radius, MD5-based MAC is currently supported, (see clause 6.9);
- Radius Key (16 bytes); and
- IP address of the local CMS (listed as the IP address of the remote gate in the Gate-Set messages).

For the gate coordination exchange to proceed, the AN MUST accept this Radius Key and ciphersuite and MUST use it to authenticate Gate Coordination messages in both directions for the specified Gate ID.

To address communication keying replay protection concerns, the same Radius key is used to authenticate only a few messages. These messages consist of the Gate-Open/Gate-Open-Ack and the Gate-Close/Gate-Close-Ack exchanges. Each message type is clearly identified and appears exactly once for each communication.

7.2.1.4.2 Gate Controller - AN COPS messages

Key management for this COPS interface MUST be implemented via IKE and use pre-shared keys. For more information on the IPCablecom use of IKE, refer to clause 6.2.2.

7.2.1.4.3 Security profile matrix summary

Table 21: Security Profile Matrix - DQoS

	COPS (AN-CMS)	**Gate Coordination (AN - CMS)**
authentication	Yes	yes (through a 3rd party)
access control	Yes	no
integrity	yes	yes
confidentiality	yes	no
non-repudiation	no	no
Security Mechanisms	IPSec with encryption and message integrity IKE w/ pre-shared keys	Radius authentication (MD5-based MAC) CMS distributes key per communication

7.2.2 Interdomain QoS

Interdomain QoS consists of two different mechanisms, Differentiated Services (DiffServ) and an admission control protocol called RSVP, see IETF RFC 2205.

7.2.2.1 Architecture overview

The overall IQoS network architecture is depicted in figure 13. The backbone consists of a general topology managed IP network that may comprise multiple administrative domains.

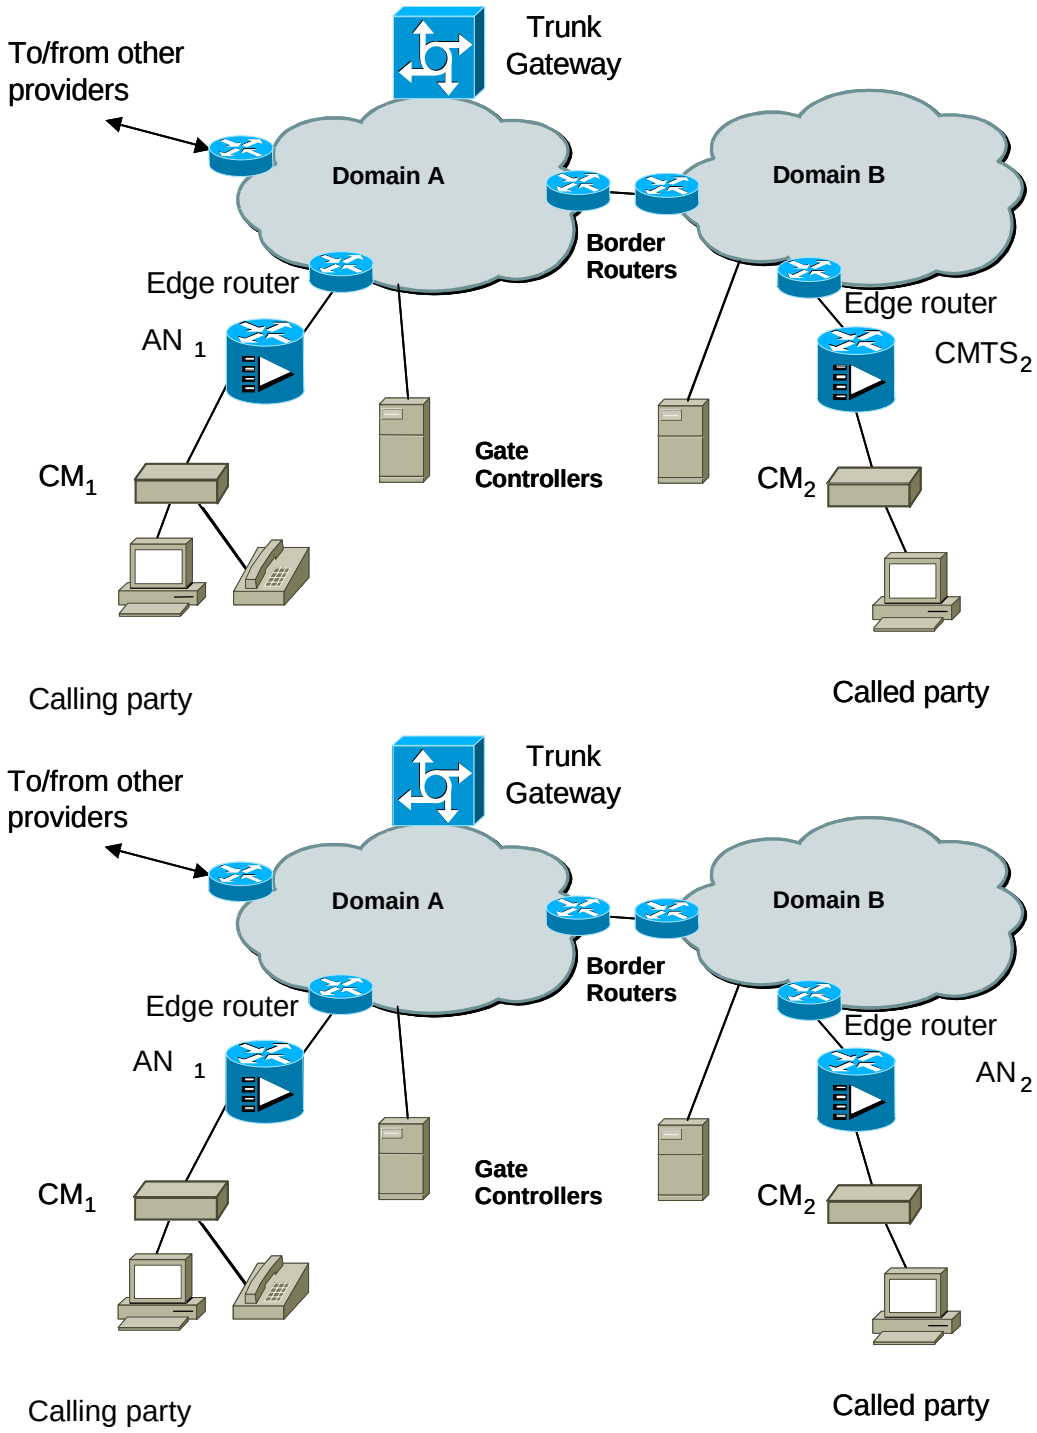

Figure 13: Interdomain QoS Architecture

In this architecture, we assume that DQoS signalling is used in the access network. At a minimum, the backbone is expected to be compliant with the DiffServ architecture. Border routers are those that sit at the boundaries between providers. They have specific roles in a DiffServ environment (such as aggregate policing and re-marking) that are discussed in more detail in the clauses that follow.

7.2.2.2 Differentiated Services (DiffServ)

DiffServ allows IP traffic to be marked with different DiffServ Code Points (DSCP) to obtain different queuing treatment on routers. Different queuing treatments in each router are called per hop behaviour (PHB) that is a mechanism for enforcing QoS for different flows in the IP Backbone. The IPCablecom Interdomain Quality of Service, TS 101 909-17 [17], defines separate per hop behaviour for media flows versus signalling flows. It also provides PHB rules of AN managing of upstream bandwidth.

The use of DiffServ Code Points (DSCP) and their associated per hop behaviour (PHB) MUST comply with requirements in TS 101 909-17 [17].

7.2.2.2.1 Security services

Authorization: This is the only security service provided by DiffServ. Based on the DSCP, each IP packet is authorized for a different level of QoS at each DiffServ router. As DiffServ does not provide any authentication, it is possible for an unauthorized router to remark an IP packet with an invalid QoS level. This DiffServ vulnerability is generally tolerated because cryptographic processing on each packet at each router is considered to be too much overhead.

7.2.2.2.2 DiffServ summary security profile matrix

Table 22: Security Profile Matrix - IQoS

	DiffServ Router- DiffServ Router
authentication	no
access control	Yes, in the form of DSCP authorization
integrity	no
confidentiality	no
non-repudiation	no
Security Mechanisms	PHB Authorization based on DSCP. DSCP marking is performed on the upstream channel of edge and border routers.

7.2.2.3 Resource reSerVation Protocol (RSVP)

RSVP is normally used to preserve bandwidth for individual media stream sessions. This type of admission control may be provided at each IPCablecom AN. When the IP traffic is entering the IP Backbone, keeping track of every individual resource reservation does not scale well enough for the expected growth of IPCablecom networks. For this reason, the IPCablecom IqoS, TS 101 909-17 [17], optionally defines the aggregation of RSVP reservations at the edge and border routers. In general, the IQoS Specification defines various RSVP mechanisms for admission control, but none of them are required for IPCablecom. This clause specifies the IPCablecom profile for RSVP security for those IPCablecom elements that employ RSVP.

The aggregation (edge and border) routers have the responsibility of creating aggregate reservations across an aggregation region, which may be the entire DiffServ cloud or a defined aggregation region within the cloud. Each aggregate reservation represents an aggregate flow of traffic from an ingress router (or aggregator) to an egress router (the de-aggregator). Aggregate reservations may be configured statically based on the expected load from an ingress to an egress router, or they may be automatically established and re-sized. Each aggregate reservation carries the traffic from a number of "end-to-end" RSVP reservations that share a common ingress/egress router pair. An end-to-end reservation represents a single microflow, and signalling for such a reservation is accomplished using standard RSVP. "End-to-end" RSVP messages are originated by the AN on behalf of the MTA. Such E2E RSVP messages are "tunneled" across the aggregation region by setting the IP protocol number in the Path message to "RSVP-E2E-IGNORE".

7.2.2.3.1 Security services

Authentication: required to prevent theft of bandwidth and denial-of-service attacks. Because RSVP messages are modified at each RSVP capable router, RSVP authentication is performed separately at each RSVP hop.

Access Control: required on this interface. The mechanism used to provide access control and policy information at RSVP routers is currently out of scope for IPCablecom.

Message Integrity: required on this interface. Without it, the same bandwidth theft scenarios are still possible, along with the denial-of-service attacks.

Confidentiality: not supported on this interface as it is not part of the IETF standard. Normally, bandwidth reservation signals via RSVP do not carry application level data that needs additional privacy protection.

7.2.2.3.2 Cryptographic mechanisms

Authentication and Message Integrity: All RSVP messages MUST include an RSVP INTEGRITY object as it is specified in IETF RFC 2747 [32]. The RSVP integrity object contains a keyed message digest over the entire RSVP message. Within IPCablecom, the keyed message digest algorithm MUST be HMAC-MD5.

In order to support secure RSVP router restart/recovery, all RSVP routers MUST support the integrity handshake mechanism as specified in IETF RFC 2747 [32].

Access Control: automatically provided with the use of pre-shared keys. Access is granted to another router when pre-shared keys are established with that router.

7.2.2.3.3 Key-management

Pre-shared keys MUST be used to secure all RSVP interfaces. IPCablecom has reviewed IETF RFC 2752, Identity Representation for RSVP where public key and Kerberos approaches are used for RSVP authentication. Within IPCablecom, public key and Kerberos approaches for RSVP authentication MUST NOT be used.

Public key authentication is considered to be too computationally expensive to be used for RSVP. Kerberos authentication, as it is specified in IETF RFC 2752, is incomplete.

7.2.2.3.4 RSVP summary security profile matrix

Table 23: Security Profile Matrix - RSVP

	RSVP Router- RSVP Router
authentication	yes
access control	yes
integrity	yes
confidentiality	no
non-repudiation	no
Security Mechanisms	HMAC-MD5 for message authentication and integrity. Pre-shared keys satisfy access control requirements.

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 insure that the billing events reported to the RKS are not falsified.

Confidentiality: required to protect subscriber information and communication patterns.

7.3.1.2 AN-RKS interface

Authentication, Access Control and Message Integrity: required to prevent service theft and denial-of-service attacks. Want to insure 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 insure 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 clause 6.1.

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, AN 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 clause 7.3.3 on how this map is maintained.

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 AN 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

CMS and RKS MUST negotiate a shared secret (CMS-RKS Secret) using IKE. IKE MUST be use one of the modes with pre-shared keys for this interface. For details, refer to clause 6.1.2.3.

IKE MUST run asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired CMS-RKS Secret. This shared secret MUST be unique to this particular CMS and RKS.

At the RKS, the CMS 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 billing event messages. For more details refer to IETF RFC 2407 [26].

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 AN - RKS interface

AN and RKS MUST negotiate a shared secret (AN-RKS Secret) using IKE. IKE may use one of the modes with pre-shared keys. For details, refer to clause 6.1.2.3.

IKE MUST be running asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired AN-RKS Secret. This shared secret SHOULD be unique to this particular AN and RKS.

At the RKS, the AN 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 billing event messages. For more details refer to IETF RFC 2407.

Later, when a billing event arrives at the RKS, it MUST be able to query the database of IPSec SAs 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.

7.3.3.3 MGC - RKS interface

MGC and RKS MUST negotiate a shared secret (MGC-RKS Secret) using IKE. IKE may use one of the modes with pre-shared keys. For details, refer to clause 6.1.2. 3.

IKE MUST be running asynchronous to the billing event generation and will guarantee that there is always a valid, non-expired MGC-RKS Secret. This shared secret SHOULD be unique to this particular MGC and RKS.

At the RKS, the 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 billing event messages. For more details refer to IETF RFC 2407 [26].

Later, when a billing event arrives at the RKS, it MUST be able to query the database of IPSec SAs 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.

7.3.4 Billing system summary security profile matrix

Table 24: Security Profile Matrix - Radius

	RADIUS Accounting (CMS-Radius Server/RKS)	RADIUS Accounting (AN - 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 with pre-shared keys	IPSec ESP with encryption and message integrity enabled key management using IKE with pre-shared keys	IPSec ESP with encryption and message integrity enabled key management using IKE with pre-shared keys

7.4 Call signalling

7.4.1 Network Call Signalling (NCS)

7.4.1.1 Reference Architecture

Figure 14 shows the network components and the various interfaces to be discussed in this clause.

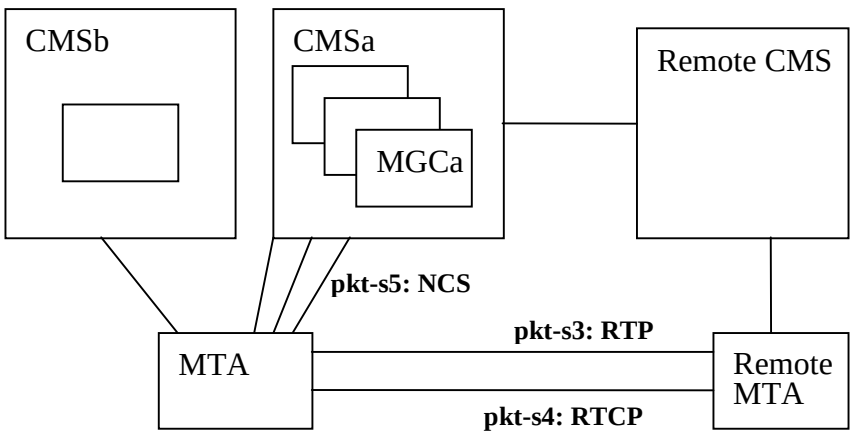

Figure 14: NCS Reference Architecture

The diagram shows the CMSs as being a cluster of several MGCs. 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 clause 6.6 of the present document. The protocol interface pkt-s16 (CMS to CMS) is SIP with IPCablecom extensions, as specified in TS 101 909-16.

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 TS 101 909-16 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.

Administratively, various SIP proxies and CMSs are grouped into Kerberos realms. If there is a signalling path between two realms, there is a trust relationship between the corresponding KDCs.

7.4.1.2 Security services

The same set of requirements applies to both CMS-MTA and CMS-CMS signalling interfaces.

Authentication: signalling messages must be authenticated, in order to prevent a third party masquerading as either an authorized MTA or CMS.

Confidentiality: NCS messages carry dialed numbers and other customer information, which must not be disclosed to a third party. Thus confidentiality of signalling messages is required.

Message integrity: must be assured in order to prevent tampering with signalling messages - e.g. changing the dialed 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 this interface. IPSec keys MUST be derived using mechanism described in clause 6.5.6.

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 <-> 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-SIP proxy and SIP Proxy - SIP Proxy interfaces

When a CMS or a SIP Proxy receives a SIP signalling message, it MUST 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 MUST utilize an IP address <-> FQDN map for all CMSs and SIP Proxies that have current IPSec security associations with this host. This map is built during key management described in the following clause and does not need to reside in permanent storage.

Whenever the SIP signalling message contains the FQDN of the previous hop (e.g. in the SIP Via header field), that FQDN MUST be identical to the FQDN in the local IP address <-> FQDN map. If they are not the same, the SIP message MUST be rejected as an invalid message.

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 clause 6.4.4). 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 clause 6.4.5).

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 clause 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 switchover from one key lifetime to another key lifetime is also defined.

The key distribution and timeout 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 timeout parameter. Similarly, the MTA will obtain the sub-key (and thus IPSec ESP keys) based on the IPSec timeout 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 6.5.6.5.3 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 single PKINIT Request/Reply sequence with the KDC and multiple AP Request/Reply sequence for each Call Agent in the cluster. The Kerberos messages are specified in clause 6.4.5

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 a IPSec SA with that MTA (see clause 6.5.4).

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 roundtrip time.

This is illustrated in figure 15:

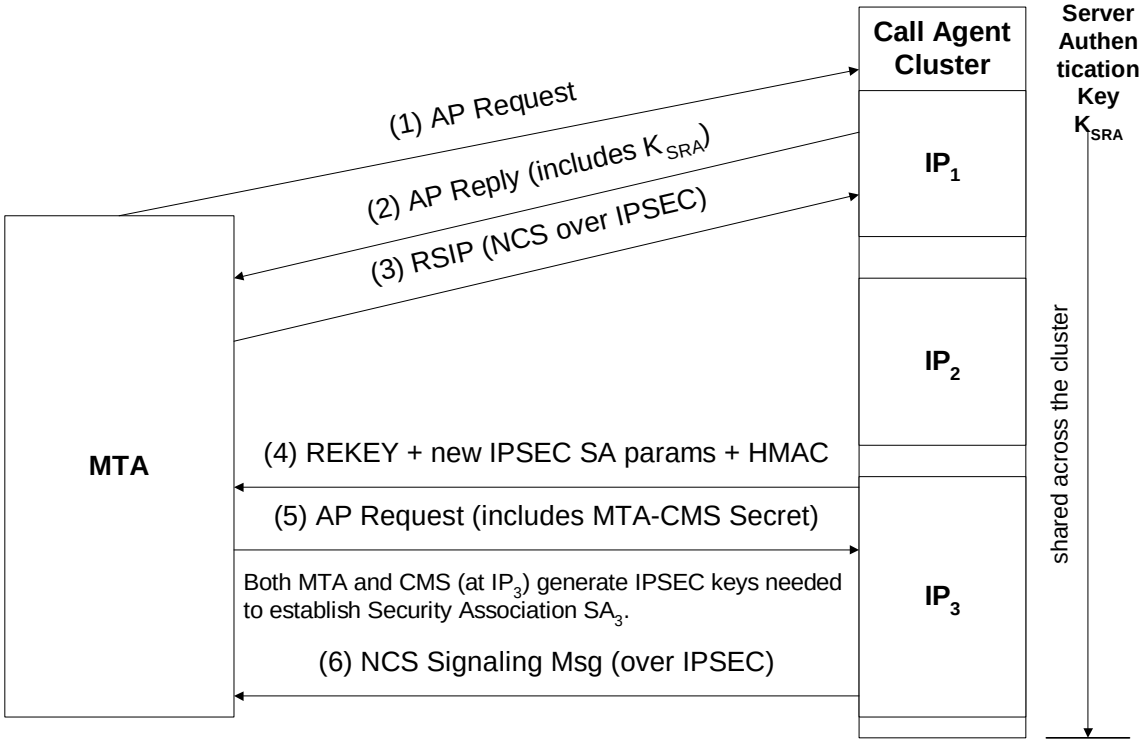

Figure 15: Key Management for NCS Clusters

In this figure, 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 KSRA for authenticating Rekey message from the CMS was derived from the session key used to encrypt the AP Reply.

When a new SA3 needs to be established between the MTA and CMS at IP3, the key management is as follows:

(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): ASD, 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 KSRA.

(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 clause 6.5.4.

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 TGS, the MTA will need to execute multiple PKINIT Request/Reply, one for each CMS and then multiple AP Request/Reply in order to create the security association with the individual MGCs.

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:

The new CMS, or other application server might not have been included in the CMS table. In that case, the corresponding table entry MUST be locally created. For a new CMS, or other application server table entry, the realm MUST be copied from the application server that previously controlled this endpoint. Default values SHOULD be used for the clock skew and AP request/AP reply adaptive backoff parameters.

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.

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-SIP Proxy, SIP Proxy-SIP Proxy key-management

When a CMS/other application server or a SIP Proxy has data to send to another application server or SIP Proxy and does not already have Security Associations with that host, it MUST utilize a Kerberized key management protocol (see clause 6.5) to establish them. In this case, any application server or SIP Proxy is responsible for obtaining a ticket for to which it wants to connect; therefore the Wake Up message MUST NOT be used.

A CMS/other application server or a SIP Proxy MUST first obtain a TGT (Ticket Granting Ticket) before it can obtain an application server ticket for another application server or SIP Proxy. In order to obtain a TGT, an application server or a SIP Proxy MUST authenticate itself to the KDC with its symmetric service key (see clause 6.4.4). Once an application server or a SIP Proxy has obtained a TGT, it uses the TGS Request/TGS Reply exchange to obtain the application server ticket (see clause 6.4.5).

7.4.1.4.2.1 Inter-domain call setup

In the case of an inter-domain calls, some of the signalling hops between the two CMSs may contain two hosts belonging to two different Kerberos realms.

In the cross-realm case, Security Associations in the form of KDC-KDC tickets are used to establish trust between the two KDCs in two different realms. These tickets are obtained with the PKCROSS protocol. After the KDC-KDC trust is established, an CMS or SIP Proxy in one realm is able to get a service ticket for an CMS or SIP Proxy in another realm (see clause 6.4.7 for details). After the service ticket is obtained, the Kerberized IPSec key management flows remain the same for the cross-realm case.

7.4.1.4.2.2 Example of inter-domain call setup with key-management flows

A CMS or a SIP Proxy SHOULD create SAs ahead of time (before they are needed) whenever possible. One such example is illustrated in the following diagram.

The diagram illustrates the beginning of an inter-domain call setup. In this example, the signalling between the two CMSs is routed through two EBPs (Exterior Border Proxies). After one roundtrip 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.

The diagram assumes that there are no prior SAs between the two CMSs and that CMS "A" does not possess a service ticket for CMS "B". It shows the key management flows necessary to establish the necessary SAs for this call.

CMS "A" sends a TGS Request to its local KDC (KDC "A") to get a ticket for CMS "B". This TGS Request is sent about the same time as the Gate-Set message to the local AN. The key management flows continue in parallel with the subsequent DQoS flows (Gate-Set ACK) and also in parallel with some NCS signalling flows (MDCX from CMS "A" to MTA "A" and the 100 response from MTA "A"). After the 100 response no more parallelism is possible. The next signalling message that CMS "A" sends out is the PRACK to CMS "B" that has to wait for the SAs with CMS "B" to be set up.

Since CMS "B" is in a different realm (and not in KDC "A's" database), the TGS Request from CMS "A" causes KDC "A" to perform a DNS lookup to retrieve CMS "B's" realm name. After the DNS lookup is complete, KDC "A" will attempt to locate a ticket for KDC "B" in its local ticket cache. If it finds that ticket, it will immediately return to CMS "A" the cross-realm TGT needed to authenticate to KDC "B". If (in a rare circumstance) KDC "A" does not currently have a ticket for KDC "B", it will first have to perform DNS lookups to locate it and then perform a PKCROSS exchange with KDC "B" to obtain the KDC ticket.

Once CMS "A" finally receives a cross-realm TGT for KDC "B", it has to send another TGS Request for KDC "B" and then finally obtain the service ticket for CMS "B". In order for CMS "A" to contact KDC "B", it will first have to perform two more DNS queries (one to get the KDC "B's" FQDN and another to get its IP address).

After CMS "A" had obtained a ticket for CMS "B", it will initiate Kerberized IPSec key management with CMS "B" in order to set up SAs. The PRACK message from CMS "A" to CMS "B" will be secured with these newly set up SAs.

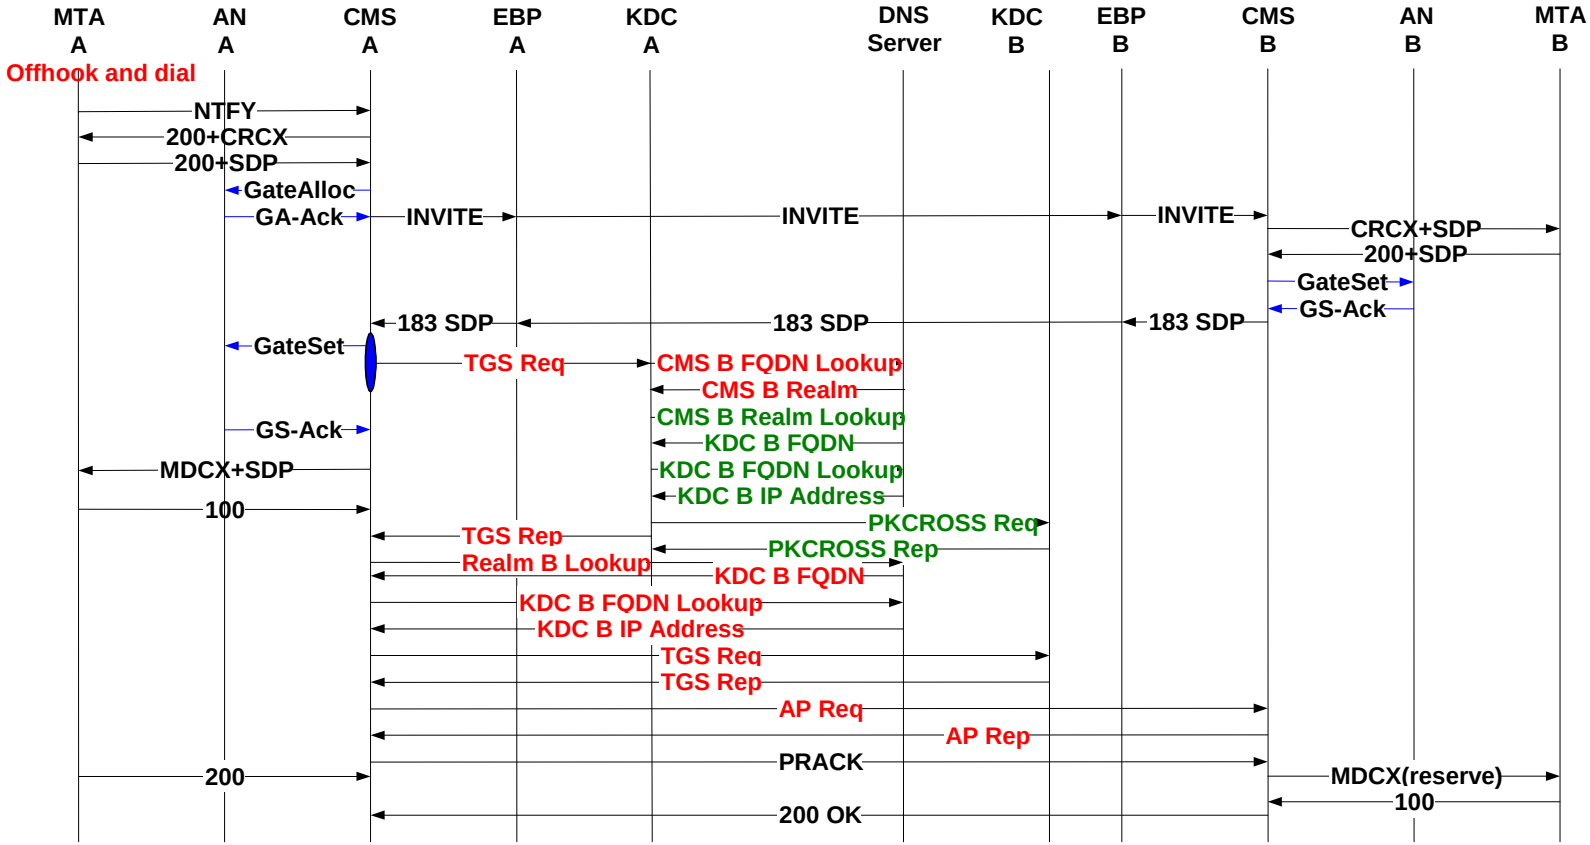

Figure 16: CMS - CMS Signalling Flow with Security

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 TS 101 909-13 [13].

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 TS 101 909-13 [13] 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 mechanism 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 dialed numbers.

Confidentiality: TGCP signalling messages carry dialed numbers and other customer information, which must not be disclosed to a third party. Thus confidentiality of the TGCP signalling messages is required.

7.5.2.2 MGC - SG Interface

Authentication: signalling messages must be authenticated, in order to prevent a third party masquerading as either an authorized MGC or SG.

Access Control: Services enable by the NCS signalling should be made available only to authorized users - thus access control is required at the MGC.

Integrity: must be assured in order to prevent tampering with the signalling messages - e.g. changing the dialed numbers.

Confidentiality: NCS messages carry dialed numbers and other customer information, which must not be disclosed to a third party. Thus confidentiality of signalling messages is required.

7.5.2.3 CMS - SG Interface

This interface is used for TCAP queries for LNP (Local Number Portability) and other voice communications services.

Authentication: TCAP queries must be authenticated, in order to prevent release of information to an unauthorized party.

Access Control: required along with the authentication, in order to prevent release of information to an unauthorized party.

Integrity: must be assured in order to prevent tampering with the TCAP queries, to prevent a class of denial-of-service attacks.

Confidentiality: TCAP queries contain dialed numbers and other subscriber information that MUST be kept private. Thus, confidentiality 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 clause 6.1 for details of how IPSec ESP is used within IPCablecom and for the list of available ciphersuites.

The ISTP protocol allows multiple redundant connections between the SG and the MGC. Multiple connections mean multiple security associations. The assumption is that the number of multiple connections is manageably small, where ahead of time we would set up a security association for each one, using IKE with pre-shared keys.

7.5.3.2 MGC - SG Interface

IPSec ESP MUST be used to both authenticate and encrypt the messages from MGC to SG and vice versa. Refer to clause 6.1 for details of how IPSec ESP is used within IPCablecom and for the list of available ciphersuites.

7.5.3.3 CMS - SG Interface

This interface is used for TCAP queries for LNP (Local Number Portability) and other voice communications services. IPSec ESP MUST be used to both authenticate and encrypt the messages from CMS to SG and vice versa. Refer to clause 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 MUST be implemented via IKE. IKE MUST use pre-shared key mode for this interface. Refer to clause 6.2 of the present document for details on the IPCablecom use of IKE.

IKE will guarantee that there is always a valid, non-expired MGC - MG Secret. This shared secret MUST be unique to this particular interface.

7.5.4.2 MGC - SG interface

Key management for MGC - SG interface MUST be implemented via IKE. IKE MUST use pre-shared key mode for this interface. Refer to clause 6.2of the present document for details on the IPCablecom use of IKE.

IKE will guarantee that there is always a valid, non-expired MGC - SG Secret. This shared secret MUST be unique to this particular interface.

7.5.4.3 CMS - SG interface

Key management for CMS - SG interface MUST be implemented via IKE. IKE MUST use pre-shared key mode for this interface. Refer to clause 6.2 of the present document for details on the IPCablecom use of IKE.

IKE will guarantee that there is always a valid, non-expired CMS - SG Secret. This shared secret MUST be unique to this particular interface.

7.5.5 MGC-MG-CMS-SG summary security profile matrix

Table 25: Security profile matrix - TCAP/IP & TGCP

	TCAP-IP, ISUP-IP (MGC-SG)	TGCP (MG - MGC)	TCAP-IP (CMS-SG)
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 IKE with pre-shared keys	IPSec IKE with pre-shared keys	IPSec IKE with pre-shared keys

7.6 Media stream

This security document 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 must 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 or 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 clause 7.6.2.1.2.1).

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 clause 7.6.1.1.

Encryption: some RTCP messages (e.g. CNAME) contain the identity of the endpoint. The requirement is to keep all RTCP message private, thus using block cipher to encrypt RTCP messages is required.

Message Integrity: RTCP signalling messages (e.g. BYE) can be manipulated to cause denial-of-service attacks. To prevent these attacks, message integrity is required for RTCP.

7.6.2 Cryptographic mechanisms

Each media RTP packet MUST be encrypted for privacy. The MTAs have an ability to negotiate a particular encryption algorithm. Encryption MUST be applied to the packet's payload. Encryption MUST NOT be applied to its header.

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 MTA as described in clause 7.6.2.2.

7.6.2.1 RTP packet format

Figure 18 shows the format of an encoded RTP packet. IPCablecom MUST adhere to the RTP packet format as defined by IETF RFC 1889 [18] and IETF RFC 1890.

The packet's header consists of 12 or more octets, as described in IETF RFC 1889. The only field of the header that is relevant to the encoding process is the timestamp field.

The RTP header has the following format (IETF RFC 1889 [18]):

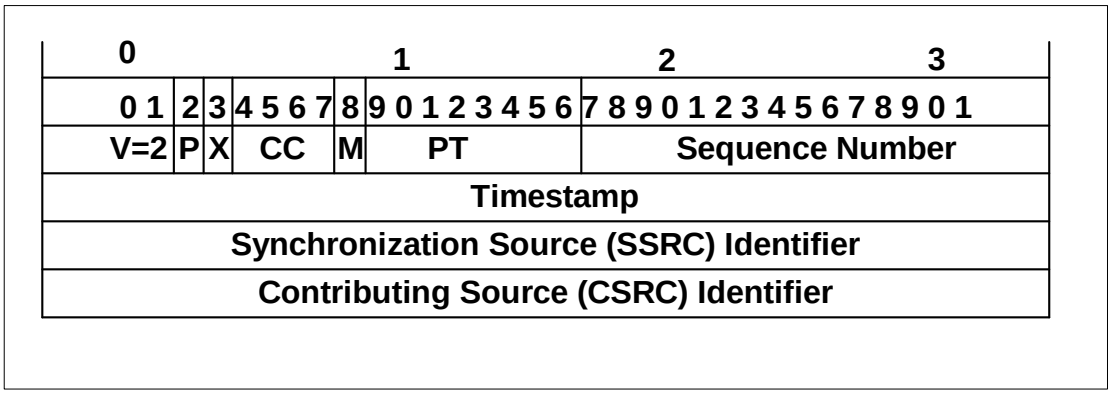

Figure 17: RTP Packet Header Format

The first twelve octets are present in every RTP packet, while the list of CSRC identifiers is present only when inserted by a mixer.

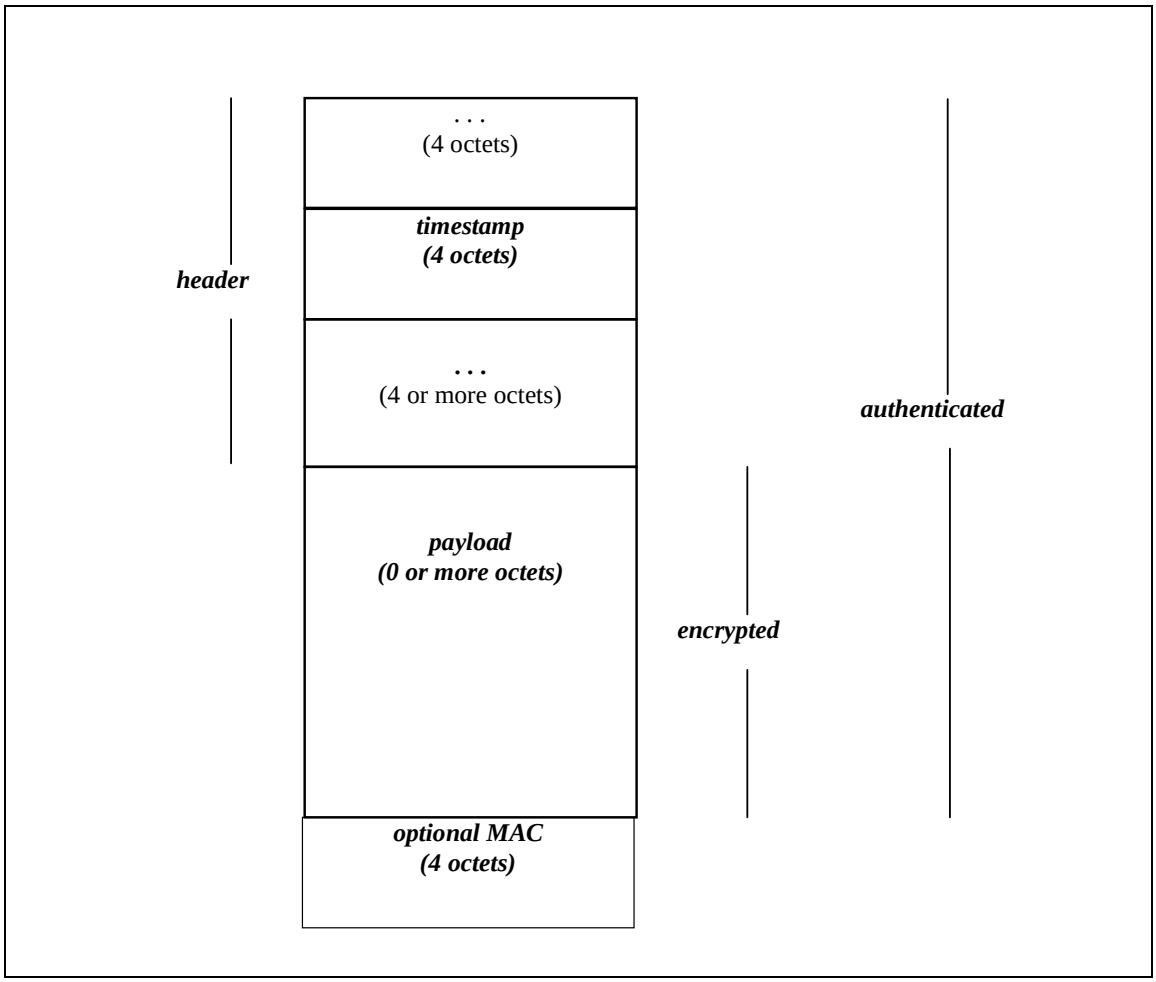

Figure 18: Format of Encoded RTP Packet

In IPCablecom, 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 the present document. 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 nonsilent 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 setup. For NCS signalling, it is a parameter in the LocalConnectionOptions - see TS 101 909-4.

For example, suppose the speech sample rate is 8,000 samples/sec, the frame rate is 10 msec, the packetization rate is 30 msec, and the compressed audio rate is 16,000 bits/sec. 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 in an event packet within the duration of one codec frame. The maximum size of the payload of an event packet is $M_f * N_e$.

NOTE: IPCablecom will use $N_e = N_c$, i.e. an event packet can have a payload as large as that of a voice packet, but no longer. Future versions of IPCablecom may provide a means for N_e to be determined in other ways.

- 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.

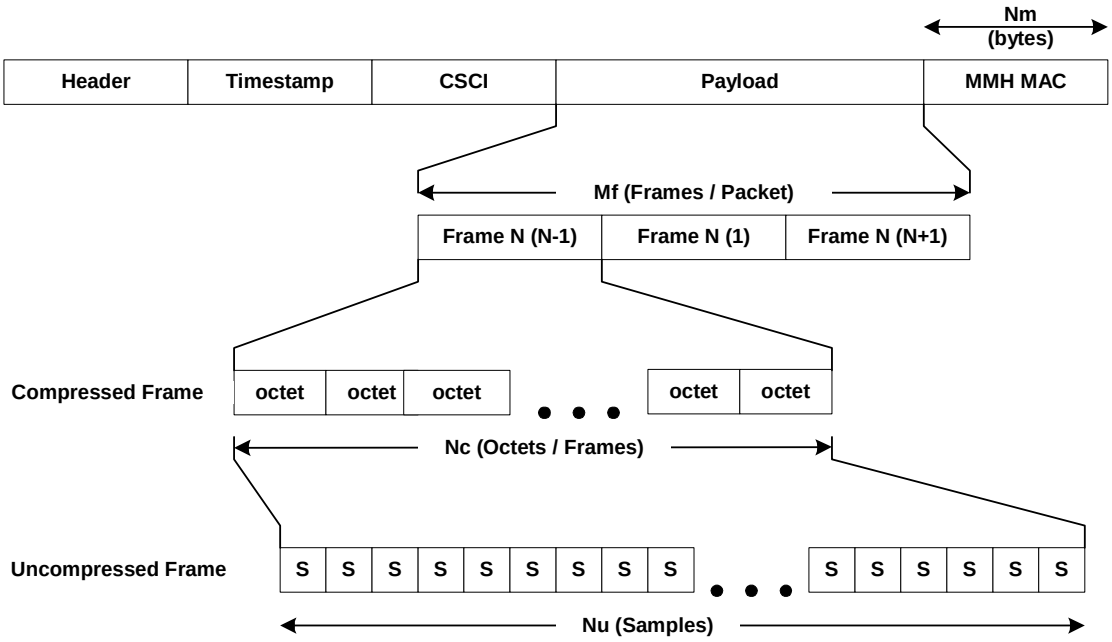

Figure 19: RTP Packet Profile Characteristics

7.6.2.1.1 RTP timestamp

According to IETF 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: $((N_f \times N_u) + (\text{RTP Initial Timestamp}))$ 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 clause 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 RC4 Encryption and MMH MAC option

7.6.2.1.2.1.1 Deriving an MMH MAC Key

Key size is ($N_h + N_e + N_m - 1$), where: N_h is the maximum number of octets in the RTP header; N_e is maximum number of octets in an event packet's payload, as defined in clause 7.6.2.1; and N_m is the number of octets in the MAC. 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 IETF 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.

7.6.2.1.2.1.2 Initializing the RC4 encryption process

The following additional parameter is defined for use with RC4:

- N_k, the state of the RC4 encryption process. "State" means the number of keystream octets that have been previously generated by the process, whether used or discarded. N_k has value 0 immediately after the RC4 process is initialized with a new key and increments with each generated octet of keystream. Prior to encoding the first packet, the following procedure MUST be used:
- The new RTP Privacy Key is used to initialize the RC4 encryption process.
- N_k is initialized to 0.
- N_f is initialized to 0.

7.6.2.1.2.1.3 Packet encoding

Each packet MUST be encoded using the following procedure:

- The timestamp is written into the timestamp field of the header. The timestamp MUST equate to the value: $((N_f * N_u)$ + (RTP Initial Timestamp)) modulo 2^{32}, where N_f is the frame number of the first frame included in the packet.
- All other fields of the header are set to values prescribed in IETF RFC 1889 [18].
- The RC4 encryption state N_k is set to the value $N_f*(N_e + N_m)$.
- The octets of the packet's payload are encrypted using the RC4 encryption process and inserted into the payload field. If there are B octets to be encrypted, then they are encrypted using octets $N_k + N_m$ to $N_k + N_m + B - 1$, inclusive and in order in the RC4 keystream.
- If the MAC option is enabled, the MAC digest is computed using the MMH algorithm (see annex G) with the RTP MAC Key. The digest calculation begins with the first octet of the unencrypted header and ends with the last octet of the encrypted payload. The computed digest is inserted into the MAC field. The digest calculation requires N_m octets of keystream from the RC4 process. These N_m octets are taken from the octets N_k to $N_k + N_m - 1$, inclusive and in order in the RC4 keystream.

Not all of the keystream octets generated by the RC4 process are necessarily used. If a packet contains m frames, then the RC4 state is advanced by $m*(N_e + N_m)$ prior to encoding the next packet. However, only $m*N_c + N_m$ keystream octets are actually used to encode the current packet; the remaining $(m - 1)*N_m + m*(N_e - N_c)$ keystream octets are unused. The RC4 encryption process is advanced for silent codec frames that are not actually transmitted, since the value of N_f increments even for silent frames. For each dropped silent frame, $(N_e + N_m)$ keystream octets are unused. Instead of dropping a silent frame, a codec might encode it using a short frame containing s octets, where $s < N_c$. For such a short frame, $(N_e - s)$ keystream octets are unused.

7.6.2.1.2.1.4 Codec change

During a codec change that is explicitly signalled by one of the endpoints via the NCS protocol, that endpoint MUST increment the N_{REKEY} counter and then re-derive a new set of RTP keys and a new RTP timestamp for both directions of traffic, according to clause 7.6.2.3.3.1. When that endpoint sends the SDP information (see IETF RFC 2327) describing the codec change to the other endpoint, it must also include the new value of the N_{REKEY} counter. (As usual, the codec change would be signalled with the NCS protocol, where all signalling messages have to first transit through the CMS.)

After the codec change the new keys MUST be used on all outbound traffic. For inbound traffic, each endpoint MAY save the old keys for some period of time - in order to decrypt RTP packets that are still using the old codec. Alternatively, endpoints MAY elect to drop incoming packets that are using the old codec.

7.6.2.1.2.1.5 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 $0 < \xi_1 <= \Delta T_{MAX}$, 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 $0 <= \xi_2 < \Delta T_{MAX}$.

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. When re-keying occurs during an existing media stream connection, N_{WRAP} MUST be re-initialized to 0. Note that this counter is unrelated to the count of key changes specified in clause 7.6.2.2.3.2.

The encryption and decryption procedures do not change after the wrap-around. However, the RTP timestamp value is no longer the RTP timestamp in the RTP header. After one or more wrap-arounds the RTP timestamp used in the encryption and decryption procedures MUST be calculated as:

$$\text{timestamp} + (N_{WRAP} * 2^{32})$$

This adjustment of the timestamp allows for a continuation of the same RC4 key stream after the wrap-around.

7.6.2.1.2.1.6 RTP SSRC collisions

According to IETF RFC 1889 [18], it is possible that an endpoint (e.g. Media Gateway) finds two different RTP sessions with different endpoints with the same RTP Synchronization Source Identifier (SSRC). IETF RFC 1889 [18] requires that when an SSRC collision is detected, one of the RTP sessions is shut down (via an RTCP BYE command) and is restarted with another SSRC value.

During this SSRC change, the sender MUST re-derive a new set of RTP keys and a new initial timestamp for the new SSRC value as specified in clause 7.6.2.2.3.1. The first RTP packet sent out with the new SSRC value MUST be encrypted and authenticated with the new set of keys.

As soon as the receiver becomes aware of the SSRC collision it MUST also re-derive a new set of RTP keys and a new initial timestamp as specified in clause 7.6.2.2.3.2. When it starts receiving RTP packets with the new SSRC value it MUST start using this new set of RTP keys and a new initial timestamp value for the decryption and validation of the inbound traffic. (In the case when the RTCP BYE packet got lost and the sender never generated packets with the new SSRC value, the receiver would continue decrypting the RTP packets with the original set of keys.)

Note that the SSRC change is uni-directional. Thus, the RTP key and timestamp re-derivation is applied to only one direction of traffic. The RTP keys and timestamp used for the other direction remain unchanged.

7.6.2.1.2.2 Block Cipher Encryption of RTP packets

7.6.2.1.2.2.1 Block termination

If an implementation supports block ciphers, the residual block termination (RBT) MUST be used to terminate streams that end with less than a full block of data to encrypt. (see clause 9.8.2.2).

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 Rijndael 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 clause 7.6.2.1.2.1.5. 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(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 0's 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 clause 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

A method for deriving the MMH-MAC pad when using RC4 was explained in an earlier clause. When using a block cipher, the pad is calculated using the block cipher. When the block cipher requires an IV, the IV value is calculated according to clause 7.2.1.2.2.2. This value will serve as the basis of the MMH-MAC pad when using a block cipher. If the MMH-MAC is used with a block cipher that does not require an IV, a corresponding value MUST be calculated according to clause 7.2.1.2.2.2 and used as the basis of the MMH-MAC pad according to this clause.

The pad is subsequently calculated by performing the MMH digest function on the resulting IV and then using the appropriate number of most-significant-bytes for the MMH-MAC pad.

Additional keying material needed to calculate the digest for the pad is derived with the MMH-MAC key from the shared secret as specified in clause 7.6.2.1.2.2.3.

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 clause 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) through (4) 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.

3) If the timestamp value in the RTP headers from a sender never comes back within the acceptable range, the receiver discontinues the session.

4) 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.

7.6.2.1.3.3 RC4 decryption and MMH MAC

Prior to decoding the first packet, the RTP Privacy Key is used to initialize the RC4 decryption process state N_k to zero.

Each packet MUST be decoded using the following procedure:

- The frame number for the first frame in the packet, N_f, is computed from the value of the timestamp field in the header as follows:

$$N_f = (\text{timestamp} - (\text{RTP Initial Timestamp})) / N_u$$

- Note that when the timestamp wraps-around its value is adjusted by adding 2^{32} (see clause 7.6.2.1.2.1.5). Thus timestamp is always >= RTP Initial Timestamp.

- The frame number for the first frame in the packet, N_f, is computed from the value of the timestamp field in the header as follows:

  - if the value of the timestamp is greater than or equal to the value of RTP Initial Timestamp, then $N_f = (\text{timestamp} - (\text{RTP Initial Timestamp})) / N_u$;

  - otherwise, $N_f = (\text{timestamp} + 2^{32} - (\text{RTP Initial Timestamp})) / N_u$.

- If the computed value of N_f is not an integer value, the packet is discarded; this indicates an invalid timestamp.

- The RC4 decryption state N_k is set to the value $N_f*(N_e + N_m)$.

- If the MAC option is enabled, a MAC digest is computed using the MMH algorithm with the RTP MAC key. The digest calculation begins with the first octet of the unencrypted header and ends with the last octet of the encrypted payload. The digest calculation requires N_m octets of keystream from the RC4 process. These N_m octets are taken from the octets N_k to $N_k + N_m - 1$, inclusive and in order in the RC4 keystream. The computed digest is compared to the value in the MAC field. If the computed digest does not match the value in the MAC field, the packet is discarded.

- The octets of the packet's payload are decrypted using the RC4 decryption process. If there are B octets to be decrypted, then they are decrypted using octets $N_k + N_m$ to $N_k + N_m + B - 1$, inclusive and in order in the RC4 keystream.

- All other fields of the header are processed as prescribed in IETF RFC 1889.

Note that the state of the RC4 decryption process is adjusted to match the state of the sender's RC4 encryption process prior to decrypting the packet's payload or verifying its MAC digest. If packets arrive out of order, the receive must, in principle, push the RC4 process backwards, as well as forwards, in order to match the state of the sender's RC4 process. In practice, this can be accomplished by having the receiver run its RC4 process in the forward direction only and synchronized to real time, thus making keystream available to decode packets in whatever order they arrive.

7.6.2.2 RTCP messages

7.6.2.2.1 RTCP format

IETF RFC 1889 [18] defines the packet format of RTCP messages.

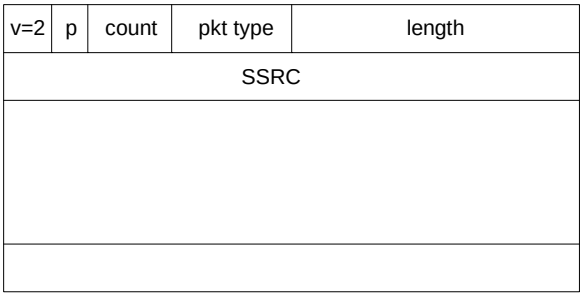

Figure 20: RTCP Packet Format

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 are always MUST be encrypted in their entirety by a block cipher in CBC mode. After the message is encrypted, an additional header and MAC (Message Authentication Code) are added. The result packet has the format in the following diagram.

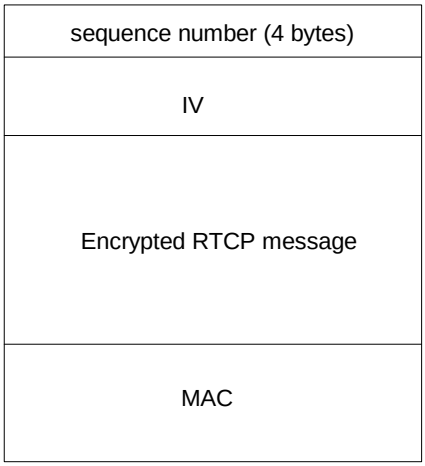

Figure 21: RTCP Encrypted Packet Format

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 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, 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 WRTCP depends on the reliability of the UDP transport and is locally configured at each endpoint. WRTCP 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 WRTCP - 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 (SRIGHT - WRTCP + 1, SRIGHT), sequence number SNEW is received and SNEW > SRIGHT, then the new window becomes:

(SNEW - WRTCP + 1, SNEW)

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 clause 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, 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 c) 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 to 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 clause 7.6.3.2.3. If the sequence number is rejected, the message is dropped. The sliding window is also updated as specified in clause 7.6.3.2.3.

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 end point, which can be an MTA or a MG. In the case that the end point is a MG, it is controlled by an MGC instead of a CMS.

During call setup MTA_0 (the initiating MTA) and MTA_1 (the terminating MTA) exchange randomly generated keying material, carried inside the call signalling messages. Call signalling messages are themselves protected by IPSec ESP 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 $Secret_0$ (46-bytes) and Pad_1 (46-bytes)

MTA_1 generates two randomly generated values: End-End $Secret_1$ (46-bytes) and Pad_0 (46-bytes).

MTA_0 uses End-End $Secret_1$ and Pad_1 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 $Secret_0$ and Pad_0 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 TS 101 909-4 and is described in the following clauses.

7.6.2.3.1 Key-management over NCS

The diagram below 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 labeled 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 labeled as SIP+ messages. However, NCS has not yet defined the CMS-CMS protocol and SIP+ is only one possible choice.

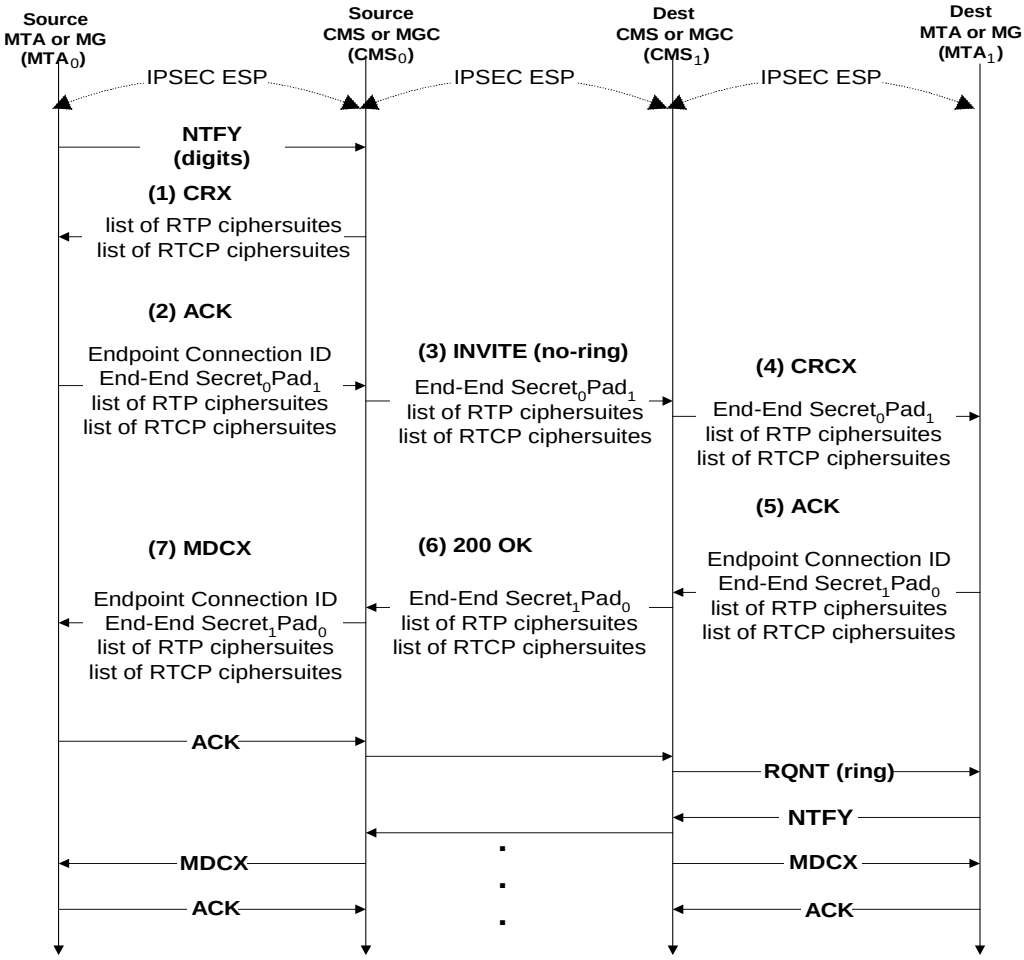

Figure 22: End-End Secret Distribution over NCS

Figure 22 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.

Each of the numbered flows in the above diagram is described below:

(1) CMS_0 -> MTA_0

CMS_0 MUST send the allowable lists of ciphersuites for the new communication to MTA_0 in the CreateConnection (CRCX) command, inside the LocalConnectionOptions parameter. There are two lists of ciphersuites, one list for RTP security and one for RTCP.

(2) MTA_0 -> CMS_0

From the list of ciphersuites in the LocalConnectionOptions, MTA_0 MUST select a (non-empty) subset for RTP and a (non-empty) subset for RTCP and return them in the subsequent ACK message, in the LocalConnectionDescriptor, in the form of SDP attributes. The resulting lists of ciphersuites are in a prioritized order - the preferred ciphersuites MUST be listed first.

MTA_0 MUST generate the End-End $Secret_0$ and Pad_1 value and return them along with the ciphersuites in the LocalConnectionDescriptor. For further details on the NCS message syntax, refer to TS 101 909-4.

The ACK also includes the Connection ID and the Endpoint for MTA_0 as described in TS 101 909-4. The pair (Connection ID, Endpoint) uniquely identifies this connection, where the Endpoint is an NCS identifier for MTA_0.

Before sending the ACK, MTA_0 MUST:

1) Establish inbound RTP security based on its preferred (first) RTP ciphersuite and End-End Secret0, as described in clause 7.6 of the present document.

2) Establish inbound RTCP key, based on End-End Secret0 and the preferred (first) RTCP ciphersuite in the list.

3) Be ready to receive RTP and RTCP messages, which may arrive any time after this message is sent by the MTA.

If MTA_1 decides to use an alternate ciphersuite listed by MTA_0, MTA_0 will later have to update its RTP and RTCP keys. If MTA_1 decides to send MTA_0 packets before ciphersuite negotiation had completed, processing on those packets at MTA_0 will fail (since it assumed a different ciphersuite).

(3) CMS_0 -> CMS_1

CMS_0 MUST send End-End $Secret_0$, Pad_1 and the list of RTP and RTCP ciphersuites to CMS_1 (local to MTA_1) as selected by MTA_0. CMS_1 will later forward this information to MTA_1.

(4) CMS_1 -> MTA_1

CMS_1 MUST send to MTA_1 in the CreateConnection (CRCX) command, inside the LocalConnectionOptions parameter, the lists of available RTP and RTCP ciphersuites. These ciphersuites SHOULD be what CMS_1 policy allows.

The RemoteConnectionDescriptor MUST be included in this CRCX command. It MUST contain End-End $Secret_0$ and Pad_1 received from MTA_0 (via CMS_0). It MUST also contain the ciphersuites preferred by MTA_0.

(5) MTA_1 -> CMS_1

MTA_1 MUST perform an interclause of the RTP ciphersuites contained in the LocalConnectionOptions with the ciphersuites listed in the RemoteConnectionDescriptor. If the interclause contains at least one RTP ciphersuite supported by MTA_1, then MTA_1 MUST select a non-empty subset and return it in the subsequent ACK message, in the LocalConnectionDescriptor, in the form of SDP attributes.

MTA_1 MUST also perform an interclause of the RTCP ciphersuites contained in the LocalConnectionOptions with the ciphersuites listed in the RemoteConnectionDescriptor. If the interclause contains at least one RTCP ciphersuite supported by MTA_1, then MTA_1 MUST select a non-empty subset and return it in the subsequent ACK message, in the LocalConnectionDescriptor, in the form of SDP attributes.

Whenever possible, MTA_1 SHOULD select the first supported ciphersuite in the RemoteConnectionDescriptor for RTP and the first supported for RTCP. This allows MTA_1 to immediately start sending RTP and RTCP packets to MTA_0. MTA_1 MAY instead select alternate ciphersuites specified by MTA_0. In that case, MTA_1 SHOULD NOT try to send any packets to MTA_0 until MTA_0 had been informed of the selected ciphersuites, as specified by TS 101 909-4.

MTA_1 MUST send back an ACK message, which includes lists of the selected ciphersuites inside the LocalConnectionDescriptor, 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 MTA_1. Additional ciphersuites in each list are alternatives in a prioritized order. If at any time, MTA_0 wants to switch to one of the alternatives that were selected by MTA_1, it would have to go through a new key negotiation. The ACK must also include the Connection ID (generated by MTA_1) as specified in TS 101 909-4. Thus, both End-End $Secret_0$ and End-End $Secret_1$ are now associated with a pair (Endpoint, Connection ID).

MTA_1 MUST generate the End-End $Secret_1$ for the incoming RTCP packets, and return it along with the ciphersuites in the LocalConnectionDescriptor. MTA_1 SHOULD also generate Pad_0 and return it in the same LocalConnectionDescriptor.

Although the option of not generating Pad_0 is provided in order to better support early media flows from MTA_1, it results in MTA_1 using a send key that is completely dependent on a random value generated by MTA_0. In other words, privacy of the media stream generated by MTA_1 in this case depends on the strength of MTA_0's random number generator.

Before sending the ACK, MTA_1 MUST:

1) Establish inbound RTP security based on its selected RTP ciphersuite, End-End $Secret_1$ and Pad_1, as described in clause 7.6.2.3.3.1 of the present document.

2) Establish inbound RTCP keys, based on End-End $Secret_1$, Pad_1 and the selected RTCP ciphersuite in the list.

3) Establish outbound RTP security based on its selected RTP ciphersuite and End-End $Secret_0$, as described in clause 7.6.2.3.3.1 of the present document. If Pad_0 was generated by MTA_1, the outbound RTP security will also be based on Pad_0.

In a specific case, where both the encryption algorithm originally selected by MTA_0 and the encryption algorithm selected by MTA_1 is AES, but the MAC size has changed (due to a change in the MAC algorithm), increment the N_{REKEY} counter before generating the outbound RTP security parameters. In this case, as specified in clause 7.6.2.3.3.1, MTA_0 must return the new value of the N_{REKEY} counter inside the ACK message in the LocalConnectionDescriptor.

1) Establish outbound RTCP keys, based on End-End Secret0 and the selected RTCP ciphersuite in the list. If Pad0 was generated by MTA1, the outbound RTCP Key Association will also be based on Pad0.

2) Be ready to receive RTP and RTCP messages, which may arrive any time after this message is sent by the MTA.

Any time after sending this ACK to the CMS1, MTA_1 MAY begin sending RTP and RTCP packets to MTA_0. However, in the case that MTA_1 generated Pad_0 or selected a different ciphersuite from the one preferred by MTA_0, MTA_0 will not be able to decrypt packets from MTA_1, until MTA_0 has received MTA_1's SDP.

(6) CMS_1 -> CMS_0

CMS_1 MUST forward the acknowledgement, End-End $Secret_1$, Pad_0 and the selected ciphersuites sent from MTA_1 and forward to CMS_0.

(7) CMS_0 -> MTA_0

CMS_0 MAY send to MTA_0 in the ModifyConnection command, inside the LocalConnectionOptions parameter, the lists of available RTP and RTCP ciphersuites. These ciphersuites SHOULD be what CMS_0 policy allows. (The reason that CMS_0 is not required to send the lists of ciphersuites is because it already sent them to MTA_0 in a CreateConnection command. CMS_0 would send the ciphersuites again for consistency.)

In the event that CMS_0 received SDP from the remote MTA, RemoteConnectionDescriptor MUST be included in this ModifyConnection command. If present, it MUST contain End-End $Secret_1$ and the ciphersuites (and alternatives) selected by MTA_1. If LocalConnectionOptions also contain a list of ciphersuites, MTA_0 MUST verify that the ciphersuites selected by MTA_1 are in that list; if not, it MUST abort the connection. RemoteConnectionDescriptor may also contain Pad_0.

If the RemoteConnectionDescriptor is not sent in this MDCX command, MTA_0 will still be able to receive RTP and RTCP messages but will be unable to send anything to MTA_1.

After receiving this message, MTA_0 MUST:

1) If Pad_0 was received, remove its inbound RTP and RTCP keys and replace them with new ones, based on the keys that are generated from both End-End $Secret_0$ and Pad_0. Re-initialize both the RTP timestamp and the RTCP sequence numbers for the newkeys. The ciphersuites used for these inbound keys are taken from the RemoteConnectionDescriptor just received from CMS_0. (Note that this change to the inbound keys due to the presense of Pad_0 does not increment the N_{REKEY} counter - its value remains 0. The N_{REKEY} counter is incremented only when explicitly stated in the present document.)

2) If the RemoteConnectionDescriptor was received without Pad_0, check if the first RTP ciphersuite in the RemoteConnectionDescriptor differs from the one that MTA_0 selected in step (2). If they differ, perform the following steps:

   a) Remove the inbound RTP key.

   b) If the N_{REKEY} counter is present in the RemoteConnectionDescriptor, update the local copy of the N_{REKEY} counter as specified in clause 7.6.2.3.3.2. This may happen when the encryption algorithm originally selected by MTA_0 and the new encryption algorithm are both RC4, but the MAC size has changed.

   c) Generate new inbound RTP keys and RTP timestamp from the same End-End $Secret_0$ as the last time, as specified in clause 4.7.2.2.3.2.

3) If the RemoteConnectionDescriptor was received without Pad_0, check if the first RTCP ciphersuite in the RemoteConnectionDescriptor differs from the one that MTA0 selected in step (2). If they differ, perform the following steps:

   a) Remove the inbound RTCP key.

   b) Generate a new key based on the key generated from the same End-End $Secret_0$ as the last time, but for the new authentication and/or encryption algorithms.

4) If the RemoteConnectionDescriptor was received, establish outbound RTP and RTCP keys, based on End-End $Secret^1$ and Pad^1.

5) Be ready to send and receive RTP and RTCP messages with MTA^1.

For full syntax of the NCS messages, please refer to the NCS signalling document, the TS 101 909-4.

7.6.2.3.2 Ciphersuite format

Each ciphersuite for both bearer channel and signalling security (via IPSec) 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 clause 6.1 for IPSec, clause 6.6 for RTP security and 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 TS 101 909-4 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 clause 7.6.2.2.1.

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, "End-End RTP Security Association"). Here, S is concatenation of the following binary values, each in MSB-first order:

   a) 4-byte SSRC identifier for this RTP session;

   b) 2-byte N_{REKEY} counter;

   c) End-End Secret; and

   d) Pad (optional, if it was negotiated through signalling).

The N_{REKEY} counter MUST be initialized to 0 at the start of an RTP connection and is incremented by one of the endpoints during various events requiring RTP rekeying (e.g. codec changes), as specified in the present document. The same N_{REKEY} counter is used in the generation of keys for both directions of traffic.

Whenever the value of the N_{REKEY} counter is updated by one endpoint, that endpoint MUST relay its value to the other endpoint inside the SDP parameters. If the other endpoint does not receive the value of this counter, it MUST use its current value of the N_{REKEY} counter.

When an endpoint receives a new value of the N_{REKEY} counter inside the SDP, it MUST make sure that it is greater than or equal to its locally stored value of this counter. If that check fails, instead of accepting the value of the N_{REKEY} counter from remote SDP, the endpoint MUST increment its local value of the counter. The endpoint MUST also return its updated value of the N_{REKEY} counter inside SDP in the ACK message back to the CMS. In the absence of error conditions, the CMS MUST in turn forward this ACK message to the other endpoint (possibly via another CMS or SIP Proxy).

In the case that the new value of the N_{REKEY} counter inside the SDP is equal to the locally stored value of this counter, if the received signalling message does not require re-keying (e.g. it is a codec change where the new codec is the same as the current one), the endpoint MUST NOT update its N_{REKEY} counter and MUST NOT perform rekeying. If the received signalling message does require rekeying (e.g. it is a codec change where the new codec is different from the current one), the endpoint MUST increment the local value of the N_{REKEY} counter. The endpoint MUST also return its updated value of the N_{REKEY} counter inside SDP in the ACK message back to the CMS. In the absence of error conditions, the CMS MUST in turn forward this ACK message to the other endpoint (possibly via another CMS or SIP Proxy).

The string "End-End RTP Security Association" is taken without quotes and without a terminating null character. Function F (specified in clause 9.7) 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 clause 7.6.3.1.2.2.2. 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 clause 7.6.3.1.2.1.1.

2) RTCP security. Derive a set of the following keys in the specified order with the derivation function F(S, "End-End RTP Control Protocol Security Association"). Here, S is concatentation of the following binary values:

   a) End-End Secret.

   b) Pad (optional, if it was negotiated through signalling).

Function F (specified in clause 9.7) 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.

7.6.2.3.3.2 End-to-End Rekey Derivation

Any time that either the N_{REKEY} counter or the SSRC value in the RTP header changes, the key derivation procedure in clause 7.6.2.3.3.1 MUST be used to derive a new set of keys and other RTP parameters for the media stream connection. In the case of the change in the N_{REKEY} counter, both directions of traffic are affected, while an SSRC change affects only one direction (the one with that SSRC value). This rekeying does not affect the keys established for RTCP.

7.6.2.4 RTP-RTCP summary security profile matrix

Table 26: Security profile matrix - RTP & RTCP

	RTP (MTA - MTA, MTA - MG)	**RTCP (MTA - MTA, MTA - MG, MG - MG)**
authentication	yes (indirect) (see note)	yes (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 End-to-End Secret distributed over secured MTA-CMS links. Final keys derived from this secret. AES-128 encryption algorithm Optional 2-byte or 4-byte MAC based on MMH algorithm IPCablecom requires support for ciphersuite negotiation.	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. Keys are derived from the end-end secret using the same mechanism as used for RTP encryption.
NOTE: MTAs do not authenticate directly. Authentication refers to the authentication of identity.		

7.7 Audio server services

7.7.1 Reference architecture

Figure 23 shows the network components and the various interfaces to be discussed in this clause, see TS 101 909-19 [16].

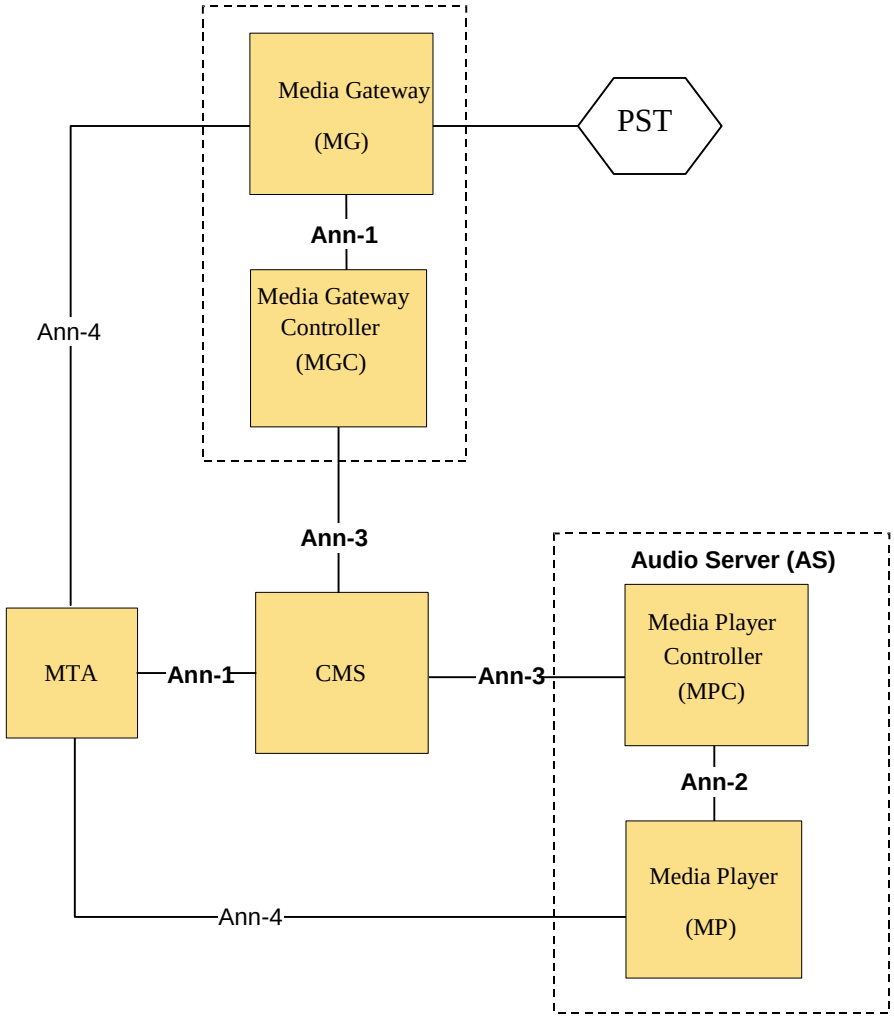

Figure 23: Audio Server Components and Interfaces

Figure 21 shows a network-based Media Player (MP). It has an optional TGCP 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 clause 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 clause 6.6 of the present document.

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 the NCS signalling clause 7.4.1.2 of the present document.

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 TGCP 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 playout 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 insures 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 clause 7.6.1.

7.7.3 Cryptographic mechanisms

7.7.3.1 MTA-CMS NCS signalling (Ann-1)

Refer to the cryptographic mechanisms in the NCS signalling clause 7.4.1.3 of the present document.

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 clause 6.1for details of how IPSec ESP is used within IPCablecom and for the list of available ciphersuites. The MPC MUST verify that the MP domain name (included in the Endpoint ID) correctly corresponds to its IP address for each signalling message received from the MP. This check is done via a lookup into a map of IP addresses to MP domain names. Also, in reverse, when the MPC performs a lookup of the MP IP address based on its domain name (in order to send a message to the MP), it consults this same map instead of performing a DNS query. Refer to clause 5.1.3.8 on how this map is maintained.

7.7.3.3 MTA-MP (Ann-4)

Cryptographic mechanisms on this media packet interface are specified in clause 7.6.2.

7.7.4 Key-management

7.7.4.1 MTA-CMS NCS Signalling (Ann-1)

Refer to the key management in the NCS signalling clause 7.4.1.4.1.

7.7.4.2 MPC-MP signalling (Ann-2)

MPC and MP MUST negotiate a shared secret (MPC-MP Secret) using IKE. IKE MUST use one of the modes with pre-shared keys for this interface. For details, refer to clause 6.1.2.3.

IKE 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.

At the MPC, MP domain names MUST be associated with the corresponding IP addresses. One possibility is to associate each pre-shared key directly with the domain name. 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 MUST be the MP domain name. For more details refer to IETF RFC 2409 [27].

Later, when a TGCP signalling message arrives at the MPC, it MUST be able to query the database of IPSec Security Associations and retrieve a source IP address, based on the MP domain name. The MPC MUST ensure that it is the same as the source IP address in the IP packet header. One way to query this database is through SNMP. Similarly, rather than doing a DNS query to find an IP address of the correct MP, MPC will look it up in the database of IPSec SAs.

7.7.4.3 MTA-MP (Ann-4)

Key Management on the media packet interface is specified in clause 7.6.2.2. 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 the present document. The corresponding column in the following matrix provides only the security requirements on that interface. Security specifications on that interface will be added in future revisions of the present document.

Table 27: Security profile matrix - Audio server services

	Ann-1: NCS (MTA - CMS) & (MG - MGC)	Ann-2: TGCP (MPC-MP)	Ann-3: unspecified (CMS-MPC) & (CMS - MGC) interface security requirements (see 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	IPSec IKE with pre-shared keys		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, the following are security requirements for the CMS-MPC interface.					

7.8 Third party interfaces

7.8.1 Reference architecture

The IPCablecom system for third party connections consists of the following elements and interfaces:

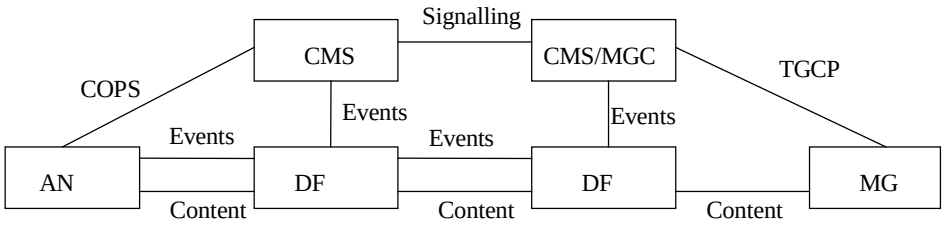

Figure 24: Security for third party Interfaces

The COPS interface between the CMS and the AN is used to signal the AN 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 clauses 7.2.1.2.3, 7.2.1.3.3 and 7.2.1.4.2.

The TGCP signalling interface between the CMS/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 setup on the PSTN Gateway side. For the corresponding security services, refer to clauses 7.8.2.1, 7.8.3.1 and 7.8.4.1.

The event interface between the AN 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 AN.

The interface between the AN and DF for call content is where the ANencapsulates 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.

The event interface between the two DFs is used to forward call information in the case where a call is forwarded to another location that is 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 call is forwarded to another location that is 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, AN-DF and DF-DF

Authentication, Access Control and Message Integrity: required to prevent service theft and denial-of-service attacks. Want to insure that the DF (law enforcement) has the right parameters to prevent denial-of-service. Also, want 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 AN-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 clause 7.3.3.1.

Also the same 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 160 bytes.

7.8.3.2 Interface between AN 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 AN-RKS interface specified in clause 7.3.1.2.

Also the same as with the AN-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 0-bytes.

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 clause 7.3.3.1.

Also the same 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 0-bytes.

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 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. The corresponding IPSec keys MUST be unique to this particular CMS and DF.

At the DF, CMS 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.

7.8.4.2 Interface between AN and DF

AN and DF MUST negotiate a pair of 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. The corresponding IPSec keys MUST be unique to this particular AN and DF.

At the DF, AN 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.

7.8.4.3 Interface between DF and DF

AN and DF MUST negotiate a shared secret (AN-DF Secret) using IKE with certificates. The IPCablecom profile for IKE with certificates is specified in clause 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 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 a 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.

8 IPCablecom certificates

IPCablecom uses digital certificates, which comply with the X.509 [14] and IETF RFC 2459 [29].

8.1 Generic structure

8.1.1 Version

The Version of the certificates MUST be V3. All certificates MUST comply with IETF RFC 2459 [29] except where the non-compliance with the RFC is explicitly stated in this clause of the present document.

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 extension MUST be used as specified in the clauses below. Any other certificate extensions MAY also be included as non-critical. The encoding tags are [c:critical, n:non-critical; m:mandatory, o:optional] and these are identified in the table for each certificate.

8.1.3.1 subjectKeyIdentifier

The subjectKeyIdentifier extension included in all IPCablecom CA certificates as required by IETF RFC 2459 [29] (e.g. all certificates except the device and ancillary certificates) MUST include the keyIdentifier value composed of the 160-bit SHA1 hash of the value of the BIT STRING subjectPublicKey (excluding the tag, length and number of unused bits from the ASN1 encoding) (see IETF RFC 2459 [29]).

8.1.3.2 authorityKeyIdentifier

The authorityKeyIdentifier extension included in all IPCablecom certificates as required by IETF RFC 2459 [29] MUST include the keyIdentifier value composed of the 160-bit SHA1 hash of the value of the BIT STRING subjectPublicKey (excluding the tag, length and number of unused bits from the ASN1 encoding) (see IETF RFC 2459 [29]).

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. The end-entity certificates may use the key usage extension as listed in IETF RFC 2459 [29].

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:

- Each RelativeDistinguishedName (RDN) MUST contain only a single element in the set of X.500 attributes.
- The order of the RDNs in an X.500 name MUST be the same as the order in which they are presented in the present document.

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.

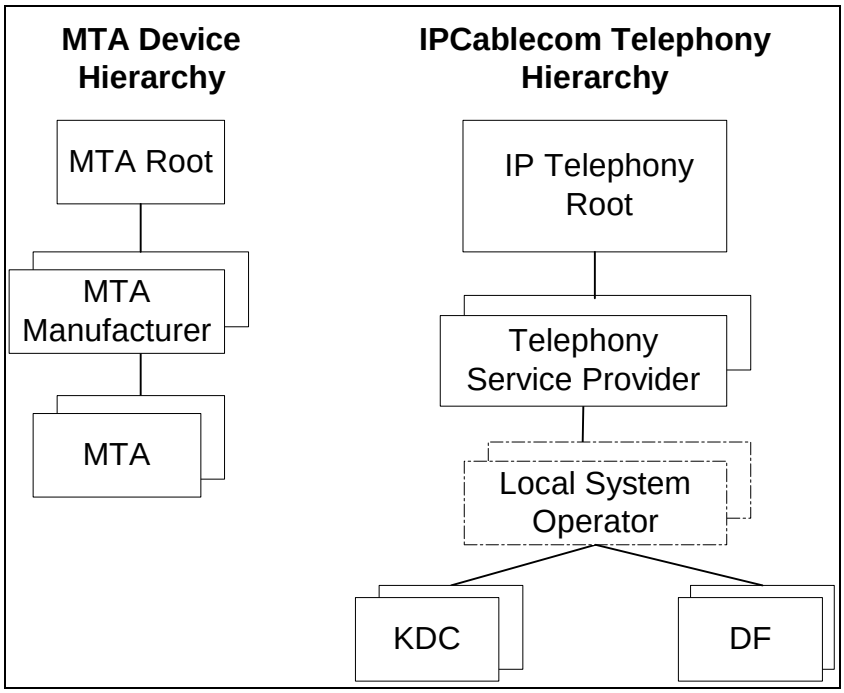

Figure 25: IPCablecom Certificate Hierarchy

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 then the certificate serial number, validity period and the value of the signature, exist in the MTA Root certificate that was passed over the wire in comparison to the known MTA Root certificate, the device making the comparison MUST fail the certificate verification. 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 IETF RFC 2459 [29], 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. IETF 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 the present document. 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 25.

The IPCablecom validation of validity periods for nesting is not checked and intentionally not enforced, which is compliant with current standards. At the time of issuance, the validity start date for any end-entity certificate MUST be the same as or later than the start date of the issuing CA certificate validity period. In some cases, however, IPCablecom CA certificates are renewed more often than end-entity certificates. After a CA certificate is renewed, the start dates of end-entity certificates MAY be earlier than the start date of the issuing CA certificate. The validity end date for entities may be before, the same as or after the validity end date for the issuing CA as specified in the IPCablecom Certificate tables.

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 the following table contains the IPCablecom specific values for the required fields according to IETF RFC 2459. These IPCablecom specific values MUST be followed according to the table below. If a required field is not specifically listed for IPCablecom then the guidelines in IETF RFC 2459 MUST be followed. The generic extensions for IPCablecom must also be included as specified in clause 8.1.3.

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 28: MTA Root Certificate

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 is long enough that this certificate is never re-issued.
Modulus Length	2 048
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 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 included in the MTA secure code download as specified by the IPCablecom Security document and 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 IETF RFC 2459.

Table 29: MTA Manufacturer Certificate

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 the IPCablecom Security document (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	2 048
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.

The state/province, city and manufacturer's facility are optional attributes.

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.

An MTA device certificate is permanently installed and not renewable or replaceable. Therefore, it must have a validity period greater than the operational lifetime of the MTA.

Table 30: MTA Device Certificate

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	1 024, 1 536 or 2 048
Extensions	keyUsage[n,o](digitalSignature, keyEncipherment), authorityKeyIdentifier The keyUsage tag is optional. When it is used it SHOULD be marked as critical.

8.2.2.4 MTA Manufacturer code verification certificate

Code Verification Certificate (CVC) specification for embedded MTAs MUST be identical to the J.112 CVC.

8.2.3 IPCablecom telephony certificate hierarchy

The Service Provider Certificate Hierarchy is rooted in a 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 the following table contains the IPCablecom specific values for the required fields according to IETF RFC 2459. These IPCablecom specific values MUST be followed according to table 31. If a required field is not specifically listed for IPCablecom then the guidelines in IETF RFC 2459 MUST be followed. The generic extensions for IPCablecom MUST also be included as specified in clause 8.1.3

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 31: IP Telephony root certificate

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 the IPCablecom Security document 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+. It is intended that the validity period is long enough that this certificate is never re-issued.
Modulus Length	2 048
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 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 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 Certificate changes, its signature MUST be verified with the IP Telephony Root Certificate. However, new certificate for the same Service Provider MUST preserve the same value of the OrganizationName attribute in the SubjectName.

Table 32: Telephony service provider Certificate

Telephony service provider certificate	
Subject Name Form	C=<country>, O=<Company>, OU=PacketCable, CN=<Company> PacketCable System Operator CA
Intended Usage	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. 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. 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.
Signed By	Signed by IP Telephony Root Certificate
Validity Period	20 years
Modulus Length	2 048
Extensions	keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier, authorityKeyIdentifier, basicConstraints[c,m](cA=true, pathLenConstraint=1)

8.2.3.3 Local system certificate

This is the certificate 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 Certificate, Service Provider Certificate, Local System Certificate and an end-entity certificate.

Table 33: Local system certificate

Local System Certificate	
Subject Name Form	C=<Country>, O=<Company>, OU=PacketCable, OU=<Local System Name>, CN=<Company> PacketCable Local System CA
Intended Usage	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). 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).
Signed By	Telephony Service Provider Certificate
Validity Period	20 years.
Modulus Length	1 024, 1 536, 2 048
Extensions	keyUsage[c,m](keyCertSign, cRLSign), subjectKeyIdentifier, authorityKeyIdentifier, basicConstraints[c,m](cA=true, pathLenConstraint=0)

8.2.4 Operational ancillary certificates

All of these are signed by the either the Local System CA or by the Telephony Service Provider CA. Other ancillary certificates may be added to the present document at a later time.

The information contained in the following table contains the IPCablecom specific values for the required fields according to IETF RFC 2459. These IPCablecom specific values MUST be followed according to table 34. If a required field is not specifically listed for IPCablecom then the guidelines in IETF RFC 2459 must be followed. The generic extensions for IPCablecom MUST also be included as specified in clause 8.1.3.

8.2.4.1 Key Distribution Center certificate

This certificate 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 annex C requires the KDC certificate to include the subjectAltName v.3 certificate extension, the value of which must be the Kerberos principal name of the KDC.

The encoding of the Kerberos PKINIT subjectAltName is:

```
SEQUENCE {              -- subjectAltName
    [0] SEQUENCE {      -- otherName

      OBJECT IDENTIFIER,   --  1.3.6.1.5.2.2, which is { krb5 2}

      [0] EXPLICIT SEQUENCE {          -- KerberosName

        [0] Realm,          -- Realm: a realm for which this Ticket Granting Service

                            -- issues tickets.

        [1] PrincipalName

     }

   }

 }
```

where PrincipalName is as defined by annex B:

```
      PrincipalName ::= SEQUENCE {

                  name-type[0]     INTEGER,

                  name-string[1]   SEQUENCE OF GeneralString

  -- namestring has two components:

  -- the first element is the string "krbtgt", without the quotation marks

  -- the second element is the name of the local realm - it has the same

  -- value as the Realm element of KerberosName

      }
```

Refer to annex C for a more detailed specification of the subjectAltName syntax.

Table 34: Key Distribution Center certificate

Key Distribution Center 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	1 024, 1 536 or 2 048
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)

8.2.4.2 Distribution Function (DF)

This certificate 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 betweenDFs. 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.

The extendedKeyUsage field (commonly called "EKU") MUST be present and MUST contain only the object identifier iKEIntermediate: iso.org.dod.internet.security.mechanisms.ipsec.certificate.2 (1.3.6.1.5.5.8.2.2).

Table 35: DF certificate

DF certificate	
Subject Name Form	C=<Country>, O=<Company>, OU=PacketCable, OU=[<Local System Name>], OU=Lawful Intercept, 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 forwards the call and event messages containing call info have to be forwarded to a new server (DF). For intra-domain Security Associations, these certificates are trusted only if the trust extends to the same Telephony Service Provider Certificate. The inter-domain case is not covered by this version of the document.
Signed By	Telephony Service Provider Certificate or Local System Certificate
Validity Period	20 years
Modulus Length	2 048
Extensions	KeyUsage[n,o](digitalSignature), authorityKeyIdentifier The keyUsage tag is optional. When it is used it SHOULD be marked as critical. subjectAltName[n,m](dNSName=<DNSName>) extendedKeyUsage[n,m](iKEIntermediate)

Note that this particular OID is not part of the IETF RFC 2459 certificate profile.

8.2.4.3 Operator Code Verification Certificate

Code Verification Certificate (CVC) specification for embedded MTAs MUST be identical to the J.112 CVC.

8.2.5 Certificate revocation

For future study.

9 Cryptographic algorithms

This clause describes the cryptographic algorithms used in the IPCablecom security document. When a particular algorithm is used, the algorithm MUST follow the corresponding specification.

9.1 AES

AES-Rijndael is a 128-bit block cipher that MUST be implemented according to the AES (Advanced Encryption Standard) proposed submission specified in AES-The Rijndael Block Cipher.

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 analyzing 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-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 clause 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-81. The CBC mode of DESX-XEX is shown a figure below, where DESX-XEX is executed within the block called "block cipher." Inside the block, DESX-XEX is performed as shown in a figure below using a 192-bit key. K1 is the first 8-bytes of the key, and K2 represents the second 8-bytes of key; and K3 the third 8-bytes of 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-81, and an additional 64-bits of keying material is 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 XOR-ed 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 clause 9.3), the left-most-bits of the IV are used in the final XOR padding operation. This mode of DES-CBC is shown a figure 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 a figure below, where 3DES-EDE is executed within the block called "block cipher." Inside the block, 3DES-EDE is performed as shown in a figure 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 a figure below. Residual block termination (RBT) is executed as follows:

Given a final block having n bits, where n is less than block size, the n bits are padded up to a block by appending (block size - n) bites 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 informative reference, 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 the figure below for both encryption and decryption operations.

BC Encryption Architecture

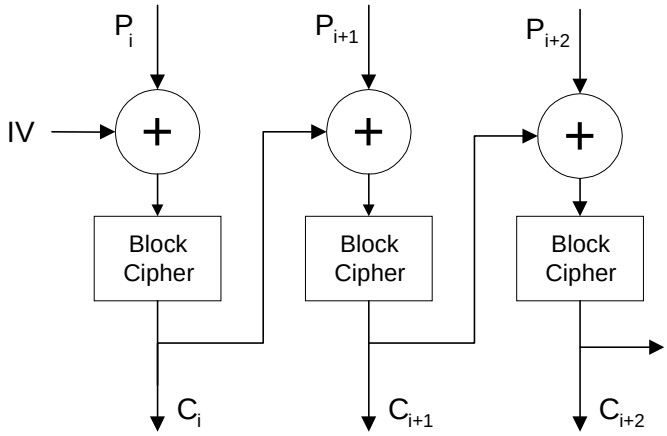

BC Decryption Architecture

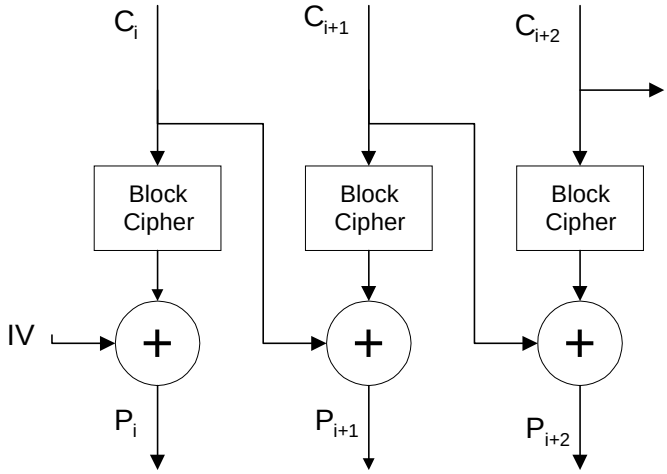

Figure 26: CBC Mode

BC-PAD Encryption Architecture

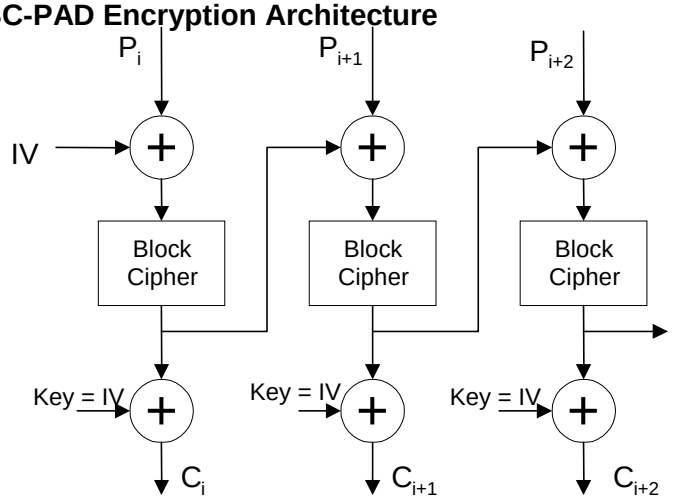

CBC-PAD Decryption Architecture

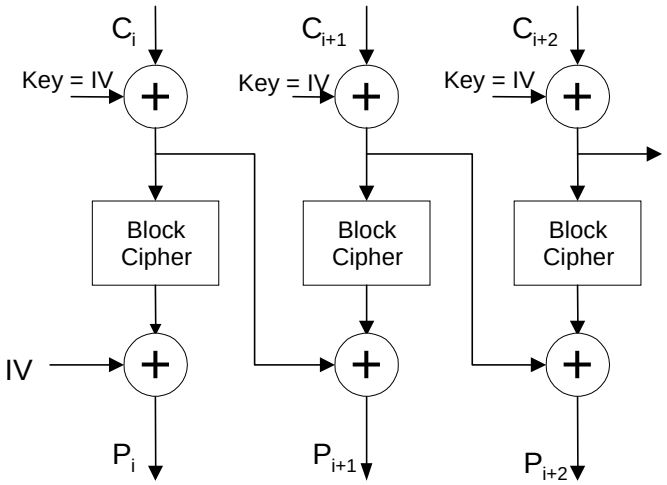

Figure 27: CBC Pad Mode

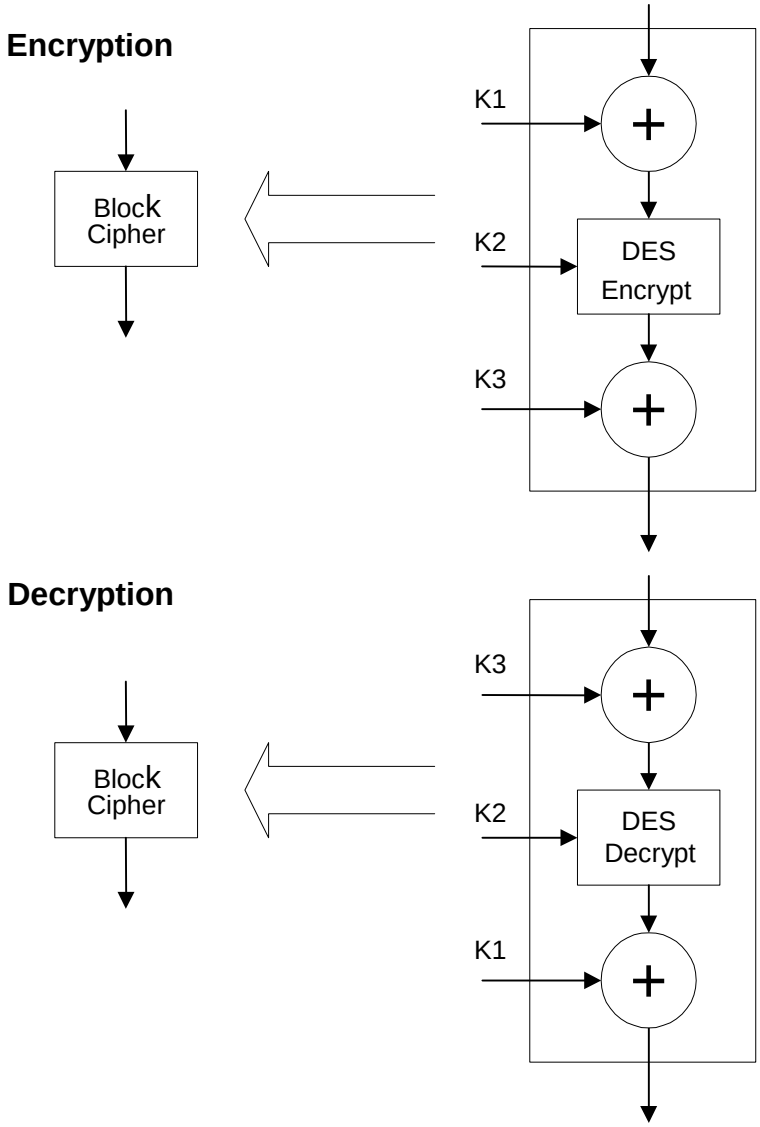

Figure 28: DESX-XEX as Block Cipher

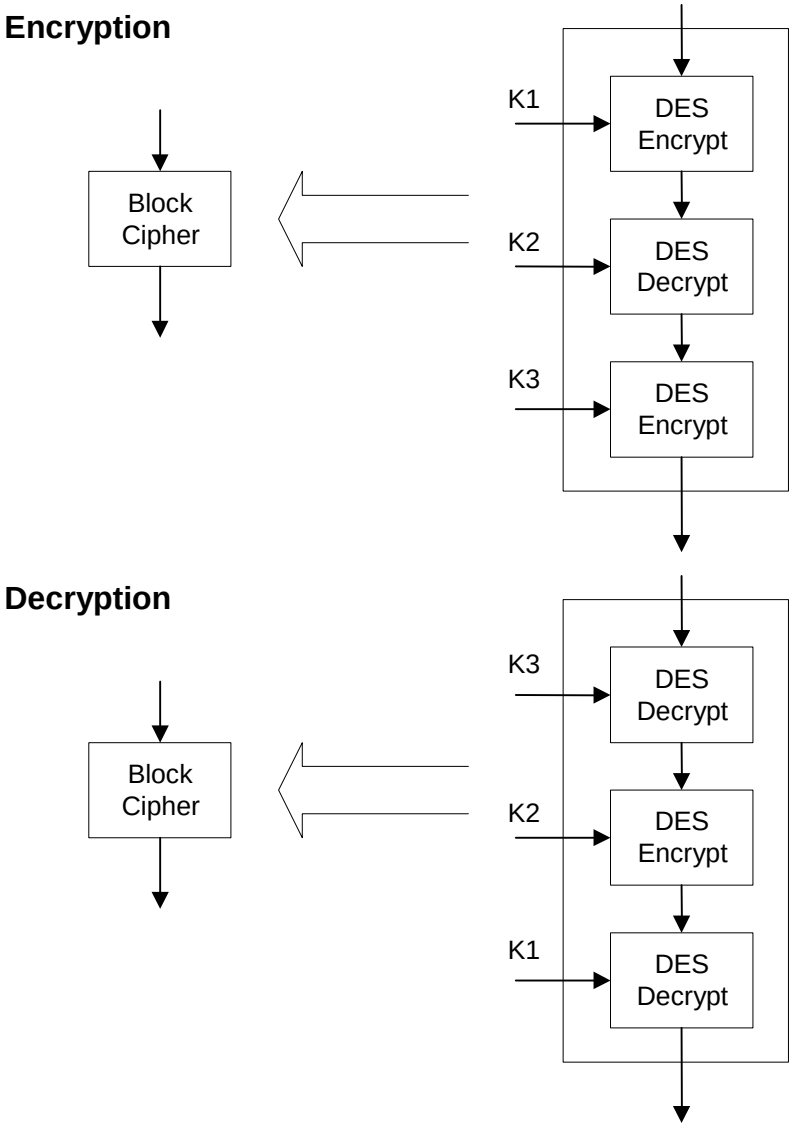

Figure 29: 3DES-EDE as Block Cipher

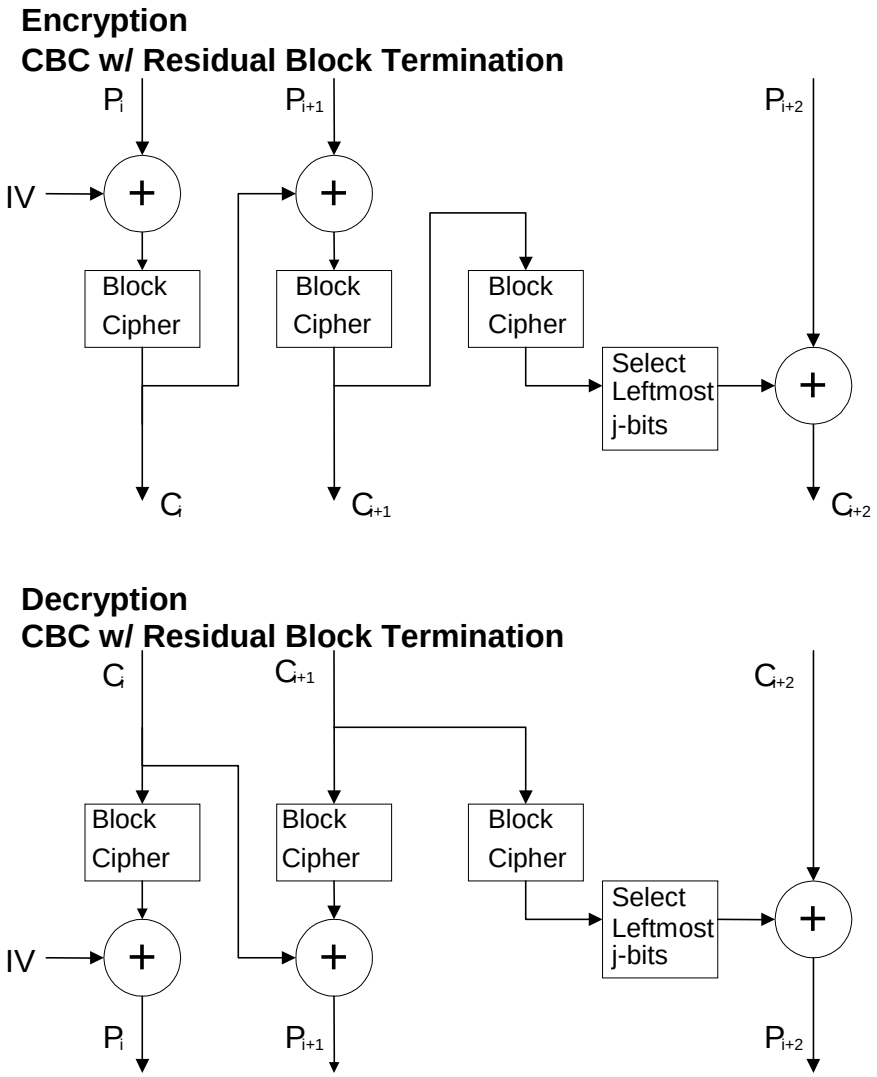

Figure 30: CBC with Residual Block Termination

9.4 RC4

RC4 [1] is a very efficient symmetric cipher. RC4 is used in clause 7.6 to encrypt media flows. Key management is described in clause 7.6.2.2. The algorithm uses variable length keys. For IPCablecom the key length MUST be set to 128 bits. The generation of this 128-bit key from the session key is described in clause 9.4. RC4 is a pseudo-random number generator in output feedback mode. A stream is generated from the key and XORed with the plaintext. There are no integrity protections on the data. RC4 uses a 256-entry substitution box (Sbox), which must be initialized. The entries in the Sbox are represented as bytes S_0, S_1, …, S_{255}. To initialize the Sbox, it is first filled with each entry matching its index. So,

$S_0 = 0$, $S_1 = 1$, … $S_{255} = 255$.

Another 256-byte array is filled with the key, repeating as necessary to fill the array, K_0, K_1, … , K_{255}. An index, j, is set to 0. Then, the Sbox is filled as follows:

for i = 0 to 255

$j = (j + S_i + K_i) \bmod 256$

swap S_i and S_j.

After the loop completes, the Sbox is initialized. The Sbox is dependent on the key, so a new one must be initialized for each key. The Sbox can now be used to generate a pseudo-random stream. i and j are initialized to 0, and a random byte is produced as follows:

$i = (i + 1) \bmod 256$

$j = (j + S_i) \bmod 256$

swap S_i and S_j

$t = (S_i + S_j) \bmod 256$

random_byte = S_t

To generate more bytes, the process is repeated using the values for i and j that result from the previous iteration.

9.5 RSA signature

All public key signatures for IPCablecom MUST be generated and verified using the RSA signature algorithm described in IETF RFC 2437 [33]. The format for all IPCablecom RSA signatures MUST be compliant with the Cryptographic Message Syntax of IETF RFC 2630 [31].

9.6 HMAC-SHA1

The keyed hash employed by the HMAC-Digest Attribute MUST use the HMAC message authentication method per IETF RFC 2104 with the SHA-1 hash algorithm per FIPS 180-1 [37].

HMAC-SHA1 is used to authenticate AN-CMS UDP messages for DQoS, as described in clause 7.2.1.2.2. The key is 160 bits long. The key management for the HMAC keys is described in clause 7.2.1.4.1.

9.7 Key derivation

Key derivation clauses in the present document refer to a function F(S, 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, seed), borrowed from TLS (IETF RFC 2246):

F(S, seed) = HMAC_SHA-1(S, A(1) + seed) +

HMAC_SHA-1(S, A(2) + seed) +

HMAC_SHA-1(S, A(3) + seed) + …

where + indicates concatenation.

A() is defined as: A(0) = seed

A(i) = HMAC_SHA-1(S, A(i-1))

F(S, 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.

9.8 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 clause 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.

9.8.1 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 Bbit/sec Rates.* Some of the computations described below use signed arithmetic whereas the computations in *MMH:Software Message Authentication in Bbit/sec 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 Bbit/sec 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. MMH[ω,*s*,*t*] specifies the hash function with word size ω, *s* input words and *t* output words. For IPCablecom the word size is fixed to 16 bits: ω=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.

9.8.1.1 MMH[16,s,1]

For the remainder of this clause 9.8, 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 clause 9.8.2. The key is denoted by k and the ith word of the key by ki: k=k1,k2,…,ks. Likewise the input message is denoted by m and the ith word of the input message by mi: m = m1, m2,…, ms.

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,…,0,…,(*n*/2)-1}. For example, $S_{2^{16}}$ = {-2^{15},…,0,…,2^{15}-1} is the set of signed 16 bit integers. For any integer *z*, *z* 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 twos complement representation, then *z* smod 2^{16} can be computed by taking the 16 least significant bits of *z* and interpreting those bits in 16 bit twos complement representation.

For any positive integer *q*, Z_q denotes the following set of *q* integers: {0, 1, …, *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 $H_1(k,m) = \sum_{i=1}^{s} k_i \cdot m_i \; s \bmod 2^{32}$.
2) Define H_2 as $H_2(k,m) = H_1(k,m) \bmod p$ 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,

$$H(k,m) = \left(\left(\left(\sum_{i=1}^{s} k_i \cdot m_i \right) s \bmod 2^{32} \right) \bmod p \right) \bmod 2^{16}$$

Each step is discussed in detail below.

Step1. $H_1(k,m)$ is the inner product of two vectors each of *s* 16 bit signed integers. The result of the inner product is taken smod 2^{32} to yield an element of $S_{2^{32}}$.

NOTE: The entire sum need not be computed before performing the smod 2^{32} operation. The smod 2^{32} operation can be computed on partial sums since (x + y) smod 2^{32} = (*x* smod 2^{32} + *y* smod 2^{32}) smod 2^{32}.

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 mod 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 1. If x is non-negative then $x = a2^{16}+b$ where $a \in \{0,\ldots,2^{15}-1\}$ and $b \in \{0,\ldots,2^{16}-1\}$. 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 $\{-2^{15}+1,\ldots,2^{16}-1\}$. 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 2. If x is negative then $x = a2^{16} + b - 2^{32}$ where $a \in \{2^{15},\ldots,2^{16}-1\}$ and $b \in \{0,\ldots,2^{16} - 1\}$. 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 $\{2^{15} + 1,\ldots,2^{17} - 1\}$. 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 mod 2^{16}. This is equivalent to taking the 16 least significant bits.

9.8.1.2 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' = MMH[16,s,2]. H' takes a key of s+1 words. Let $k = k_1,\ldots,k_{s+1}$. Furthermore, define $k^{(1)}$ to be the s words of k starting with k_1, i.e. $k^{(1)} = k_1,\ldots,k_s$. Define $k^{(2)}$ to be the s words of k, starting with k_2, i.e. $k^{(2)}=k_2,\ldots,k_{s+1}$. For any $k \in S_{2^{16}}^{s+1}$ and any $m \in S_{2^{16}}^{s}$ m, $H'(k,m)$ is computed by first computing $H(k^{(1)},m)$ and then H($k^{(2)},m$) and concatenating the results. This can be written as: $H'(k,m) = H(k^{(1)},m) \circ H(k^{(2)},m)$.

9.8.2 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 wordsize 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 = 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 ith message sent and received. For the ith message, $m^{(i)}$, the message authentication code is computed as:

$$M(k, r^{(i)}, m^{(i)}) = \mathrm{H}(k, m^{(i)}) + r^{(i)}$$

Here H = MMH[16,s,t], $r^{(i)}$ is in $Z_{2^{16}}$ and addition is mod 2^{16}.

9.8.2.1 MMH-MAC when using RC-4

When calculating the MMH-MAC for use with RC4, the sequence of key streams is generated by an RC4 key stream as described in clause 7.6.2. The 2(s + t - 1)-byte key for MMH-MAC[16,s,t] are randomly generated as described in clause 7.6.2 from a session key for the media flows that is generated by the key agreement protocol give in clause 7.6.2.2.

9.8.2.2 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 clause 7.6.2.1.2.2.3.

9.8.2.3 Odd payload sizes

If a message m is not of length s words, but rather of length $v < s$ words, then the input to M is a new message m' given by $m' = m_1,\ldots m_v, e_{v+1},\ldots e_s$ where $e_{v+1} = \ldots = e_s$ is the all zeroes word.

9.9 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. IETF 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

The IPCablecom security specification requires that an embedded MTA (MTA-E) and a standalone MTA (MTA-S) maintain persistent IPSec encryption and authentication keys and Kerberos session keys. An MTA MUST also 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 the IPCablecom security specification for an MTA is specified in terms of the security levels defined in the FIPS PUBS 140-1, Security Requirements for Cryptographic Modules, standard. An MTA-E or MTA-S SHOULD, at a minimum meet FIPS PUBS 140-1 Security Level 1 requirements.

The IPCablecom Security specification'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.

The IPCablecom Security specification requirements 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 MTA's, 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 masters 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 AN could limit any access based upon 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 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 documents 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 AN 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 that 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 the IPCablecom security specification.

10.2 MTA key Encapsulation

As stated in the previous clause, FIPS PUB 140-1 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.

Two of the (many) requirements of FIPS PUB 140-1 Security Level 3 recommends "data ports for critical security parameters be physically separated from other data ports" and, entry/exit of keys in encrypted form or direct entry/exit with split knowledge procedures". As also mentioned in the previous clause, the IPCablecom security specification is not requiring compliance with any of the FIPS PUB 140-1 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

The scope of IPCablecom includes only Embedded MTAs. Therefore IPCablecom MTAs MUST be embedded with the J.112 cable modem that MUST implement the enhancements specified in clause 6.8. IPCablecom Embedded MTAs MUST have their software upgraded according to the J.112 requirements. The cable modem will verify the code file using J.112 parameters that include the J.112 root key and the Code Verification Certificate (CVC). Requirements for J.112 software upgrade are specified in ITU-T Recommendation J.112 [1].

The future implementation of secure software upgrades for Standalone MTAs is expected to utilize a similar method for secure software upgrades. The details of these requirements are for future study.

Annex A (normative): Security events

Table A.1 shows the recommended security events.

Table A.1: Security events

Event Name	Default Severity for Event Raised	Default Display String	Comments	Reference
SEC-EV-1	Major	"Unable to retrieve TGT %Realm %IpAddress"	AS Request/AS Response timed out because TGT ticket was not retrieved during SEC-5/SEC-6	7.1.1.2.5
SEC-EV-2	Major	"Unable to retrieve CMS Ticket %Realm %IpAddress"	AS Request/AS Response or TGS Request/TGS Response timed out because CMS ticket was not retrieved during SEC-5/SEC-6 or SEC-7/SEC-8 respectively	7.1.1.2.5
SEC-EV-3	Major	"Unable to MTA/CMS Security Association %Realm %Fqdn %IpAddress"	AP Request/AP Response timed out because IP Sec Security Associated was not established during SEC-9/SEC-10	7.1.1.2.5
SEC-EV-4	Major	"Name expired %Client"	Kerberos Error 1 - Client's entry in database has expired	Annex F
SEC-EV-5	Major	"Service expired %Service"	Kerberos Error 2 - Server's entry in database has expired	Annex F
SEC-EV-6	Critical	"Bad protocol version number %Version"	Kerberos Error 3 - Requested protocol version number not supported	Annex F
SEC-EV-7	Major	"Client Unknown %Client"	Kerberos Error 6 - Client not found in Kerberos database	Annex F
SEC-EV-8	Critical	"Server Unknown %Server"	Kerberos Error 7 - Server not found in Kerberos database	Annex F
SEC-EV-9	Minor	"Principal Not Unique %PrincileName"	Kerberos Error 8 - Multiple principal entries in database	Annex F
SEC-EV-10	Major	"Client Null Key %Client"	Kerberos Error 9 - Kerberos Error 1 - The client has a null key	Annex F
SEC-EV-11	Critical	"Server Null Key %Server"	Kerberos Error 9 - Kerberos Error 1 - The server has a null key	Annex F
SEC-EV-12	Major	"Policy %Client %Reason"	Kerberos Error 12 - KDC policy rejects request	Annex F
SEC-EV-13	Major	"Bad Option %client %Option"	Kerberos Error 13 - KDC cannot accommodate requested option	Annex F
SEC-EV-14	Critical	"Encryption type not supported %Client %EncryptionType"	Kerberos Error 14 - KDC has no support for encryption type	Annex F
SEC-EV-15	Critical	"Checksum type not supported %Client %Type"	Kerberos Error 15 - KDC has no support for checksum type	Annex F
SEC-EV-16	Major	"PAD type not supported %Client PadType"	Kerberos Error 16 - KDC has no support for padata type	Annex F
SEC-EV-17	Major	"Client Credentials revoked %Client"	Kerberos Error 18 - Clients credentials have been revoked	Annex F
SEC-EV-18	Critical	"Server Credentials revoked %Server":	Kerberos Error 19 - Credentials for server have been revoked	Annex F
SEC-EV-19	Major	""TGT Revoked %Client"	Kerberos Error 20 - TGT has been revoked	Annex F
SEC-EV-20	Minor	"Client not yet valid %Client"	Kerberos Error 21 - Client not yet valid try again later	Annex F
SEC-EV-21	Minor	"Pre authorization info invalid %server %client %information"	Kerberos Error 24 - Pre authorization information invalid	Annex F

Event Name	Default Severity for Event Raised	Default Display String	Comments	Reference
SEC-EV-22	Minor	"Require additional pre authorization info %server %client"	Kerberos Error 25 - Additional pre authorization information required	Annex F
SEC-EV-23	Minor	"Integrity check failure %client"	Kerberos Error 31 - Integrity check on decrypted field failed	Annex F
SEC-EV-24	Minor	"Iticket Expired %client"	Kerberos Error 32 - Ticket expired	Annex F
SEC-EV-25	Major	"Ticket not yet valid %client %ticket"	Kerberos Error 33 - Ticket not yet valid	Annex F
SEC-EV-26	Minor	"Replay request %client"	Kerberos Error 34 - Request is a replay	Annex F
SEC-EV-27	Major	"Bad ticket destination %server %ticket"	Kerberos Error 35 - Ticket not for us	Annex F
SEC-EV-28	Minor	"Ticket/authenticator mismatch %server %client	Kerberos Error 36 - Ticket and authenticator do not match	Annex F
SEC-EV-29	Minor	"Clock skew error %client, %server"	Kerberos Error 37 - Clock skew too great	Annex F
SEC-EV-30	Critical	"Protocol Version Mismatch %realm %server"	Kerberos Error 39 - Protocol Version Mismatch	Annex F
SEC-EV-31	Minor	"Invalid message type %client %server %realm"	Kerberos Error 40 - Invalid message type	Annex F
SEC-EV-32	Minor	"Bad key version %client %version"	Kerberos Error 44 - Specified version of key is not available	Annex F
SEC-EV-33	Major	"No service key %server"	Kerberos Error 45 - Service key not available	Annex F
SEC-EV-34	Minor	"Bad check sum type %client %type"	Kerberos Error 50 - Inappropriate type of checksum in message	Annex F
SEC-EV-35	Minor	"Generic %error %client"	Kerberos Error 60 - Generic Error	Annex F
SEC-EV-36	Minor	"Field too long %client"	Kerberos Error 61 - Field is too long for this implementation	Annex F
SEC-EV-37	Major	"Client not trusted %client"	Kerberos Error 62 - Client not trusted	Annex F
SEC-EV-38	Minor	"Invalid Signature %Client"	Kerberos Error 64 - Invalid Signature	Annex F
SEC-EV-39	Minor	"Weak Key %Client"	Kerberos Error 65 - Key chosen is very weak	Annex F
SEC-EV-40	Critical	"Cannot parse manufacturer certificate %Client, %certindex"	Kerberos Error 70 - Not possible to parse the manufacturer certificate	Annex F
SEC-EV-41	Major	"Cannot parse device certificate %Client, %certindex"	Kerberos Error 70 - Not possible to parse the device certificate	Annex F
SEC-EV-42	Major	"Invalid manufacturer certificate %Client, %certindex"	Kerberos Error 71 - Invalid manufacturer certificate	Annex F
SEC-EV-43	Major	"Invalid device certificate %Client, %certindex"	Kerberos Error 71 - Invalid device certificate	Annex F
SEC-EV-44	Critical	"Revoked manufacturer certificate %Client, %certindex"	Kerberos Error 72 - Manufacturer certificate made void	Annex F
SEC-EV-45	Major	"Revoked device certificate %Client, %certindex"	Kerberos Error 72 - Device certificate made void	Annex F
SEC-EV-46	Critical	"Manufacturer revocation status unknown %Client, %certindex"	Kerberos Error 73 - Manufacturer revocation status unknown	Annex F
SEC-EV-47	Major	"Device revocation status unknown %Client, %certindex"	Kerberos Error 73 - Device revocation status unknown	Annex F
SEC-EV-48	Critical	"Revocation status not available %Client, %certindex"	Kerberos Error 74 - Revocation status not available	Annex F

Event Name	Default Severity for Event Raised	Default Display String	Comments	Reference
SEC-EV-49	Minor	"Client name mismatch %Client"	Kerberos Error 75 - Client name mismatch	Annex F

Annex B (normative): Kerberos network authentication service

The Kerberos Network Authentication Service specification is currently still an IETF draft. The present document complies only with the version of the draft that is listed in this clause. The IPCablecom security experts will continue to track progress of the Kerberos Network Authentication Service draft through the IETF and will advise the Study Group concerning how to best deal with the situation as the document becomes an RFC. The goal is to incorporate by reference and eliminate this annex.

The Kerberos Network Authentication Service

INTERNET-DRAFT

Clifford Neuman
John Kohl
Theodore Ts'o

November 24, 2000

Expires May 24, 2001

The Kerberos Network Authentication Service (V5)

draft-ietf-cat-kerberos-revisions-07.txt.

STATUS OF THIS MEMO

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC 2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at
http://www.ietf.org/shadow.html.

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.ietf.org (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

The distribution of this memo is unlimited. It is filed as draft-ietf-cat-kerberos-revisions-07.txt, and expires May 24, 2001. Please send comments to: ietf-krb-wg@anl.gov

ABSTRACT

This document provides an overview and specification of Version 5 of the Kerberos protocol, and updates RFC1510 to clarify aspects of the protocol and its intended use that require more detailed or clearer explanation than was provided in RFC1510. This document is intended to provide a detailed description of the protocol, suitable for implementation, together with descriptions of the appropriate use of protocol messages and fields within those messages.

This document is not intended to describe Kerberos to the end user, system administrator, or application developer. Higher level papers describing Version 5 of the Kerberos system [NT94] and documenting version 4 [SNS88], are available elsewhere.

OVERVIEW

This INTERNET-DRAFT describes the concepts and model upon which the Kerberos network authentication system is based. It also specifies Version 5 of the Kerberos protocol.

The motivations, goals, assumptions, and rationale behind most design decisions are treated cursorily; they are more fully described in a paper available in IEEE communications [NT94] and earlier in the Kerberos portion of the Athena Technical Plan [MNSS87]. The protocols have been a proposed standard and are being considered for advancement for draft standard through the IETF standard process. Comments are encouraged on the presentation, but only minor refinements to the protocol as implemented or extensions that fit within current protocol framework will be considered at this time.

Requests for addition to an electronic mailing list for discussion of Kerberos, kerberos@MIT.EDU, may be addressed to kerberos-request@MIT.EDU. This mailing list is gatewayed onto the Usenet as the group comp.protocols.kerberos. Requests for further information, including documents and code availability, may be sent to info-kerberos@MIT.EDU.

BACKGROUND

The Kerberos model is based in part on Needham and Schroeder's trusted third-party authentication protocol [NS78] and on modifications suggested by Denning and Sacco [DS81]. The original design and implementation of Kerberos Versions 1 through 4 was the work of two former Project Athena staff members, Steve Miller of Digital Equipment Corporation and Clifford Neuman (now at the Information Sciences Institute of the University of Southern California), along with Jerome Saltzer, Technical Director of Project Athena, and Jeffrey Schiller, MIT Campus Network Manager. Many other members of Project Athena have also contributed to the work on Kerberos.

Version 5 of the Kerberos protocol (described in this document) has evolved from Version 4 based on new requirements and desires for features not available in Version 4. The design of Version 5 of the Kerberos protocol was led by Clifford Neuman and John Kohl with much input from the community. The development of the MIT reference implementation was led at MIT by John Kohl and Theodore T'so, with help and contributed code from many others. Since RFC1510 was issued, extensions and revisions to the protocol have been proposed by many individuals. Some of these proposals are reflected in this document. Where such changes involved significant effort, the document cites the contribution of the proposer.

Reference implementations of both version 4 and version 5 of Kerberos are publicly available and commercial implementations have been developed and are widely used. Details on the differences between Kerberos Versions 4 and 5 can be found in [KNT92].

1. Introduction

Kerberos provides a means of verifying the identities of principals, (e.g. a workstation user or a network server) on an open (unprotected) network. This is accomplished without relying on assertions by the host operating system, without basing trust on host addresses, without requiring physical security of all the hosts on the network, and under the assumption that packets traveling along the network can be read, modified, and inserted at will[1.1]. Kerberos performs authentication under these conditions as a trusted third-party authentication service by using conventional (shared secret key [1.2]) cryptography. Kerberos extensions described in [PKINIT reference as RFC] provide for the use of public key cryptography during certain phases of the authentication protocol. These extensions allow authentication of users registered with public key certification authorities, and provide certain benefits of public key cryptography in situations where they are needed.

The basic Kerberos authentication process proceeds as follows: A client sends a request to the authentication server (AS) requesting 'credentials' for a given server. The AS responds with these credentials, encrypted in the client's key. The credentials consist of 1) a 'ticket' for the server and 2) a temporary encryption key (often called a "session key"). The client transmits the ticket (which contains the client's identity and a copy of the session key, all encrypted in the server's key) to the server. The session key (now shared by the client and server) is used to authenticate the client, and may optionally be used to authenticate the server. It may also be used to encrypt further communication between the two parties or to exchange a separate sub-session key to be used to encrypt further communication.

Implementation of the basic protocol consists of one or more authentication servers running on physically secure hosts. The authentication servers maintain a database of principals (i.e. users and servers) and their secret keys. Code libraries provide encryption and implement the Kerberos protocol. In order to add authentication to its transactions, a typical network application adds one or two calls to the Kerberos library directly or through the Generic Security Services Application Programming Interface, GSSAPI, described in separate document [ref to GSSAPI RFC]. These calls result in the transmission of the necessary messages to achieve authentication.

The Kerberos protocol consists of several sub-protocols (or exchanges). There are two basic methods by which a client can ask a Kerberos server for credentials. In the first approach, the client sends a cleartext request for a ticket for the desired server to the AS. The reply is sent encrypted in the client's secret key. Usually this request is for a ticket-granting ticket (TGT) which can later be used with the ticket-granting server (TGS). In the second method, the client sends a request to the TGS. The client uses the TGT to authenticate itself to the TGS in the same manner as if it were contacting any other application server that requires Kerberos authentication. The reply is encrypted in the session key from the TGT. Though the protocol specification describes the AS and the TGS as separate servers, they are implemented in practice as different protocol entry points within a single Kerberos server.

Once obtained, credentials may be used to verify the identity of the principals in a transaction, to ensure the integrity of messages exchanged between them, or to preserve privacy of the messages. The application is free to choose whatever protection may be necessary.

To verify the identities of the principals in a transaction, the client transmits the ticket to the application server. Since the ticket is sent "in the clear" (parts of it are encrypted, but this encryption doesn't thwart replay) and might be intercepted and reused by an attacker, additional information is sent to prove that the message originated with the principal to whom the ticket was issued. This information (called the authenticator) is encrypted in the session key, and includes a timestamp. The timestamp proves that the message was recently generated and is not a replay. Encrypting the authenticator in the session key proves that it was generated by a party possessing the session key. Since no one except the requesting principal and the server know the session key (it is never sent over the network in the clear) this guarantees the identity of the client.

The integrity of the messages exchanged between principals can also be guaranteed using the session key (passed in the ticket and contained in the credentials). This approach provides detection of both replay attacks and message stream modification attacks. It is accomplished by generating and transmitting a collision-proof checksum (elsewhere called a hash or digest function) of the client's message, keyed with the session key. Privacy and integrity of the messages exchanged between principals can be secured by encrypting the data to be passed using the session key contained in the ticket or the sub-session key found in the authenticator.

The authentication exchanges mentioned above require read-only access to the Kerberos database. Sometimes, however, the entries in the database must be modified, such as when adding new principals or changing a principal's key. This is done using a protocol between a client and a third Kerberos server, the Kerberos Administration Server (KADM). There is also a protocol for maintaining multiple copies of the Kerberos database. Neither of these protocols are described in this document.

1.1. Cross-realm operation

The Kerberos protocol is designed to operate across organizational boundaries. A client in one organization can be authenticated to a server in another. Each organization wishing to run a Kerberos server establishes its own 'realm'. The name of the realm in which a client is registered is part of the client's name, and can be used by the end-service to decide whether to honor a request.

By establishing 'inter-realm' keys, the administrators of two realms can allow a client authenticated in the local realm to prove its identity to servers in other realms[1.3]. The exchange of inter-realm keys (a separate key may be used for each direction) registers the ticket-granting service of each realm as a principal in the other realm. A client is then able to obtain a ticket-granting ticket for the remote realm's ticket-granting service from its local realm. When that ticket-granting ticket is used, the remote ticket-granting service uses the inter-realm key (which usually differs from its own normal TGS key) to decrypt the ticket-granting ticket, and is thus certain that it was issued by the client's own TGS. Tickets issued by the remote ticket-granting service will indicate to the end-service that the client was authenticated from another realm.

A realm is said to communicate with another realm if the two realms share an inter-realm key, or if the local realm shares an inter-realm key with an intermediate realm that communicates with the remote realm. An authentication path is the sequence of intermediate realms that are transited in communicating from one realm to another.

Realms are typically organized hierarchically. Each realm shares a key with its parent and a different key with each child. If an inter-realm key is not directly shared by two realms, the hierarchical organization allows an authentication path to be easily constructed. If a hierarchical organization is not used, it may be necessary to consult a database in order to construct an authentication path between realms.

Although realms are typically hierarchical, intermediate realms may be bypassed to achieve cross-realm authentication through alternate authentication paths (these might be established to make communication between two realms more efficient). It is important for the end-service to know which realms were transited when deciding how much faith to place in the authentication process. To facilitate this decision, a field in each ticket contains the names of the realms that were involved in authenticating the client.

The application server is ultimately responsible for accepting or rejecting authentication and should check the transited field. The application server may choose to rely on the KDC for the application server's realm to check the transited field. The application server's KDC will set the TRANSITED-POLICY-CHECKED flag in this case. The KDC's for intermediate realms may also check the transited field as they issue ticket-granting-tickets for other realms, but they are encouraged not to do so. A client may request that the KDC's not check the transited field by setting the DISABLE-TRANSITED-CHECK flag. KDC's are encouraged but not required to honor this flag.

1.2. Choosing a principal with which to communicate

The Kerberos protocol provides the means for verifying (subject to the assumptions in 1.4) that the entity with which one communicates is the same entity that was registered with the KDC using the claimed identity (principal name). It is still necessary to determine whether that identity corresponds to the entity with which one intends to communicate.

When appropriate data has been exchanged in advance, this determination may be performed syntactically by the application based on the application protocol specification, information provided by the user, and configuration files. For example, the server principal name (including realm) for a telnet server might be derived from the user specified host name (from the telnet command line), the "host/" prefix specified in the application protocol specification, and a mapping to a Kerberos realm derived syntactically from the domain part of the specified hostname and information from the local Kerberos realms database.

One can also rely on trusted third parties to make this determination, but only when the data obtained from the third party is suitably integrity protected wile resident on the third party server and when transmitted. Thus, for example, one should not rely on an unprotected domain name system record to map a host alias to the primary name of a server, accepting the primary name as the party one intends to contact since an attacker can modify the mapping and impersonate the party with which one intended to communicate.

If a Kerberos server supports name canonicalization, it may be relied upon as a third party to aid in this determination. When utilizing the name canonicalization function provided by the Kerberos server, a client, having already located the instance of a service it wishes to contact, makes a request to the KDC using the server's name information as specified by the user. The Kerberos server will attempt to locate a service principal in its database that corresponds to the requested name and return a ticket for the appropriate server principal to the client. If the KDC determines that the correct server principal is registered in another realm, the KDC will provide a referral to the Kerberos realm that is known to contain the requested service principal. The name canonicalization function supports identity mapping only, and it may not be used as a general name service to locate service instances. There is no guarantee that the returned server principal name (identity) will embed the name of the host on which the server resides.

1.3. Authorization

As an authentication service, Kerberos provides a means of verifying the identity of principals on a network. Authentication is usually useful primarily as a first step in the process of authorization, determining whether a client may use a service, which objects the client is allowed to access, and the type of access allowed for each. Kerberos does not, by itself, provide authorization. Possession of a client ticket for a service provides only for authentication of the client to that service, and in the absence of a separate authorization procedure, it should not be considered by an application as authorizing the use of that service.

Such separate authorization methods may be implemented as application specific access control functions and may utilize files on the application server, or on separately issued authorization credentials such as those based on proxies [Neu93], or on other authorization services. Separately authenticated authorization credentials may be embedded in a tickets authorization data when encapsulated by the kdc-issued authorization data element.

Applications should not accept the mere issuance of a service ticket by the Kerberos server (even by a modified Kerberos server) as granting authority to use the service, since such applications may become vulnerable to the bypass of this authorization check in an environment if they interoperate with other KDCs or where other options for application authentication (e.g. the PKTAPP proposal) are provided.

1.4. Environmental assumptions

Kerberos imposes a few assumptions on the environment in which it can properly function:

* 'Denial of service' attacks are not solved with Kerberos. There are places in the protocols where an intruder can prevent an application from participating in the proper authentication steps. Detection and solution of such attacks (some of which can appear to be not-uncommon 'normal' failure modes for the system) is usually best left to the human administrators and users.
* Principals must keep their secret keys secret. If an intruder somehow steals a principal's key, it will be able to masquerade as that principal or impersonate any server to the legitimate principal.

* 'Password guessing' attacks are not solved by Kerberos. If a user chooses a poor password, it is possible for an attacker to successfully mount an offline dictionary attack by repeatedly attempting to decrypt, with successive entries from a dictionary, messages obtained which are encrypted under a key derived from the user's password.
* Each host on the network must have a clock which is 'loosely synchronized' to the time of the other hosts; this synchronization is used to reduce the bookkeeping needs of application servers when they do replay detection. The degree of "looseness" can be configured on a per-server basis, but is typically on the order of 5 minutes. If the clocks are synchronized over the network, the clock synchronization protocol must itself be secured from network attackers.
* Principal identifiers are not recycled on a short-term basis. A typical mode of access control will use access control lists (ACLs) to grant permissions to particular principals. If a stale ACL entry remains for a deleted principal and the principal identifier is reused, the new principal will inherit rights specified in the stale ACL entry. By not re-using principal identifiers, the danger of inadvertent access is removed.

1.5. Glossary of terms

Below is a list of terms used throughout this document.

Authentication

Verifying the claimed identity of a principal.

Authentication header

A record containing a Ticket and an Authenticator to be presented to a server as part of the authentication process.

Authentication path

A sequence of intermediate realms transited in the authentication process when communicating from one realm to another.

Authenticator

A record containing information that can be shown to have been recently generated using the session key known only by the client and server.

Authorization

The process of determining whether a client may use a service, which objects the client is allowed to access, and the type of access allowed for each.

Capability

A token that grants the bearer permission to access an object or service. In Kerberos, this might be a ticket whose use is restricted by the contents of the authorization data field, but which lists no network addresses, together with the session key necessary to use the ticket.

Ciphertext

The output of an encryption function. Encryption transforms plaintext into ciphertext.

Client

A process that makes use of a network service on behalf of a user. Note that in some cases a Server may itself be a client of some other server (e.g. a print server may be a client of a file server).

Credentials

A ticket plus the secret session key necessary to successfully use that ticket in an authentication exchange.

KDC

Key Distribution Center, a network service that supplies tickets and temporary session keys; or an instance of that service or the host on which it runs. The KDC services both initial ticket and ticket-granting ticket requests. The initial ticket portion is sometimes referred to as the Authentication Server (or service). The ticket-granting ticket portion is sometimes referred to as the ticket-granting server (or service).

Kerberos

Aside from the 3-headed dog guarding Hades, the name given to Project Athena's authentication service, the protocol used by that service, or the code used to implement the authentication service.

Plaintext

The input to an encryption function or the output of a decryption function. Decryption transforms ciphertext into plaintext.

Principal

A named client or server entity that participates in a network communication, with one name that is considered canonical.

Principal identifier

The canonical name used to uniquely identify each different principal.

Seal

To encipher a record containing several fields in such a way that the fields cannot be individually replaced without either knowledge of the encryption key or leaving evidence of tampering.

Secret key

An encryption key shared by a principal and the KDC, distributed outside the bounds of the system, with a long lifetime. In the case of a human user's principal, the secret key may be derived from a password.

Server

A particular Principal which provides a resource to network clients. The server is sometimes referred to as the Application Server.

Service

A resource provided to network clients; often provided by more than one server (for example, remote file service).

Session key

A temporary encryption key used between two principals, with a lifetime limited to the duration of a single login "session".

Sub-session key

A temporary encryption key used between two principals, selected and exchanged by the principals using the session key, and with a lifetime limited to the duration of a single association.

Ticket

A record that helps a client authenticate itself to a server; it contains the client's identity, a session key, a timestamp, and other information, all sealed using the server's secret key. It only serves to authenticate a client when presented along with a fresh Authenticator.

2. Ticket flag uses and requests

Each Kerberos ticket contains a set of flags which are used to indicate attributes of that ticket. Most flags may be requested by a client when the ticket is obtained; some are automatically turned on and off by a Kerberos server as required. The following sections explain what the various flags mean, and gives examples of reasons to use such a flag.

2.1. Initial, pre-authenticated, and hardware authenticated tickets

The INITIAL flag indicates that a ticket was issued using the AS protocol and not issued based on a ticket-granting ticket. Application servers that want to require the demonstrated knowledge of a client's secret key (e.g. a password-changing program) can insist that this flag be set in any tickets they accept, and thus be assured that the client's key was recently presented to the application client.

The PRE-AUTHENT and HW-AUTHENT flags provide additional information about the initial authentication, regardless of whether the current ticket was issued directly (in which case INITIAL will also be set) or issued on the basis of a ticket-granting ticket (in which case the INITIAL flag is clear, but the PRE-AUTHENT and HW-AUTHENT flags are carried forward from the ticket-granting ticket).

2.2. Invalid tickets

The INVALID flag indicates that a ticket is invalid. Application servers must reject tickets which have this flag set. A postdated ticket will usually be issued in this form. Invalid tickets must be validated by the KDC before use, by presenting them to the KDC in a TGS request with the VALIDATE option specified. The KDC will only validate tickets after their starttime has passed. The validation is required so that postdated tickets which have been stolen before their starttime can be rendered permanently invalid (through a hot-list mechanism) (see section 3.3.3.1).

2.3. Renewable tickets

Applications may desire to hold tickets which can be valid for long periods of time. However, this can expose their credentials to potential theft for equally long periods, and those stolen credentials would be valid until the expiration time of the ticket(s). Simply using short-lived tickets and obtaining new ones periodically would require the client to have long-term access to its secret key, an even greater risk. Renewable tickets can be used to mitigate the consequences of theft. Renewable tickets have two "expiration times": the first is when the current instance of the ticket expires, and the second is the latest permissible value for an individual expiration time. An application client must periodically (i.e. before it expires) present a renewable ticket to the KDC, with the RENEW option set in the KDC request. The KDC will issue a new ticket with a new session key and a later expiration time. All other fields of the ticket are left unmodified by the renewal process. When the latest permissible expiration time arrives, the ticket expires permanently. At each renewal, the KDC may consult a hot-list to determine if the ticket had been reported stolen since its last renewal; it will refuse to renew such stolen tickets, and thus the usable lifetime of stolen tickets is reduced.

The RENEWABLE flag in a ticket is normally only interpreted by the ticket-granting service (discussed below in section 3.3). It can usually be ignored by application servers. However, some particularly careful application servers may wish to disallow renewable tickets.

If a renewable ticket is not renewed by its expiration time, the KDC will not renew the ticket. The RENEWABLE flag is reset by default, but a client may request it be set by setting the RENEWABLE option in the KRB_AS_REQ message. If it is set, then the renew-till field in the ticket contains the time after which the ticket may not be renewed.

2.4. Postdated tickets

Applications may occasionally need to obtain tickets for use much later, e.g. a batch submission system would need tickets to be valid at the time the batch job is serviced. However, it is dangerous to hold valid tickets in a batch queue, since they will be on-line longer and more prone to theft. Postdated tickets provide a way to obtain these tickets from the KDC at job submission time, but to leave them "dormant" until they are activated and validated by a further request of the KDC. If a ticket theft were reported in the interim, the KDC would refuse to validate the ticket, and the thief would be foiled.

The MAY-POSTDATE flag in a ticket is normally only interpreted by the ticket-granting service. It can be ignored by application servers. This flag must be set in a ticket-granting ticket in order to issue a postdated ticket based on the presented ticket. It is reset by default; it may be requested by a client by setting the ALLOW-POSTDATE option in the KRB_AS_REQ message. This flag does not allow a client to obtain a postdated ticket-granting ticket; postdated ticket-granting tickets can only by obtained by requesting the postdating in the KRB_AS_REQ message. The life (endtime-starttime) of a postdated ticket will be the remaining life of the ticket-granting ticket at the time of the request, unless the RENEWABLE option is also set, in which case it can be the full life (endtime-starttime) of the ticket-granting ticket. The KDC may limit how far in the future a ticket may be postdated.

The POSTDATED flag indicates that a ticket has been postdated. The application server can check the authtime field in the ticket to see when the original authentication occurred. Some services may choose to reject postdated tickets, or they may only accept them within a certain period after the original authentication. When the KDC issues a POSTDATED ticket, it will also be marked as INVALID, so that the application client must present the ticket to the KDC to be validated before use.

2.5. Proxiable and proxy tickets

At times it may be necessary for a principal to allow a service to perform an operation on its behalf. The service must be able to take on the identity of the client, but only for a particular purpose. A principal can allow a service to take on the principal's identity for a particular purpose by granting it a proxy.

The process of granting a proxy using the proxy and proxiable flags is used to provide credentials for use with specific services. Though conceptually also a proxy, user's wishing to delegate their identity for ANY purpose must use the ticket forwarding mechanism described in the next section to forward a ticket granting ticket.

The PROXIABLE flag in a ticket is normally only interpreted by the ticket-granting service. It can be ignored by application servers. When set, this flag tells the ticket-granting server that it is OK to issue a new ticket (but not a ticket-granting ticket) with a different network address based on this ticket. This flag is set if requested by the client on initial authentication. By default, the client will request that it be set when requesting a ticket granting ticket, and reset when requesting any other ticket.

This flag allows a client to pass a proxy to a server to perform a remote request on its behalf, e.g. a print service client can give the print server a proxy to access the client's files on a particular file server in order to satisfy a print request.

In order to complicate the use of stolen credentials, Kerberos tickets are usually valid from only those network addresses specifically included in the ticket[2.1]. When granting a proxy, the client must specify the new network address from which the proxy is to be used, or indicate that the proxy is to be issued for use from any address.

The PROXY flag is set in a ticket by the TGS when it issues a proxy ticket. Application servers may check this flag and at their option they may require additional authentication from the agent presenting the proxy in order to provide an audit trail.

2.6. Forwardable tickets

Authentication forwarding is an instance of a proxy where the service granted is complete use of the client's identity. An example where it might be used is when a user logs in to a remote system and wants authentication to work from that system as if the login were local.

The FORWARDABLE flag in a ticket is normally only interpreted by the ticket-granting service. It can be ignored by application servers. The FORWARDABLE flag has an interpretation similar to that of the PROXIABLE flag, except ticket-granting tickets may also be issued with different network addresses. This flag is reset by default, but users may request that it be set by setting the FORWARDABLE option in the AS request when they request their initial ticket-granting ticket.

This flag allows for authentication forwarding without requiring the user to enter a password again. If the flag is not set, then authentication forwarding is not permitted, but the same result can still be achieved if the user engages in the AS exchange specifying the requested network addresses and supplies a password.

The FORWARDED flag is set by the TGS when a client presents a ticket with the FORWARDABLE flag set and requests a forwarded ticket by specifying the FORWARDED KDC option and supplying a set of addresses for the new ticket. It is also set in all tickets issued based on tickets with the FORWARDED flag set. Application servers may choose to process FORWARDED tickets differently than non-FORWARDED tickets.

2.7 Transited Policy Checking

While the application server is ultimately responsible for accepting or rejecting authentication and should check the transited field, a KDC may apply a realm specific policy for validating the transited field and accepting credentials for cross-realm authentication. When the KDC applies such checks and accepts such cross-realm authentication it will set the TRANSITED-POLICY-CHECKED flag in the service tickets it issues based on the cross-realm TGT. A client may request that the KDC's not check the transited field by setting the DISABLE-TRANSITED-CHECK flag. KDC's are encouraged but not required to honor this flag.

2.8 Anonymous Tickets

When policy allows, a KDC may issue anonymous tickets for the purpose of enabling encrypted communication between a client and server without identifying the client to the server. Such anonymous tickets are issued with a generic principal name configured on the KDC (e.g. "anonymous@") and will have the ANONYMOUS flag set. A server accepting such a ticket may assume that subsequent requests using the same ticket and session key originate from the same user. Requests with the same username but different tickets are likely to originate from different users. Users request anonymous ticket by setting the REQUEST-ANONYMOUS option in an AS or TGS request.

2.9. Other KDC options

There are three additional options which may be set in a client's request of the KDC.

2.9.1 Name canonicalization [JBrezak]

The NAME-CANONICALIZATION option allows the KDC to replace the name of the client or server requested by the client with the canonical form of the principal's name, if known, or to refer the client to a KDC for the realm with which the requested principal is registered.

Where name cannonicalization is supported a client who can identify a principal but does not know the full principal name can request that the Kerberos server attempt to lookup the name in its database and use the canonical name of the requested principal or return a referral to a realm that has the requested principal in its namespace. Use of name canonicalization supports the case where a principal has multiple common names (names typed by a user[2.2]), all of which are known to the KDC, but only one Kerberos identity (the canonical name is the Kerberos principal name). Name canonicalization is intended solely to provide a secure mapping from the name known by a user to its principal identifier. It is not intended for use as a general purpose nameserver or to identify instances of a service.

The CANONICALIZE flag in a ticket request is used to indicate to the Kerberos server that the client will accept an alternative name to the principal in the request or a referral to another realm. When name cannonicalization is supported in a realm, all instances of the AS and TGS for the realm must be able to interpret requests with this flag. In realms where name cannonicalization is not supported, this flag may be ignored. By using this flag, the client can avoid extensive configuration needed to map specific host names to a particular realm.

2.9.2 Renewable-OK

The RENEWABLE-OK option indicates that the client will accept a renewable ticket if a ticket with the requested life cannot otherwise be provided. If a ticket with the requested life cannot be provided, then the KDC may issue a renewable ticket with a renew-till equal to the requested endtime. The value of the renew-till field may still be adjusted by site-determined limits or limits imposed by the individual principal or server.

2.9.3 ENC-TKT-IN-SKEY

The ENC-TKT-IN-SKEY option supports user-to-user authentication. It allows the KDC to issue a service ticket encrypted using the session key from a ticket granting ticket issued to another user. This is needed to support peer-to-peer authentication since the long term key of the user does not remain on the workstation after initial login. The ENC-TKT-IN-SKEY option is honored only by the ticket-granting service. It indicates that the ticket to be issued for the end server is to be encrypted in the session key from the additional second ticket-granting ticket provided with the request. See section 3.3.3 for specific details.

3. Message Exchanges

The following sections describe the interactions between network clients and servers and the messages involved in those exchanges.

3.1. The Authentication Service Exchange

Summary

Message direction	Message type	Section
1. Client to Kerberos	KRB_AS_REQ	5.4.1
2. Kerberos to client	KRB_AS_REP or KRB_ERROR	5.4.2 5.9.1

The Authentication Service (AS) Exchange between the client and the Kerberos Authentication Server is initiated by a client when it wishes to obtain authentication credentials for a given server but currently holds no credentials. In its basic form, the client's secret key is used for encryption and decryption. This exchange is typically used at the initiation of a login session to obtain credentials for a Ticket-Granting Server which will subsequently be used to obtain credentials for other servers (see section 3.3) without requiring further use of the client's secret key. This exchange is also used to request credentials for services which must not be mediated through the Ticket-Granting Service, but rather require a principal's secret key, such as the password-changing service[3.1]. This exchange does not by itself provide any assurance of the identity of the user[3.2].

The exchange consists of two messages: KRB_AS_REQ from the client to Kerberos, and KRB_AS_REP or KRB_ERROR in reply. The formats for these messages are described in sections 5.4.1, 5.4.2, and 5.9.1.

In the request, the client sends (in cleartext) its own identity and the identity of the server for which it is requesting credentials. The response, KRB_AS_REP, contains a ticket for the client to present to the server, and a session key that will be shared by the client and the server. The session key and additional information are encrypted in the client's secret key. The KRB_AS_REP message contains information which can be used to detect replays, and to associate it with the message to which it replies.

Without pre-authentication, the authentication server does not know whether the client is actually the principal named in the request. It simply sends a reply without knowing or caring whether they are the same. This is acceptable because nobody but the principal whose identity was given in the request will be able to use the reply. Its critical information is encrypted in that principal's key. The initial request supports an optional field that can be used to pass additional information that might be needed for the initial exchange. This field may be used for pre-authentication as described in section 3.1.1.

Various errors can occur; these are indicated by an error response (KRB_ERROR) instead of the KRB_AS_REP response. The error message is not encrypted. The KRB_ERROR message contains information which can be used to associate it with the message to which it replies. If suitable preauthentication has occurred, an optional checksum may be included in the KRB_ERROR message to prevent fabrication or modification of the KRB_ERROR message. When a checksum is not present, the lack of integrity protection precludes the ability to detect replays, fabrications, or modifications of the message, and the client must not depend on information in the KRB_ERROR message for security critical operations.

3.1.1. Generation of KRB_AS_REQ message

The client may specify a number of options in the initial request. Among these options are whether pre-authentication is to be performed; whether the requested ticket is to be renewable, proxiable, or forwardable; whether it should be postdated or allow postdating of derivative tickets; whether the client requests name-canonicalization or an anonymous ticket; and whether a renewable ticket will be accepted in lieu of a non-renewable ticket if the requested ticket expiration date cannot be satisfied by a non-renewable ticket (due to configuration constraints; see section 4). See section A.1 for pseudocode.

The client prepares the KRB_AS_REQ message and sends it to the KDC.

3.1.2. Receipt of KRB_AS_REQ message

If all goes well, processing the KRB_AS_REQ message will result in the creation of a ticket for the client to present to the server. The format for the ticket is described in section 5.3.1. The contents of the ticket are determined as follows.

3.1.3. Generation of KRB_AS_REP message

The authentication server looks up the client and server principals named in the KRB_AS_REQ in its database, extracting their respective keys. If the requested client principal named in the request is not known because it doesn't exist in the KDC's principal database and if an acceptable canonical name of the client is not known, then an error message with a KDC_ERR_C_PRINCIPAL_UNKNOWN is returned.

If the request had the CANONICALIZE option set and if the AS finds the canonical name for the client and it is in another realm, then an error message with a KDC_ERR_WRONG_REALM error code and the cname and crealm in the error message will contain the true client principal name and realm. In this case, since no key is shared with the client, the response from the KDC is not integrity protected and the referral can only be considered a hint; the validity of the referral is validated upon successful completion of initial authentication with the correct AS using the appropriate user key.

If required, the server pre-authenticates the request, and if the pre-authentication check fails, an error message with the code KDC_ERR_PREAUTH_FAILED is returned. If pre-authentication is required, but was not present in the request, an error message with the code KDC_ERR_PREAUTH_FAILED is returned and the PA-ETYPE-INFO pre-authentication field will be included in the KRB-ERROR message. If the server cannot accommodate an encryption type requested by the client, an error message with code KDC_ERR_ETYPE_NOSUPP is returned. Otherwise the KDC generates a 'random' session key[3.3].

When responding to an AS request, if there are multiple encryption keys registered for a client in the Kerberos database (or if the key registered supports multiple encryption types; e.g. DES3-CBC-SHA1 and DES3-CBC-SHA1-KD), then the etype field from the AS request is used by the KDC to select the encryption method to be used to protect the encrypted part of the KRB_AS_REP message which is sent to the client. If there is more than one supported strong encryption type in the etype list, the first valid etype for which an encryption key is available is used. The encryption method used to protect the encrypted part of the KRB_TGS_REP message is the keytype of the session key found in the ticket granting ticket presented in the KRB_TGS_REQ.

If the user's key was generated using an alternate string to key function than that used by the selected encryption type, information needed by the string to key function will be returned to the client in the padata field of the KRB_AS_REP message using the PA-PW-SALT, PA-AFS3-SALT, or similar pre-authentication typed values. This does not affect the encryption performed by the KDC since the key stored in the principal database already has the string to key transformation applied.

When the etype field is present in a KDC request, whether an AS or TGS request, the KDC will attempt to assign the type of the random session key from the list of methods in the etype field. The KDC will select the appropriate type using the list of methods provided together with information from the Kerberos database indicating acceptable encryption methods for the application server. The KDC will not issue tickets with a weak session key encryption type.

If the requested start time is absent, indicates a time in the past, or is within the window of acceptable clock skew for the KDC and the POSTDATE option has not been specified, then the start time of the ticket is set to the authentication server's current time. If it indicates a time in the future beyond the acceptable clock skew, but the POSTDATED option has not been specified then the error KDC_ERR_CANNOT_POSTDATE is returned. Otherwise the requested start time is checked against the policy of the local realm (the administrator might decide to prohibit certain types or ranges of postdated tickets), and if acceptable, the ticket's start time is set as requested and the INVALID flag is set in the new ticket. The postdated ticket must be validated before use by presenting it to the KDC after the start time has been reached.

The expiration time of the ticket will be set to the earlier of the requested endtime and a time determined by local policy, possibly determined using realm or principal specific factors. For example, the expiration time may be set to the minimum of the following:

* The expiration time (endtime) requested in the KRB_AS_REQ message.

* The ticket's start time plus the maximum allowable lifetime associated with the client principal from the authentication server's database (see section 4).
* The ticket's start time plus the maximum allowable lifetime associated with the server principal.
* The ticket's start time plus the maximum lifetime set by the policy of the local realm.

If the requested expiration time minus the start time (as determined above) is less than a site-determined minimum lifetime, an error message with code KDC_ERR_NEVER_VALID is returned. If the requested expiration time for the ticket exceeds what was determined as above, and if the 'RENEWABLE-OK' option was requested, then the 'RENEWABLE' flag is set in the new ticket, and the renew-till value is set as if the 'RENEWABLE' option were requested (the field and option names are described fully in section 5.4.1).

If the RENEWABLE option has been requested or if the RENEWABLE-OK option has been set and a renewable ticket is to be issued, then the renew-till field is set to the minimum of:

* Its requested value.
* The start time of the ticket plus the minimum of the two maximum renewable lifetimes associated with the principals' database entries.
* The start time of the ticket plus the maximum renewable lifetime set by the policy of the local realm.

The flags field of the new ticket will have the following options set if they have been requested and if the policy of the local realm allows: FORWARDABLE, MAY-POSTDATE, POSTDATED, PROXIABLE, RENEWABLE, ANONYMOUS. If the new ticket is post-dated (the start time is in the future), its INVALID flag will also be set.

If all of the above succeed, the server will encrypt ciphertext part of the ticket using the encryption key extracted from the server principal's record in the Kerberos database using the encryption type associated with the server principal's key (this choice is NOT affected by the etype field in the request). It then formats a KRB_AS_REP message (see section 5.4.2), copying the addresses in the request into the caddr of the response, placing any required pre-authentication data into the padata of the response, and encrypts the ciphertext part in the client's key using an acceptable encryption method requested in the etype field of the request, and sends the message to the client. See section A.2 for pseudocode.

3.1.4. Generation of KRB_ERROR message

Several errors can occur, and the Authentication Server responds by returning an error message, KRB_ERROR, to the client, with the error-code, e-text, and optional e-cksum fields set to appropriate values. The error message contents and details are described in Section 5.9.1.

3.1.5. Receipt of KRB_AS_REP message

If the reply message type is KRB_AS_REP, then the client verifies that the cname and crealm fields in the cleartext portion of the reply match what it requested. If any padata fields are present, they may be used to derive the proper secret key to decrypt the message. The client decrypts the encrypted part of the response using its secret key, verifies that the nonce in the encrypted part matches the nonce it supplied in its request (to detect replays). It also verifies that the sname and srealm in the response match those in the request (or are otherwise expected values), and that the host address field is also correct. It then stores the ticket, session key, start and expiration times, and other information for later use. The key-expiration field from the encrypted part of the response may be checked to notify the user of impending key expiration (the client program could then suggest remedial action, such as a password change). See section A.3 for pseudocode.

Proper decryption of the KRB_AS_REP message is not sufficient for the host to verify the identity of the user; the user and an attacker could cooperate to generate a KRB_AS_REP format message which decrypts properly but is not from the proper KDC. If the host wishes to verify the identity of the user, it must require the user to present application credentials which can be verified using a securely-stored secret key for the host. If those credentials can be verified, then the identity of the user can be assured.

3.1.6. Receipt of KRB_ERROR message

If the reply message type is KRB_ERROR, then the client interprets it as an error and performs whatever application-specific tasks are necessary to recover. If the client set the CANONICALIZE option and a KDC_ERR_WRONG_REALM error was returned, the AS request should be retried to the realm and client principal name specified in the error message crealm and cname field respectively.

3.2. The Client/Server Authentication Exchange

Summary

Message direction	Message type	Section
Client to Application server	KRB_AP_REQ	5.5.1
[optional] Application server to client	KRB_AP_REP or KRB_ERROR	5.5.2 5.9.1

The client/server authentication (CS) exchange is used by network applications to authenticate the client to the server and vice versa. The client must have already acquired credentials for the server using the AS or TGS exchange.

3.2.1. The KRB_AP_REQ message

The KRB_AP_REQ contains authentication information which should be part of the first message in an authenticated transaction. It contains a ticket, an authenticator, and some additional bookkeeping information (see section 5.5.1 for the exact format). The ticket by itself is insufficient to authenticate a client, since tickets are passed across the network in cleartext[3.4], so the authenticator is used to prevent invalid replay of tickets by proving to the server that the client knows the session key of the ticket and thus is entitled to use the ticket. The KRB_AP_REQ message is referred to elsewhere as the 'authentication header.'

3.2.2. Generation of a KRB_AP_REQ message

When a client wishes to initiate authentication to a server, it obtains (either through a credentials cache, the AS exchange, or the TGS exchange) a ticket and session key for the desired service. The client may re-use any tickets it holds until they expire. To use a ticket the client constructs a new Authenticator from the system time, its name, and optionally an application specific checksum, an initial sequence number to be used in KRB_SAFE or KRB_PRIV messages, and/or a session subkey to be used in negotiations for a session key unique to this particular session. Authenticators may not be re-used and will be rejected if replayed to a server[3.5]. If a sequence number is to be included, it should be randomly chosen so that even after many messages have been exchanged it is not likely to collide with other sequence numbers in use.

The client may indicate a requirement of mutual authentication or the use of a session-key based ticket by setting the appropriate flag(s) in the ap-options field of the message.

The Authenticator is encrypted in the session key and combined with the ticket to form the KRB_AP_REQ message which is then sent to the end server along with any additional application-specific information. See section A.9 for pseudocode.

3.2.3. Receipt of KRB_AP_REQ message

Authentication is based on the server's current time of day (clocks must be loosely synchronized), the authenticator, and the ticket. Several errors are possible. If an error occurs, the server is expected to reply to the client with a KRB_ERROR message. This message may be encapsulated in the application protocol if its 'raw' form is not acceptable to the protocol. The format of error messages is described in section 5.9.1.

The algorithm for verifying authentication information is as follows. If the message type is not KRB_AP_REQ, the server returns the KRB_AP_ERR_MSG_TYPE error. If the key version indicated by the Ticket in the KRB_AP_REQ is not one the server can use (e.g. it indicates an old key, and the server no longer possesses a copy of the old key), the KRB_AP_ERR_BADKEYVER error is returned. If the USE-SESSION-KEY flag is set in the ap-options field, it indicates to the server that the ticket is encrypted in the session key from the server's ticket-granting ticket rather than its secret key [3.6].

Since it is possible for the server to be registered in multiple realms, with different keys in each, the srealm field in the unencrypted portion of the ticket in the KRB_AP_REQ is used to specify which secret key the server should use to decrypt that ticket. The KRB_AP_ERR_NOKEY error code is returned if the server doesn't have the proper key to decipher the ticket.

The ticket is decrypted using the version of the server's key specified by the ticket. If the decryption routines detect a modification of the ticket (each encryption system must provide safeguards to detect modified ciphertext; see section 6), the KRB_AP_ERR_BAD_INTEGRITY error is returned (chances are good that different keys were used to encrypt and decrypt).

The authenticator is decrypted using the session key extracted from the decrypted ticket. If decryption shows it to have been modified, the KRB_AP_ERR_BAD_INTEGRITY error is returned. The name and realm of the client from the ticket are compared against the same fields in the authenticator. If they don't match, the KRB_AP_ERR_BADMATCH error is returned (they might not match, for example, if the wrong session key was used to encrypt the authenticator). The addresses in the ticket (if any) are then searched for an address matching the operating-system reported address of the client. If no match is found or the server insists on ticket addresses but none are present in the ticket, the KRB_AP_ERR_BADADDR error is returned. If the local (server) time and the client time in the authenticator differ by more than the allowable clock skew (e.g. 5 minutes), the KRB_AP_ERR_SKEW error is returned.

Unless the application server provides its own suitable means to protect against replay (for example, a challenge-response sequence initiated by the server after authentication, or use of a server-generated encryption subkey), the server must utilize a replay cache to remember any authenticator presented within the allowable clock skew. Careful analysis of the application protocol and implementation is recommended before eliminating this cache. The replay cache will store the server name, along with the client name, time and microsecond fields from the recently-seen authenticators and if a matching tuple is found, the KRB_AP_ERR_REPEAT error is returned [3.7]. If a server loses track of authenticators presented within the allowable clock skew, it must reject all requests until the clock skew interval has passed, providing assurance that any lost or re-played authenticators will fall outside the allowable clock skew and can no longer be successfully replayed[3.8].

If a sequence number is provided in the authenticator, the server saves it for later use in processing KRB_SAFE and/or KRB_PRIV messages. If a subkey is present, the server either saves it for later use or uses it to help generate its own choice for a subkey to be returned in a KRB_AP_REP message.

If multiple servers (for example, different services on one machine, or a single service implemented on multiple machines) share a service principal (a practice we do not recommend in general, but acknowledge will be used in some cases), they should also share this replay cache, or the application protocol should be designed so as to eliminate the need for it. Note that this applies to all of the services, if any of the application protocols does not have replay protection built in; an authenticator used with such a service could later be replayed to a different service with the same service principal but no replay protection, if the former doesn't record the authenticator information in the common replay cache.

The server computes the age of the ticket: local (server) time minus the start time inside the Ticket. If the start time is later than the current time by more than the allowable clock skew or if the INVALID flag is set in the ticket, the KRB_AP_ERR_TKT_NYV error is returned. Otherwise, if the current time is later than end time by more than the allowable clock skew, the KRB_AP_ERR_TKT_EXPIRED error is returned.

If all these checks succeed without an error, the server is assured that the client possesses the credentials of the principal named in the ticket and thus, the client has been authenticated to the server. See section A.10 for pseudocode.

Passing these checks provides only authentication of the named principal; it does not imply authorization to use the named service. Applications must make a separate authorization decisions based upon the authenticated name of the user, the requested operation, local access control information such as that contained in a .k5login or .k5users file, and possibly a separate distributed authorization service.

3.2.4. Generation of a KRB_AP_REP message

Typically, a client's request will include both the authentication information and its initial request in the same message, and the server need not explicitly reply to the KRB_AP_REQ. However, if mutual authentication (not only authenticating the client to the server, but also the server to the client) is being performed, the KRB_AP_REQ message will have MUTUAL-REQUIRED set in its ap-options field, and a KRB_AP_REP message is required in response. As with the error message, this message may be encapsulated in the application protocol if its "raw" form is not acceptable to the application's protocol. The timestamp and microsecond field used in the reply must be the client's timestamp and microsecond field (as provided in the authenticator)[3.9]. If a sequence number is to be included, it should be randomly chosen as described above for the authenticator. A subkey may be included if the server desires to negotiate a different subkey. The KRB_AP_REP message is encrypted in the session key extracted from the ticket. See section A.11 for pseudocode.

3.2.5. Receipt of KRB_AP_REP message

If a KRB_AP_REP message is returned, the client uses the session key from the credentials obtained for the server[3.10] to decrypt the message, and verifies that the timestamp and microsecond fields match those in the Authenticator it sent to the server. If they match, then the client is assured that the server is genuine. The sequence number and subkey (if present) are retained for later use. See section A.12 for pseudocode.

3.2.6. Using the encryption key

After the KRB_AP_REQ/KRB_AP_REP exchange has occurred, the client and server share an encryption key which can be used by the application. In some cases, the use of this session key will be implicit in the protocol; in others the method of use must be chosen from several alternatives. The 'true session key' to be used for KRB_PRIV, KRB_SAFE, or other application-specific uses may be chosen by the application based on the session key from the ticket and subkeys in the KRB_AP_REP message and the authenticator[3.11]. To mitigate the effect of failures in random number generation on the client it is strongly encouraged that any key derived by an application for subsequent use include the full key entropy derived from the KDC generated session key carried in the ticket. We leave the protocol negotiations of how to use the key (e.g. selecting an encryption or checksum type) to the application programmer; the Kerberos protocol does not constrain the implementation options, but an example of how this might be done follows.

One way that an application may choose to negotiate a key to be used for subsequent integrity and privacy protection is for the client to propose a key in the subkey field of the authenticator. The server can then choose a key using the proposed key from the client as input, returning the new subkey in the subkey field of the application reply. This key could then be used for subsequent communication.

To make this example more concrete, if the communication patterns of an application dictates the use of encryption modes of operation incompatible with the encryption system used for the authenticator, then a key compatible with the required encryption system may be generated by either the client, the server, or collaboratively by both and exchanged using the subkey field. This generation might involve the use of a random number as a pre-key,

initially generated by either party, which could then be encrypted using the session key from the ticket, and the result exchanged and used for subsequent encryption. By encrypting the pre-key with the session key from the ticket, randomness from the KDC generated key is assured of being present in the negotiated key. Application developers must be careful however, to use a means of introducing this entropy that does not allow an attacker to learn the session key from the ticket if it learns the key generated and used for subsequent communication. The reader should note that this is only an example, and that an analysis of the particular cryptosystem to be used, must be made before deciding how to generate values for the subkey fields, and the key to be used for subsequent communication.

With both the one-way and mutual authentication exchanges, the peers should take care not to send sensitive information to each other without proper assurances. In particular, applications that require privacy or integrity should use the KRB_AP_REP response from the server to client to assure both client and server of their peer's identity. If an application protocol requires privacy of its messages, it can use the KRB_PRIV message (section 3.5). The KRB_SAFE message (section 3.4) can be used to assure integrity.

3.3. The Ticket-Granting Service (TGS) Exchange

Summary

Message direction	Message type	Section
1. Client to Kerberos	KRB_TGS_REQ	5.4.1
2. Kerberos to client	KRB_TGS_REP or KRB_ERROR	5.4.2 5.9.1

The TGS exchange between a client and the Kerberos Ticket-Granting Server is initiated by a client when it wishes to obtain authentication credentials for a given server (which might be registered in a remote realm), when it wishes to renew or validate an existing ticket, or when it wishes to obtain a proxy ticket. In the first case, the client must already have acquired a ticket for the Ticket-Granting Service using the AS exchange (the ticket-granting ticket is usually obtained when a client initially authenticates to the system, such as when a user logs in). The message format for the TGS exchange is almost identical to that for the AS exchange.

The primary difference is that encryption and decryption in the TGS exchange does not take place under the client's key. Instead, the session key from the ticket-granting ticket or renewable ticket, or sub-session key from an Authenticator is used. As is the case for all application servers, expired tickets are not accepted by the TGS, so once a renewable or ticket-granting ticket expires, the client must use a separate exchange to obtain valid tickets.

The TGS exchange consists of two messages: A request (KRB_TGS_REQ) from the client to the Kerberos Ticket-Granting Server, and a reply (KRB_TGS_REP or KRB_ERROR). The KRB_TGS_REQ message includes information authenticating the client plus a request for credentials. The authentication information consists of the authentication header (KRB_AP_REQ) which includes the client's previously obtained ticket-granting, renewable, or invalid ticket. In the ticket-granting ticket and proxy cases, the request may include one or more of: a list of network addresses, a collection of typed authorization data to be sealed in the ticket for authorization use by the application server, or additional tickets (the use of which are described later). The TGS reply (KRB_TGS_REP) contains the requested credentials, encrypted in the session key from the ticket-granting ticket or renewable ticket, or if present, in the sub-session key from the Authenticator (part of the authentication header). The KRB_ERROR message contains an error code and text explaining what went wrong. The KRB_ERROR message is not encrypted. The KRB_TGS_REP message contains information which can be used to detect replays, and to associate it with the message to which it replies. The KRB_ERROR message also contains information which can be used to associate it with the message to which it replies, but except when an optional checksum is included in the KRB_ERROR message, it is not possible to detect replays or fabrications of such messages.

3.3.1. Generation of KRB_TGS_REQ message

Before sending a request to the ticket-granting service, the client must determine in which realm the application server is believed to be registered[3.12]. If the client knows the service principal name and realm and it does not already possess a ticket-granting ticket for the appropriate realm, then one must be obtained. This is first attempted by requesting a

ticket-granting ticket for the destination realm from a Kerberos server for which the client possesses a ticket-granting ticket (using the KRB_TGS_REQ message recursively). The Kerberos server may return a TGT for the desired realm in which case one can proceed. Alternatively, the Kerberos server may return a TGT for a realm which is 'closer' to the desired realm (further along the standard hierarchical path between the client's realm and the requested realm server's realm).

If the client does not know the realm of the service or the true service principal name, then the CANONICALIZE option must be used in the request. This will cause the TGS to locate the service principal based on the target service name in the ticket and return the service principal name in the response. This function allows the KDC to inform the user of the registered Kerberos principal name and registered KDC for a server that may have more than one host name or whose registered realm can not be determined from the name of the host, but it is not to be used to locate the application server.

If the server name determined by a TGS supporting name canonicalization is with a remote KDC, then the response will include the principal name determined by the KDC, and will include a TGT for the remote realm or a realm 'closer' to the realm with which the server principal is registered. In this case, the canonicalization request must be repeated with a Kerberos server in the realm specified in the returned TGT. If neither are returned, then the request may be retried with a Kerberos server for a realm higher in the hierarchy. This request will itself require a ticket-granting ticket for the higher realm which must be obtained by recursively applying these directions.

Once the client obtains a ticket-granting ticket for the appropriate realm, it determines which Kerberos servers serve that realm, and contacts one. The list might be obtained through a configuration file or network service or it may be generated from the name of the realm; as long as the secret keys exchanged by realms are kept secret, only denial of service results from using a false Kerberos server.

As in the AS exchange, the client may specify a number of options in the KRB_TGS_REQ message. The client prepares the KRB_TGS_REQ message, providing an authentication header as an element of the padata field, and including the same fields as used in the KRB_AS_REQ message along with several optional fields: the enc-authorization-data field for application server use and additional tickets required by some options.

In preparing the authentication header, the client can select a sub-session key under which the response from the Kerberos server will be encrypted[3.13]. If the sub-session key is not specified, the session key from the ticket-granting ticket will be used. If the enc-authorization-data is present, it must be encrypted in the sub-session key, if present, from the authenticator portion of the authentication header, or if not present, using the session key from the ticket-granting ticket.

Once prepared, the message is sent to a Kerberos server for the destination realm. See section A.5 for pseudocode.

3.3.2. Receipt of KRB_TGS_REQ message

The KRB_TGS_REQ message is processed in a manner similar to the KRB_AS_REQ message, but there are many additional checks to be performed. First, the Kerberos server must determine which server the accompanying ticket is for and it must select the appropriate key to decrypt it. For a normal KRB_TGS_REQ message, it will be for the ticket granting service, and the TGS's key will be used. If the TGT was issued by another realm, then the appropriate inter-realm key must be used. If the accompanying ticket is not a ticket granting ticket for the current realm, but is for an application server in the current realm, the RENEW, VALIDATE, or PROXY options are specified in the request, and the server for which a ticket is requested is the server named in the accompanying ticket, then the KDC will decrypt the ticket in the authentication header using the key of the server for which it was issued. If no ticket can be found in the padata field, the KDC_ERR_PADATA_TYPE_NOSUPP error is returned.

Once the accompanying ticket has been decrypted, the user-supplied checksum in the Authenticator must be verified against the contents of the request, and the message rejected if the checksums do not match (with an error code of KRB_AP_ERR_MODIFIED) or if the checksum is not keyed or not collision-proof (with an error code of KRB_AP_ERR_INAPP_CKSUM). If the checksum type is not supported, the KDC_ERR_SUMTYPE_NOSUPP error is returned. If the authorization-data are present, they are decrypted using the sub-session key from the Authenticator.

If any of the decryptions indicate failed integrity checks, the KRB_AP_ERR_BAD_INTEGRITY error is returned. If the CANONICALIZE option is set in the KRB_TGS_REQ, then the requested service name might not be the true principal name or the service might not be in the TGS realm and the correct name must be determined.

3.3.3. Generation of KRB_TGS_REP message

The KRB_TGS_REP message shares its format with the KRB_AS_REP (KRB_KDC_REP), but with its type field set to KRB_TGS_REP. The detailed specification is in section 5.4.2.

The response will include a ticket for the requested server or for a ticket granting server of an intermediate KDC to be contacted to obtain the requested ticket. The Kerberos database is queried to retrieve the record for the appropriate server (including the key with which the ticket will be encrypted). If the request is for a ticket granting ticket for a remote realm, and if no key is shared with the requested realm, then the Kerberos server will select the realm 'closest' to the requested realm with which it does share a key, and use that realm instead. If the CANONICALIZE option is set, the TGS may return a ticket containing the server name of the true service principal. If the requested server cannot be found in the TGS database, then a TGT for another trusted realm may be returned instead of a ticket for the service. This TGT is a referral mechanism to cause the client to retry the request to the realm of the TGT. These are the only cases where the response for the KDC will be for a different server than that requested by the client.

By default, the address field, the client's name and realm, the list of transited realms, the time of initial authentication, the expiration time, and the authorization data of the newly-issued ticket will be copied from the ticket-granting ticket (TGT) or renewable ticket. If the transited field needs to be updated, but the transited type is not supported, the KDC_ERR_TRTYPE_NOSUPP error is returned.

If the request specifies an endtime, then the endtime of the new ticket is set to the minimum of (a) that request, (b) the endtime from the TGT, and (c) the starttime of the TGT plus the minimum of the maximum life for the application server and the maximum life for the local realm (the maximum life for the requesting principal was already applied when the TGT was issued). If the new ticket is to be a renewal, then the endtime above is replaced by the minimum of (a) the value of the renew_till field of the ticket and (b) the starttime for the new ticket plus the life (endtime-starttime) of the old ticket.

If the FORWARDED option has been requested, then the resulting ticket will contain the addresses specified by the client. This option will only be honored if the FORWARDABLE flag is set in the TGT. The PROXY option is similar; the resulting ticket will contain the addresses specified by the client. It will be honored only if the PROXIABLE flag in the TGT is set. The PROXY option will not be honored on requests for additional ticket-granting tickets.

If the requested start time is absent, indicates a time in the past, or is within the window of acceptable clock skew for the KDC and the POSTDATE option has not been specified, then the start time of the ticket is set to the authentication server's current time. If it indicates a time in the future beyond the acceptable clock skew, but the POSTDATED option has not been specified or the MAY-POSTDATE flag is not set in the TGT, then the error KDC_ERR_CANNOT_POSTDATE is returned. Otherwise, if the ticket-granting ticket has the MAY-POSTDATE flag set, then the resulting ticket will be postdated and the requested starttime is checked against the policy of the local realm. If acceptable, the ticket's start time is set as requested, and the INVALID flag is set. The postdated ticket must be validated before use by presenting it to the KDC after the starttime has been reached. However,

in no case may the starttime, endtime, or renew-till time of a newly-issued postdated ticket extend beyond the renew-till time of the ticket-granting ticket.

If the ENC-TKT-IN-SKEY option has been specified and an additional ticket has been included in the request, the KDC will decrypt the additional ticket using the key for the server to which the additional ticket was issued and verify that it is a ticket-granting ticket. If the name of the requested server is missing from the request, the name of the client in the additional ticket will be used. Otherwise the name of the requested server will be compared to the name of the client in the additional ticket and if different, the request will be rejected. If the request succeeds, the session key from the additional ticket will be used to encrypt the new ticket that is issued instead of using the key of the server for which the new ticket will be used.

If the name of the server in the ticket that is presented to the KDC as part of the authentication header is not that of the ticket-granting server itself, the server is registered in the realm of the KDC, and the RENEW option is requested, then the KDC will verify that the RENEWABLE flag is set in the ticket, that the INVALID flag is not set in the ticket, and that the renew_till time is still in the future. If the VALIDATE option is requested, the KDC will check that the starttime has passed and the INVALID flag is set. If the PROXY option is requested, then the KDC will check that the PROXIABLE flag is set in the ticket. If the tests succeed, and the ticket passes the hotlist check described in the next section, the KDC will issue the appropriate new ticket.

The ciphertext part of the response in the KRB_TGS_REP message is encrypted in the sub-session key from the Authenticator, if present, or the session key from the ticket-granting ticket. It is not encrypted using the client's secret key. Furthermore, the client's key's expiration date and the key version number fields are left out since these values are stored along with the client's database record, and that record is not needed to satisfy a request based on a ticket-granting ticket. See section A.6 for pseudocode.

3.3.3.1. Checking for revoked tickets

Whenever a request is made to the ticket-granting server, the presented ticket(s) is(are) checked against a hot-list of tickets which have been canceled. This hot-list might be implemented by storing a range of issue timestamps for 'suspect tickets'; if a presented ticket had an authtime in that range, it would be rejected. In this way, a stolen ticket-granting ticket or renewable ticket cannot be used to gain additional tickets (renewals or otherwise) once the theft has been reported to the KDC for the realm in which the server resides. Any normal ticket obtained before it was reported stolen will still be valid (because they require no interaction with the KDC), but only until their normal expiration time. If TGT's have been issued for cross-realm authentication, use of the cross-realm TGT will not be affected unless the hot-list is propagated to the KDC's for the realms for which such cross-realm tickets were issued.

3.3.3.2. Encoding the transited field

If the identity of the server in the TGT that is presented to the KDC as part of the authentication header is that of the ticket-granting service, but the TGT was issued from another realm, the KDC will look up the inter-realm key shared with that realm and use that key to decrypt the ticket. If the ticket is valid, then the KDC will honor the request, subject to the constraints outlined above in the section describing the AS exchange. The realm part of the client's identity will be taken from the ticket-granting ticket. The name of the realm that issued the ticket-granting ticket will be added to the transited field of the ticket to be issued. This is accomplished by reading the transited field from the ticket-granting ticket (which is treated as an unordered set of realm names), adding the new realm to the set, then constructing and writing out its encoded (shorthand) form (this may involve a rearrangement of the existing encoding).

Note that the ticket-granting service does not add the name of its own realm. Instead, its responsibility is to add the name of the previous realm. This prevents a malicious Kerberos server from intentionally leaving out its own name (it could, however, omit other realms' names).

The names of neither the local realm nor the principal's realm are to be included in the transited field. They appear elsewhere in the ticket and both are known to have taken part in authenticating the principal. Since the endpoints are not included, both local and single-hop inter-realm authentication result in a transited field that is empty.

Because the name of each realm transited is added to this field, it might potentially be very long. To decrease the length of this field, its contents are encoded. The initially supported encoding is optimized for the normal case of inter-realm communication: a hierarchical arrangement of realms using either domain or X.500 style realm names. This encoding (called DOMAIN-X500-COMPRESS) is now described.

Realm names in the transited field are separated by a ",". The ",", "\", trailing "."s, and leading spaces (" ") are special characters, and if they are part of a realm name, they must be quoted in the transited field by preceding them with a "\".

A realm name ending with a "." is interpreted as being prepended to the previous realm. For example, we can encode traversal of EDU, MIT.EDU, ATHENA.MIT.EDU, WASHINGTON.EDU, and CS.WASHINGTON.EDU as:

"EDU,MIT.,ATHENA.,WASHINGTON.EDU,CS.".

Note that if ATHENA.MIT.EDU, or CS.WASHINGTON.EDU were end-points, that they would not be included in this field, and we would have:

"EDU,MIT.,WASHINGTON.EDU"

A realm name beginning with a "/" is interpreted as being appended to the previous realm[18]. If it is to stand by itself, then it should be preceded by a space (" "). For example, we can encode traversal of /COM/HP/APOLLO, /COM/HP, /COM, and /COM/DEC as:

"/COM,/HP,/APOLLO, /COM/DEC".

Like the example above, if /COM/HP/APOLLO and /COM/DEC are endpoints, they they would not be included in this field, and we would have:

"/COM,/HP"

A null subfield preceding or following a "," indicates that all realms between the previous realm and the next realm have been traversed[19]. Thus, "," means that all realms along the path between the client and the server

have been traversed. ",EDU, /COM," means that that all realms from the client's realm up to EDU (in a domain style hierarchy) have been traversed, and that everything from /COM down to the server's realm in an X.500 style has also been traversed. This could occur if the EDU realm in one hierarchy shares an inter-realm key directly with the /COM realm in another hierarchy.

3.3.4. Receipt of KRB_TGS_REP message

When the KRB_TGS_REP is received by the client, it is processed in the same manner as the KRB_AS_REP processing described above. The primary difference is that the ciphertext part of the response must be decrypted using the session key from the ticket-granting ticket rather than the client's secret key. The server name returned in the reply is the true principal name of the service. See section A.7 for pseudocode.

3.4. The KRB_SAFE Exchange

The KRB_SAFE message may be used by clients requiring the ability to detect modifications of messages they exchange. It achieves this by including a keyed collision-proof checksum of the user data and some control information. The checksum is keyed with an encryption key (usually the last key negotiated via subkeys, or the session key if no negotiation has occurred).

3.4.1. Generation of a KRB_SAFE message

When an application wishes to send a KRB_SAFE message, it collects its data and the appropriate control information and computes a checksum over them. The checksum algorithm should be a keyed one-way hash function (such as the

RSA- MD5-DES checksum algorithm specified in section 6.4.5, or the DES MAC), generated using the sub-session key if present, or the session key. Different algorithms may be selected by changing the checksum type in the message. Unkeyed or non-collision-proof checksums are not suitable for this use.

The control information for the KRB_SAFE message includes both a timestamp and a sequence number. The designer of an application using the KRB_SAFE message must choose at least one of the two mechanisms. This choice should be based on the needs of the application protocol.

Sequence numbers are useful when all messages sent will be received by one's peer. Connection state is presently required to maintain the session key, so maintaining the next sequence number should not present an additional problem.

If the application protocol is expected to tolerate lost messages without them being resent, the use of the timestamp is the appropriate replay detection mechanism. Using timestamps is also the appropriate mechanism for multi-cast protocols where all of one's peers share a common sub-session key, but some messages will be sent to a subset of one's peers.

After computing the checksum, the client then transmits the information and checksum to the recipient in the message format specified in section 5.6.1.

3.4.2. Receipt of KRB_SAFE message

When an application receives a KRB_SAFE message, it verifies it as follows. If any error occurs, an error code is reported for use by the application.

The message is first checked by verifying that the protocol version and type fields match the current version and KRB_SAFE, respectively. A mismatch generates a KRB_AP_ERR_BADVERSION or KRB_AP_ERR_MSG_TYPE error. The application verifies that the checksum used is a collision-proof keyed checksum, and if it is not, a KRB_AP_ERR_INAPP_CKSUM error is generated. If the sender's address was included in the control information, the recipient verifies that the operating system's report of the sender's address matches the sender's address in the message, and (if a recipient address is specified or the recipient requires an address) that one of the recipient's addresses appears as the recipient's address in the message. A failed match for either case generates a KRB_AP_ERR_BADADDR error. Then the timestamp and usec and/or the sequence number fields are checked. If timestamp and usec are expected and not present, or they are present but not current, the KRB_AP_ERR_SKEW error is generated. If the server name, along with the client name, time and microsecond fields from the Authenticator match any recently-seen (sent or received[20]) such tuples, the KRB_AP_ERR_REPEAT error is generated. If an incorrect sequence number is included, or a sequence number is expected but not present, the KRB_AP_ERR_BADORDER error is generated. If neither a time-stamp and usec or a sequence number is present, a KRB_AP_ERR_MODIFIED error is generated. Finally, the checksum is computed over the data and control information, and if it doesn't match the received checksum, a KRB_AP_ERR_MODIFIED error is generated.

If all the checks succeed, the application is assured that the message was generated by its peer and was not modified in transit.

3.5. The KRB_PRIV Exchange

The KRB_PRIV message may be used by clients requiring confidentiality and the ability to detect modifications of exchanged messages. It achieves this by encrypting the messages and adding control information.

3.5.1. Generation of a KRB_PRIV message

When an application wishes to send a KRB_PRIV message, it collects its data and the appropriate control information (specified in section 5.7.1) and encrypts them under an encryption key (usually the last key negotiated via subkeys, or the session key if no negotiation has occurred). As part of the control information, the client must choose to use either a timestamp or a sequence number (or both); see the discussion in section 3.4.1 for guidelines on which to use. After the user data and control information are encrypted, the client transmits the ciphertext and some 'envelope' information to the recipient.

3.5.2. Receipt of KRB_PRIV message

When an application receives a KRB_PRIV message, it verifies it as follows. If any error occurs, an error code is reported for use by the application.

The message is first checked by verifying that the protocol version and type fields match the current version and KRB_PRIV, respectively. A mismatch generates a KRB_AP_ERR_BADVERSION or KRB_AP_ERR_MSG_TYPE error. The application then decrypts the ciphertext and processes the resultant plaintext. If decryption shows the data to have been modified, a KRB_AP_ERR_BAD_INTEGRITY error is generated. If the sender's address was included in the control information, the recipient verifies that the operating system's report of the sender's address matches the sender's address in the message, and (if a recipient address is specified or the recipient requires an address) that one of the recipient's addresses appears as the recipient's address in the message. A failed match for either case generates a KRB_AP_ERR_BADADDR error. Then the timestamp and usec and/or the sequence number fields are checked. If timestamp and usec are expected and not present, or they are present but not current, the KRB_AP_ERR_SKEW error is generated. If the server name, along with the client name, time and microsecond fields from the Authenticator match any recently-seen such tuples, the KRB_AP_ERR_REPEAT error is generated. If an incorrect sequence number is included, or a sequence number is expected but not present, the KRB_AP_ERR_BADORDER error is generated. If neither a time-stamp and usec or a sequence number is present, a KRB_AP_ERR_MODIFIED error is generated.

If all the checks succeed, the application can assume the message was generated by its peer, and was securely transmitted (without intruders able to see the unencrypted contents).

3.6. The KRB_CRED Exchange

The KRB_CRED message may be used by clients requiring the ability to send Kerberos credentials from one host to another. It achieves this by sending the tickets together with encrypted data containing the session keys and other information associated with the tickets.

3.6.1. Generation of a KRB_CRED message

When an application wishes to send a KRB_CRED message it first (using the KRB_TGS exchange) obtains credentials to be sent to the remote host. It then constructs a KRB_CRED message using the ticket or tickets so obtained, placing the session key needed to use each ticket in the key field of the corresponding KrbCredInfo sequence of the encrypted part of the KRB_CRED message.

Other information associated with each ticket and obtained during the KRB_TGS exchange is also placed in the corresponding KrbCredInfo sequence in the encrypted part of the KRB_CRED message. The current time and, if specifically required by the application the nonce, s-address, and r-address fields, are placed in the encrypted part of the KRB_CRED message which is then encrypted under an encryption key previously exchanged in the KRB_AP exchange (usually the last key negotiated via subkeys, or the session key if no negotiation has occurred).

3.6.2. Receipt of KRB_CRED message

When an application receives a KRB_CRED message, it verifies it. If any error occurs, an error code is reported for use by the application. The message is verified by checking that the protocol version and type fields match the current version and KRB_CRED, respectively. A mismatch generates a KRB_AP_ERR_BADVERSION or KRB_AP_ERR_MSG_TYPE error. The application then decrypts the ciphertext and processes the resultant plaintext. If decryption shows the data to have been modified, a KRB_AP_ERR_BAD_INTEGRITY error is generated.

If present or required, the recipient verifies that the operating system's report of the sender's address matches the sender's address in the message, and that one of the recipient's addresses appears as the recipient's address in the message. A failed match for either case generates a KRB_AP_ERR_BADADDR error. The timestamp and usec fields (and the nonce field if required) are checked next. If the timestamp and usec are not present, or they are present but not current, the KRB_AP_ERR_SKEW error is generated.

If all the checks succeed, the application stores each of the new tickets in its ticket cache together with the session key and other information in the corresponding KrbCredInfo sequence from the encrypted part of the KRB_CRED message.

4. The Kerberos Database

The Kerberos server must have access to a database containing the principal identifiers and secret keys of any principals to be authenticated[4.1] using such secret keys. The keying material in the database must be protected so that they are only accessible to the Kerberos server and administrative functions specifically authorized to access such material. Specific implementations may handle the storage of keying material separate from the Kerberos database (e.g. in hardware) or by encrypting the keying material before placing it in the Kerberos database. Some implementations might provide a means for using long term secret keys, but not for retrieving them from the Kerberos database.

4.1. Database contents

A database entry will typically contain the following fields, though in some instances a KDC may obtain these values through other means:

Field	Value
name	Principal's identifier
key	Principal's secret key
p_kvno	Principal's key version
max_life	Maximum lifetime for Tickets
max_renewable_life	Maximum total lifetime for renewable Tickets

The name field is an encoding of the principal's identifier. The key field contains an encryption key. This key is the principal's secret key. (The key can be encrypted before storage under a Kerberos "master key" to protect it in case the database is compromised but the master key is not. In that case, an extra field must be added to indicate the master key version used, see below.) The p_kvno field is the key version number of the principal's secret key. The max_life field contains the maximum allowable lifetime (endtime - starttime) for any Ticket issued for this principal. The max_renewable_life field contains the maximum allowable total lifetime for any renewable Ticket issued for this principal. (See section 3.1 for a description of how these lifetimes are used in determining the lifetime of a given Ticket.)

A server may provide KDC service to several realms, as long as the database representation provides a mechanism to distinguish between principal records with identifiers which differ only in the realm name.

When an application server's key changes, if the change is routine (i.e. not the result of disclosure of the old key), the old key should be retained by the server until all tickets that had been issued using that key have expired. Because of this, it is possible for several keys to be active for a single principal. Ciphertext encrypted in a principal's key is always tagged with the version of the key that was used for encryption, to help the recipient find the proper key for decryption.

When more than one key is active for a particular principal, the principal will have more than one record in the Kerberos database. The keys and key version numbers will differ between the records (the rest of the fields may or may not be the same). Whenever Kerberos issues a ticket, or responds to a request for initial authentication, the most recent key (known by the Kerberos server) will be used for encryption. This is the key with the highest key version number.

4.2. Additional fields

Project Athena's KDC implementation uses additional fields in its database:

Field	Value
K_kvno	Kerberos' key version
expiration	Expiration date for entry
attributes	Bit field of attributes
mod_date	Timestamp of last modification
mod_name	Modifying principal's identifier

The K_kvno field indicates the key version of the Kerberos master key under which the principal's secret key is encrypted.

After an entry's expiration date has passed, the KDC will return an error to any client attempting to gain tickets as or for the principal. (A database may want to maintain two expiration dates: one for the principal, and one for the principal's current key. This allows password aging to work independently of the principal's expiration date. However, due to the limited space in the responses, the KDC combines the key expiration and principal expiration date into a single value called 'key_exp', which is used as a hint to the user to take administrative action.)

The attributes field is a bitfield used to govern the operations involving the principal. This field might be useful in conjunction with user registration procedures, for site-specific policy implementations (Project Athena currently uses it for their user registration process controlled by the system-wide database service, Moira [LGDSR87]), to identify whether a principal can play the role of a client or server or both, to note whether a server is appropriately trusted to receive credentials delegated by a client, or to identify the 'string to key' conversion algorithm used for a principal's key[4.2]. Other bits are used to indicate that certain ticket options should not be allowed in tickets encrypted under a principal's key (one bit each): Disallow issuing postdated tickets, disallow issuing forwardable tickets, disallow issuing tickets based on TGT authentication, disallow issuing renewable tickets, disallow issuing proxiable tickets, and disallow issuing tickets for which the principal is the server.

The mod_date field contains the time of last modification of the entry, and the mod_name field contains the name of the principal which last modified the entry.

4.3. Frequently Changing Fields

Some KDC implementations may wish to maintain the last time that a request was made by a particular principal. Information that might be maintained includes the time of the last request, the time of the last request for a ticket-granting ticket, the time of the last use of a ticket-granting ticket, or other times. This information can then be returned to the user in the last-req field (see section 5.2).

Other frequently changing information that can be maintained is the latest expiration time for any tickets that have been issued using each key. This field would be used to indicate how long old keys must remain valid to allow the continued use of outstanding tickets.

4.4. Site Constants

The KDC implementation should have the following configurable constants or options, to allow an administrator to make and enforce policy decisions:

* The minimum supported lifetime (used to determine whether the KDC_ERR_NEVER_VALID error should be returned). This constant should reflect reasonable expectations of round-trip time to the KDC, encryption/decryption time, and processing time by the client and target server, and it should allow for a minimum 'useful' lifetime.
* The maximum allowable total (renewable) lifetime of a ticket (renew_till - starttime).
* The maximum allowable lifetime of a ticket (endtime - starttime).
* Whether to allow the issue of tickets with empty address fields (including the ability to specify that such tickets may only be issued if the request specifies some authorization_data).
* Whether proxiable, forwardable, renewable or post-datable tickets are to be issued.

5. Message Specifications

This section (5) still has revisions that are pending based on comments by Tom Yu. Please see http://www.isi.edu/people/bcn/krb-revisions for the latest versions. There will be additional updates prior to the San Diego IETF meeting.

The following sections describe the exact contents and encoding of protocol messages and objects. The ASN.1 base definitions are presented in the first subsection. The remaining subsections specify the protocol objects (tickets and authenticators) and messages. Specification of encryption and checksum techniques, and the fields related to them, appear in section 6.

Optional field in ASN.1 sequences

For optional integer value and date fields in ASN.1 sequences where a default value has been specified, certain default values will not be allowed in the encoding because these values will always be represented through defaulting by the absence of the optional field. For example, one will not send a microsecond zero value because one must make sure that there is only one way to encode this value.

Additional fields in ASN.1 sequences

Implementations receiving Kerberos messages with additional fields present in ASN.1 sequences should carry those fields through, unmodified, when the message is forwarded. Implementations should not drop such fields if the sequence is re-encoded.

5.1. ASN.1 Distinguished Encoding Representation

All uses of ASN.1 in Kerberos shall use the Distinguished Encoding Representation of the data elements as described in the X.509 specification, section 8.7 [X509-88].

5.2. ASN.1 Base Definitions

The following ASN.1 base definitions are used in the rest of this section. Note that since the underscore character (_) is not permitted in ASN.1 names, the hyphen (-) is used in its place for the purposes of ASN.1 names.

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

Kerberos realms are encoded as GeneralStrings. Realms shall not contain a character with the code 0 (the ASCII NUL). Most realms will usually consist of several components separated by periods (.), in the style of Internet Domain Names, or separated by slashes (/) in the style of X.500 names. Acceptable forms for realm names are specified in section 7. A PrincipalName is a typed sequence of components consisting of the following sub-fields:

name-type

This field specifies the type of name that follows. Pre-defined values for this field are specified in section 7.2. The name-type should be treated as a hint. Ignoring the name type, no two names can be the same (i.e. at least one of the components, or the realm, must be different). This constraint may be eliminated in the future.

name-string

This field encodes a sequence of components that form a name, each component encoded as a GeneralString. Taken together, a PrincipalName and a Realm form a principal identifier. Most PrincipalNames will have only a few components (typically one or two).

```
KerberosTime ::=   GeneralizedTime
                   -- Specifying UTC time zone (Z)
```

The timestamps used in Kerberos are encoded as GeneralizedTimes. An encoding shall specify the UTC time zone (Z) and shall not include any fractional portions of the seconds. It further shall not include any separators. Example: The only valid format for UTC time 6 minutes, 27 seconds after 9 pm on 6 November 1985 is 19851106210627Z.

```
HostAddress ::=     SEQUENCE  {
                    addr-type[0]            INTEGER,
                    address[1]              OCTET STRING
}

HostAddresses ::=   SEQUENCE OF HostAddress
```

The host address encodings consists of two fields:

addr-type
     This field specifies the type of address that follows. Pre-defined values for this field are specified in section 8.1.

address
     This field encodes a single address of type addr-type.

The two forms differ slightly. HostAddress contains exactly one address; HostAddresses contains a sequence of possibly many addresses.

```
AuthorizationData ::=   SEQUENCE OF SEQUENCE {
                        ad-type[0]              INTEGER,
                        ad-data[1]              OCTET STRING
}
```

ad-data
     This field contains authorization data to be interpreted according to the value of the corresponding ad-type field.

ad-type
     This field specifies the format for the ad-data subfield. All negative values are reserved for local use. Non-negative values are reserved for registered use.

Each sequence of type and data is referred to as an authorization element. Elements may be application specific, however, there is a common set of recursive elements that should be understood by all implementations. These elements contain other elements embedded within them, and the interpretation of the encapsulating element determines which of the embedded elements must be interpreted, and which may be ignored. Definitions for these common elements may be found in Appendix B.

```
TicketExtensions ::= SEQUENCE OF SEQUENCE {
           te-type[0]       INTEGER,
           te-data[1]       OCTET STRING
}
```

```
te-data
```

This field contains opaque data that must be carried with the ticket to support extensions to the Kerberos protocol including but not limited to some forms of inter-realm key exchange and plaintext authorization data. See appendix C for some common uses of this field.

```
te-type
```

This field specifies the format for the te-data subfield. All negative values are reserved for local use. Non-negative values are reserved for registered use.

```
APOptions ::=   BIT STRING
                  -- reserved(0),
                  -- use-session-key(1),
                  -- mutual-required(2)

TicketFlags ::= BIT STRING
                  -- reserved(0),
                  -- forwardable(1),
                  -- forwarded(2),
                  -- proxiable(3),
                  -- proxy(4),
                  -- may-postdate(5),
                  -- postdated(6),
                  -- invalid(7),
                  -- renewable(8),
                  -- initial(9),
                  -- pre-authent(10),
                  -- hw-authent(11),
                  -- transited-policy-checked(12),
                  -- ok-as-delegate(13)
```

```
KDCOptions ::=   BIT STRING io
                  -- reserved(0),
                  -- forwardable(1),
                  -- forwarded(2),
                  -- proxiable(3),
                  -- proxy(4),
                  -- allow-postdate(5),
                  -- postdated(6),
                  -- unused7(7),
                  -- renewable(8),
                  -- unused9(9),
                  -- unused10(10),
                  -- unused11(11),
                  -- unused12(12),
                  -- unused13(13),
                  -- requestanonymous(14),
                  -- canonicalize(15),
                  -- disable-transited-check(26),
                  -- renewable-ok(27),
                  -- enc-tkt-in-skey(28),
                  -- renew(30),
                  -- validate(31)
```

ASN.1 Bit strings have a length and a value. When used in Kerberos for the APOptions, TicketFlags, and KDCOptions, the length of the bit string on generated values should be the smallest number of bits needed to include the highest order bit that is set (1), but in no case less than 32 bits. The ASN.1 representation of the bit strings uses unnamed bits, with the meaning of the individual bits defined by the comments in the specification above. Implementations should accept values of bit strings of any length and treat the value of flags corresponding to bits beyond the end of the bit string as if the bit were reset (0). Comparison of bit strings of different length should treat the smaller string as if it were padded with zeros beyond the high order bits to the length of the longer string[23].

```
LastReq ::=   SEQUENCE OF SEQUENCE {
               lr-type[0]               INTEGER,
               lr-value[1]              KerberosTime
}
```

lr-type

This field indicates how the following lr-value field is to be interpreted. Negative values indicate that the information pertains only to the responding server. Non-negative values pertain to all servers for the realm. If the lr-type field is zero (0), then no information is conveyed by the lr-value subfield. If the absolute value of the lr-type field is one (1), then the lr-value subfield is the time of last initial request for a TGT. If it is two (2), then the lr-value subfield is the time of last initial request. If it is three (3), then the lr-value subfield is the time of issue for the newest ticket-granting ticket used. If it is four (4), then the lr-value subfield is the time of the last renewal. If it is five (5), then the lr-value subfield is the time of last request (of any type). If it is (6), then the lr-value subfield is the time when the password will expire.

lr-value

This field contains the time of the last request. the time must be interpreted according to the contents of the accompanying lr-type subfield.

See section 6 for the definitions of Checksum, ChecksumType, EncryptedData, EncryptionKey, EncryptionType, and KeyType.

5.3. Tickets and Authenticators

This section describes the format and encryption parameters for tickets and authenticators. When a ticket or authenticator is included in a protocol message it is treated as an opaque object.

5.3.1. Tickets

A ticket is a record that helps a client authenticate to a service. A Ticket contains the following information:

```
Ticket ::=        [APPLICATION 1] SEQUENCE {
                  tkt-vno[0]                    INTEGER,
                  realm[1]                      Realm,
                  sname[2]                      PrincipalName,
                  enc-part[3]                   EncryptedData, -- EncTicketPart
                  extensions[4]                 TicketExtensions OPTIONAL
}

-- Encrypted part of ticket
EncTicketPart ::= [APPLICATION 3] SEQUENCE {
                  flags[0]                      TicketFlags,
                  key[1]                        EncryptionKey,
                  crealm[2]                     Realm,
                  cname[3]                      PrincipalName,
                  transited[4]                  TransitedEncoding,
                  authtime[5]                   KerberosTime,
                  starttime[6]                  KerberosTime OPTIONAL,
                  endtime[7]                    KerberosTime,
                  renew-till[8]                 KerberosTime OPTIONAL,
                  caddr[9]                      HostAddresses OPTIONAL,
                  authorization-data[10]        AuthorizationData OPTIONAL
}
-- encoded Transited field
TransitedEncoding ::=   SEQUENCE {
                        tr-type[0]              INTEGER, -- must be registered
                        contents[1]             OCTET STRING
}
```

The encoding of EncTicketPart is encrypted in the key shared by Kerberos and the end server (the server's secret key). See section 6 for the format of the ciphertext.

tkt-vno

This field specifies the version number for the ticket format. This document describes version number 5.

realm

This field specifies the realm that issued a ticket. It also serves to identify the realm part of the server's principal identifier. Since a Kerberos server can only issue tickets for servers within its realm, the two will always be identical.

sname

This field specifies all components of the name part of the server's identity, including those parts that identify a specific instance of a service.

enc-part

This field holds the encrypted encoding of the EncTicketPart sequence.

extensions

This optional field contains a sequence of extensions that may be used to carry information that must be carried with the ticket to support several extensions, including but not limited to plaintext authorization data, tokens for exchanging inter-realm keys, and other information that must be associated with a ticket for use by the application server. See Appendix C for definitions of common extensions.

Note that some older versions of Kerberos did not support this field. Because this is an optional field it will not break older clients, but older clients might strip this field from the ticket before sending it to the application server. This limits the usefulness of this ticket field to environments where the ticket will not be parsed and reconstructed by these older Kerberos clients.

If it is known that the client will strip this field from the ticket, as an interim measure the KDC may append this field to the end of the enc-part of the ticket and append a trailer indicating the length of the appended extensions field.

flags

This field indicates which of various options were used or requested when the ticket was issued. It is a bit-field, where the selected options are indicated by the bit being set (1), and the unselected options and reserved fields being reset (0). Bit 0 is the most significant bit. The encoding of the bits is specified in section 5.2. The flags are described in more detail above in section 2. The meanings of the flags are:

Bits	Name	Description
0	RESERVED	Reserved for future expansion of this field.
1	FORWARDABLE	The FORWARDABLE flag is normally only interpreted by the TGS, and can be ignored by end servers. When set, this flag tells the ticket-granting server that it is OK to issue a new ticket-granting ticket with a different network address based on the presented ticket.
2	FORWARDED	When set, this flag indicates that the ticket has either been forwarded or was issued based on authentication involving a forwarded ticket-granting ticket.

3	PROXIABLE	The PROXIABLE flag is normally only interpreted by the TGS, and can be ignored by end servers. The PROXIABLE flag has an interpretation identical to that of the FORWARDABLE flag, except that the PROXIABLE flag tells the ticket-granting server that only non-ticket-granting tickets may be issued with different network addresses.
4	PROXY	When set, this flag indicates that a ticket is a proxy.
5	MAY-POSTDATE	The MAY-POSTDATE flag is normally only interpreted by the TGS, and can be ignored by end servers. This flag tells the ticket-granting server that a post-dated ticket may be issued based on this ticket-granting ticket.
6	POSTDATED	This flag indicates that this ticket has been postdated. The end-service can check the authtime field to see when the original authentication occurred.
7	INVALID	This flag indicates that a ticket is invalid, and it must be validated by the KDC before use. Application servers must reject tickets which have this flag set.

8	RENEWABLE	The RENEWABLE flag is normally only interpreted by the TGS, and can usually be ignored by end servers (some particularly careful servers may wish to disallow renewable tickets). A renewable ticket can be used to obtain a replacement ticket that expires at a later date.
9	INITIAL	This flag indicates that this ticket was issued using the AS protocol, and not issued based on a ticket-granting ticket.
10	PRE-AUTHENT	This flag indicates that during initial authentication, the client was authenticated by the KDC before a ticket was issued. The strength of the preauthentication method is not indicated, but is acceptable to the KDC.
11	HW-AUTHENT	This flag indicates that the protocol employed for initial authentication required the use of hardware expected to be possessed solely by the named client. The hardware authentication method is selected by the KDC and the strength of the method is not indicated.

		This flag indicates that the KDC for the realm has checked the transited field against a realm defined policy for trusted certifiers. If this flag is reset (0), then the application server must check the transited field itself, and if unable to do so it must
12	TRANSITED-POLICY-CHECKED	reject the authentication. If the flag is set (1) then the application server may skip its own validation of the transited field, relying on the validation performed by the KDC. At its option the application server may still apply its own validation based on a separate policy for acceptance.
13	OK-AS-DELEGATE	This flag indicates that the server (not the client) specified in the ticket has been determined by policy of the realm to be a suitable recipient of delegation. A client can use the presence of this flag to help it make a decision whether to delegate credentials (either grant a proxy or a forwarded ticket granting ticket) to this server. The client is free to ignore the value of this flag. When setting this flag, an administrator should consider the Security and placement of the server on which the service will run, as well as whether the service requires the use of delegated credentials.

		This flag indicates that the principal named in the ticket is a generic principal for the realm and does not identify the individual using the ticket. The purpose of the ticket is only to securely distribute a session
14	ANONYMOUS	key, and not to identify the user. Subsequent requests using the same ticket and session may be considered as originating from the same user, but requests with the same username but a different ticket are likely to originate from different users.
15-31	RESERVED	Reserved for future use.

key

This field exists in the ticket and the KDC response and is used to pass the session key from Kerberos to the application server and the client. The field's encoding is described in section 6.2.

crealm

This field contains the name of the realm in which the client is registered and in which initial authentication took place.

cname

This field contains the name part of the client's principal identifier.

transited

This field lists the names of the Kerberos realms that took part in authenticating the user to whom this ticket was issued. It does not specify the order in which the realms were transited. See section 3.3.3.2 for details on how this field encodes the traversed realms. When the names of CA's are to be embedded in the transited field (as specified for some extensions to the protocol), the X.500 names of the CA's should be mapped into items in the transited field using the mapping defined by RFC2253.

authtime

This field indicates the time of initial authentication for the named principal. It is the time of issue for the original ticket on which this ticket is based. It is included in the ticket to provide additional information to the end service, and to provide the necessary information for implementation of a `hot list' service at the KDC. An end service that is particularly paranoid could refuse to accept tickets for which the initial authentication occurred "too far" in the past. This field is also returned as part of the response from the KDC. When returned as part of the response to initial authentication (KRB_AS_REP), this is the current time on the Kerberos server[24].

starttime

This field in the ticket specifies the time after which the ticket is valid. Together with endtime, this field specifies the life of the ticket. If it is absent from the ticket, its value should be treated as that of the authtime field.

endtime

This field contains the time after which the ticket will not be honored (its expiration time). Note that individual services may place their own limits on the life of a ticket and may reject tickets which have not yet expired. As such, this is really an upper bound on the expiration time for the ticket.

renew-till

This field is only present in tickets that have the RENEWABLE flag set in the flags field. It indicates the maximum endtime that may be included in a renewal. It can be thought of as the absolute expiration time for the ticket, including all renewals.

caddr

This field in a ticket contains zero (if omitted) or more (if present) host addresses. These are the addresses from which the ticket can be used. If there are no addresses, the ticket can be used from any location. The decision by the KDC to issue or by the end server to accept zero-address tickets is a policy decision and is left to the Kerberos and end-service administrators; they may refuse to issue or accept such tickets. The suggested and default policy, however, is that such tickets will only be issued or accepted when additional information that can be used to restrict the use of the ticket is included in the authorization_data field. Such a ticket is a capability.

Network addresses are included in the ticket to make it harder for an attacker to use stolen credentials. Because the session key is not sent over the network in cleartext, credentials can't be stolen simply by listening to the network; an attacker has to gain access to the session key (perhaps through operating system security breaches or a careless user's unattended session) to make use of stolen tickets.

It is important to note that the network address from which a connection is received cannot be reliably determined. Even if it could be, an attacker who has compromised the client's workstation could use the credentials from there. Including the network addresses only makes it more difficult, not impossible, for an attacker to walk off with stolen credentials and then use them from a "safe" location.

authorization-data

The authorization-data field is used to pass authorization data from the principal on whose behalf a ticket was issued to the application service. If no authorization data is included, this field will be left out. Experience has shown that the name of this field is confusing, and that a better name for this field would be restrictions. Unfortunately, it is not possible to change the name of this field at this time.

This field contains restrictions on any authority obtained on the basis of authentication using the ticket. It is possible for any principal in posession of credentials to add entries to the authorization data field since these entries further restrict what can be done with the ticket. Such additions can be made by specifying the additional entries when a new ticket is obtained during the TGS exchange, or they may be added during chained delegation using the authorization data field of the authenticator.

Because entries may be added to this field by the holder of credentials, except when an entry is separately authenticated by encapsulation in the kdc-issued element, it is not allowable for the presence of an entry in the authorization data field of a ticket to amplify the privileges one would obtain from using a ticket.

The data in this field may be specific to the end service; the field will contain the names of service specific objects, and the rights to those objects. The format for this field is described in section 5.2. Although Kerberos is not concerned with the format of the contents of the sub-fields, it does carry type information (ad-type).

By using the authorization_data field, a principal is able to issue a proxy that is valid for a specific purpose. For example, a client wishing to print a file can obtain a file server proxy to be passed to the print server. By specifying the name of the file in the authorization_data field, the file server knows that the print server can only use the client's rights when accessing the particular file to be printed.

A separate service providing authorization or certifying group membership may be built using the authorization-data field. In this case, the entity granting authorization (not the authorized entity), may obtain a ticket in its own name (e.g. the ticket is issued in the name of a privilege server), and this entity adds restrictions on its own authority and delegates the restricted authority through a proxy to the client. The client would then present this authorization credential to the application server separately from the authentication exchange.

Alternatively, such authorization credentials may be embedded in the ticket authenticating the authorized entity, when the authorization is separately authenticated using the kdc-issued authorization data element (see B.4).

Similarly, if one specifies the authorization-data field of a proxy and leaves the host addresses blank, the resulting ticket and session key can be treated as a capability. See [Neu93] for some suggested uses of this field.

The authorization-data field is optional and does not have to be included in a ticket.

5.3.2. Authenticators

An authenticator is a record sent with a ticket to a server to certify the client's knowledge of the encryption key in the ticket, to help the server detect replays, and to help choose a "true session key" to use with the particular session. The encoding is encrypted in the ticket's session key shared by the client and the server:

```
-- Unencrypted authenticator
Authenticator ::= [APPLICATION 2] SEQUENCE  {
                  authenticator-vno[0]          INTEGER,
                  crealm[1]                     Realm,
                  cname[2]                      PrincipalName,
                  cksum[3]                      Checksum OPTIONAL,
                  cusec[4]                      INTEGER,
                  ctime[5]                      KerberosTime,
                  subkey[6]                     EncryptionKey OPTIONAL,
                  seq-number[7]                 INTEGER OPTIONAL,
                  authorization-data[8]         AuthorizationData OPTIONAL
}
```

authenticator-vno

This field specifies the version number for the format of the authenticator. This document specifies version 5.

crealm and cname

These fields are the same as those described for the ticket in section 5.3.1.

cksum

This field contains a checksum of the the application data that accompanies the KRB_AP_REQ.

cusec

This field contains the microsecond part of the client's timestamp. Its value (before encryption) ranges from 0 to 999999. It often appears along with ctime. The two fields are used together to specify a reasonably accurate timestamp.

ctime

This field contains the current time on the client's host.

subkey

This field contains the client's choice for an encryption key which is to be used to protect this specific application session. Unless an application specifies otherwise, if this field is left out the session key from the ticket will be used.

seq-number

This optional field includes the initial sequence number to be used by the KRB_PRIV or KRB_SAFE messages when sequence numbers are used to detect replays (It may also be used by application specific messages). When included in the authenticator this field specifies the initial sequence number for messages from the client to the server. When included in the AP-REP message, the initial sequence number is that for messages from the server to the client. When used in KRB_PRIV or KRB_SAFE messages, it is incremented by one after each message is sent. Sequence numbers fall in the range of 0 through 2^32 - 1 and wrap to zero following the value 2^32 - 1.

For sequence numbers to adequately support the detection of replays they should be non-repeating, even across connection boundaries. The initial sequence number should be random and uniformly distributed across the full space of possible sequence numbers, so that it cannot be guessed by an attacker and so that it and the successive sequence numbers do not repeat other sequences.

authorization-data

This field is the same as described for the ticket in section 5.3.1. It is optional and will only appear when additional restrictions are to be placed on the use of a ticket, beyond those carried in the ticket itself.

5.4. Specifications for the AS and TGS exchanges

This section specifies the format of the messages used in the exchange between the client and the Kerberos server. The format of possible error messages appears in section 5.9.1.

5.4.1. KRB_KDC_REQ definition

The KRB_KDC_REQ message has no type of its own. Instead, its type is one of KRB_AS_REQ or KRB_TGS_REQ depending on whether the request is for an initial ticket or an additional ticket. In either case, the message is sent from the client to the Authentication Server to request credentials for a service.

The message fields are:

```
AS-REQ ::=          [APPLICATION 10] KDC-REQ
TGS-REQ ::=         [APPLICATION 12] KDC-REQ

KDC-REQ ::=         SEQUENCE {
                    pvno[1]             INTEGER,
                    msg-type[2]         INTEGER,
                    padata[3]           SEQUENCE OF PA-DATA OPTIONAL,
                    req-body[4]         KDC-REQ-BODY
}
```

```
PA-DATA ::=         SEQUENCE {
                    padata-type[1]     INTEGER,
                    padata-value[2]    OCTET STRING,
                                       -- might be encoded AP-REQ
}

KDC-REQ-BODY ::=    SEQUENCE {
                     kdc-options[0]         KDCOptions,
                     cname[1]               PrincipalName OPTIONAL,
                                            -- Used only in AS-REQ
                     realm[2]               Realm, -- Server's realm
                                            -- Also client's in AS-REQ
                     sname[3]               PrincipalName OPTIONAL,
                     from[4]                KerberosTime OPTIONAL,
                     till[5]                KerberosTime OPTIONAL,
                     rtime[6]               KerberosTime OPTIONAL,
                     nonce[7]               INTEGER,
                     etype[8]               SEQUENCE OF INTEGER,
                                            -- EncryptionType,
                                            -- in preference order
                     addresses[9]           HostAddresses OPTIONAL,
                 enc-authorization-data[10] EncryptedData OPTIONAL,
                                            -- Encrypted AuthorizationData
                                            -- encoding
                     additional-tickets[11] SEQUENCE OF Ticket OPTIONAL
}
```

The fields in this message are:

pvno

This field is included in each message, and specifies the protocol version number. This document specifies protocol version 5.

msg-type

This field indicates the type of a protocol message. It will almost always be the same as the application identifier associated with a message. It is included to make the identifier more readily accessible to the application. For the KDC-REQ message, this type will be KRB_AS_REQ or KRB_TGS_REQ.

padata

The padata (pre-authentication data) field contains a sequence of authentication information which may be needed before credentials can be issued or decrypted. In the case of requests for additional tickets (KRB_TGS_REQ), this field will include an element with padata-type of PA-TGS-REQ and data of an authentication header (ticket-granting ticket and authenticator). The checksum in the authenticator (which must be collision-proof) is to be computed over the KDC-REQ-BODY encoding. In most requests for initial authentication (KRB_AS_REQ) and most replies (KDC-REP), the padata field will be left out.

This field may also contain information needed by certain extensions to the Kerberos protocol. For example, it might be used to initially verify the identity of a client before any response is returned. When this field is used to authenticate or pre-authenticate a request, it should contain a keyed checksum over the KDC-REQ-BODY to bind the pre-authentication data to rest of the request. The KDC, as a matter of policy, may decide whether to honor a KDC-REQ which includes any pre-authentication data that does not contain the checksum field.

PA-ENC-TIMESTAMP defines a pre-authentication data type that is used for authenticating a client by way of an encrypted timestamp. This is accomplished with a padata field with padata-type equal to PA-ENC-TIMESTAMP and padata-value defined as follows (query: the checksum is new in this definition. If the optional field will break things we can keep the old PA-ENC-TS-ENC, and define a new alternate form that includes the checksum). :

```
padata-type     ::= PA-ENC-TIMESTAMP
padata-value    ::= EncryptedData -- PA-ENC-TS-ENC

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
}
```

with patimestamp containing the client's time and pausec containing the microseconds which may be omitted if a client will not generate more than one request per second. The ciphertext (padata-value) consists of the PA-ENC-TS-ENC sequence, encrypted using the client's secret key.

It may also be used by the client to specify the version of a key that is being used for accompanying preauthentication, and/or which should be used to encrypt the reply from the KDC.

```
padata-type     ::= PA-USE-SPECIFIED-KVNO
padata-value    ::= Integer
}
```

The KDC should only accept and abide by the value of the use-specified-kvno preauthentication data field when the specified key is still valid and until use of a new key is confirmed. This situation is likely to occur primarily during the period during which an updated key is propagating to other KDC's in a realm.

The padata field can also contain information needed to help the KDC or the client select the key needed for generating or decrypting the response. This form of the padata is useful for supporting the use of certain token cards with Kerberos. The details of such extensions are specified in separate documents. See [Pat92] for additional uses of this field.

padata-type

The padata-type element of the padata field indicates the way that the padata-value element is to be interpreted. Negative values of padata-type are reserved for unregistered use; non-negative values are used for a registered interpretation of the element type.

req-body

This field is a placeholder delimiting the extent of the remaining fields. If a checksum is to be calculated over the request, it is calculated over an encoding of the KDC-REQ-BODY sequence which is enclosed within the req-body field.

kdc-options

This field appears in the KRB_AS_REQ and KRB_TGS_REQ requests to the KDC and indicates the flags that the client wants set on the tickets as well as other information that is to modify the behaviour of the KDC. Where appropriate, the name of an option may be the same as the flag that is set by that option. Although in most case, the bit in the options field will be the same as that in the flags field, this is not guaranteed, so it is not acceptable to simply copy the options field to the flags field. There are various checks that must be made before honoring an option anyway.

The kdc_options field is a bit-field, where the selected options are indicated by the bit being set (1), and the unselected options and reserved fields being reset (0). The encoding of the bits is specified in section 5.2. The options are described in more detail above in section 2. The meanings of the options are:

Bits	Name	Description
0	RESERVED	Reserved for future expansion of this field.

1	FORWARDABLE	The FORWARDABLE option indicates that the ticket to be issued is to have its forwardable flag set. It may only be set on the initial request, or in a subsequent request if the ticket-granting ticket on which it is based is also forwardable. The FORWARDED option is only specified in a request to the ticket-granting server and will only be honored if the ticket-granting ticket in the request has its
2	FORWARDED	FORWARDABLE bit set. This option indicates that this is a request for forwarding. The address(es) of the host from which the resulting ticket is to be valid are included in the addresses field of the request. The PROXIABLE option indicates that the ticket to be issued is to have its proxiable flag set. It may only
3	PROXIABLE	be set on the initial request, or in a subsequent request if the ticket-granting ticket on which it is based is also proxiable.

4	PROXY	The PROXY option indicates that this is a request for a proxy. This option will only be honored if the ticket-granting ticket in the request has its PROXIABLE bit set. The address(es) of the host from which the resulting ticket is to be valid are included in the addresses field of the request.
5	ALLOW-POSTDATE	The ALLOW-POSTDATE option indicates that the ticket to be issued is to have its MAY-POSTDATE flag set. It may only be set on the initial request, or in a subsequent request if the ticket-granting ticket on which it is based also has its MAY-POSTDATE flag set.
6	POSTDATED	The POSTDATED option indicates that this is a request for a postdated ticket. This option will only be honored if the ticket-granting ticket on which it is based has its MAY-POSTDATE flag set. The resulting ticket will also have its INVALID flag set, and that flag may be reset by a subsequent request to the KDC after the starttime in the ticket has been reached.

7	UNUSED	This option is presently unused.
8	RENEWABLE	The RENEWABLE option indicates that the ticket to be issued is to have its RENEWABLE flag set. It may only be set on the initial request, or when the ticket-granting ticket on which the request is based is also renewable. If this option is requested, then the rtime field in the request contains the desired absolute expiration time for the ticket.
9	RESERVED	Reserved for PK-Cross
10-13	UNUSED	These options are presently unused.
14	REQUEST-ANONYMOUS	The REQUEST-ANONYMOUS option indicates that the ticket to be issued is not to identify the user to which it was issued. Instead, the principal identifier is to be generic, as specified by the policy of the realm (e.g. usually anonymous@realm). The purpose of the ticket is only to securely distribute a session key, and not to identify the user. The ANONYMOUS flag on the ticket to be returned should be set. If the local realms policy does not permit anonymous credentials, the request is to be rejected.

15	CANONICALIZE	The CANONICALIZE option indicates that the client will accept the return of a true server name instead of the name specified in the request. In addition the client will be able to process any TGT referrals that will direct the client to another realm to locate the requested server. If a KDC does not support name- canonicalization, the option is ignored and the appropriate KDC_ERR_C_PRINCIPAL_UNKNOWN or KDC_ERR_S_PRINCIPAL_UNKNOWN error is returned. [JBrezak]
16-25	RESERVED	Reserved for future use.
26	DISABLE-TRANSITED-CHECK	By default the KDC will check the transited field of a ticket-granting-ticket against the policy of the local realm before it will issue derivative tickets based on the ticket granting ticket. If this flag is set in the request, checking of the transited field is disabled. Tickets issued without the performance of this check will be noted by the reset (0) value of the TRANSITED-POLICY-CHECKED flag, indicating to the application server that the tranisted field must be checked locally. KDC's are encouraged but not required to honor the DISABLE-TRANSITED-CHECK option.

27	RENEWABLE-OK	The RENEWABLE-OK option indicates that a renewable ticket will be acceptable if a ticket with the requested life cannot otherwise be provided. If a ticket with the requested life cannot be provided, then a renewable ticket may be issued with a renew-till equal to the requested endtime. The value of the renew-till field may still be limited by local limits, or limits selected by the individual principal or server.
28	ENC-TKT-IN-SKEY	This option is used only by the ticket-granting service. The ENC-TKT-IN-SKEY option indicates that the ticket for the end server is to be encrypted in the session key from the additional ticket-granting ticket provided.
29	RESERVED	Reserved for future use.

30	RENEW	This option is used only by the ticket-granting service. The RENEW option indicates that the present request is for a renewal. The ticket provided is encrypted in the secret key for the server on which it is valid. This option will only be honored if the ticket to be renewed has its RENEWABLE flag set and if the time in its renew-till field has not passed. The ticket to be renewed is passed in the padata field as part of the authentication header.
31	VALIDATE	This option is used only by the ticket-granting service. The VALIDATE option indicates that the request is to validate a postdated ticket. It will only be honored if the ticket presented is postdated, presently has its INVALID flag set, and would be otherwise usable at this time. A ticket cannot be validated before its starttime. The ticket presented for validation is encrypted in the key of the server for which it is valid and is passed in the padata field as part of the authentication header.

cname and sname

These fields are the same as those described for the ticket in section 5.3.1. sname may only be absent when the ENC-TKT-IN-SKEY option is specified. If absent, the name of the server is taken from the name of the client in the ticket passed as additional-tickets.

enc-authorization-data

The enc-authorization-data, if present (and it can only be present in the TGS_REQ form), is an encoding of the desired authorization-data encrypted under the sub-session key if present in the Authenticator, or alternatively from the session key in the ticket-granting ticket, both from the padata field in the KRB_AP_REQ.

realm

This field specifies the realm part of the server's principal identifier. In the AS exchange, this is also the realm part of the client's principal identifier. If the CANONICALIZE option is set, the realm is used as a hint to the KDC for its database lookup.

from

This field is included in the KRB_AS_REQ and KRB_TGS_REQ ticket requests when the requested ticket is to be postdated. It specifies the desired start time for the requested ticket. If this field is omitted then the KDC should use the current time instead.

till

This field contains the expiration date requested by the client in a ticket request. It is optional and if omitted the requested ticket is to have the maximum endtime permitted according to KDC policy for the parties to the authentication exchange as limited by expiration date of the ticket granting ticket or other preauthentication credentials.

rtime

This field is the requested renew-till time sent from a client to the KDC in a ticket request. It is optional.

nonce

This field is part of the KDC request and response. It it intended to hold a random number generated by the client. If the same number is included in the encrypted response from the KDC, it provides evidence that the response is fresh and has not been replayed by an attacker. Nonces must never be re-used. Ideally, it should be generated randomly, but if the correct time is known, it may suffice[25].

etype

This field specifies the desired encryption algorithm to be used in the response.

addresses

This field is included in the initial request for tickets, and optionally included in requests for additional tickets from the ticket-granting server. It specifies the addresses from which the requested ticket is to be valid. Normally it includes the addresses for the client's host. If a proxy is requested, this field will contain other addresses. The contents of this field are usually copied by the KDC into the caddr field of the resulting ticket.

additional-tickets

Additional tickets may be optionally included in a request to the ticket-granting server. If the ENC-TKT-IN-SKEY option has been specified, then the session key from the additional ticket will be used in place of the server's key to encrypt the new ticket. When he ENC-TKT-IN-SKEY option is used for user-to-user authentication, this addional ticket may be a TGT issued by the local realm or an inter-realm TGT issued for the current KDC's realm by a remote KDC. If more than one option which requires additional tickets has been specified, then the additional tickets are used in the order specified by the ordering of the options bits (see kdc-options, above).

The application code will be either ten (10) or twelve (12) depending on whether the request is for an initial ticket (AS-REQ) or for an additional ticket (TGS-REQ).

The optional fields (addresses, authorization-data and additional-tickets) are only included if necessary to perform the operation specified in the kdc-options field.

It should be noted that in KRB_TGS_REQ, the protocol version number appears twice and two different message types appear: the KRB_TGS_REQ message contains these fields as does the authentication header (KRB_AP_REQ) that is passed in the padata field.

5.4.2. KRB_KDC_REP definition

The KRB_KDC_REP message format is used for the reply from the KDC for either an initial (AS) request or a subsequent (TGS) request. There is no message type for KRB_KDC_REP. Instead, the type will be either KRB_AS_REP or KRB_TGS_REP. The key used to encrypt the ciphertext part of the reply depends on the message type. For KRB_AS_REP, the ciphertext is encrypted in the client's secret key, and the client's key version number is included in the key version number for the encrypted data. For KRB_TGS_REP, the ciphertext is encrypted in the sub-session key from the Authenticator, or if absent, the session key from the ticket-granting ticket used in the request. In that case, no version number will be present in the EncryptedData sequence.

The KRB_KDC_REP message contains the following fields:

```
AS-REP ::=    [APPLICATION 11] KDC-REP
TGS-REP ::=   [APPLICATION 13] KDC-REP

KDC-REP ::=   SEQUENCE {
              pvno[0]                    INTEGER,
              msg-type[1]                INTEGER,
              padata[2]                  SEQUENCE OF PA-DATA OPTIONAL,
              crealm[3]                  Realm,
              cname[4]                   PrincipalName,
              ticket[5]                  Ticket,
              enc-part[6]                EncryptedData
                                         -- EncASREpPart or EncTGSReoOart
}

EncASRepPart ::=    [APPLICATION 25[27]] EncKDCRepPart
EncTGSRepPart ::=   [APPLICATION 26] EncKDCRepPart
```

```
EncKDCRepPart ::=   SEQUENCE {
                    key[0]               EncryptionKey,
                    last-req[1]          LastReq,
                    nonce[2]             INTEGER,
                    key-expiration[3]    KerberosTime OPTIONAL,
                    flags[4]             TicketFlags,
                    authtime[5]          KerberosTime,
                    starttime[6]         KerberosTime OPTIONAL,
                    endtime[7]           KerberosTime,
                    renew-till[8]        KerberosTime OPTIONAL,
                    srealm[9]            Realm,
                    sname[10]            PrincipalName,
                    caddr[11]            HostAddresses OPTIONAL
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is either KRB_AS_REP or KRB_TGS_REP.

padata

This field is described in detail in section 5.4.1. One possible use for this field is to encode an alternate "mix-in" string to be used with a string-to-key algorithm (such as is described in section 6.3.2). This ability is useful to ease transitions if a realm name needs to change (e.g. when a company is acquired); in such a case all existing password-derived entries in the KDC database would be flagged as needing a special mix-in string until the next password change.

crealm, cname, srealm and sname

These fields are the same as those described for the ticket in section 5.3.1.

ticket

The newly-issued ticket, from section 5.3.1.

enc-part

This field is a place holder for the ciphertext and related information that forms the encrypted part of a message. The description of the encrypted part of the message follows each appearance of this field. The encrypted part is encoded as described in section 6.1.

key

This field is the same as described for the ticket in section 5.3.1.

last-req

This field is returned by the KDC and specifies the time(s) of the last request by a principal. Depending on what information is available, this might be the last time that a request for a ticket-granting ticket was made, or the last time that a request based on a ticket-granting ticket was successful. It also might cover all servers for a realm, or just the particular server. Some implementations may display this information to the user to aid in discovering unauthorized use of one's identity. It is similar in spirit to the last login time displayed when logging into timesharing systems.

nonce

This field is described above in section 5.4.1.

key-expiration

The key-expiration field is part of the response from the KDC and specifies the time that the client's secret key is due to expire. The expiration might be the result of password aging or an account expiration. This field will usually be left out of the TGS reply since the response to the TGS request is encrypted in a session key and no client information need be retrieved from the KDC database. It is up to the application client (usually the login program) to take appropriate action (such as notifying the user) if the expiration time is imminent.

flags, authtime, starttime, endtime, renew-till and caddr

These fields are duplicates of those found in the encrypted portion of the attached ticket (see section 5.3.1), provided so the client may verify they match the intended request and to assist in proper ticket caching. If the message is of type KRB_TGS_REP, the caddr field will only be filled in if the request was for a proxy or forwarded ticket, or if the user is substituting a subset of the addresses from the ticket granting ticket. If the client-requested addresses are not present or not used, then the addresses contained in the ticket will be the same as those included in the ticket-granting ticket.

5.5. Client/Server (CS) message specifications

This section specifies the format of the messages used for the authentication of the client to the application server.

5.5.1. KRB_AP_REQ definition

The KRB_AP_REQ message contains the Kerberos protocol version number, the message type KRB_AP_REQ, an options field to indicate any options in use, and the ticket and authenticator themselves. The KRB_AP_REQ message is often referred to as the 'authentication header'.

```
AP-REQ ::=      [APPLICATION 14] SEQUENCE {
                pvno[0]                         INTEGER,
                msg-type[1]                     INTEGER,
                ap-options[2]                   APOptions,
                ticket[3]                       Ticket,
                authenticator[4]                EncryptedData
                                                -- Authenticator from 5.3.2
}

APOptions ::=   BIT STRING {
                reserved(0),
                use-session-key(1),
                mutual-required(2)
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_AP_REQ.

ap-options

This field appears in the application request (KRB_AP_REQ) and affects the way the request is processed. It is a bit-field, where the selected options are indicated by the bit being set (1), and the unselected options and reserved fields being reset (0). The encoding of the bits is specified in section 5.2. The meanings of the options are:

Bit(s)	Name	Description
0	RESERVED	Reserved for future expansion of this field.
1	USE-SESSION-KEY	The USE-SESSION-KEY option indicates that the ticket the client is presenting to a server is encrypted in the session key from the server's ticket-granting ticket. When this option is not specified, the ticket is encrypted in the server's secret key.
2	MUTUAL-REQUIRED	The MUTUAL-REQUIRED option tells the server that the client requires mutual authentication, and that it must respond with a KRB_AP_REP message.
3-31	RESERVED	Reserved for future use.

ticket

This field is a ticket authenticating the client to the server.

authenticator

This contains the authenticator, which includes the client's choice of a subkey. Its encoding is described in section 5.3.2.

5.5.2. KRB_AP_REP definition

The KRB_AP_REP message contains the Kerberos protocol version number, the message type, and an encrypted time- stamp. The message is sent in in response to an application request (KRB_AP_REQ) where the mutual authentication option has been selected in the ap-options field.

```
AP-REP ::=          [APPLICATION 15] SEQUENCE {
                    pvno[0]                          INTEGER,
                    msg-type[1]                      INTEGER,
                    enc-part[2]                      EncryptedData
                                                     -- EncAPRepPart
}


EncAPRepPart ::=    [APPLICATION 27[29]] SEQUENCE {
                    ctime[0]                         KerberosTime,
                    cusec[1]                         INTEGER,
                    subkey[2]                        EncryptionKey OPTIONAL,
                    seq-number[3]                    INTEGER OPTIONAL
}
```

The encoded EncAPRepPart is encrypted in the shared session key of the ticket. The optional subkey field can be used in an application-arranged negotiation to choose a per association session key.

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_AP_REP.

enc-part

This field is described above in section 5.4.2.

ctime

This field contains the current time on the client's host.

cusec

This field contains the microsecond part of the client's timestamp.

subkey

This field contains an encryption key which is to be used to protect this specific application session. See section 3.2.6 for specifics on how this field is used to negotiate a key. Unless an application specifies otherwise, if this field is left out, the sub-session key from the authenticator, or if also left out, the session key from the ticket will be used.

seq-number

This field is described above in section 5.3.2.

5.5.3. Error message reply

If an error occurs while processing the application request, the KRB_ERROR message will be sent in response. See section 5.9.1 for the format of the error message. The cname and crealm fields may be left out if the server cannot determine their appropriate values from the corresponding KRB_AP_REQ message. If the authenticator was decipherable, the ctime and cusec fields will contain the values from it.

5.6. KRB_SAFE message specification

This section specifies the format of a message that can be used by either side (client or server) of an application to send a tamper-proof message to its peer. It presumes that a session key has previously been exchanged (for example, by using the KRB_AP_REQ/KRB_AP_REP messages).

5.6.1. KRB_SAFE definition

The KRB_SAFE message contains user data along with a collision-proof checksum keyed with the last encryption key negotiated via subkeys, or the session key if no negotiation has occurred. The message fields are:

```
KRB-SAFE ::=        [APPLICATION 20] SEQUENCE {
                    pvno[0]                     INTEGER,
                    msg-type[1]                 INTEGER,
                    safe-body[2]                KRB-SAFE-BODY,
                    cksum[3]                    Checksum
```

```
}
KRB-SAFE-BODY ::=   SEQUENCE {
                    user-data[0]                OCTET STRING,
                    timestamp[1]                KerberosTime OPTIONAL,
                    usec[2]                     INTEGER OPTIONAL,
                    seq-number[3]               INTEGER OPTIONAL,
                    s-address[4]                HostAddress OPTIONAL,
                    r-address[5]                HostAddress OPTIONAL
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_SAFE.

safe-body

This field is a placeholder for the body of the KRB-SAFE message.

cksum

This field contains the checksum of the application data. Checksum details are described in section 6.4. The checksum is computed over the encoding of the KRB-SAFE sequence. First, the cksum is zeroed and the checksum is computed over the encoding of the KRB-SAFE sequence, then the checksum is set to the result of that computation, and finally the KRB-SAFE sequence is encoded again.

user-data

This field is part of the KRB_SAFE and KRB_PRIV messages and contain the application specific data that is being passed from the sender to the recipient.

timestamp

This field is part of the KRB_SAFE and KRB_PRIV messages. Its contents are the current time as known by the sender of the message. By checking the timestamp, the recipient of the message is able to make sure that it was recently generated, and is not a replay.

usec

This field is part of the KRB_SAFE and KRB_PRIV headers. It contains the microsecond part of the timestamp.

seq-number

This field is described above in section 5.3.2.

s-address

This field specifies the address in use by the sender of the message. It may be omitted if not required by the application protocol. The application designer considering omission of this field is warned, that the inclusion of this address prevents some kinds of replay attacks (e.g. reflection attacks) and that it is only acceptable to omit this address if there is sufficient information in the integrity protected part of the application message for the recipient to unambiguously determine if it was the intended recipient.

r-address

This field specifies the address in use by the recipient of the message. It may be omitted for some uses (such as broadcast protocols), but the recipient may arbitrarily reject such messages. This field along with s-address can be used to help detect messages which have been incorrectly or maliciously delivered to the wrong recipient.

5.7. KRB_PRIV message specification

This section specifies the format of a message that can be used by either side (client or server) of an application to securely and privately send a message to its peer. It presumes that a session key has previously been exchanged (for example, by using the KRB_AP_REQ/KRB_AP_REP messages).

5.7.1. KRB_PRIV definition

The KRB_PRIV message contains user data encrypted in the Session Key. The message fields are:

```
KRB-PRIV ::=        [APPLICATION 21] SEQUENCE {
                    pvno[0]                         INTEGER,
                    msg-type[1]                     INTEGER,
                    enc-part[3]                     EncryptedData
                                                    -- EncKrbPrivPart
}
```

```
EncKrbPrivPart ::=   [APPLICATION 28[31]] SEQUENCE {
                     user-data[0]        OCTET STRING,
                     timestamp[1]        KerberosTime OPTIONAL,
                     usec[2]             INTEGER OPTIONAL,
                     seq-number[3]       INTEGER OPTIONAL,
                     s-address[4]        HostAddress OPTIONAL, -- sender's addr
                     r-address[5]        HostAddress OPTIONAL -- recip's addr
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_PRIV.

enc-part

This field holds an encoding of the EncKrbPrivPart sequence encrypted under the session key[32]. This encrypted encoding is used for the enc-part field of the KRB-PRIV message. See section 6 for the format of the ciphertext.

user-data, timestamp, usec, s-address and r-address

These fields are described above in section 5.6.1.

seq-number

This field is described above in section 5.3.2.

5.8. KRB_CRED message specification

This section specifies the format of a message that can be used to send Kerberos credentials from one principal to another. It is presented here to encourage a common mechanism to be used by applications when forwarding tickets or providing proxies to subordinate servers. It presumes that a session key has already been exchanged perhaps by using the KRB_AP_REQ/KRB_AP_REP messages.

5.8.1. KRB_CRED definition

The KRB_CRED message contains a sequence of tickets to be sent and information needed to use the tickets, including the session key from each. The information needed to use the tickets is encrypted under an encryption key previously exchanged or transferred alongside the KRB_CRED message. The message fields are:

```
KRB-CRED         ::= [APPLICATION 22]   SEQUENCE {
                 pvno[0]                INTEGER,
                 msg-type[1]            INTEGER, -- KRB_CRED
                 tickets[2]             SEQUENCE OF Ticket,
                 enc-part[3]            EncryptedData -- EncKrbCredPart
}



EncKrbCredPart   ::= [APPLICATION 29]   SEQUENCE {
                 ticket-info[0]         SEQUENCE OF KrbCredInfo,
                 nonce[1]               INTEGER OPTIONAL,
                 timestamp[2]           KerberosTime OPTIONAL,
                 usec[3]                INTEGER OPTIONAL,
                 s-address[4]           HostAddress OPTIONAL,
                 r-address[5]           HostAddress OPTIONAL
}

KrbCredInfo      ::=                    SEQUENCE {
                 key[0]                 EncryptionKey,
                 prealm[1]              Realm OPTIONAL,
                 pname[2]               PrincipalName OPTIONAL,
                 flags[3]               TicketFlags OPTIONAL,
                 authtime[4]            KerberosTime OPTIONAL,
                 starttime[5]           KerberosTime OPTIONAL,
                 endtime[6]             KerberosTime OPTIONAL
                 renew-till[7]          KerberosTime OPTIONAL,
                 srealm[8]              Realm OPTIONAL,
```

```
                sname[9]            PrincipalName OPTIONAL,
                caddr[10]           HostAddresses OPTIONAL
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_CRED.

tickets

These are the tickets obtained from the KDC specifically for use by the intended recipient. Successive tickets are paired with the corresponding KrbCredInfo sequence from the enc-part of the KRB-CRED message.

enc-part

This field holds an encoding of the EncKrbCredPart sequence encrypted under the session key shared between the sender and the intended recipient. This encrypted encoding is used for the enc-part field of the KRB-CRED message. See section 6 for the format of the ciphertext.

nonce

If practical, an application may require the inclusion of a nonce generated by the recipient of the message. If the same value is included as the nonce in the message, it provides evidence that the message is fresh and has not been replayed by an attacker. A nonce must never be re-used; it should be generated randomly by the recipient of the message and provided to the sender of the message in an application specific manner.

timestamp and usec

These fields specify the time that the KRB-CRED message was generated. The time is used to provide assurance that the message is fresh.

s-address and r-address

These fields are described above in section 5.6.1. They are used optionally to provide additional assurance of the integrity of the KRB-CRED message.

key

This field exists in the corresponding ticket passed by the KRB-CRED message and is used to pass the session key from the sender to the intended recipient. The field's encoding is described in section 6.2.

The following fields are optional. If present, they can be associated with the credentials in the remote ticket file. If left out, then it is assumed that the recipient of the credentials already knows their value.

prealm and pname

The name and realm of the delegated principal identity.

flags, authtime, starttime, endtime, renew-till, srealm, sname, and caddr

These fields contain the values of the corresponding fields from the ticket found in the ticket field. Descriptions of the fields are identical to the descriptions in the KDC-REP message.

5.9. Error message specification

This section specifies the format for the KRB_ERROR message. The fields included in the message are intended to return as much information as possible about an error. It is not expected that all the information required by the fields will be available for all types of errors. If the appropriate information is not available when the message is composed, the corresponding field will be left out of the message.

Note that since the KRB_ERROR message is only optionally integrity protected, it is quite possible for an intruder to synthesize or modify such a message. In particular, this means that unless appropriate integrity protection mechanisms have been applied to the KRB_ERROR message, the client should not use any fields in this message for security-critical purposes, such as setting a system clock or generating a fresh authenticator. The message can be useful, however, for advising a user on the reason for some failure.

5.9.1. KRB_ERROR definition

The KRB_ERROR message consists of the following fields:

```
KRB-ERROR ::=   [APPLICATION 30] SEQUENCE {
                pvno[0]                         INTEGER,
                msg-type[1]                     INTEGER,
                ctime[2]                        KerberosTime OPTIONAL,
                cusec[3]                        INTEGER OPTIONAL,
                stime[4]                        KerberosTime,
                susec[5]                        INTEGER,
                error-code[6]                   INTEGER,
                crealm[7]                       Realm OPTIONAL,
                cname[8]                        PrincipalName OPTIONAL,
                realm[9]                        Realm, -- Correct realm
                sname[10]                       PrincipalName, -- Correct name
                e-text[11]                      GeneralString OPTIONAL,
                e-data[12]                      OCTET STRING OPTIONAL,
                e-cksum[13]                     Checksum OPTIONAL,
}
```

pvno and msg-type

These fields are described above in section 5.4.1. msg-type is KRB_ERROR.

ctime

This field is described above in section 5.4.1.

cusec

This field is described above in section 5.5.2.

stime

This field contains the current time on the server. It is of type KerberosTime.

susec

This field contains the microsecond part of the server's timestamp. Its value ranges from 0 to 999999. It appears along with stime. The two fields are used in conjunction to specify a reasonably accurate timestamp.

error-code

This field contains the error code returned by Kerberos or the server when a request fails. To interpret the value of this field see the list of error codes in section 8. Implementations are encouraged to provide for national language support in the display of error messages.

crealm, cname, srealm and sname

These fields are described above in section 5.3.1.

e-text

This field contains additional text to help explain the error code associated with the failed request (for example, it might include a principal name which was unknown).

e-data

This field contains additional data about the error for use by the application to help it recover from or handle the error. If present, this field will contain the encoding of a sequence of TypedData (TYPED-DATA below), unless the errorcode is KDC_ERR_PREAUTH_REQUIRED, in which case it will contain the encoding of a sequence of of padata fields (METHOD-DATA below), each corresponding to an acceptable pre-authentication method and optionally containing data for the method:

```
TYPED-DATA   ::=   SEQUENCE of TypedData
METHOD-DATA  ::=   SEQUENCE of PA-DATA

TypedData ::=   SEQUENCE {
                    data-type[0]   INTEGER,
                    data-value[1]  OCTET STRING OPTIONAL
}
```

Note that the padata-type field in the PA-DATA structure and the data-type field in the TypedData structure share a common range of allocated values which are coordinated to avoid conflicts. One Kerberos error message, KDC_ERR_PREAUTH_REQUIRED, embeds elements of type PA-DATA, while all other error messages embed TypedData.

While preauthentication methods of type PA-DATA should be encapsulated within a TypedData element of type TD-PADATA, for compatibility with old clients, the KDC should include PA-DATA types below 22 directly as method-data. All new implementations interpreting the METHOD-DATA field for the KDC_ERR_PREAUTH_REQUIRED message must accept a type of TD-PADATA, extract the typed data field and interpret the use any elements encapsulated in the TD-PADATA elements as if they were present in the METHOD-DATA sequence.

Unless otherwise specified, unrecognized TypedData elements within the KRB-ERROR message MAY be ignored by implementations that do not support them. Note that all TypedData MAY be bound to the KRB-ERROR message via the checksum field.

An application may use the TD-APP-DEFINED-ERROR typed data type for data carried in a Kerberos error message that is specific to the application. TD-APP-SPECIFIC must set the data-type value of TypedData to TD-APP-SPECIFIC and the data-value field to

AppSpecificTypedData as follows:

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

The TD-REQ-NONCE TypedData MAY be used to bind a KRB-ERROR to a KDC-REQ. The data-value is an INTEGER that is equivalent to the nonce in a KDC-REQ.

The TD-REQ-SEQ TypedData MAY be used for binding a KRB-ERROR to the sequence number from an authenticator. The data-value is an INTEGER, and it is identical to sequence number sent in the authenticator.

The data-value for TD-KRB-PRINCIPAL is the Kerberos-defined PrincipalName. The data-value for TD-KRB-REALM is the Kerberos-defined Realm. These TypedData types MAY be used to indicate principal and realm name when appropriate.

e-cksum

This field contains an optional checksum for the KRB-ERROR message. The checksum is calculated over the Kerberos ASN.1 encoding of the KRB-ERROR message with the checksum absent. The checksum is then added to the KRB-ERROR structure and the message is re-encoded. The Checksum should be calculated using the session key from the ticket granting ticket or service ticket, where available. If the error is in response to a TGS or AP request, the checksum should be calculated using the the session key from the client's ticket. If the error is in response to an AS request, then the checksum should be calulated using the client's secret key ONLY if there has been suitable preauthentication to prove knowledge of the secret key by the client[33]. If a checksum can not be computed because the key to be used is not available, no checksum will be included.

6. Encryption and Checksum Specifications

This section is undergoing major revision to include rijndael support based on the Internet Draft by Ken Raeburn (draft-raeburn-krb-rijndael-krb-00.txt). The discussions of 3DES are also undergoing revision. Please see http://www.isi.edu/people/bcn/krb-revisions for the latest versions of this section when it becomes available.

7. Naming Constraints

7.1. Realm Names

Although realm names are encoded as GeneralStrings and although a realm can technically select any name it chooses, interoperability across realm boundaries requires agreement on how realm names are to be assigned, and what information they imply.
To enforce these conventions, each realm must conform to the conventions itself, and it must require that any realms with which inter-realm keys are shared also conform to the conventions and require the same from its neighbors.

Kerberos realm names are case sensitive. Realm names that differ only in the case of the characters are not equivalent. There are presently four styles of realm names: domain, X500, other, and reserved. Examples of each style follow:

```
     domain:   ATHENA.MIT.EDU (example)
       X500:   C=US/O=OSF (example)
      other:   NAMETYPE:rest/of.name=without-restrictions (example)
   reserved:   reserved, but will not conflict with above
```

Domain names must look like domain names: they consist of components separated by periods (.) and they contain neither colons (:) nor slashes (/). Though domain names themselves are case insensitive, in order for realms to match, the case must match as well. When establishing a new realm name based on an internet domain name it is recommended by convention that the characters be converted to upper case.

X.500 names contain an equal (=) and cannot contain a colon (:) before the equal. The realm names for X.500 names will be string representations of the names with components separated by slashes. Leading and trailing slashes will not be included. Note that the slash separator is consistent with Kerberos implementations based on RFC1510, but it is different from the separator recommended in RFC2253.

Names that fall into the other category must begin with a prefix that contains no equal (=) or period (.) and the prefix must be followed by a colon (:) and the rest of the name. All prefixes must be assigned before they may be used. Presently none are assigned.

The reserved category includes strings which do not fall into the first three categories. All names in this category are reserved. It is unlikely that names will be assigned to this category unless there is a very strong argument for not using the 'other' category.

These rules guarantee that there will be no conflicts between the various name styles. The following additional constraints apply to the assignment of realm names in the domain and X.500 categories: the name of a realm for the domain or X.500 formats must either be used by the organization owning (to whom it was assigned) an Internet domain name or X.500 name, or in the case that no such names are registered, authority to use a realm name may be derived from the authority of the parent realm. For example, if there is no domain name for E40.MIT.EDU, then the administrator of the MIT.EDU realm can authorize the creation of a realm with that name.

This is acceptable because the organization to which the parent is assigned is presumably the organization authorized to assign names to its children in the X.500 and domain name systems as well. If the parent assigns a realm name without also registering it in the domain name or X.500 hierarchy, it is the parent's responsibility to make sure that there will not in the future exist a name identical to the realm name of the child unless it is assigned to the same entity as the realm name.

7.2. Principal Names

As was the case for realm names, conventions are needed to ensure that all agree on what information is implied by a principal name. The name-type field that is part of the principal name indicates the kind of information implied by the name. The name-type should be treated as a hint. Ignoring the name type, no two names can be the same (i.e. at least one of the components, or the realm, must be different). The following name types are defined:

name-type	value	meaning
NT-UNKNOWN	0	Name type not known
NT-PRINCIPAL	1	General principal name (e.g. username, or DCE principal)
NT-SRV-INST	2	Service and other unique instance (krbtgt)
NT-SRV-HST	3	Service with host name as instance (telnet, rcommands)
NT-SRV-XHST	4	Service with slash-separated host name components
NT-UID	5	Unique ID
NT-X500-PRINCIPAL	6	Encoded X.509 Distingished name [RFC 1779]
NT-SMTP-NAME	7	Name in form of SMTP email name (e.g. user@foo.com)

When a name implies no information other than its uniqueness at a particular time the name type PRINCIPAL should be used. The principal name type should be used for users, and it might also be used for a unique server. If the name is a unique machine generated ID that is guaranteed never to be reassigned then the name type of UID should be used (note that it is generally a bad idea to reassign names of any type since stale entries might remain in access control lists).

If the first component of a name identifies a service and the remaining components identify an instance of the service in a server specified manner, then the name type of SRV-INST should be used. An example of this name type is the Kerberos ticket-granting service whose name has a first component of krbtgt and a second component identifying the realm for which the ticket is valid.

If instance is a single component following the service name and the instance identifies the host on which the server is running, then the name type SRV-HST should be used. This type is typically used for Internet services such as telnet and the Berkeley R commands. If the separate components of the host name appear as successive components following the name of the service, then the name type SRV-XHST should be used. This type might be used to identify servers on hosts with X.500 names where the slash (/) might otherwise be ambiguous.

A name type of NT-X500-PRINCIPAL should be used when a name from an X.509 certificate is translated into a Kerberos name. The encoding of the X.509 name as a Kerberos principal shall conform to the encoding rules specified in RFC 2253.
A name type of SMTP allows a name to be of a form that resembles a SMTP email name. This name, including an "@" and a domain name, is used as the one component of the principal name. This name type can be used in conjunction with name-canonicalization to allow a free-form of email address to be specified as a client name and allow the KDC to determine the Kerberos principal name for the requested name. [JBrezak, Raeburn]

A name type of UNKNOWN should be used when the form of the name is not known. When comparing names, a name of type UNKNOWN will match principals authenticated with names of any type. A principal authenticated with a name of type UNKNOWN, however, will only match other names of type UNKNOWN.

Names of any type with an initial component of 'krbtgt' are reserved for the Kerberos ticket granting service. See section 8.2.3 for the form of such names.

7.2.1. Name of server principals

The principal identifier for a server on a host will generally be composed of two parts: (1) the realm of the KDC with which the server is registered, and (2) a two-component name of type NT-SRV-HST if the host name is an Internet domain name or a multi-component name of type NT-SRV-XHST if the name of the host is of a form such as X.500 that allows slash (/) separators. The first component of the two- or multi-component name will identify the service and the latter components will identify the host. Where the name of the host is not case sensitive (for example, with Internet domain names) the name of the host must be lower case. If specified by the application protocol for services such as telnet and the Berkeley R commands which run with system privileges, the first component may be the string 'host' instead of a service specific identifier. When a host has an official name and one or more aliases and the official name can be reliably determined, the official name of the host should be used when constructing the name of the server principal.

8. Constants and other defined values

8.1. Host address types

All negative values for the host address type are reserved for local use. All non-negative values are reserved for officially assigned type fields and interpretations.

The values of the types for the following addresses are chosen to match the defined address family constants in the Berkeley Standard Distributions of Unix. They can be found in with symbolic names AF_xxx (where xxx is an abbreviation of the address family name).

Internet (IPv4) Addresses

Internet (IPv4) addresses are 32-bit (4-octet) quantities, encoded in MSB order. The IPv4 loopback address should not appear in a Kerberos packet. The type of IPv4 addresses is two (2).

Internet (IPv6) Addresses [Westerlund]

IPv6 addresses are 128-bit (16-octet) quantities, encoded in MSB order. The type of IPv6 addresses is twenty-four (24). [RFC1883] [RFC1884]. The following addresses (see [RFC1884]) MUST not appear in any Kerberos packet:

* the Unspecified Address
* the Loopback Address
* Link-Local addresses

IPv4-mapped IPv6 addresses MUST be represented as addresses of type 2.

CHAOSnet addresses

CHAOSnet addresses are 16-bit (2-octet) quantities, encoded in MSB order. The type of CHAOSnet addresses is five (5).

ISO addresses

ISO addresses are variable-length. The type of ISO addresses is seven (7).

Xerox Network Services (XNS) addresses

XNS addresses are 48-bit (6-octet) quantities, encoded in MSB order. The type of XNS addresses is six (6).

AppleTalk Datagram Delivery Protocol (DDP) addresses

AppleTalk DDP addresses consist of an 8-bit node number and a 16-bit network number. The first octet of the address is the node number; the remaining two octets encode the network number in MSB order. The type of AppleTalk DDP addresses is sixteen (16).

DECnet Phase IV addresses

DECnet Phase IV addresses are 16-bit addresses, encoded in LSB order. The type of DECnet Phase IV addresses is twelve (12).

Netbios addresses

Netbios addresses are 16-octet addresses typically composed of 1 to 15 characters, trailing blank (ascii char 20) filled, with a 16th octet of 0x0. The type of Netbios addresses is 20 (0x14).

8.2. KDC messages

8.2.1. UDP/IP transport

When contacting a Kerberos server (KDC) for a KRB_KDC_REQ request using UDP IP transport, the client shall send a UDP datagram containing only an encoding of the request to port 88 (decimal) at the KDC's IP address; the KDC will respond with a reply datagram containing only an encoding of the reply message (either a KRB_ERROR or a KRB_KDC_REP) to the sending port at the sender's IP address. Kerberos servers supporting IP transport must accept UDP requests on port 88 (decimal). The response to a request made through UDP/IP transport must also use UDP/IP transport.

8.2.2. TCP/IP transport [Westerlund,Danielsson]

Kerberos servers (KDC's) should accept TCP requests on port 88 (decimal) and clients should support the sending of TCP requests on port 88 (decimal). When the KRB_KDC_REQ message is sent to the KDC over a TCP stream, a new connection will be established for each authentication exchange (request and response). The KRB_KDC_REP or KRB_ERROR message will be returned to the client on the same TCP stream that was established for the request. The response to a request made through TCP/IP transport must also use TCP/IP transport. Implementors should note that some extensions to the Kerberos protocol will not work if any implementation not supporting the TCP transport is involved (client or KDC). Implementors are strongly urged to support the TCP transport on both the client and server and are advised that the current notation of "should" support will likely change in the future to must support. The KDC may close the TCP stream after sending a response, but may leave the stream open if it expects a followup - in which case it may close the stream at any time if resource constraints or other factors make it desirable to do so. Care must be taken in managing TCP/IP connections with the KDC to prevent denial of service attacks based on the number of TCP/IP connections with the KDC that remain open. If multiple exchanges with the KDC are needed for certain forms of preauthentication, multiple TCP connections may be required. A client may close the stream after receiving response, and should close the stream if it does not expect to send followup messages. The client must be prepared to have the stream closed by the KDC at anytime, in which case it must simply connect again when it is ready to send subsequent messages.

The first four octets of the TCP stream used to transmit the request request will encode in network byte order the length of the request (KRB_KDC_REQ), and the length will be followed by the request itself. The response will similarly be preceded by a 4 octet encoding in network byte order of the length of the KRB_KDC_REP or the KRB_ERROR message and will be followed by the KRB_KDC_REP or the KRB_ERROR response. If the sign bit is set on the integer represented by the first 4 octets, then the next 4 octets will be read, extending the length of the field by another 4 octets (less the sign bit of the additional four octets which is reserved for future expansion and which at present must be zero).

8.2.3. OSI transport

During authentication of an OSI client to an OSI server, the mutual authentication of an OSI server to an OSI client, the transfer of credentials from an OSI client to an OSI server, or during exchange of private or integrity checked messages, Kerberos protocol messages may be treated as opaque objects and the type of the authentication mechanism will be:

OBJECT IDENTIFIER ::= {iso (1), org(3), dod(6),internet(1), security(5),kerberosv5(2)}

Depending on the situation, the opaque object will be an authentication header (KRB_AP_REQ), an authentication reply (KRB_AP_REP), a safe message (KRB_SAFE), a private message (KRB_PRIV), or a credentials message (KRB_CRED). The opaque data contains an application code as specified in the ASN.1 description for each message. The application code may be used by Kerberos to determine the message type.

8.2.3. Name of the TGS

The principal identifier of the ticket-granting service shall be composed of three parts: (1) the realm of the KDC issuing the TGS ticket (2) a two-part name of type NT-SRV-INST, with the first part "krbtgt" and the second part the name of the realm which will accept the ticket-granting ticket. For example, a ticket-granting ticket issued by the ATHENA.MIT.EDU realm to be used to get tickets from the ATHENA.MIT.EDU KDC has a principal identifier of "ATHENA.MIT.EDU" (realm), ("krbtgt", "ATHENA.MIT.EDU") (name). A ticket-granting ticket issued by the ATHENA.MIT.EDU realm to be used to get tickets from the MIT.EDU realm has a principal identifier of "ATHENA.MIT.EDU" (realm), ("krbtgt", "MIT.EDU") (name).

8.3. Protocol constants and associated values

The following tables list constants used in the protocol and define their meanings. Ranges are specified in the "specification" section that limit the values of constants for which values are defined here. This allows implementations to make assumptions about the maximum values that will be received for these constants. Implementation receiving values outside the range specified in the "specification" section may reject the request, but they must recover cleanly.

Encryption type	etype value	block size	minimum pad size	confounder size
NULL	0	1	0	0
des-cbc-crc	1	8	4	8
des-cbc-md4	2	8	0	8
des-cbc-md5	3	8	0	8
[reserved]	4			
des3-cbc-md5	5	8	0	8
[reserved]	6			
des3-cbc-sha1	7	8	0	8
dsaWithSHA1-CmsOID	9			(pkinit)
md5WithRSAEncryption-CmsOID	10			(pkinit)
sha1WithRSAEncryption-CmsOID	11			(pkinit)
rc2CBC-EnvOID	12			(pkinit)
rsaEncryption-EnvOID	13			(pkinit from PKCS#1 v1.5)
rsaES-OAEP-ENV-OID	14			(pkinit from PKCS#1 v2.0)
des-ede3-cbc-Env-OID	15			(pkinit)
des3-cbc-sha1-kd	16			(Tom Yu)
rc4-hmac	23			(swift)
rc4-hmac-exp	24			(swift)
subkey-keynaterial	65			(opaque mhur)

[reserved]	0x8003

Checksum type	sumtype value	checksum size
CRC32	1	4
rsa-md4	2	16
rsa-md4-des	3	24
des-mac	4	16
des-mac-k	5	8
rsa-md4-des-k	6	16 (drop rsa ?)
rsa-md5	7	16 (drop rsa ?)
rsa-md5-des	8	24 (drop rsa ?)
rsa-md5-des3	9	24 (drop rsa ?)
hmac-sha1-des3-kd	12	20
hmac-sha1-des3	13	20
sha1 (unkeyed)	14	20

padata and data types	padata-type value	comment
PA-TGS-REQ	1	
PA-ENC-TIMESTAMP	2	
PA-PW-SALT	3	
[reserved]	4	
PA-ENC-UNIX-TIME	5	(depricated)
PA-SANDIA-SECUREID	6	
PA-SESAME	7	
PA-OSF-DCE	8	
PA-CYBERSAFE-SECUREID	9	
PA-AFS3-SALT	10	
PA-ETYPE-INFO	11	
PA-SAM-CHALLENGE	12	(sam/otp)
PA-SAM-RESPONSE	13	(sam/otp)
PA-PK-AS-REQ	14	(pkinit)
PA-PK-AS-REP	15	(pkinit)
PA-USE-SPECIFIED-KVNO	20	
PA-SAM-REDIRECT	21	(sam/otp)
PA-GET-FROM-TYPED-DATA	22	(embedded in typed data)
TD-PADATA	22	(embeds padata)
PA-SAM-ETYPE-INFO	23	(sam/otp)
TD-PKINIT-CMS-CERTIFICATES	101	CertificateSet from CMS

TD-KRB-PRINCIPAL	102	PrincipalName (see Sec.5.9.1)
TD-KRB-REALM	103	Realm (see Sec.5.9.1)
TD-TRUSTED-CERTIFIERS	104	from PKINIT
TD-CERTIFICATE-INDEX	105	from PKINIT
TD-APP-DEFINED-ERROR	106	application specific (see Sec.5.9.1)
TD-REQ-NONCE	107	INTEGER (see Sec.5.9.1)
TD-REQ-SEQ	108	INTEGER (see Sec.5.9.1)

authorization data type	ad-type value	
AD-IF-RELEVANT	1	
AD-INTENDED-FOR-SERVER	2	
AD-INTENDED-FOR-APPLICATION-CLASS	3	
AD-KDC-ISSUED	4	
AD-OR	5	
AD-MANDATORY-TICKET-EXTENSIONS	6	
AD-IN-TICKET-EXTENSIONS	7	
reserved values	8-63	
OSF-DCE	64	
SESAME	65	
AD-OSF-DCE-PKI-CERTID	66	(hemsath@us.ibm.com)
AD-WIN200-PAC	128	(jbrezak@exchange.microsoft.com)

Ticket Extension Types

TE-TYPE-NULL	0	Null ticket extension
TE-TYPE-EXTERNAL-ADATA	1	Integrity protected authorization data
[reserved]	2	TE-TYPE-PKCROSS-KDC (I have reservations)
TE-TYPE-PKCROSS-CLIENT	3	PKCROSS cross realm key ticket
TE-TYPE-CYBERSAFE-EXT	4	Assigned to CyberSafe Corp
[reserved]	5	TE-TYPE-DEST-HOST (I have reservations)

alternate authentication type	method-type value
reserved values	0-63
ATT-CHALLENGE-RESPONSE	64

transited encoding type	tr-type value
DOMAIN-X500-COMPRESS	1
reserved values	all others

Label	Value	Meaning or MIT code
pvno	5	current Kerberos protocol version number
message types		
KRB_AS_REQ	10	Request for initial authentication
KRB_AS_REP	11	Response to KRB_AS_REQ request
KRB_TGS_REQ	12	Request for authentication based on TGT
KRB_TGS_REP	13	Response to KRB_TGS_REQ request
KRB_AP_REQ	14	application request to server
KRB_AP_REP	15	Response to KRB_AP_REQ_MUTUAL
KRB_SAFE	20	Safe (checksummed) application message
KRB_PRIV	21	Private (encrypted) application message
KRB_CRED	22	Private (encrypted) message to forward credentials
KRB_ERROR	30	Error response
name types		
KRB_NT_UNKNOWN	0	Name type not known
KRB_NT_PRINCIPAL	1	Just the name of the principal as in DCE, or for users
KRB_NT_SRV_INST	2	Service and other unique instance (krbtgt)
KRB_NT_SRV_HST	3	Service with host name as instance (telnet, rcommands)
KRB_NT_SRV_XHST	4	Service with host as remaining components
KRB_NT_UID	5	Unique ID
KRB_NT_X500_PRINCIPAL	6	Encoded X.509 Distingished name [RFC 2253]

error codes

KDC_ERR_NONE	0	No error
KDC_ERR_NAME_EXP	1	Client's entry in database has expired
KDC_ERR_SERVICE_EXP	2	Server's entry in database has expired
KDC_ERR_BAD_PVNO	3	Requested protocol version number not supported
KDC_ERR_C_OLD_MAST_KVNO	4	Client's key encrypted in old master key
KDC_ERR_S_OLD_MAST_KVNO	5	Server's key encrypted in old master key
KDC_ERR_C_PRINCIPAL_UNKNOWN	6	Client not found in Kerberos database
KDC_ERR_S_PRINCIPAL_UNKNOWN	7	Server not found in Kerberos database
KDC_ERR_PRINCIPAL_NOT_UNIQUE	8	Multiple principal entries in database
KDC_ERR_NULL_KEY	9	The client or server has a null key
KDC_ERR_CANNOT_POSTDATE	10	Ticket not eligible for postdating
KDC_ERR_NEVER_VALID	11	Requested start time is later than end time
KDC_ERR_POLICY	12	KDC policy rejects request
KDC_ERR_BADOPTION	13	KDC cannot accommodate requested option
KDC_ERR_ETYPE_NOSUPP	14	KDC has no support for encryption type
KDC_ERR_SUMTYPE_NOSUPP	15	KDC has no support for checksum type
KDC_ERR_PADATA_TYPE_NOSUPP	16	KDC has no support for padata type
KDC_ERR_TRTYPE_NOSUPP	17	KDC has no support for transited type
KDC_ERR_CLIENT_REVOKED	18	Clients credentials have been revoked
KDC_ERR_SERVICE_REVOKED	19	Credentials for server have been revoked
KDC_ERR_TGT_REVOKED	20	TGT has been revoked
KDC_ERR_CLIENT_NOTYET	21	Client not yet valid - try again later
KDC_ERR_SERVICE_NOTYET	22	Server not yet valid - try again later
KDC_ERR_KEY_EXPIRED	23	Password has expired - change password to reset
KDC_ERR_PREAUTH_FAILED	24	Pre-authentication information was invalid
KDC_ERR_PREAUTH_REQUIRED	25	Additional pre-authenticationrequired [40]
KDC_ERR_SERVER_NOMATCH	26	Requested server and ticket don't match
KDC_ERR_MUST_USE_USER2USER	27	Server principal valid for user2user only
KDC_ERR_PATH_NOT_ACCPETED	28	KDC Policy rejects transited path
KDC_ERR_SVC_UNAVAILABLE	29	A service is not available
KRB_AP_ERR_BAD_INTEGRITY	31	Integrity check on decrypted field failed
KRB_AP_ERR_TKT_EXPIRED	32	Ticket expired
KRB_AP_ERR_TKT_NYV	33	Ticket not yet valid
KRB_AP_ERR_REPEAT	34	Request is a replay
KRB_AP_ERR_NOT_US	35	The ticket isn't for us

```
KRB_AP_ERR_BADMATCH            36   Ticket and authenticator don't match
KRB_AP_ERR_SKEW                37   Clock skew too great
KRB_AP_ERR_BADADDR             38   Incorrect net address
KRB_AP_ERR_BADVERSION          39   Protocol version mismatch
KRB_AP_ERR_MSG_TYPE            40   Invalid msg type
KRB_AP_ERR_MODIFIED            41   Message stream modified
KRB_AP_ERR_BADORDER            42   Message out of order
KRB_AP_ERR_BADKEYVER           44   Specified version of key is not available
KRB_AP_ERR_NOKEY               45   Service key not available
KRB_AP_ERR_MUT_FAIL            46   Mutual authentication failed
KRB_AP_ERR_BADDIRECTION        47   Incorrect message direction
KRB_AP_ERR_METHOD              48   Alternative authentication method required
KRB_AP_ERR_BADSEQ              49   Incorrect sequence number in message
KRB_AP_ERR_INAPP_CKSUM         50   Inappropriate type of checksum in message
KRB_AP_PATH_NOT_ACCEPTED       51   Policy rejects transited path
KRB_ERR_RESPONSE_TOO_BIG       52   Response too big for UDP, retry with TCP
KRB_ERR_GENERIC                60   Generic error (description in e-text)
KRB_ERR_FIELD_TOOLONG          61   Field is too long for this implementation
KDC_ERROR_CLIENT_NOT_TRUSTED        62 (pkinit)
KDC_ERROR_KDC_NOT_TRUSTED           63 (pkinit)


KDC_ERROR_INVALID_SIG               64 (pkinit)
KDC_ERR_KEY_TOO_WEAK                65 (pkinit)
KDC_ERR_CERTIFICATE_MISMATCH        66 (pkinit)
KRB_AP_ERR_NO_TGT                   67 (user-to-user)
KDC_ERR_WRONG_REALM                 68 (user-to-user)
KRB_AP_ERR_USER_TO_USER_REQUIRED    69 (user-to-user)
KDC_ERR_CANT_VERIFY_CERTIFICATE     70 (pkinit)
KDC_ERR_INVALID_CERTIFICATE         71 (pkinit)
KDC_ERR_REVOKED_CERTIFICATE         72 (pkinit)
KDC_ERR_REVOCATION_STATUS_UNKNOWN   73 (pkinit)
KDC_ERR_REVOCATION_STATUS_UNAVAILABLE   74 (pkinit)
KDC_ERR_CLIENT_NAME_MISMATCH        75 (pkinit)
KDC_ERR_KDC_NAME_MISMATCH           76 (pkinit)
```

9. Interoperability requirements

Version 5 of the Kerberos protocol supports a myriad of options. Among these are multiple encryption and checksum types, alternative encoding schemes for the transited field, optional mechanisms for pre-authentication, the handling of tickets with no addresses, options for mutual authentication, user to user authentication, support for proxies, forwarding, postdating, and renewing tickets, the format of realm names, and the handling of authorization data.

In order to ensure the interoperability of realms, it is necessary to define a minimal configuration which must be supported by all implementations. This minimal configuration is subject to change as technology does. For example, if at some later date it is discovered that one of the required encryption or checksum algorithms is not secure, it will be replaced.

9.1. Specification 2

This section defines the second specification of these options. Implementations which are configured in this way can be said to support Kerberos Version 5 Specification 2 (5.1). Specification 1 (deprecated) may be found in RFC1510.

Transport

TCP/IP and UDP/IP transport must be supported by KDCs claiming conformance to specification 2. Kerberos clients claiming conformance to specification 2 must support UDP/IP transport for messages with the KDC and should support TCP/IP transport.

Encryption and checksum methods

The following encryption and checksum mechanisms must be supported. Implementations may support other mechanisms as well, but the additional mechanisms may only be used when communicating with principals known to also support them: This list is to be determined.

Encryption: DES-CBC-MD5, DES3-CBC-SHA1-KD, RIJNDAEL(decide identifier)
Checksums: CRC-32, DES-MAC, DES-MAC-K, DES-MD5, HMAC-SHA1-DES3-KD

Realm Names

All implementations must understand hierarchical realms in both the Internet Domain and the X.500 style. When a ticket granting ticket for an unknown realm is requested, the KDC must be able to determine the names of the intermediate realms between the KDCs realm and the requested realm.

Transited field encoding

DOMAIN-X500-COMPRESS (described in section 3.3.3.2) must be supported. Alternative encodings may be supported, but they may be used only when that encoding is supported by ALL intermediate realms.

Pre-authentication methods

The TGS-REQ method must be supported. The TGS-REQ method is not used on the initial request. The PA-ENC-TIMESTAMP method must be supported by clients but whether it is enabled by default may be determined on a realm by realm basis. If not used in the initial request and the error KDC_ERR_PREAUTH_REQUIRED is returned specifying PA-ENC-TIMESTAMP as an acceptable method, the client should retry the initial request using the PA-ENC-TIMESTAMP preauthentication method. Servers need not support the PA-ENC-TIMESTAMP method, but if not supported the server should ignore the presence of PA-ENC-TIMESTAMP pre-authentication in a request.

Mutual authentication

Mutual authentication (via the KRB_AP_REP message) must be supported.

Ticket addresses and flags

All KDC's must pass through tickets that carry no addresses (i.e. if a TGT contains no addresses, the KDC will return derivative tickets), but each realm may set its own policy for issuing such tickets, and each application server will set its own policy with respect to accepting them.

Proxies and forwarded tickets must be supported. Individual realms and application servers can set their own policy on when such tickets will be accepted.

All implementations must recognize renewable and postdated tickets, but need not actually implement them. If these options are not supported, the starttime and endtime in the ticket shall specify a ticket's entire useful life. When a postdated ticket is decoded by a server, all implementations shall make the presence of the postdated flag visible to the calling server.

User-to-user authentication

Support for user to user authentication (via the ENC-TKT-IN-SKEY KDC option) must be provided by implementations, but individual realms may decide as a matter of policy to reject such requests on a per-principal or realm-wide basis.

Authorization data

Implementations must pass all authorization data subfields from ticket-granting tickets to any derivative tickets unless directed to suppress a subfield as part of the definition of that registered subfield type (it is never incorrect to pass on a subfield, and no registered subfield types presently specify suppression at the KDC).

Implementations must make the contents of any authorization data subfields available to the server when a ticket is used. Implementations are not required to allow clients to specify the contents of the authorization data fields.

Constant ranges

All protocol constants are constrained to 32 bit (signed) values unless further constrained by the protocol definition. This limit is provided to allow implementations to make assumptions about the maximum values that will be received for these constants. Implementation receiving values outside this range may reject the request, but they must recover cleanly.

9.2. Recommended KDC values

Following is a list of recommended values for a KDC implementation, based on the list of suggested configuration constants (see section 4.4).

minimum lifetime	5 minutes
maximum renewable lifetime	1 week
maximum ticket lifetime	1 day
empty addresses	only when suitable restrictions appear in authorization data
proxiable, etc.	Allowed.

10. REFERENCES

[NT94] B. Clifford Neuman and Theodore Y. Ts'o, "An Authentication Service for Computer Networks," IEEE Communications Magazine, Vol. 32(9), pp. 33-38 (September 1994).

[MNSS87] S. P. Miller, B. C. Neuman, J. I. Schiller, and J. H. Saltzer, Section E.2.1: Kerberos Authentication and Authorization System, M.I.T. Project Athena, Cambridge, Massachusetts (December 21, 1987).

[SNS88] J. G. Steiner, B. C. Neuman, and J. I. Schiller, "Kerberos: An Authentication Service for Open Network Systems," pp. 191-202 in Usenix Conference Proceedings, Dallas, Texas (February, 1988).

[NS78] Roger M. Needham and Michael D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," Communications of the ACM, Vol. 21(12), pp. 993-999 (December, 1978).

[DS81] Dorothy E. Denning and Giovanni Maria Sacco, "Timestamps in Key Distribution Protocols," Communications of the ACM, Vol. 24(8), pp. 533-536 (August 1981).

[KNT92] John T. Kohl, B. Clifford Neuman, and Theodore Y. Ts'o, "The Evolution of the Kerberos Authentication Service," in an IEEE Computer Society Text soon to be published (June 1992).

[Neu93] B. Clifford Neuman, "Proxy-Based Authorization and Accounting for Distributed Systems," in Proceedings of the 13th International Conference on Distributed Computing Systems, Pittsburgh, PA (May, 1993).

[DS90] Don Davis and Ralph Swick, "Workstation Services and Kerberos Authentication at Project Athena," Technical Memorandum TM-424, MIT Laboratory for Computer Science (February 1990).

[LGDSR87] P. J. Levine, M. R. Gretzinger, J. M. Diaz, W. E. Sommerfeld, and K. Raeburn, Section E.1: Service Management System, M.I.T. Project Athena, Cambridge, Massachusetts (1987).

[X509-88] CCITT, Recommendation X.509: The Directory Authentication Framework, December 1988.

[Pat92]. J. Pato, Using Pre-Authentication to Avoid Password Guessing Attacks, Open Software Foundation DCE Request for Comments 26 (December 1992).

[DES77] National Bureau of Standards, U.S. Department of Commerce, "Data Encryption Standard," Federal Information Processing Standards Publication 46, Washington, DC (1977).

[DESM80] National Bureau of Standards, U.S. Department of Commerce, "DES Modes of Operation," Federal Information Processing Standards Publication 81, Springfield, VA (December 1980).

[SG92] Stuart G. Stubblebine and Virgil D. Gligor, "On Message Integrity in Cryptographic Protocols," in Proceedings of the IEEE Symposium on Research in Security and Privacy, Oakland, California (May 1992).

[IS3309] International Organization for Standardization, "ISO Information Processing Systems - Data Communication - High-Level Data Link Control Procedure - Frame Structure," IS 3309 (October 1984). 3rd Edition.

[MD4-92] R. Rivest, "The MD4 Message Digest Algorithm," RFC 1320, MIT Laboratory for Computer Science (April 1992).

[MD5-92] R. Rivest, "The MD5 Message Digest Algorithm," RFC 1321, MIT Laboratory for Computer Science (April 1992).

[KBC96] H. Krawczyk, M. Bellare, and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication," Working Draft draft-ietf-ipsec-hmac-md5-01.txt, (August 1996).

[Horowitz96] Horowitz, M., "Key Derivation for Authentication, Integrity, and Privacy", draft-horowitz-key-derivation-02.txt, August 1998.

[HorowitzB96] Horowitz, M., "Key Derivation for Kerberos V5", draft-
horowitz-kerb-key-derivation-01.txt, September 1998.

[Krawczyk96] Krawczyk, H., Bellare, and M., Canetti, R., "HMAC:
Keyed-Hashing for Message Authentication", =
draft-ietf-ipsec-hmac-
md5-01.txt, August, 1996.

A. Pseudo-code for protocol processing

This appendix provides pseudo-code describing how the messages are to be constructed and interpreted by clients and servers.

A.1. KRB_AS_REQ generation

```
        request.pvno :=3D protocol version; /* pvno =3D 5 */
        request.msg-type :=3D message type; /* type =3D KRB_AS_REQ */

        if(pa_enc_timestamp_required) then
                request.padata.padata-type =3D PA-ENC-TIMESTAMP;
                get system_time;
                padata-body.patimestamp,pausec =3D system_time;
                encrypt padata-body into request.padata.padata-value
                        using client.key; /* derived from password */
        endif

        body.kdc-options :=3D users's preferences;
        body.cname :=3D user's name;
        body.realm :=3D user's realm;
        body.sname :=3D service's name; /* usually "krbtgt",  =
"localrealm" */
        if (body.kdc-options.POSTDATED is set) then
                body.from :=3D requested starting time;
        else
                omit body.from;
```

```
        endif
        body.till :=3D requested end time;
        if (body.kdc-options.RENEWABLE is set) then
                body.rtime :=3D requested final renewal time;
        endif
        body.nonce :=3D random_nonce();
        body.etype :=3D requested etypes;
        if (user supplied addresses) then
                body.addresses :=3D user's addresses;
        else
                omit body.addresses;
        endif
        omit body.enc-authorization-data;
        request.req-body :=3D body;

        kerberos :=3D lookup(name of local kerberos server (or =
servers));
        send(packet,kerberos);

        wait(for response);
        if (timed_out) then
                retry or use alternate server;
        endif
```

A.2. KRB_AS_REQ verification and KRB_AS_REP generation

```
        decode message into req;

        client :=3D lookup(req.cname,req.realm);
        server :=3D lookup(req.sname,req.realm);


        get system_time;
        kdc_time :=3D system_time.seconds;

        if (!client) then
                /* no client in Database */
```

```
                error_out(KDC_ERR_C_PRINCIPAL_UNKNOWN);
        endif
        if (!server) then
                /* no server in Database */
                error_out(KDC_ERR_S_PRINCIPAL_UNKNOWN);
        endif

        if(client.pa_enc_timestamp_required and
           pa_enc_timestamp not present) then
                error_out(KDC_ERR_PREAUTH_REQUIRED(PA_ENC_TIMESTAMP));
        endif

        if(pa_enc_timestamp present) then
                decrypt req.padata-value into decrypted_enc_timestamp
                        using client.key;
                        using auth_hdr.authenticator.subkey;
                if (decrypt_error()) then
                        error_out(KRB_AP_ERR_BAD_INTEGRITY);
                if(decrypted_enc_timestamp is not within allowable skew) =
then
                        error_out(KDC_ERR_PREAUTH_FAILED);
                endif
                if(decrypted_enc_timestamp and usec is replay)
                        error_out(KDC_ERR_PREAUTH_FAILED);
                endif
                add decrypted_enc_timestamp and usec to replay cache;
        endif

        use_etype :=3D first supported etype in req.etypes;

        if (no support for req.etypes) then
                error_out(KDC_ERR_ETYPE_NOSUPP);
        endif

        new_tkt.vno :=3D ticket version; /* =3D 5 */
        new_tkt.sname :=3D req.sname;
        new_tkt.srealm :=3D req.srealm;
```

```
reset all flags in new_tkt.flags;

/* It should be noted that local policy may affect the  */
/* processing of any of these flags.  For example, some */
/* realms may refuse to issue renewable tickets         */

if (req.kdc-options.FORWARDABLE is set) then
        set new_tkt.flags.FORWARDABLE;
endif
if (req.kdc-options.PROXIABLE is set) then
        set new_tkt.flags.PROXIABLE;
endif

if (req.kdc-options.ALLOW-POSTDATE is set) then
        set new_tkt.flags.MAY-POSTDATE;
endif



if ((req.kdc-options.RENEW is set) or
    (req.kdc-options.VALIDATE is set) or
    (req.kdc-options.PROXY is set) or
    (req.kdc-options.FORWARDED is set) or
    (req.kdc-options.ENC-TKT-IN-SKEY is set)) then
        error_out(KDC_ERR_BADOPTION);
endif

new_tkt.session :=3D random_session_key();
new_tkt.cname :=3D req.cname;
new_tkt.crealm :=3D req.crealm;
new_tkt.transited :=3D empty_transited_field();

new_tkt.authtime :=3D kdc_time;

if (req.kdc-options.POSTDATED is set) then
   if (against_postdate_policy(req.from)) then
        error_out(KDC_ERR_POLICY);
```

```
            endif
            set new_tkt.flags.POSTDATED;
            set new_tkt.flags.INVALID;
            new_tkt.starttime :=3D req.from;
         else
            omit new_tkt.starttime; /* treated as authtime when omitted =
*/
         endif
         if (req.till =3D 0) then
                 till :=3D infinity;
         else
                 till :=3D req.till;
         endif

         new_tkt.endtime :=3D min(till,
                              new_tkt.starttime+client.max_life,
                              new_tkt.starttime+server.max_life,
                              new_tkt.starttime+max_life_for_realm);

         if ((req.kdc-options.RENEWABLE-OK is set) and
             (new_tkt.endtime < req.till)) then
                 /* we set the RENEWABLE option for later processing */
                 set req.kdc-options.RENEWABLE;
                 req.rtime :=3D req.till;
         endif

         if (req.rtime =3D 0) then
                 rtime :=3D infinity;
         else
                 rtime :=3D req.rtime;
         endif

         if (req.kdc-options.RENEWABLE is set) then
                 set new_tkt.flags.RENEWABLE;
                 new_tkt.renew-till :=3D min(rtime,
                                          new_tkt.starttime+client.max_rlife,
                                          new_tkt.starttime+server.max_rlife,
```

```
                                        new_tkt.starttime+max_rlife_for_realm);


        else
                omit new_tkt.renew-till; /* only present if RENEWABLE */
        endif

        if (req.addresses) then
                new_tkt.caddr :=3D req.addresses;
        else
                omit new_tkt.caddr;
        endif

        new_tkt.authorization_data :=3D empty_authorization_data();

        encode to-be-encrypted part of ticket into OCTET STRING;
        new_tkt.enc-part :=3D encrypt OCTET STRING
                using etype_for_key(server.key), server.key, =
server.p_kvno;

        /* Start processing the response */

        resp.pvno :=3D 5;
        resp.msg-type :=3D KRB_AS_REP;
        resp.cname :=3D req.cname;
        resp.crealm :=3D req.realm;
        resp.ticket :=3D new_tkt;

        resp.key :=3D new_tkt.session;
        resp.last-req :=3D fetch_last_request_info(client);
        resp.nonce :=3D req.nonce;
        resp.key-expiration :=3D client.expiration;
        resp.flags :=3D new_tkt.flags;
```

```
        resp.authtime :=3D new_tkt.authtime;
        resp.starttime :=3D new_tkt.starttime;
        resp.endtime :=3D new_tkt.endtime;

        if (new_tkt.flags.RENEWABLE) then
                resp.renew-till :=3D new_tkt.renew-till;
        endif

        resp.realm :=3D new_tkt.realm;
        resp.sname :=3D new_tkt.sname;

        resp.caddr :=3D new_tkt.caddr;

        encode body of reply into OCTET STRING;

        resp.enc-part :=3D encrypt OCTET STRING
                         using use_etype, client.key, client.p_kvno;
        send(resp);
```

A.3. KRB_AS_REP verification

```
        decode response into resp;

        if (resp.msg-type =3D KRB_ERROR) then
                if(error =3D KDC_ERR_PREAUTH_REQUIRED(PA_ENC_TIMESTAMP)) =
then
                        set pa_enc_timestamp_required;
                        goto KRB_AS_REQ;
                endif
                process_error(resp);
                return;
        endif

        /* On error, discard the response, and zero the session key */
        /* from the response immediately */
```

```
        key =3D get_decryption_key(resp.enc-part.kvno, =
resp.enc-part.etype,
                                 resp.padata);
        unencrypted part of resp :=3D decode of decrypt of resp.enc-part
                                using resp.enc-part.etype and key;
        zero(key);

        if (common_as_rep_tgs_rep_checks fail) then
                destroy resp.key;
                return error;
        endif

        if near(resp.princ_exp) then
                print(warning message);
        endif
        save_for_later(ticket,session,client,server,times,flags);
```

A.4. KRB_AS_REP and KRB_TGS_REP common checks

```
        if (decryption_error() or
            (req.cname !=3D resp.cname) or
            (req.realm !=3D resp.crealm) or
            (req.sname !=3D resp.sname) or
            (req.realm !=3D resp.realm) or
            (req.nonce !=3D resp.nonce) or
            (req.addresses !=3D resp.caddr)) then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
        endif

        /* make sure no flags are set that shouldn't be, and that all =
that */
        /* should be are set                                          =
  */
        if (!check_flags_for_compatability(req.kdc-options,resp.flags)) =
```

```
then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
        endif


        if ((req.from =3D 0) and
            (resp.starttime is not within allowable skew)) then
                destroy resp.key;
                return KRB_AP_ERR_SKEW;
        endif
        if ((req.from !=3D 0) and (req.from !=3D resp.starttime)) then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
        endif
        if ((req.till !=3D 0) and (resp.endtime > req.till)) then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
        endif

        if ((req.kdc-options.RENEWABLE is set) and
            (req.rtime !=3D 0) and (resp.renew-till > req.rtime)) then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
        endif
        if ((req.kdc-options.RENEWABLE-OK is set) and
            (resp.flags.RENEWABLE) and
            (req.till !=3D 0) and
            (resp.renew-till > req.till)) then
                destroy resp.key;
                return KRB_AP_ERR_MODIFIED;
```

```
        endif
```

A.5. KRB_TGS_REQ generation

```
        /* Note that make_application_request might have to recursivly   =
  */
        /* call this routine to get the appropriate ticket-granting =
ticket */

        request.pvno :=3D protocol version; /* pvno =3D 5 */
        request.msg-type :=3D message type; /* type =3D KRB_TGS_REQ */

        body.kdc-options :=3D users's preferences;
        /* If the TGT is not for the realm of the end-server  */
        /* then the sname will be for a TGT for the end-realm */
        /* and the realm of the requested ticket (body.realm) */
        /* will be that of the TGS to which the TGT we are    */
        /* sending applies                                    */
        body.sname :=3D service's name;
        body.realm :=3D service's realm;

        if (body.kdc-options.POSTDATED is set) then
                body.from :=3D requested starting time;
        else
                omit body.from;
        endif
        body.till :=3D requested end time;
        if (body.kdc-options.RENEWABLE is set) then
                body.rtime :=3D requested final renewal time;
        endif


        body.nonce :=3D random_nonce();
        body.etype :=3D requested etypes;
        if (user supplied addresses) then
                body.addresses :=3D user's addresses;
```

```
        else
                omit body.addresses;
        endif

        body.enc-authorization-data :=3D user-supplied data;
        if (body.kdc-options.ENC-TKT-IN-SKEY) then
                body.additional-tickets_ticket :=3D second TGT;
        endif

        request.req-body :=3D body;
        check :=3D generate_checksum (req.body,checksumtype);

        request.padata[0].padata-type :=3D PA-TGS-REQ;
        request.padata[0].padata-value :=3D create a KRB_AP_REQ using
                                        the TGT and checksum

        /* add in any other padata as required/supplied */

        kerberos :=3D lookup(name of local kerberose server (or =
servers));
        send(packet,kerberos);

        wait(for response);
        if (timed_out) then
                retry or use alternate server;
        endif
```

A.6. KRB_TGS_REQ verification and KRB_TGS_REP generation

```
        /* note that reading the application request requires first
        determining the server for which a ticket was issued, and =
choosing the
        correct key for decryption.  The name of the server appears in =
the
        plaintext part of the ticket. */
```

```
        if (no KRB_AP_REQ in req.padata) then
                error_out(KDC_ERR_PADATA_TYPE_NOSUPP);
        endif
        verify KRB_AP_REQ in req.padata;

        /* Note that the realm in which the Kerberos server is operating =
is
        determined by the instance from the ticket-granting ticket.  The =
realm
        in the ticket-granting ticket is the realm under which the =
ticket
        granting ticket was issued.  It is possible for a single =
Kerberos
        server to support more than one realm. */

        auth_hdr :=3D KRB_AP_REQ;
        tgt :=3D auth_hdr.ticket;


        if (tgt.sname is not a TGT for local realm and is not req.sname) =
then
                error_out(KRB_AP_ERR_NOT_US);

        realm :=3D realm_tgt_is_for(tgt);

        decode remainder of request;

        if (auth_hdr.authenticator.cksum is missing) then
                error_out(KRB_AP_ERR_INAPP_CKSUM);
        endif

        if (auth_hdr.authenticator.cksum type is not supported) then
                error_out(KDC_ERR_SUMTYPE_NOSUPP);
        endif
        if (auth_hdr.authenticator.cksum is not both collision-proof and =
keyed) then
                error_out(KRB_AP_ERR_INAPP_CKSUM);
```

```
endif

set computed_checksum :=3D checksum(req);
if (computed_checksum !=3D auth_hdr.authenticatory.cksum) then
        error_out(KRB_AP_ERR_MODIFIED);
endif

server :=3D lookup(req.sname,realm);

if (!server) then
        if (is_foreign_tgt_name(req.sname)) then
                server :=3D best_intermediate_tgs(req.sname);
        else
                /* no server in Database */
                error_out(KDC_ERR_S_PRINCIPAL_UNKNOWN);
        endif
endif

session :=3D generate_random_session_key();

use_etype :=3D first supported etype in req.etypes;

if (no support for req.etypes) then
        error_out(KDC_ERR_ETYPE_NOSUPP);
endif

new_tkt.vno :=3D ticket version; /* =3D 5 */
new_tkt.sname :=3D req.sname;
new_tkt.srealm :=3D realm;
reset all flags in new_tkt.flags;

/* It should be noted that local policy may affect the  */
/* processing of any of these flags.  For example, some */
/* realms may refuse to issue renewable tickets         */
```

```
new_tkt.caddr :=3D tgt.caddr;
resp.caddr :=3D NULL; /* We only include this if they change */
if (req.kdc-options.FORWARDABLE is set) then
        if (tgt.flags.FORWARDABLE is reset) then
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.FORWARDABLE;
endif
if (req.kdc-options.FORWARDED is set) then
        if (tgt.flags.FORWARDABLE is reset) then
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.FORWARDED;
        new_tkt.caddr :=3D req.addresses;
        resp.caddr :=3D req.addresses;
endif
if (tgt.flags.FORWARDED is set) then
        set new_tkt.flags.FORWARDED;
endif

if (req.kdc-options.PROXIABLE is set) then
        if (tgt.flags.PROXIABLE is reset)
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.PROXIABLE;
endif
if (req.kdc-options.PROXY is set) then
        if (tgt.flags.PROXIABLE is reset) then
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.PROXY;
        new_tkt.caddr :=3D req.addresses;
        resp.caddr :=3D req.addresses;
```

```
endif

if (req.kdc-options.ALLOW-POSTDATE is set) then
        if (tgt.flags.MAY-POSTDATE is reset)
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.MAY-POSTDATE;
endif
if (req.kdc-options.POSTDATED is set) then
        if (tgt.flags.MAY-POSTDATE is reset) then
                error_out(KDC_ERR_BADOPTION);
        endif
        set new_tkt.flags.POSTDATED;
        set new_tkt.flags.INVALID;
        if (against_postdate_policy(req.from)) then
                error_out(KDC_ERR_POLICY);
        endif
        new_tkt.starttime :=3D req.from;
endif


if (req.kdc-options.VALIDATE is set) then
        if (tgt.flags.INVALID is reset) then
                error_out(KDC_ERR_POLICY);
        endif
        if (tgt.starttime > kdc_time) then
                error_out(KRB_AP_ERR_NYV);
        endif
        if (check_hot_list(tgt)) then
                error_out(KRB_AP_ERR_REPEAT);
        endif
        tkt :=3D tgt;
        reset new_tkt.flags.INVALID;
```

```
        endif

        if (req.kdc-options.(any flag except ENC-TKT-IN-SKEY, RENEW,
                             and those already processed) is set) then
                error_out(KDC_ERR_BADOPTION);
        endif

        new_tkt.authtime :=3D tgt.authtime;

        if (req.kdc-options.RENEW is set) then
          /* Note that if the endtime has already passed, the ticket =
would  */
          /* have been rejected in the initial authentication stage, so  =
   */
          /* there is no need to check again here                        =
   */
                if (tgt.flags.RENEWABLE is reset) then
                        error_out(KDC_ERR_BADOPTION);
                endif
                if (tgt.renew-till < kdc_time) then
                        error_out(KRB_AP_ERR_TKT_EXPIRED);
                endif
                tkt :=3D tgt;
                new_tkt.starttime :=3D kdc_time;
                old_life :=3D tgt.endttime - tgt.starttime;
                new_tkt.endtime :=3D min(tgt.renew-till,
                                        new_tkt.starttime + old_life);
        else
                new_tkt.starttime :=3D kdc_time;
                if (req.till =3D 0) then
                        till :=3D infinity;
                else
                        till :=3D req.till;
                endif
                new_tkt.endtime :=3D min(till,
                                        =
```

```
new_tkt.starttime+client.max_life,
                                        =
new_tkt.starttime+server.max_life,
                                        =
new_tkt.starttime+max_life_for_realm,
                                        tgt.endtime);



                if ((req.kdc-options.RENEWABLE-OK is set) and
                    (new_tkt.endtime < req.till) and
                    (tgt.flags.RENEWABLE is set) then
                        /* we set the RENEWABLE option for later =
processing */
                        set req.kdc-options.RENEWABLE;
                        req.rtime :=3D min(req.till, tgt.renew-till);
                endif
        endif

        if (req.rtime =3D 0) then
                rtime :=3D infinity;
        else
                rtime :=3D req.rtime;
        endif

        if ((req.kdc-options.RENEWABLE is set) and
            (tgt.flags.RENEWABLE is set)) then
                set new_tkt.flags.RENEWABLE;
                new_tkt.renew-till :=3D min(rtime,
                                           =
new_tkt.starttime+client.max_rlife,
                                           =
new_tkt.starttime+server.max_rlife,
                                           =
new_tkt.starttime+max_rlife_for_realm,
                                           tgt.renew-till);
```

```
        else
                new_tkt.renew-till :=3D OMIT; /* leave the renew-till =
field out */
        endif
        if (req.enc-authorization-data is present) then
                decrypt req.enc-authorization-data into =
decrypted_authorization_data
                        using auth_hdr.authenticator.subkey;
                if (decrypt_error()) then
                        error_out(KRB_AP_ERR_BAD_INTEGRITY);
                endif
        endif
        new_tkt.authorization_data :=3D =
req.auth_hdr.ticket.authorization_data +
                                decrypted_authorization_data;

        new_tkt.key :=3D session;
        new_tkt.crealm :=3D tgt.crealm;
        new_tkt.cname :=3D req.auth_hdr.ticket.cname;

        if (realm_tgt_is_for(tgt) :=3D tgt.realm) then
                /* tgt issued by local realm */
                new_tkt.transited :=3D tgt.transited;
        else
                /* was issued for this realm by some other realm */
                if (tgt.transited.tr-type not supported) then
                        error_out(KDC_ERR_TRTYPE_NOSUPP);
                endif
                new_tkt.transited :=3D compress_transited(tgt.transited =
+ tgt.realm)




                /* Don't check tranited field if TGT for foreign realm,=20
                 * or requested not to check */
                if (is_not_foreign_tgt_name(new_tkt.server)=20
```

```
                    && req.kdc-options.DISABLE-TRANSITED-CHECK not set) =
then
                        /* Check it, so end-server does not have to=20
                         * but don't fail, end-server may still accept =
it */
                        if (check_transited_field(new_tkt.transited) =
=3D=3D OK)
                                set =
new_tkt.flags.TRANSITED-POLICY-CHECKED;
                        endif
                endif
        endif

        encode encrypted part of new_tkt into OCTET STRING;
        if (req.kdc-options.ENC-TKT-IN-SKEY is set) then
                if (server not specified) then
                        server =3D req.second_ticket.client;
                endif
                if ((req.second_ticket is not a TGT) or
                    (req.second_ticket.client !=3D server)) then
                        error_out(KDC_ERR_POLICY);
                endif

                new_tkt.enc-part :=3D encrypt OCTET STRING using
                        using etype_for_key(second-ticket.key), =
second-ticket.key;
        else
                new_tkt.enc-part :=3D encrypt OCTET STRING
                        using etype_for_key(server.key), server.key, =
server.p_kvno;
```

```
        endif

        resp.pvno :=3D 5;
        resp.msg-type :=3D KRB_TGS_REP;
        resp.crealm :=3D tgt.crealm;
        resp.cname :=3D tgt.cname;
        resp.ticket :=3D new_tkt;

        resp.key :=3D session;
        resp.nonce :=3D req.nonce;
        resp.last-req :=3D fetch_last_request_info(client);
        resp.flags :=3D new_tkt.flags;

        resp.authtime :=3D new_tkt.authtime;
        resp.starttime :=3D new_tkt.starttime;
        resp.endtime :=3D new_tkt.endtime;

        omit resp.key-expiration;

        resp.sname :=3D new_tkt.sname;
        resp.realm :=3D new_tkt.realm;

        if (new_tkt.flags.RENEWABLE) then
                resp.renew-till :=3D new_tkt.renew-till;
        endif



        encode body of reply into OCTET STRING;

        if (req.padata.authenticator.subkey)
                resp.enc-part :=3D encrypt OCTET STRING using use_etype,
                        req.padata.authenticator.subkey;
        else resp.enc-part :=3D encrypt OCTET STRING using use_etype, =
tgt.key;
```

```
        send(resp);

=09
```

A.7. KRB_TGS_REP verification

```
        decode response into resp;

        if (resp.msg-type =3D KRB_ERROR) then
                process_error(resp);
                return;
        endif

        /* On error, discard the response, and zero the session key from
        the response immediately */

        if (req.padata.authenticator.subkey)
                unencrypted part of resp :=3D decode of decrypt of =
resp.enc-part
                        using resp.enc-part.etype and subkey;
        else unencrypted part of resp :=3D decode of decrypt of =
resp.enc-part
                                using resp.enc-part.etype and tgt's =
session key;
        if (common_as_rep_tgs_rep_checks fail) then
                destroy resp.key;
                return error;
        endif

        check authorization_data as necessary;
        save_for_later(ticket,session,client,server,times,flags);
```

A.8. Authenticator generation

```
body.authenticator-vno :=3D authenticator vno; /* =3D 5 */
body.cname, body.crealm :=3D client name;
if (supplying checksum) then
        body.cksum :=3D checksum;
endif
get system_time;
body.ctime, body.cusec :=3D system_time;
if (selecting sub-session key) then
        select sub-session key;
        body.subkey :=3D sub-session key;
endif
if (using sequence numbers) then
        select initial sequence number;
        body.seq-number :=3D initial sequence;
endif
```

A.9. KRB_AP_REQ generation

```
obtain ticket and session_key from cache;

packet.pvno :=3D protocol version; /* 5 */
packet.msg-type :=3D message type; /* KRB_AP_REQ */

if (desired(MUTUAL_AUTHENTICATION)) then
        set packet.ap-options.MUTUAL-REQUIRED;
else
        reset packet.ap-options.MUTUAL-REQUIRED;
endif
if (using session key for ticket) then
        set packet.ap-options.USE-SESSION-KEY;
else
        reset packet.ap-options.USE-SESSION-KEY;
endif
packet.ticket :=3D ticket; /* ticket */
generate authenticator;
```

```
        encode authenticator into OCTET STRING;
        encrypt OCTET STRING into packet.authenticator using =
session_key;
```

A.10. KRB_AP_REQ verification

```
        receive packet;
        if (packet.pvno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.msg-type !=3D KRB_AP_REQ) then
                error_out(KRB_AP_ERR_MSG_TYPE);
        endif
        if (packet.ticket.tkt_vno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.ap_options.USE-SESSION-KEY is set) then
                retrieve session key from ticket-granting ticket for
                 packet.ticket.{sname,srealm,enc-part.etype};
        else
                retrieve service key for
                 =
packet.ticket.{sname,srealm,enc-part.etype,enc-part.skvno};
        endif
        if (no_key_available) then
                if (cannot_find_specified_skvno) then
                        error_out(KRB_AP_ERR_BADKEYVER);
                else
                        error_out(KRB_AP_ERR_NOKEY);
                endif
        endif
        decrypt packet.ticket.enc-part into decr_ticket using retrieved =
key;
        if (decryption_error()) then
                error_out(KRB_AP_ERR_BAD_INTEGRITY);
```

```
        endif



        decrypt packet.authenticator into decr_authenticator
                using decr_ticket.key;
        if (decryption_error()) then
                error_out(KRB_AP_ERR_BAD_INTEGRITY);
        endif
        if (decr_authenticator.{cname,crealm} !=3D
            decr_ticket.{cname,crealm}) then
                error_out(KRB_AP_ERR_BADMATCH);
        endif
        if (decr_ticket.caddr is present) then
                if (sender_address(packet) is not in decr_ticket.caddr) =
then
                        error_out(KRB_AP_ERR_BADADDR);
                endif
        elseif (application requires addresses) then
                error_out(KRB_AP_ERR_BADADDR);
        endif
        if (not in_clock_skew(decr_authenticator.ctime,
                              decr_authenticator.cusec)) then
                error_out(KRB_AP_ERR_SKEW);
        endif
        if (repeated(decr_authenticator.{ctime,cusec,cname,crealm})) =
then
                error_out(KRB_AP_ERR_REPEAT);
        endif
        save_identifier(decr_authenticator.{ctime,cusec,cname,crealm});
        get system_time;
        if ((decr_ticket.starttime-system_time > CLOCK_SKEW) or
            (decr_ticket.flags.INVALID is set)) then
                /* it hasn't yet become valid */
                error_out(KRB_AP_ERR_TKT_NYV);
        endif
        if (system_time-decr_ticket.endtime > CLOCK_SKEW) then
                error_out(KRB_AP_ERR_TKT_EXPIRED);
```

```
        endif
        if (decr_ticket.transited) then
            /* caller may ignore the TRANSITED-POLICY-CHECKED and do
             * check anyway */
            if (decr_ticket.flags.TRANSITED-POLICY-CHECKED not set) then
                 if (check_transited_field(decr_ticket.transited) then
                      error_out(KDC_AP_PATH_NOT_ACCPETED);
                 endif
            endif
        endif
        /* caller must check decr_ticket.flags for any pertinent details =
*/
        return(OK, decr_ticket, packet.ap_options.MUTUAL-REQUIRED);
```

A.11. KRB_AP_REP generation

```
        packet.pvno :=3D protocol version; /* 5 */
        packet.msg-type :=3D message type; /* KRB_AP_REP */

        body.ctime :=3D packet.ctime;
        body.cusec :=3D packet.cusec;
        if (selecting sub-session key) then
                select sub-session key;
                body.subkey :=3D sub-session key;
        endif

        if (using sequence numbers) then
                select initial sequence number;
                body.seq-number :=3D initial sequence;
        endif

        encode body into OCTET STRING;

        select encryption type;
        encrypt OCTET STRING into packet.enc-part;
```

A.12. KRB_AP_REP verification

```
        receive packet;
        if (packet.pvno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.msg-type !=3D KRB_AP_REP) then
                error_out(KRB_AP_ERR_MSG_TYPE);
        endif
        cleartext :=3D decrypt(packet.enc-part) using ticket's session =
key;
        if (decryption_error()) then
                error_out(KRB_AP_ERR_BAD_INTEGRITY);
        endif
        if (cleartext.ctime !=3D authenticator.ctime) then
                error_out(KRB_AP_ERR_MUT_FAIL);
        endif
        if (cleartext.cusec !=3D authenticator.cusec) then
                error_out(KRB_AP_ERR_MUT_FAIL);
        endif
        if (cleartext.subkey is present) then
                save cleartext.subkey for future use;
        endif
        if (cleartext.seq-number is present) then
                save cleartext.seq-number for future verifications;
        endif
        return(AUTHENTICATION_SUCCEEDED);
```

A.13. KRB_SAFE generation

```
        collect user data in buffer;

        /* assemble packet: */
        packet.pvno :=3D protocol version; /* 5 */
        packet.msg-type :=3D message type; /* KRB_SAFE */
```

```
        body.user-data :=3D buffer; /* DATA */
        if (using timestamp) then
                get system_time;
                body.timestamp, body.usec :=3D system_time;
        endif
        if (using sequence numbers) then
                body.seq-number :=3D sequence number;
        endif
        body.s-address :=3D sender host addresses;
        if (only one recipient) then
                body.r-address :=3D recipient host address;
        endif




        checksum.cksumtype :=3D checksum type;
        compute checksum over body;
        checksum.checksum :=3D checksum value; /* checksum.checksum */
        packet.cksum :=3D checksum;
        packet.safe-body :=3D body;
```

A.14. KRB_SAFE verification

```
        receive packet;
        if (packet.pvno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.msg-type !=3D KRB_SAFE) then
                error_out(KRB_AP_ERR_MSG_TYPE);
        endif
        if (packet.checksum.cksumtype is not both collision-proof and =
keyed) then
                error_out(KRB_AP_ERR_INAPP_CKSUM);
        endif
        if (safe_priv_common_checks_ok(packet)) then
                set computed_checksum :=3D checksum(packet.body);
```

```
                if (computed_checksum !=3D packet.checksum) then
                        error_out(KRB_AP_ERR_MODIFIED);
                endif
                return (packet, PACKET_IS_GENUINE);
        else
                return common_checks_error;
        endif
```

A.15. KRB_SAFE and KRB_PRIV common checks

```
        if (packet.s-address !=3D O/S_sender(packet)) then
                /* O/S report of sender not who claims to have sent it =
*/
                error_out(KRB_AP_ERR_BADADDR);
        endif
        if ((packet.r-address is present) and
            (packet.r-address !=3D local_host_address)) then
                /* was not sent to proper place */
                error_out(KRB_AP_ERR_BADADDR);
        endif
        if (((packet.timestamp is present) and
             (not in_clock_skew(packet.timestamp,packet.usec))) or
            (packet.timestamp is not present and timestamp expected)) =
then
                error_out(KRB_AP_ERR_SKEW);
        endif
        if (repeated(packet.timestamp,packet.usec,packet.s-address)) =
then
                error_out(KRB_AP_ERR_REPEAT);
        endif

        if (((packet.seq-number is present) and
             ((not in_sequence(packet.seq-number)))) or
            (packet.seq-number is not present and sequence expected)) =
then
                error_out(KRB_AP_ERR_BADORDER);
```

```
        endif



        if (packet.timestamp not present and packet.seq-number not =
present) then
                error_out(KRB_AP_ERR_MODIFIED);
        endif

        save_identifier(packet.{timestamp,usec,s-address},
                        sender_principal(packet));

        return PACKET_IS_OK;
```

A.16. KRB_PRIV generation

```
        collect user data in buffer;

        /* assemble packet: */
        packet.pvno :=3D protocol version; /* 5 */
        packet.msg-type :=3D message type; /* KRB_PRIV */

        packet.enc-part.etype :=3D encryption type;

        body.user-data :=3D buffer;
        if (using timestamp) then
                get system_time;
                body.timestamp, body.usec :=3D system_time;
        endif
        if (using sequence numbers) then
                body.seq-number :=3D sequence number;
        endif
        body.s-address :=3D sender host addresses;
        if (only one recipient) then
                body.r-address :=3D recipient host address;
```

```
        endif

        encode body into OCTET STRING;

        select encryption type;
        encrypt OCTET STRING into packet.enc-part.cipher;
```

A.17. KRB_PRIV verification

```
        receive packet;
        if (packet.pvno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.msg-type !=3D KRB_PRIV) then
                error_out(KRB_AP_ERR_MSG_TYPE);
        endif

        cleartext :=3D decrypt(packet.enc-part) using negotiated key;
        if (decryption_error()) then
                error_out(KRB_AP_ERR_BAD_INTEGRITY);
        endif

        if (safe_priv_common_checks_ok(cleartext)) then
                return(cleartext.DATA, =
PACKET_IS_GENUINE_AND_UNMODIFIED);
        else
                return common_checks_error;
        endif
```

A.18. KRB_CRED generation

```
        invoke KRB_TGS; /* obtain tickets to be provided to peer */

        /* assemble packet: */
        packet.pvno :=3D protocol version; /* 5 */
```

```
        packet.msg-type :=3D message type; /* KRB_CRED */

        for (tickets[n] in tickets to be forwarded) do
                packet.tickets[n] =3D tickets[n].ticket;
        done

        packet.enc-part.etype :=3D encryption type;

        for (ticket[n] in tickets to be forwarded) do
                body.ticket-info[n].key =3D tickets[n].session;
                body.ticket-info[n].prealm =3D tickets[n].crealm;
                body.ticket-info[n].pname =3D tickets[n].cname;
                body.ticket-info[n].flags =3D tickets[n].flags;
                body.ticket-info[n].authtime =3D tickets[n].authtime;
                body.ticket-info[n].starttime =3D tickets[n].starttime;
                body.ticket-info[n].endtime =3D tickets[n].endtime;
                body.ticket-info[n].renew-till =3D =
tickets[n].renew-till;
                body.ticket-info[n].srealm =3D tickets[n].srealm;
                body.ticket-info[n].sname =3D tickets[n].sname;
                body.ticket-info[n].caddr =3D tickets[n].caddr;
        done

        get system_time;
        body.timestamp, body.usec :=3D system_time;

        if (using nonce) then
                body.nonce :=3D nonce;
        endif

        if (using s-address) then
                body.s-address :=3D sender host addresses;
        endif
        if (limited recipients) then
                body.r-address :=3D recipient host address;
```

```
        endif

        encode body into OCTET STRING;

        select encryption type;
        encrypt OCTET STRING into packet.enc-part.cipher
               using negotiated encryption key;
```

A.19. KRB_CRED verification

```
        receive packet;
        if (packet.pvno !=3D 5) then
                either process using other protocol spec
                or error_out(KRB_AP_ERR_BADVERSION);
        endif
        if (packet.msg-type !=3D KRB_CRED) then
                error_out(KRB_AP_ERR_MSG_TYPE);
        endif

        cleartext :=3D decrypt(packet.enc-part) using negotiated key;
        if (decryption_error()) then
                error_out(KRB_AP_ERR_BAD_INTEGRITY);
        endif
        if ((packet.r-address is present or required) and
           (packet.s-address !=3D O/S_sender(packet)) then
                /* O/S report of sender not who claims to have sent it =
*/
                error_out(KRB_AP_ERR_BADADDR);
        endif
        if ((packet.r-address is present) and
            (packet.r-address !=3D local_host_address)) then
                /* was not sent to proper place */
                error_out(KRB_AP_ERR_BADADDR);
        endif
        if (not in_clock_skew(packet.timestamp,packet.usec)) then
                error_out(KRB_AP_ERR_SKEW);
```

```
        endif
        if (repeated(packet.timestamp,packet.usec,packet.s-address)) =
then
                error_out(KRB_AP_ERR_REPEAT);
        endif
        if (packet.nonce is required or present) and
           (packet.nonce !=3D expected-nonce) then
                error_out(KRB_AP_ERR_MODIFIED);
        endif

        for (ticket[n] in tickets that were forwarded) do
                save_for_later(ticket[n],key[n],principal[n],
                                server[n],times[n],flags[n]);
        return
```

A.20. KRB_ERROR generation

```
        /* assemble packet: */
        packet.pvno :=3D protocol version; /* 5 */
        packet.msg-type :=3D message type; /* KRB_ERROR */

        get system_time;
        packet.stime, packet.susec :=3D system_time;
        packet.realm, packet.sname :=3D server name;

        if (client time available) then
                packet.ctime, packet.cusec :=3D client_time;
        endif




        packet.error-code :=3D error code;
        if (client name available) then
                packet.cname, packet.crealm :=3D client name;
        endif
        if (error text available) then
                packet.e-text :=3D error text;
```

```
        endif
        if (error data available) then
                packet.e-data :=3D error data;
        endif
```

B. Definition of common authorization data elements

This appendix contains the definitions of common authorization data elements. These common authorization data elements are recursivly defined, meaning the ad-data for these types will itself contain a sequence of authorization data whose interpretation is affected by the encapsulating element. Depending on the meaning of the encapsulating element, the encapsulated elements may be ignored, might be interpreted as issued directly by the KDC, or they might be stored in a separate plaintext part of the ticket. The types of the encapsulating elements are specified as part of the Kerberos specification because the behaviour based on these values should be understood across implementations whereas other elements need only be understood by the applications which they affect.

In the definitions that follow, the value of the ad-type for the element will be specified in the subsection number, and the value of the ad-data will be as shown in the ASN.1 structure that follows the subsection heading.

B.1. If relevant

AD-IF-RELEVANT AuthorizationData

AD elements encapsulated within the if-relevant element are intended for interpretation only by application servers that understand the particular ad-type of the embedded element. Application servers that do not understand the type of an element embedded within the if-relevant element may ignore the uninterpretable element. This element promotes interoperability across implementations which may have local extensions for authorization.

B.2. Intended for server

```
AD-INTENDED-FOR-SERVER   SEQUENCE {
         intended-server[0]     SEQUENCE OF PrincipalName
         elements[1]            AuthorizationData
}
```

AD elements encapsulated within the intended-for-server element may be ignored if the application server is not in the list of principal names of intended servers. Further, a KDC issuing a ticket for an application server can remove this element if the application server is not in the list of intended servers.

Application servers should check for their principal name in the intended-server field of this element. If their principal name is not found, this element should be ignored. If found, then the encapsulated elements should be evaluated in the same manner as if they were present in the top level authorization data field. Applications and application servers that do not implement this element should reject tickets that contain authorization data elements of this type.

B.3. Intended for application class

AD-INTENDED-FOR-APPLICATION-CLASS SEQUENCE { intended-application-class[0] SEQUENCE OF GeneralString elements[1] AuthorizationData } AD elements encapsulated within the intended-for-application-class element may be ignored if the application server is not in one of the named classes of application servers. Examples of application server classes include "FILESYSTEM", and other kinds of servers.=20
This element and the elements it encapulates may be safely ignored by applications, application servers, and KDCs that do not implement this element.

B.4. KDC Issued

```
AD-KDCIssued   SEQUENCE {
               ad-checksum[0]    Checksum,
                i-realm[1]       Realm OPTIONAL,
                i-sname[2]       PrincipalName OPTIONAL,
               elements[3]       AuthorizationData.
}
```

ad-checksum

A checksum over the elements field using a cryptographic checksum method that is identical to the checksum used to protect the ticket itself (i.e. using the same hash function and the same encryption algorithm used to encrypt the ticket) and using a key derived from the same key used to protect the ticket.

i-realm, i-sname

The name of the issuing principal if different from the KDC itself. This field would be used when the KDC can verify the authenticity of elements signed by the issuing principal and it allows this KDC to notify the application server of the validity of those elements.

elements

A sequence of authorization data elements issued by the KDC.

The KDC-issued ad-data field is intended to provide a means for Kerberos principal credentials to embed within themselves privilege attributes and other mechanisms for positive authorization, amplifying the priveleges of the principal beyond what can be done using a credentials without such an a-data element.

This can not be provided without this element because the definition of the authorization-data field allows elements to be added at will by the bearer of a TGT at the time that they request service tickets and elements may also be added to a delegated ticket by inclusion in the authenticator.

For KDC-issued elements this is prevented because the elements are signed by the KDC by including a checksum encrypted using the server's key (the same key used to encrypt the ticket - or a key derived from that key). Elements encapsulated with in the KDC-issued element will be ignored by the application server if this "signature" is not present. Further, elements encapsulated within this element from a ticket granting ticket may be interpreted by the KDC, and used as a basis according to policy for including new signed elements within derivative tickets, but they will not be copied to a derivative ticket directly. If they are copied directly to a derivative ticket by a KDC that is not aware of this element, the signature will not be correct for the application ticket elements, and the field will be ignored by the application server.

This element and the elements it encapulates may be safely ignored by applications, application servers, and KDCs that do not implement this element.

B.5. And-Or

```
AD-AND-OR           SEQUENCE {
                        condition-count[0]    INTEGER,
                        elements[1]           AuthorizationData
}=20
```

When restrictive AD elements encapsulated within the and-or element are encountered, only the number specified in condition-count of the encapsulated conditions must be met in order to satisfy this element. This element may be used to implement an "or" operation by setting the condition-count field to 1, and it may specify an "and" operation by setting the condition count to the number of embedded elements. Application servers that do not implement this element must reject tickets that contain authorization data elements of this type.

B.6. Mandatory ticket extensions

```
AD-Mandatory-Ticket-Extensions          SEQUENCE {
```

```
                    te-type[0]       INTEGER,
                    te-checksum[0]    Checksum
}=20
```

An authorization data element of type mandatory-ticket-extensions specifies the type and a collision-proof checksum using the same hash algorithm used to protect the integrity of the ticket itself. This checksum will be calculated over an individual extension field of the type indicated. If there are more than one extension, multiple Mandatory-Ticket-Extensions authorization data elements may be present, each with a checksum for a different extension field. This restriction indicates that the ticket should not be accepted if a ticket extension is not present in the ticket for which the type and checksum do not match that checksum specified in the authorization data element. Note that although the type is redundant for the purposes of the comparison, it makes the comparison easier when multiple extensions are present. Application servers that do not implement this element must reject tickets that contain authorization data elements of this type.

B.7. Authorization Data in ticket extensions

AD-IN-Ticket-Extensions Checksum

An authorization data element of type in-ticket-extensions specifies a collision-proof checksum using the same hash algorithm used to protect the integrity of the ticket itself. This checksum is calculated over a separate external AuthorizationData field carried in the ticket extensions. Application servers that do not implement this element must reject tickets that contain authorization data elements of this type. Application servers that do implement this element will search the ticket extensions for authorization data fields, calculate the specified checksum over each authorization data field and look for one matching the checksum in this in-ticket-extensions element. If not found, then the ticket must be rejected. If found, the corresponding authorization data elements will be interpreted in the same manner as if they were contained in the top level authorization data field.

Note that if multiple external authorization data fields are present in a ticket, each will have a corresponding element of type in-ticket-extensions in the top level authorization data field, and the external entries will be linked to the corresponding element by their checksums.

C. Definition of common ticket extensions

This appendix contains the definitions of common ticket extensions. Support for these extensions is optional. However, certain extensions have associated authorization data elements that may require rejection of a ticket containing an extension by application servers that do not implement the particular extension. Other extensions have been defined beyond those described in this specification. Such extensions are described elswhere and for some of those extensions the reserved number may be found in the list of constants.

It is known that older versions of Kerberos did not support this field, and that some clients will strip this field from a ticket when they parse and then reassemble a ticket as it is passed to the application servers. The presence of the extension will not break such clients, but any functionaly dependent on the extensions will not work when such tickets are handled by old clients. In such situations, some implementation may use alternate methods to transmit the information in the extensions field.

C.1. Null ticket extension

TE-NullExtension OctetString -- The empty Octet String

The te-data field in the null ticket extension is an octet string of lenght zero. This extension may be included in a ticket granting ticket so that the KDC can determine on presentation of the ticket granting ticket whether the client software will strip the extensions field. =20

C.2. External Authorization Data

TE-ExternalAuthorizationData AuthorizationData

The te-data field in the external authorization data ticket extension is field of type AuthorizationData containing one or more authorization data elements. If present, a corresponding authorization data element will be present in the primary authorization data for the ticket and that element will contain a checksum of the external authorization data ticket extension.

D. Significant changes since RFC 1510

Commentary

Section 1: The preamble and introduction does not define the protocol, mention is made in the introduction regarding the ability to rely on the KDC to check the transited field, and on the inclusion of a flag in a ticket indicating that this check has occurred. This is a new capability not present in RFC1510. Pre-existing implementation may ignore or not set this flag without negative security implications.

The definition of the secret key says that in the case of a user the key may be derived from a password. In 1510, it said that the key was derived from the password. This change was made to accommodate situations where the user key might be stored on a smart-card, or otherwise obtained independent of a password.

The introduction also mentions the use of public key for initial authentication in Kerberos by reference. RFC1510 did not include such a reference.

Section 1.2 was added to explain that while Kerberos provides authentication of a named principal, it is still the responsibility of the application to ensure that the authenticated name is the entity with which the application wishes to communicate. Because section 1.2 is completely new, I am particularly interested in suggestions to improve the wording of this section. Sections 1.2-4 were renumbered.

Section 2: No changes were made to existing options and flags specified in RFC1510, though some of the sections in the specification were renumbered, and text was revised to make the description and intent of existing options clearer, especially with respect to the ENC-TKT-IN-SKEY option (now section 2.9.3) which is used for user-to-user authentication. New options and ticket flags added since RFC1510 include transited policy checking (section 2.7), anonymous tickets (section 2.8) and name canonicalization (section 2.9.1).

Section 3: Added mention of the optional checksum field in the KRB-ERROR message. Added mention of name canonicalization and anonymous tickets in exposition on KDC options. Mention of the name canonicalization case is included in the description of the KDC reply (3.1.3). A warning regarding generation of session keys for application use was added, urging the inclusion of key entropy from the KDC generated session key in the ticket. An example regarding use of the subsession key was added to section 3.2.6. Descriptions of the pa-etype-info, and pa-pw-salt preauthentication data items were added.

Changes to section 4: Added language about who has access to the keys in the Kerberos database. Also made it clear that KDC's may obtain the information from some database field through other means - for example, one form of pkinit may extract some of these fields from a certificate.

Regarding the discussion on the list regarding the use of tamper resistant hardware to store keys, I was not able to determine specific suggested changes to the text in the RFC regarding this. Much of this discussion centers around particular implementations. I did however loosen the wording about the database so as not to preclude keys that can not be extracted in the clear from such hardware.

Section 5: A statement regarding the carrying of unrecognized additional fields in ASN.1 encoding through in tickets was added (still waiting on some better text regarding this).

Ticket flags and KDC options were added to support the new functions described elsewhere in this document. The encoding of the options flags are now described to be no less than 32 bits, and the smallest number of bits beyond 32 needed to encode any set bits. It also describes the encoding of the bitstring as using "unnamed" bits.

An optional ticket extensions field was added to support the carrying of auxiliary data that allows the passing of auxiliary that is to accompany a ticket to the verifier.

(I would like to drop the part about optionally appending it of the opaque part of the ciphertext. We are still waiting on some text regarding how to assure backward compatibility).

(Still pending, Tom Yu's request to change the application codes on KDC message to indicate which minor rev of the protocol - I think this might break things, but am not sure).

Definition of the PA-USE-SPECIFIED-KVNO preauthentication data field was added.

The optional e-cksum field was added to the KRB-ERROR message and the e-data filed was generalized for use in other than the KDC_ERR_PREAUTH_REQUIRED error. The TypedData structure was defined. Type tags for TypedData are defined in the same sequence as the PA-DATA type space to avoid confusion with the use of the PA-DATA namespace previously used for the e-data field for the KDC_ERR_PREAUTH_REQUIRED error.

Section 7: Words were added describing the convention that domain based realm names for newly created realms should be specified as upper case. This recommendation does not make lower case realm names illegal. Words were added highlighting that the slash separated components in the X500 style of realm names is consistent with existing RFC1510 based implementations, but that it conflicts with the general recommendation of X.500 name representation specified in RFC2253.

There were suggestions on the list regarding extensions to or new name types. These require discussion at the IETF meeting. My own feeling at this point is that in the absence of a strong consensus for for adding new types at this time, I would rather not add new name types in the current draft, but leave things open for additions later.

Section 8: Since RFC1510, the definition of the TCP transport for Kerberos messages was added.

Section 9: Requirements for supporting DES3-CBC-SHA1-KD encryption and HMAC-SHA1-DES3-KD checksums were added.
I would like to make support for Rijndael mandatory and for us to have a SINGLE standard for use of Rijndale in these revisions.

Discussion

Section 8: Regarding the suggestion of weakening the requirement for use of port 88 for cases where the port can be looked up elsewhere - I did not incorporate this suggestion because cross realm authentication requires the ability to contact the appropriate KDC, and unless ALL implementations of Kerberos include support for finding such alternate port numbers, use of such KDC's would be non-interoperable.

[TM] Project Athena, Athena, and Kerberos are trademarks of the Massachusetts Institute of Technology (MIT). No commercial use of these trademarks may be made without prior written permission of MIT.

[1.1] Note, however, that many applications use Kerberos' functions only upon the initiation of a stream-based network connection. Unless an application subsequently provides integrity protection for the data stream, the identity verification applies only to the initiation of the connection, and does not guarantee that subsequent messages on the connection originate from the same principal.

[1.2] Secret and private are often used interchangeably in the literature. In our usage, it takes two (or more) to share a secret, thus a shared DES key is a secret key. Something is only private when no one but its owner knows it. Thus, in public key cryptosystems, one has a public and a private key.

[1.3] Of course, with appropriate permission the client could arrange registration of a separately-named prin- cipal in a remote realm, and engage in normal exchanges with that realm's services. However, for even small numbers of clients this becomes cumbersome, and more automatic methods as described here are necessary.
[2.1] Though it is permissible to request or issue tick- ets with no network addresses specified.

[2.2] It is important that the KDC be sent the name as typed by the user, and not only the canonical form of the name. If the domain name system was used to find the canonical name on the client side, the mapping is vulnerable. [3.1] The password-changing request must not be honored unless the requester can provide the old password (the user's current secret key). Otherwise, it would be possible for someone to walk up to an unattended session and change another user's password.

[3.2] To authenticate a user logging on to a local system, the credentials obtained in the AS exchange may first be used in a TGS exchange to obtain credentials for a local server. Those credentials must then be verified by a local server through successful completion of the Client/Server exchange.

[3.3] "Random" means that, among other things, it should be impossible to guess the next session key based on knowledge of past session keys. This can only be achieved in a pseudo-random number generator if it is based on cryptographic principles. It is more desirable to use a truly random number generator, such as one based on measurements of random physical phenomena.

[3.4] Tickets contain both an encrypted and unencrypted portion, so cleartext here refers to the entire unit, which can be copied from one message and replayed in another without any cryptographic skill.

[3.5] Note that this can make applications based on unreliable transports difficult to code correctly. If the transport might deliver duplicated messages, either a new authenticator must be generated for each retry, or the application server must match requests and replies and replay the first reply in response to a detected duplicate.

[3.6] This allows easy implementation of user-to-user authentication [8], which uses ticket-granting ticket session keys in lieu of secret server keys in situations where such secret keys could be easily compromised.

[3.7]Note also that the rejection here is restricted to authenticators from the same principal to the same server. Other client principals communicating with the same server principal should not be have their authenticators rejected if the time and microsecond fields happen to match some other client's authenticator.

[3.8] If this is not done, an attacker could subvert the authentication by recording the ticket and authenticator sent over the network to a server and replaying them following an event that caused the server to lose track of recently seen authenticators.

[3.9] In the Kerberos version 4 protocol, the timestamp in the reply was the client's timestamp plus one. This is not necessary in version 5 because version 5 messages are formatted in such a way that it is not possible to create the reply by judicious message surgery (even in encrypted form) without knowledge of the appropriate encryption keys.

[3.10] Note that for encrypting the KRB_AP_REP message, the sub-session key is not used, even if present in the Authenticator.

[3.11] Implementations of the protocol may wish to provide routines to choose subkeys based on session keys and random numbers and to generate a negotiated key to be returned in the KRB_AP_REP message.

[3.12]This can be accomplished in several ways. It might be known beforehand (since the realm is part of the principal identifier), it might be stored in a nameserver, or it might be obtained from a configuration file. If the realm to be used is obtained from a nameserver, there is a danger of being spoofed if the nameservice providing the realm name is not authenticated. This might result in the use of a realm which has been compromised, and would result in an attacker's ability to compromise the authentication of the application server to the client.

[3.13] If the client selects a sub-session key, care must be taken to ensure the randomness of the selected sub-session key. One approach would be to generate a random number and XOR it with the session key from the ticket-granting ticket.

[4.1] The implementation of the Kerberos server need not combine the database and the server on the same machine; it is feasible to store the principal database in, say, a network name service, as long as the entries stored therein are protected from disclosure to and modification by unauthorized parties. However, we recommend against such strategies, as they can make system management and threat analysis quite complex.

[4.2] See the discussion of the padata field in section 5.4.2 for details on why this can be useful.

Annex C (normative): PKINIT specification

The PKINIT specification is currently still an IETF draft. The present document complies only with the version of the draft that is listed in this annex. The IPCablecom security experts will continue to track progress of the PKINIT draft through the IETF and will advise the Study Group concerning how to best deal with the situation as the document becomes an RFC. The goal is to incorporate by reference and eliminate this annex.

INTERNET-DRAFT Brian Tung
draft-ietf-cat-kerberos-pk-init-14.txt Clifford Neuman
Updates: RFC 1510bis USC/ISI
expires January 15, 2002 Matthew Hur
Cisco
Ari Medvinsky
Keen.com, Inc.
Sasha Medvinsky
Motorola
John Wray
Iris Associates, Inc.
Jonathan Trostle
Cisco

Public Key Cryptography for Initial Authentication in Kerberos

0. Status Of This Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC 2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.ietf.org (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

The distribution of this memo is unlimited. It is filed as draft-ietf-cat-kerberos-pk-init-14.txt, and expires January 15, 2002. Please send comments to the authors.

1. Abstract

This document defines extensions (PKINIT) to the Kerberos protocol specification (RFC 1510bis [1]) to provide a method for using public key cryptography during initial authentication. The methods defined specify the ways in which preauthentication data fields and error data fields in Kerberos messages are to be used to transport public key data.

2. Introduction

The popularity of public key cryptography has produced a desire for its support in Kerberos [2]. The advantages provided by public key cryptography include simplified key management (from the Kerberos perspective) and the ability to leverage existing and developing public key certification infrastructures.

Public key cryptography can be integrated into Kerberos in a number of ways. One is to associate a key pair with each realm, which can then be used to facilitate cross-realm authentication; this is the topic of another draft proposal. Another way is to allow users with public key certificates to use them in initial authentication. This is the concern of the current document.

PKINIT utilizes ephemeral-ephemeral Diffie-Hellman keys in combination with DSA keys as the primary, required mechanism. Note that PKINIT supports the use of separate signature and encryption keys.

PKINIT enables access to Kerberos-secured services based on initial authentication utilizing public key cryptography. PKINIT utilizes standard public key signature and encryption data formats within the standard Kerberos messages. The basic mechanism is as follows: The user sends an AS-REQ message to the KDC as before, except that if that user is to use public key cryptography in the initial authentication step, his certificate and a signature accompany the initial request in the preauthentication fields. Upon receipt of this request, the KDC verifies the certificate and issues a ticket granting ticket (TGT) as before, except that the encPart from the AS-REP message carrying the TGT is now encrypted utilizing either a Diffie-Hellman derived key or the user's public key. This message is authenticated utilizing the public key signature of the KDC.

Note that PKINIT does not require the use of certificates. A KDC may store the public key of a principal as part of that principal's record. In this scenario, the KDC is the trusted party that vouches for the principal (as in a standard, non-cross realm, Kerberos environment). Thus, for any principal, the KDC may maintain a symmetric key, a public key, or both.

The PKINIT specification may also be used as a building block for other specifications. PKINIT may be utilized to establish inter-realm keys for the purposes of issuing cross-realm service tickets. It may also be used to issue anonymous Kerberos tickets using the Diffie-Hellman option. Efforts are under way to draft specifications for these two application protocols.

Additionally, the PKINIT specification may be used for direct peer to peer authentication without contacting a central KDC. This application of PKINIT is based on concepts introduced in [6, 7]. For direct client-to-server authentication, the client uses PKINIT to authenticate to the end server (instead of a central KDC), which then issues a ticket for itself. This approach has an advantage over TLS [5] in that the server does not need to save state (cache session keys). Furthermore, an additional benefit is that Kerberos tickets can facilitate delegation (see [6]).

3. Proposed Extensions

This section describes extensions to RFC 1510bis for supporting the use of public key cryptography in the initial request for a ticket granting ticket (TGT).

In summary, the following change to RFC 1510bis is proposed:

* Users may authenticate using either a public key pair or a conventional (symmetric) key. If public key cryptography is used, public key data is transported in preauthentication data fields to help establish identity. The user presents a public key certificate and obtains an ordinary TGT that may be used for subsequent authentication, with such authentication using only conventional cryptography.

Section 3.1 provides definitions to help specify message formats. Section 3.2 describes the extensions for the initial authentication method.

3.1. Definitions

The extensions involve new preauthentication fields; we introduce the following preauthentication types:

PA-PK-AS-REQ	14
PA-PK-AS-REP	15

The extensions also involve new error types; we introduce the following types:

KDC_ERR_CLIENT_NOT_TRUSTED	62
KDC_ERR_KDC_NOT_TRUSTED	63
KDC_ERR_INVALID_SIG	64
KDC_ERR_KEY_TOO_WEAK	65
KDC_ERR_CERTIFICATE_MISMATCH	66
KDC_ERR_CANT_VERIFY_CERTIFICATE	70
KDC_ERR_INVALID_CERTIFICATE	71
KDC_ERR_REVOKED_CERTIFICATE	72
KDC_ERR_REVOCATION_STATUS_UNKNOWN	73
KDC_ERR_REVOCATION_STATUS_UNAVAILABLE	74
KDC_ERR_CLIENT_NAME_MISMATCH	75
KDC_ERR_KDC_NAME_MISMATCH	76

We utilize the following typed data for errors:

TD-PKINIT-CMS-CERTIFICATES	101
TD-KRB-PRINCIPAL	102
TD-KRB-REALM	103
TD-TRUSTED-CERTIFIERS	104
TD-CERTIFICATE-INDEX	105

We utilize the following encryption types (which map directly to OIDs):

dsaWithSHA1-CmsOID	9
md5WithRSAEncryption-CmsOID	10
sha1WithRSAEncryption-CmsOID	11
rc2CBC-EnvOID	12
rsaEncryption-EnvOID (PKCS#1 v1.5)	13
rsaES-OAEP-ENV-OID (PKCS#1 v2.0)	14
des-ede3-cbc-Env-OID	15

These mappings are provided so that a client may send the appropriate enctypes in the AS-REQ message in order to indicate support for the corresponding OIDs (for performing PKINIT).

In many cases, PKINIT requires the encoding of the X.500 name of a certificate authority as a Realm. When such a name appears as a realm it will be represented using the "Other" form of the realm name as specified in the naming constraints section of RFC 1510bis. For a realm derived from an X.500 name, NAMETYPE will have the value X500-RFC2253. The full realm name will appear as follows:

```
<nametype> + ":" + <string>
```

where nametype is "X500-RFC2253" and string is the result of doing an RFC2253 encoding of the distinguished name, i.e.

```
    "X500-RFC2253:" + RFC2253Encode(DistinguishedName)
```

where DistinguishedName is an X.500 name, and RFC2253Encode is a function returing a readable UTF encoding of an X.500 name, as defined by RFC 2253 [11] (part of LDAPv3 [15]).

To ensure that this encoding is unique, we add the following rule to those specified by RFC 2253:

> The order in which the attributes appear in the RFC 2253 encoding MUST be the reverse of the order in the ASN.1 encoding of the X.500 name that appears in the public key certificate. The order of the relative distinguished names (RDNs), as well as the order of the AttributeTypeAndValues within each RDN, will be reversed. (This is despite the fact that an RDN is defined as a SET of AttributeTypeAndValues, where an order is normally not important.)

Similarly, in cases where the KDC does not provide a specific policy-based mapping from the X.500 name or X.509 Version 3 SubjectAltName extension in the user's certificate to a Kerberos principal name, PKINIT requires the direct encoding of the X.500 name as a PrincipalName. In this case, the name-type of the principal name MUST be set to KRB_NT-X500-PRINCIPAL. This new name type is defined in RFC 1510bis as:

```
    KRB_NT_X500_PRINCIPAL    6
```

For this type, the name-string MUST be set as follows:

```
    RFC2253Encode(DistinguishedName)
```

as described above. When this name type is used, the principal's realm MUST be set to the certificate authority's distinguished name using the X500-RFC2253 realm name format described earlier in this section.

RFC 1510bis specifies the ASN.1 structure for PrincipalName as follows:

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

The following rules relate to the matching of PrincipalNames with regard to the PKI name constraints for CAs as laid out in RFC 2459 [12]. In order to be regarded as a match (for permitted and excluded name trees), the following MUST be satisfied.

1. If the constraint is given as a user plus realm name, or as a client principal name plus realm name (as specified in RFC 1510bis), the realm name MUST be valid (see 2.a-d below) and the match MUST be exact, byte for byte.

2. If the constraint is given only as a realm name, matching depends on the type of the realm:

   a. If the realm contains a colon (':') before any equal sign ('='), it is treated as a realm of type Other, and MUST match exactly, byte for byte.

   b. Otherwise, if the realm name conforms to rules regarding the format of DNS names, it is considered a realm name of type Domain. The constraint may be given as a realm name 'FOO.BAR', which matches any PrincipalName within the realm 'FOO.BAR' but not those in subrealms such as 'CAR.FOO.BAR'. A constraint of the form '.FOO.BAR' matches PrincipalNames in subrealms of the form 'CAR.FOO.BAR' but not the realm 'FOO.BAR' itself.

   c. Otherwise, the realm name is invalid and does not match under any conditions.

3.1.1. Encryption and Key Formats

In the exposition below, we use the terms public key and private key generically. It should be understood that the term "public key" may be used to refer to either a public encryption key or a signature verification key, and that the term "private key" may be used to refer to either a private decryption key or a signature generation key. The fact that these are logically distinct does not preclude the assignment of bitwise identical keys for RSA keys.

In the case of Diffie-Hellman, the key is produced from the agreed bit string as follows:

* Truncate the bit string to the appropriate length.
* Rectify parity in each byte (if necessary) to obtain the key.

For instance, in the case of a DES key, we take the first eight bytes of the bit stream, and then adjust the least significant bit of each byte to ensure that each byte has odd parity. Appropriate key constraints for each valid cryptosystem are given in RFC 1510bis.

3.1.2. Algorithm Identifiers

PKINIT does not define, but does permit, the algorithm identifiers listed below.

3.1.2.1. Signature Algorithm Identifiers

The following signature algorithm identifiers specified in [8] and in [12] are used with PKINIT:

id-dsa-with-sha1 (DSA with SHA1)
md5WithRSAEncryption (RSA with MD5)
sha-1WithRSAEncryption (RSA with SHA1)

3.1.2.2 Diffie-Hellman Key Agreement Algorithm Identifier

The following algorithm identifier shall be used within the SubjectPublicKeyInfo data structure: dhpublicnumber

This identifier and the associated algorithm parameters are specified in RFC 2459 [12].

3.1.2.3. Algorithm Identifiers for RSA Encryption

These algorithm identifiers are used inside the EnvelopedData data structure, for encrypting the temporary key with a public key:

rsaEncryption (RSA encryption, PKCS#1 v1.5)
id-RSAES-OAEP (RSA encryption, PKCS#1 v2.0)

Both of the above RSA encryption schemes are specified in [13]. Currently, only PKCS#1 v1.5 is specified by CMS [8], although the CMS specification says that it will likely include PKCS#1 v2.0 in the future. (PKCS#1 v2.0 addresses adaptive chosen ciphertext vulnerability discovered in PKCS#1 v1.5.)

3.1.2.4. Algorithm Identifiers for Encryption with Secret Keys

These algorithm identifiers are used inside the EnvelopedData data structure in the PKINIT Reply, for encrypting the reply key with the temporary key:

des-ede3-cbc (3-key 3-DES, CBC mode)
rc2-cbc (RC2, CBC mode)

The full definition of the above algorithm identifiers and their corresponding parameters (an IV for block chaining) is provided in the CMS specification [8].

3.2. Public Key Authentication

Implementation of the changes in this section is REQUIRED for compliance with PKINIT.

3.2.1. Client Request

Public keys may be signed by some certification authority (CA), or they may be maintained by the KDC in which case the KDC is the trusted authority. Note that the latter mode does not require the use of certificates.

The initial authentication request is sent as per RFC 1510bis, except that a preauthentication field containing data signed by the user's private key accompanies the request:

```
PA-PK-AS-REQ ::= SEQUENCE {
                              -- PA TYPE 14
    signedAuthPack          [0] SignedData
                                  -- Defined in CMS [8];
                                  -- AuthPack (below) defines the
                                  -- data that is signed.
    trustedCertifiers       [1] SEQUENCE OF TrustedCas OPTIONAL,
                                  -- This is a list of CAs that the
                                  -- client trusts and that certify
                                  -- KDCs.
    kdcCert                 [2] IssuerAndSerialNumber OPTIONAL
                                  -- As defined in CMS [8];
                                  -- specifies a particular KDC
                                  -- certificate if the client
                                  -- already has it.
    encryptionCert          [3] IssuerAndSerialNumber OPTIONAL
                                  -- For example, this may be the
                                  -- client's Diffie-Hellman
                                  -- certificate, or it may be the
                                  -- client's RSA encryption
                                  -- certificate.
```

```
}

TrustedCas ::= CHOICE {
    principalName         [0] KerberosName,
                              -- as defined below
    caName                [1] Name
                              -- fully qualified X.500 name
                              -- as defined by X.509
    issuerAndSerial       [2] IssuerAndSerialNumber
                              -- Since a CA may have a number of
                              -- certificates, only one of which
                              -- a client trusts
}
```

Usage of SignedData:

The SignedData data type is specified in the Cryptographic Message Syntax, a product of the S/MIME working group of the IETF. The following describes how to fill in the fields of this data:

1. The encapContentInfo field MUST contain the PKAuthenticator and, optionally, the client's Diffie Hellman public value.

    a. The eContentType field MUST contain the OID value for pkauthdata: iso (1) org (3) dod (6) internet (1) security (5) kerberosv5 (2) pkinit (3) pkauthdata (1)

    b. The eContent field is data of the type AuthPack (below).

2. The signerInfos field contains the signature of AuthPack.

3. The Certificates field, when non-empty, contains the client's certificate chain. If present, the KDC uses the public key from the client's certificate to verify the signature in the request. Note that the client may pass different certificate chains that are used for signing or for encrypting. Thus, the KDC may utilize a different client certificate for signature verification than the one it uses to encrypt the reply to the client. For example, the client may place a Diffie-Hellman certificate in this field in order to convey its static Diffie Hellman certificate to the KDC to enable static-ephemeral Diffie-Hellman mode for the reply; in this case, the client does NOT place its public value in the AuthPack (defined below). As another example, the client may place an RSA encryption certificate in this field. However, there MUST always be (at least) a signature certificate.

4. When a DH key is being used, the public exponent is provided in the subjectPublicKey field of the SubjectPublicKeyInfo and the DH parameters are supplied as a DHParameter in the AlgorithmIdentitfier parameters. The DH paramters SHOULD be chosen from the First and Second defined Oakley Groups [The Internet Key Exchange (IKE) RFC-2409], if a server will not accept either of these groups, it will respond with a krb-error of KDC_ERR_KEY_TOO_WEAK and the e_data will contain a DHParameter with appropriate parameters for the client to use.

5. The KDC may wish to use cached Diffie-Hellman parameters (see Section 3.2.2, KDC Response). To indicate acceptance of cached parameters, the client sends zero in the nonce field of the PKAuthenticator. Zero is not a valid value for this field under any other circumstances. If cached parameters are used, the client and the KDC MUST perform key derivation (for the appropriate cryptosystem) on the resulting encryption key, as specified in RFC 1510bis. (With a zero nonce, message binding is performed using the nonce in the main request, which must be encrypted using the encapsulated reply key.)

```
AuthPack ::= SEQUENCE {
    pkAuthenticator         [0] PKAuthenticator,
    clientPublicValue       [1] SubjectPublicKeyInfo OPTIONAL
                                -- if client is using Diffie-Hellman
                                -- (ephemeral-ephemeral only)
}

PKAuthenticator ::= SEQUENCE {
    cusec                   [0] INTEGER,
                                -- for replay prevention as in RFC 1510bis
    ctime                   [1] KerberosTime,
                                -- for replay prevention as in RFC 1510bis
    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
}

SubjectPublicKeyInfo ::= SEQUENCE {
    algorithm                   AlgorithmIdentifier,
                                -- dhKeyAgreement
    subjectPublicKey            BIT STRING
                                -- for DH, equals
                                -- public exponent (INTEGER encoded
                                -- as payload of BIT STRING)
}   -- as specified by the X.509 recommendation [7]

AlgorithmIdentifier ::= SEQUENCE {
    algorithm                   OBJECT IDENTIFIER,
                                -- for dhKeyAgreement, this is
                                -- { iso (1) member-body (2) US (840)
                                -- rsadsi (113459) pkcs (1) 3 1 }
                                -- from PKCS #3 [17]
```

```
    parameters                  ANY DEFINED by algorithm OPTIONAL
                                -- for dhKeyAgreement, this is
                                -- DHParameter
}   -- as specified by the X.509 recommendation [7]

DHParameter ::= SEQUENCE {
    prime                       INTEGER,
                                -- p
    base                        INTEGER,
                                -- g
    privateValueLength          INTEGER OPTIONAL
                                -- l
}   -- as defined in PKCS #3 [17]
```

If the client passes an issuer and serial number in the request, the KDC is requested to use the referred-to certificate. If none exists, then the KDC returns an error of type KDC_ERR_CERTIFICATE_MISMATCH. It also returns this error if, on the other hand, the client does not pass any trustedCertifiers, believing that it has the KDC's certificate, but the KDC has more than one certificate. The KDC should include information in the KRB-ERROR message that indicates the KDC certificate(s) that a client may utilize. This data is specified in the e-data, which is defined in RFC 1510bis revisions as a SEQUENCE of TypedData:

```
TypedData ::=  SEQUENCE {
                data-type      [0] INTEGER,
                data-value     [1] OCTET STRING,
} -- per Kerberos RFC 1510bis
```

where:

data-type = TD-PKINIT-CMS-CERTIFICATES = 101

data-value = CertificateSet // as specified by CMS [8]

The PKAuthenticator carries information to foil replay attacks, to bind the pre-authentication data to the KDC-REQ-BODY, and to bind the request and response. The PKAuthenticator is signed with the client's signature key.

3.2.2. KDC Response

Upon receipt of the AS_REQ with PA-PK-AS-REQ pre-authentication type, the KDC attempts to verify the user's certificate chain (userCert), if one is provided in the request. This is done by verifying the certification path against the KDC's policy of legitimate certifiers.

If the client's certificate chain contains no certificate signed by a CA trusted by the KDC, then the KDC sends back an error message of type KDC_ERR_CANT_VERIFY_CERTIFICATE. The accompanying e-data is a SEQUENCE of one TypedData (with type TD-TRUSTED-CERTIFIERS=104) whose data-value is an OCTET STRING which is the DER encoding of

```
TrustedCertifiers ::= SEQUENCE OF PrincipalName
                      -- X.500 name encoded as a principal name
                      -- see Section 3.1
```

If while verifying a certificate chain the KDC determines that the signature on one of the certificates in the CertificateSet from the signedAuthPack fails verification, then the KDC returns an error of type KDC_ERR_INVALID_CERTIFICATE. The accompanying e-data is a SEQUENCE of one TypedData (with type TD-CERTIFICATE-INDEX=105) whose data-value is an OCTET STRING which is the DER encoding of the index into the CertificateSet ordered as sent by the client.

```
CertificateIndex  ::= INTEGER
                      -- 0 = 1st certificate,
                      --     (in order of encoding)
                      -- 1 = 2nd certificate, etc
```

The KDC may also check whether any of the certificates in the client's chain has been revoked. If one of the certificates has been revoked, then the KDC returns an error of type KDC_ERR_REVOKED_CERTIFICATE; if such a query reveals that the certificate's revocation status is unknown or not available, then if required by policy, the KDC returns the appropriate error of type KDC_ERR_REVOCATION_STATUS_UNKNOWN or KDC_ERR_REVOCATION_STATUS_UNAVAILABLE. In any of these three cases, the affected certificate is identified by the accompanying e-data, which contains a CertificateIndex as described for KDC_ERR_INVALID_CERTIFICATE.

If the certificate chain can be verified, but the name of the client in the certificate does not match the client's name in the request, then the KDC returns an error of type KDC_ERR_CLIENT_NAME_MISMATCH. There is no accompanying e-data field in this case.

Even if all succeeds, the KDC may--for policy reasons--decide not to trust the client. In this case, the KDC returns an error message of type KDC_ERR_CLIENT_NOT_TRUSTED. One specific case of this is the presence or absence of an Enhanced Key Usage (EKU) OID within the certificate extensions. The rules regarding acceptability of an EKU sequence (or the absence of any sequence) are a matter of local policy. For the benefit of implementers, we define a PKINIT EKU OID as the following: iso (1) org (3) dod (6) internet (1) security (5) kerberosv5 (2) pkinit (3) pkekuoid (2).

If a trust relationship exists, the KDC then verifies the client's signature on AuthPack. If that fails, the KDC returns an error message of type KDC_ERR_INVALID_SIG. Otherwise, the KDC uses the timestamp (ctime and cusec) in the PKAuthenticator to assure that the request is not a replay. The KDC also verifies that its name is specified in the PKAuthenticator.

If the clientPublicValue field is filled in, indicating that the client wishes to use Diffie-Hellman key agreement, then the KDC checks to see that the parameters satisfy its policy. If they do not (e.g. the prime size is insufficient for the expected encryption type), then the KDC sends back an error message of type KDC_ERR_KEY_TOO_WEAK, with an e-data containing a structure of type DHParameter with appropriate DH parameters for the client to retry the request. Otherwise, it generates its own public and private values for the response.

The KDC also checks that the timestamp in the PKAuthenticator is within the allowable window and that the principal name and realm are correct. If the local (server) time and the client time in the authenticator differ by more than the allowable clock skew, then the KDC returns an error message of type KRB_AP_ERR_SKEW as defined in RFC 1510bis.

Assuming no errors, the KDC replies as per RFC 1510bis, except as follows. The user's name in the ticket is determined by the following decision algorithm:

1. If the KDC has a mapping from the name in the certificate to a Kerberos name, then use that name.
   Else
2. If the certificate contains the SubjectAltName extention and the local KDC policy defines a mapping from the SubjectAltName to a Kerberos name, then use that name.
   Else
3. Use the name as represented in the certificate, mapping as necessary (e.g. as per RFC 2253 for X.500 names). In this case the realm in the ticket MUST be the name of the certifier that issued the user's certificate.

Note that a principal name may be carried in the subjectAltName field of a certificate. This name may be mapped to a principal record in a security database based on local policy, for example the subjectAltName may be kerberos/principal@realm format. In this case the realm name is not that of the CA but that of the local realm doing the mapping (or some realm name chosen by that realm).

If a non-KDC X.509 certificate contains the principal name within the subjectAltName version 3 extension, that name may utilize KerberosName as defined below, or, in the case of an S/MIME certificate [14], may utilize the email address. If the KDC is presented with an S/MIME certificate, then the email address within subjectAltName will be interpreted as a principal and realm separated by the "@" sign, or as a name that needs to be mapped according to local policy. If the resulting name does not correspond to a registered principal name, then the principal name is formed as defined in section 3.1.

The trustedCertifiers field contains a list of certification authorities trusted by the client, in the case that the client does not possess the KDC's public key certificate. If the KDC has no certificate signed by any of the trustedCertifiers, then it returns an error of type KDC_ERR_KDC_NOT_TRUSTED.

KDCs should try to (in order of preference):

1. Use the KDC certificate identified by the serialNumber included in the client's request.
2. Use a certificate issued to the KDC by one of the client's trustedCertifier(s);

If the KDC is unable to comply with any of these options, then the KDC returns an error message of type KDC_ERR_KDC_NOT_TRUSTED to the client.

```
The KDC encrypts the reply not with the user's long-term key, but
with the Diffie Hellman derived key or a random key generated
for this particular response which is carried in the padata field of
the TGS-REP message.

PA-PK-AS-REP ::= CHOICE {
                        -- PA TYPE 15
    dhSignedData       [0] SignedData,
                        -- Defined in CMS and used only with
                        -- Diffie-Hellman key exchange (if the
                        -- client public value was present in the
                        -- request).
                        -- This choice MUST be supported
                        -- by compliant implementations.
    encKeyPack         [1] EnvelopedData,
                        -- Defined in CMS
                        -- The temporary key is encrypted
                        -- using the client public key
                        -- key
                        -- SignedReplyKeyPack, encrypted
                        -- with the temporary key, is also
                        -- included.
}

Usage of SignedData:

    When the Diffie-Hellman option is used, dhSignedData in
    PA-PK-AS-REP provides authenticated Diffie-Hellman parameters
    of the KDC.  The reply key used to encrypt part of the KDC reply
    message is derived from the Diffie-Hellman exchange:

    1.  Both the KDC and the client calculate a secret value
        (g^ab mod p), where a is the client's private exponent and
        b is the KDC's private exponent.
```

2. Both the KDC and the client take the first N bits of this secret value and convert it into a reply key. N depends on the reply key type.

   a. For example, if the reply key is DES, N=64 bits, where some of the bits are replaced with parity bits, according to FIPS PUB 74.

   b. As another example, if the reply key is (3-key) 3-DES, N=192 bits, where some of the bits are replaced with parity bits, according to FIPS PUB 74.

3. The encapContentInfo field MUST contain the KdcDHKeyInfo as defined below.

   a. The eContentType field MUST contain the OID value for pkdhkeydata: iso (1) org (3) dod (6) internet (1) security (5) kerberosv5 (2) pkinit (3) pkdhkeydata (2)

   b. The eContent field is data of the type KdcDHKeyInfo (below).

4. The certificates field MUST contain the certificates necessary for the client to establish trust in the KDC's certificate based on the list of trusted certifiers sent by the client in the PA-PK-AS-REQ. This field may be empty if the client did not send to the KDC a list of trusted certifiers (the trustedCertifiers field was empty, meaning that the client already possesses the KDC's certificate).

5. The signerInfos field is a SET that MUST contain at least one member, since it contains the actual signature.

6. If the client indicated acceptance of cached Diffie-Hellman parameters from the KDC, and the KDC supports such an option (for performance reasons), the KDC should return a zero in the nonce field and include the expiration time of the parameters in the dhKeyExpiration field. If this time is exceeded, the client SHOULD NOT use the reply. If the time is absent, the client SHOULD NOT use the reply and MAY resubmit a request with a non-zero nonce (thus indicating non-acceptance of cached Diffie-Hellman parameters). As indicated above in Section 3.2.1, Client Request, when the KDC uses cached parameters, the client and the KDC MUST perform key derivation (for the appropriate cryptosystem) on the resulting encryption key, as specified in RFC 1510bis.

```
KdcDHKeyInfo ::= SEQUENCE {
                          -- used only when utilizing Diffie-Hellman
  subjectPublicKey      [0] BIT STRING,
                            -- 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
}
```

Usage of EnvelopedData:

The EnvelopedData data type is specified in the Cryptographic Message Syntax, a product of the S/MIME working group of the IETF. It contains a temporary key encrypted with the PKINIT client's public key. It also contains a signed and encrypted reply key.

1. The originatorInfo field is not required, since that information may be presented in the signedData structure that is encrypted within the encryptedContentInfo field.

2. The optional unprotectedAttrs field is not required for PKINIT.

3. The recipientInfos field is a SET which MUST contain exactly one member of the KeyTransRecipientInfo type for encryption with a public key.

   a. The encryptedKey field (in KeyTransRecipientInfo) contains the temporary key which is encrypted with the PKINIT client's public key.

4. The encryptedContentInfo field contains the signed and encrypted reply key.

   a. The contentType field MUST contain the OID value for id-signedData: iso (1) member-body (2) us (840) rsadsi (113549) pkcs (1) pkcs7 (7) signedData (2)

   b. The encryptedContent field is encrypted data of the CMS type signedData as specified below.

      i. The encapContentInfo field MUST contains the ReplyKeyPack.

         * The eContentType field MUST contain the OID value for pkrkeydata: iso (1) org (3) dod (6) internet (1) security (5) kerberosv5 (2) pkinit (3) pkrkeydata (3)

         * The eContent field is data of the type ReplyKeyPack (below).

ii. The certificates field MUST contain the certificates necessary for the client to establish trust in the KDC's certificate based on the list of trusted certifiers sent by the client in the PA-PK-AS-REQ. This field may be empty if the client did not send to the KDC a list of trusted certifiers (the trustedCertifiers field was empty, meaning that the client already possesses the KDC's certificate).

iii. The signerInfos field is a SET that MUST contain at least one member, since it contains the actual signature.

```
ReplyKeyPack ::= SEQUENCE {
                         -- not used for Diffie-Hellman
    replyKey             [0] EncryptionKey,
                             -- from RFC 1510bis
                             -- used to encrypt main reply
                             -- ENCTYPE is at least as strong as
                             -- ENCTYPE of session key
    nonce                [1] INTEGER,
                             -- binds response to the request
                             -- must be same as the nonce
                             -- passed in the PKAuthenticator
}
```

3.2.2.1. Use of transited Field

Since each certifier in the certification path of a user's certificate is equivalent to a separate Kerberos realm, the name of each certifier in the certificate chain MUST be added to the transited field of the ticket. The format of these realm names is defined in Section 3.1 of this document. If applicable, the transit-policy-checked flag should be set in the issued ticket.

3.2.2.2. Kerberos Names in Certificates

The KDC's certificate(s) MUST bind the public key(s) of the KDC to a name derivable from the name of the realm for that KDC. X.509 certificates MUST contain the principal name of the KDC (defined in RFC 1510bis) as the SubjectAltName version 3 extension. Below is the definition of this version 3 extension, as specified by the X.509 standard:

```
subjectAltName EXTENSION ::= {
    SYNTAX GeneralNames
    IDENTIFIED BY id-ce-subjectAltName
}

GeneralNames ::= SEQUENCE SIZE(1..MAX) OF GeneralName

GeneralName ::= CHOICE {
    otherName       [0] OtherName,
    ...
}

OtherName ::= SEQUENCE {
    type-id         OBJECT IDENTIFIER,
    value           [0] EXPLICIT ANY DEFINED BY type-id
}
```

For the purpose of specifying a Kerberos principal name, the value in OtherName MUST be a KerberosName as defined in RFC 1510bis:

```
KerberosName ::= SEQUENCE {
    realm           [0] Realm,
    principalName   [1] PrincipalName
}
```

This specific syntax is identified within subjectAltName by setting the type-id in OtherName to krb5PrincipalName, where (from the Kerberos specification) we have

```
krb5 OBJECT IDENTIFIER ::= { iso (1)
                             org (3)
                             dod (6)
                             internet (1)
                             security (5)
                             kerberosv5 (2) }

krb5PrincipalName OBJECT IDENTIFIER ::= { krb5 2 }
```

(This specification may also be used to specify a Kerberos name within the user's certificate.) The KDC's certificate may be signed directly by a CA, or there may be intermediaries if the server resides within a large organization, or it may be unsigned if the client indicates possession (and trust) of the KDC's certificate.

Note that the KDC's principal name has the instance equal to the realm, and those fields should be appropriately set in the realm and principalName fields of the KerberosName. This is the case even when obtaining a cross-realm ticket using PKINIT.

3.2.3. Client Extraction of Reply

The client then extracts the random key used to encrypt the main reply. This random key (in encPaReply) is encrypted with either the client's public key or with a key derived from the DH values exchanged between the client and the KDC. The client uses this random key to decrypt the main reply, and subsequently proceeds as described in RFC 1510bis.

3.2.4. Required Algorithms

Not all of the algorithms in the PKINIT protocol specification have to be implemented in order to comply with the proposed standard. Below is a list of the required algorithms:

* Diffie-Hellman public/private key pairs
    * utilizing Diffie-Hellman ephemeral-ephemeral mode
* SHA1 digest and DSA for signatures
* SHA1 digest also for the Checksum in the PKAuthenticator
* 3-key triple DES keys derived from the Diffie-Hellman Exchange
* 3-key triple DES Temporary and Reply keys

4. Logistics and Policy

This section describes a way to define the policy on the use of PKINIT for each principal and request.

The KDC is not required to contain a database record for users who use public key authentication. However, if these users are registered with the KDC, it is recommended that the database record for these users be modified to an additional flag in the attributes field to indicate that the user should authenticate using PKINIT. If this flag is set and a request message does not contain the PKINIT preauthentication field, then the KDC sends back as error of type KDC_ERR_PREAUTH_REQUIRED indicating that a preauthentication field of type PA-PK-AS-REQ must be included in the request.

5. Security Considerations

PKINIT raises a few security considerations, which we will address in this section.

First of all, PKINIT introduces a new trust model, where KDCs do not (necessarily) certify the identity of those for whom they issue tickets. PKINIT does allow KDCs to act as their own CAs, in the limited capacity of self-signing their certificates, but one of the additional benefits is to align Kerberos authentication with a global public key infrastructure. Anyone using PKINIT in this way must be aware of how the certification infrastructure they are linking to works.

Also, PKINIT introduces the possibility of interactions between different cryptosystems, which may be of widely varying strengths. Many systems, for instance, allow the use of 512-bit public keys. Using such keys to wrap data encrypted under strong conventional cryptosystems, such as triple-DES, is inappropriate; it adds a weak link to a strong one at extra cost. Implementors and administrators should take care to avoid such wasteful and deceptive interactions.

Care should be taken in how certificates are choosen for the purposes of authentication using PKINIT. Some local policies require that key escrow be applied for certain certificate types. People deploying PKINIT should be aware of the implications of using certificates that have escrowed keys for the purposes of authentication.

As described in Section 3.2, PKINIT allows for the caching of the Diffie-Hellman parameters on the KDC side, for performance reasons. For similar reasons, the signed data in this case does not vary from message to message, until the cached parameters expire. Because of the persistence of these parameters, the client and the KDC are to use the appropriate key derivation measures (as described in RFC 1510bis) when using cached DH parameters.

Lastly, PKINIT calls for randomly generated keys for conventional cryptosystems. Many such systems contain systematically "weak" keys. PKINIT implementations MUST avoid use of these keys, either by discarding those keys when they are generated, or by fixing them in some way (e.g. by XORing them with a given mask). These precautions vary from system to system; it is not our intention to give an explicit recipe for them here.

6. Transport Issues

Certificate chains can potentially grow quite large and span several UDP packets; this in turn increases the probability that a Kerberos message involving PKINIT extensions will be broken in transit. In light of the possibility that the Kerberos specification will require KDCs to accept requests using TCP as a transport mechanism, we make the same recommendation with respect to the PKINIT extensions as well.

7. Bibliography

[1] J. Kohl, C. Neuman. The Kerberos Network Authentication Service (V5). Request for Comments 1510.

[2] B.C. Neuman, Theodore Ts'o. Kerberos: An Authentication Service for Computer Networks, IEEE Communications, 32(9):33-38. September 1994.

[3] M. Sirbu, J. Chuang. Distributed Authentication in Kerberos Using Public Key Cryptography. Symposium On Network and Distributed System Security, 1997.

[4] B. Cox, J.D. Tygar, M. Sirbu. NetBill Security and Transaction Protocol. In Proceedings of the USENIX Workshop on Electronic Commerce, July 1995.

[5] T. Dierks, C. Allen. The TLS Protocol, Version 1.0 Request for Comments 2246, January 1999.

[6] B.C. Neuman, Proxy-Based Authorization and Accounting for Distributed Systems. In Proceedings of the 13th International Conference on Distributed Computing Systems, May 1993.

[7] ITU-T (formerly CCITT) Information technology - Open Systems Interconnection - The Directory: Authentication Framework Recommendation X.509 ISO/IEC 9594-8

[8] R. Housley. Cryptographic Message Syntax. draft-ietf-smime-cms-13.txt, April 1999, approved for publication as RFC.

[9] PKCS #7: Cryptographic Message Syntax Standard, An RSA Laboratories Technical Note Version 1.5 Revised November 1, 1993

[10] R. Rivest, MIT Laboratory for Computer Science and RSA Data Security, Inc. A Description of the RC2(r) Encryption Algorithm March 1998. Request for Comments 2268.

[11] M. Wahl, S. Kille, T. Howes. Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names. Request for Comments 2253.

[12] R. Housley, W. Ford, W. Polk, D. Solo. Internet X.509 Public Key Infrastructure, Certificate and CRL Profile, January 1999. Request for Comments 2459.

[13] B. Kaliski, J. Staddon. PKCS #1: RSA Cryptography Specifications, October 1998. Request for Comments 2437.

[14] S. Dusse, P. Hoffman, B. Ramsdell, J. Weinstein. S/MIME Version 2 Certificate Handling, March 1998. Request for Comments 2312.

[15] M. Wahl, T. Howes, S. Kille. Lightweight Directory Access Protocol (v3), December 1997. Request for Comments 2251.

[16] ITU-T (formerly CCITT) Information Processing Systems - Open Systems Interconnection - Specification of Abstract Syntax Notation One (ASN.1) Rec. X.680 ISO/IEC 8824-1

[17] PKCS #3: Diffie-Hellman Key-Agreement Standard, An RSA Laboratories Technical Note, Version 1.4, Revised November 1, 1993.

8. Acknowledgements

Some of the ideas on which this proposal is based arose during discussions over several years between members of the SAAG, the IETF CAT working group, and the PSRG, regarding integration of Kerberos and SPX. Some ideas have also been drawn from the DASS system. These changes are by no means endorsed by these groups. This is an attempt to revive some of the goals of those groups, and this proposal approaches those goals primarily from the Kerberos perspective. Lastly, comments from groups working on similar ideas in DCE have been invaluable.

9. Expiration Date

This draft expires January 15, 2002.

10. Authors

Brian Tung
Clifford Neuman
USC Information Sciences Institute
4676 Admiralty Way Suite 1001
Marina del Rey CA 90292-6695
Phone: +1 310 822 1511
E-mail: {brian, bcn}@isi.edu

Matthew Hur

Cisco Systems

500 108th Ave. NE, Suite 500

Bellevue, WA 98004

Phone: (408) 525-0034

E-Mail: mhur@cisco.com

Ari Medvinsky

Keen.com, Inc.

150 Independence Drive

Menlo Park CA 94025

Phone: +1 650 289 3134

E-mail: ari@keen.com

Sasha Medvinsky

Motorola

6450 Sequence Drive

San Diego, CA 92121

+1 858 404 2367

E-mail: smedvinsky@gi.com

John Wray

Iris Associates, Inc.

5 Technology Park Dr.

Westford, MA 01886

E-mail: John_Wray@iris.com

Jonathan Trostle

Cisco Systems

170 W. Tasman Dr.

San Jose, CA 95134

E-mail: jtrostle@cisco.com.

Annex D (normative): PKCROSS specification

The PKCROSS specification is currently still an IETF draft. The present document complies only with the version of the PKCROSS draft that is listed in this annex. The IPCablecom security experts will continue to track progress of the PKCROSS draft through the IETF and will advise the Study Group concerning how to best deal with the situation as the document becomes an RFC. The goal is to incorporate by reference and eliminate this annex.

INTERNET-DRAFT
draft-ietf-kerberos-pk-cross-07.txt
Updates: RFC 1510
expires May 15, 2001

Matthew Hur
Cisco Systems
Brian Tung
Tatyana Ryutov
Clifford Neuman
ISI
Ari Medvinsky
Keen.com
Gene Tsudik
UC Irvine
Bill Sommerfeld
Sun Microsystems

Public Key Cryptography for Cross-Realm Authentication in Kerberos

0. Status Of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC 2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.ietf.org (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

The distribution of this memo is unlimited. It is filed as draft-ietf-kerberos-pk-cross-07.txt, and expires May 15, 2001. Please send comments to the authors.

1. Abstract

This document defines extensions to the Kerberos protocol specification [1] to provide a method for using public key cryptography to enable cross-realm authentication. The methods defined here specify the way in which message exchanges are to be used to transport cross-realm secret keys protected by encryption under public keys certified as belonging to KDCs.

2. Introduction

Symmetric and asymmetric key systems may co-exist within hybrid architectures in order to leverage the advantages and mitigate issues within the respective systems. An example of a hybrid solution that may employ both symmetric and asymmetric technologies is Kerberos ciphersuites in TLS [KERBTLS] which utilizes the Kerberos protocol [KERB] [KERB94] in conjunction with TLS [TLS] which has commonly been thought of as a public key protocol.

The Kerberos can leverage the advantages provided by public key cryptography. PKINIT [PKINIT] describes the use of public key cryptography in the initial authentication exchange in Kerberos. PKTAPP [PKTAPP] describes how an application service can essentially issue a kerberos ticket to itself after utilizing public key cryptography for authentication. This specification describes the use of public key cryptography in cross-realm authentication.

Without the use of public key cryptography, administrators must maintain separate keys for every realm that wishes to exchange authentication information with another realm (which implies n(n-1) keys), or they must utilize a hierarchical arrangement of realms, which may increase network traffic and complicate the trust model by requiring evaluation of transited realms.

Even with the multi-hop cross-realm authentication, there must be some way to locate the path by which separate realms are to be transited. The current method, which makes use of the DNS-like realm names typical to Kerberos, requires trust of the intermediate KDCs.

PKCROSS utilizes a public key infrastructure (PKI) [X509] to simplify the administrative burden of maintaining cross-realm keys. Such usage leverages a PKI for a non-centrally-administratable environment (namely, inter-realm). Thus, a shared key for cross-realm authentication can be established for a set period of time, and a remote realm is able to issue policy information that is returned to itself when a client requests cross-realm authentication. Such policy information may be in the form of restrictions [NEUMAN]. Furthermore, these methods are transparent to the client; therefore, only the KDCs need to be modified to use them. In this way, we take advantage of the distributed trust management capabilities of public key cryptography while maintaining the advantages of localized trust management provided by Kerberos.

Although this specification utilizes the protocol specified in the PKINIT specification, it is not necessary to implement client changes in order to make use of the changes in this document.

3. Objectives

The objectives of this specification are as follows:

1. Simplify the administration required to establish Kerberos cross-realm keys.

2. Avoid modification of clients and application servers.

3. Allow remote KDC to control its policy on cross-realm keys shared between KDCs, and on cross-realm tickets presented by clients.

4. Remove any need for KDCs to maintain state about keys shared with other KDCs.

5. Leverage the work done for PKINIT to provide the public key protocol for establishing symmetric cross realm keys.

4. Definitions

The following notation is used throughout this specification:

KDC_l ..... local KDC
KDC_r ..... remote KDC
XTKT_(l,r) ..... PKCROSS ticket that the remote KDC issues to the local KDC
TGT_(c,r) ..... cross-realm TGT that the local KDC issues to the client for presentation to the remote KDC

This specification defines the following new types to be added to the Kerberos specification:

```
  PKCROSS kdc-options field in the AS_REQ is bit 9
  TE-TYPE-PKCROSS-KDC       2
  TE-TYPE-PKCROSS-CLIENT    3
```

This specification defines the following ASN.1 type for conveying policy information:

```
CrossRealmTktData ::= SEQUENCE OF TypedData
```

This specification defines the following types for policy information conveyed in CrossRealmTktData:

```
  PLC_LIFETIME              1
  PLC_SET_TKT_FLAGS         2
  PLC_NOSET_TKT_FLAGS       3
```

TicketExtensions are defined per the Kerberos specification [KERB-REV]:

```
TicketExtensions ::= SEQUENCE OF TypedData
  Where
    TypedData ::=   SEQUENCE {
        data-type[0]   INTEGER,
        data-value[1]  OCTET STRING OPTIONAL
    }
```

5. Protocol Specification

We assume that the client has already obtained a TGT. To perform cross-realm authentication, the client does exactly what it does with ordinary (i.e. non-public-key-enabled) Kerberos; the only changes are in the KDC; although the ticket which the client forwards to the remote realm may be changed. This is acceptable since the client treats the ticket as opaque.

5.1. Overview of Protocol

The basic operation of the PKCROSS protocol is as follows:

1. The client submits a request to the local KDC for credentials for the remote realm. This is just a typical cross realm request that may occur with or without PKCROSS.

2. The local KDC submits a PKINIT request to the remote KDC to obtain a "special" PKCROSS ticket. This is a standard PKINIT request, except that PKCROSS flag (bit 9) is set in the kdc-options field in the AS_REQ.

3. The remote KDC responds as per PKINIT, except that the ticket contains a TicketExtension, which contains policy information such as lifetime of cross realm tickets issued by KDC_l to a client. The local KDC must reflect this policy information in the credentials it forwards to the client. Call this ticket XTKT_(l,r) to indicate that this ticket is used to authenticate the local KDC to the remote KDC.

4. The local KDC passes a ticket, TGT_(c,r) (the cross realm TGT between the client and remote KDC), to the client. This ticket contains in its TicketExtension field the ticket, XTKT_(l,r), which contains the cross-realm key. The TGT_(c,r) ticket is encrypted using the key sealed in XTKT_(l,r). (The TicketExtension field is not encrypted.) The local KDC may optionally include another TicketExtension type that indicates the hostname and/or IP address for the remote KDC.

5. The client submits the request directly to the remote KDC, as before.

6. The remote KDC extracts XTKT_(l,r) from the TicketExtension in order to decrypt the encrypted part of TGT_(c,r).

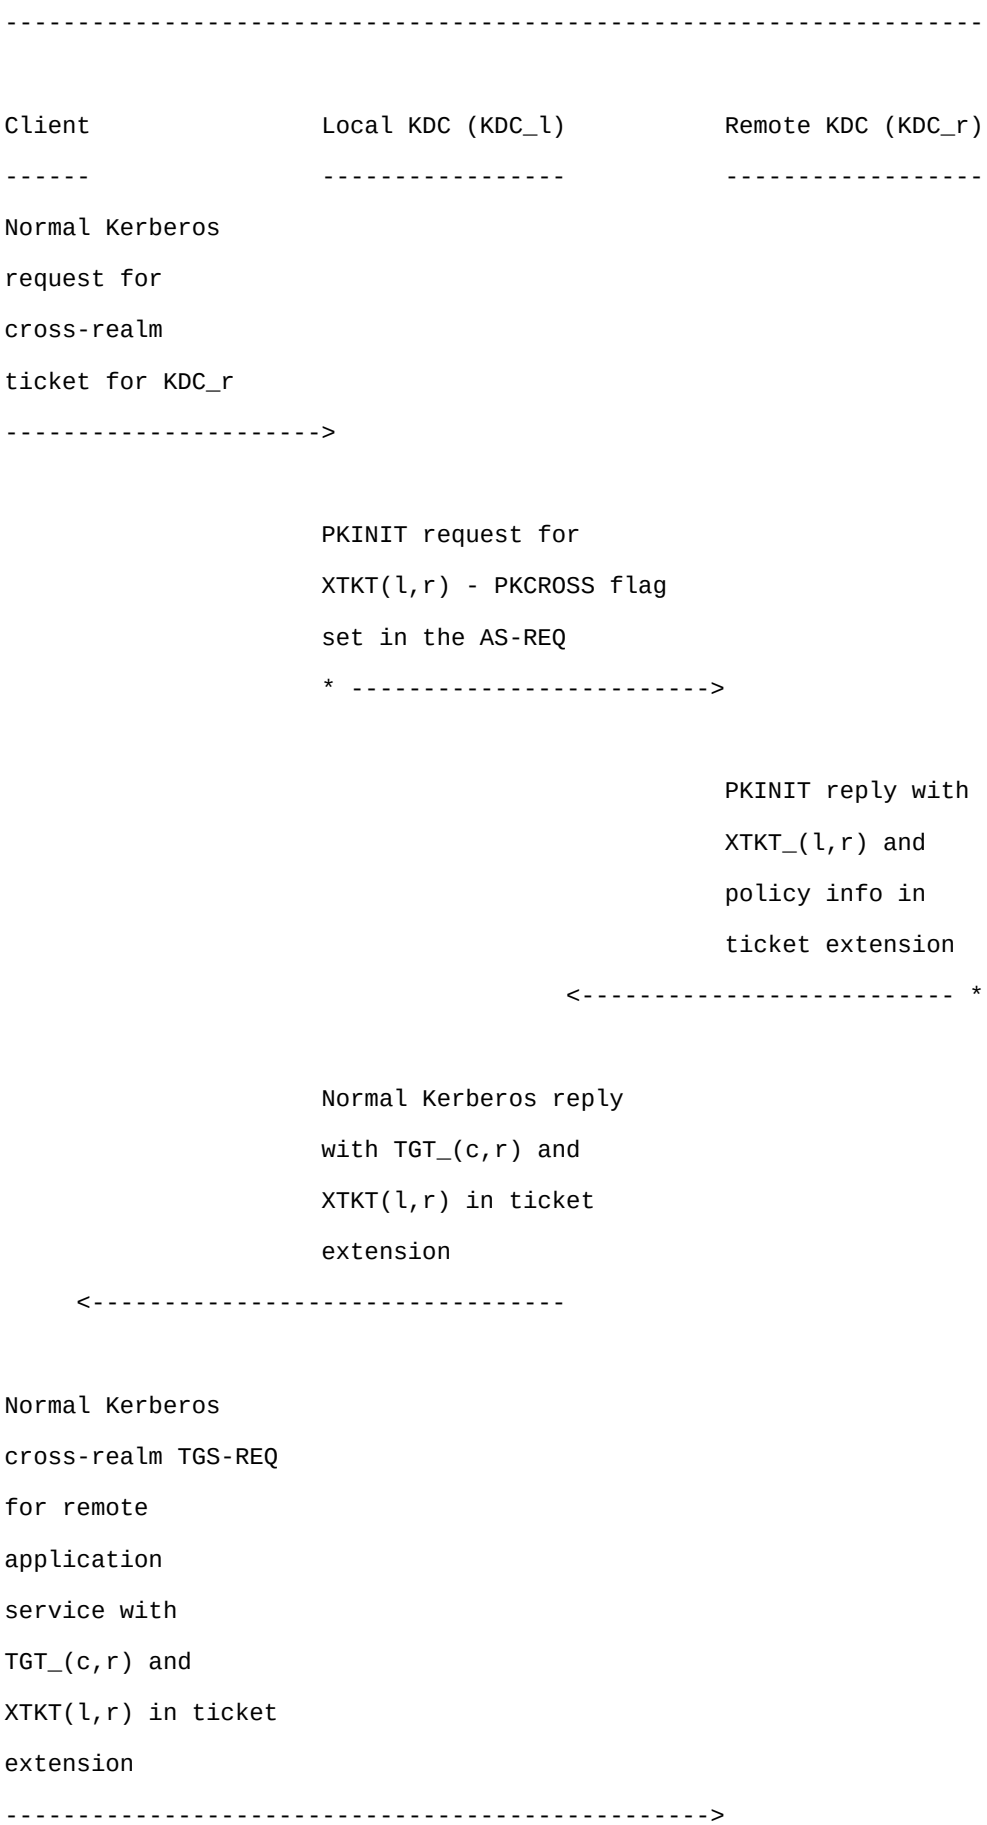
Client
Local KDC (KDC_l)
Remote KDC (KDC_r)
Normal Kerberos request for cross-realm ticket for KDC_r
PKINIT request for XTKT(l,r) - PKCROSS flag set in the AS-REQ
*
PKINIT reply with XTKT_(l,r) and policy info in ticket extension
*
Normal Kerberos reply with TGT_(c,r) and XTKT(l,r) in ticket extension
Normal Kerberos cross-realm TGS-REQ for remote application service with TGT_(c,r) and XTKT(l,r) in ticket extension

```
                                                    Normal Kerberos
                                                    cross-realm
                                                    TGS-REP
    <-------------------------------------------------------------

* Note that the KDC to KDC messages occur only periodically, since
  the local KDC caches the XTKT_(l,r).
-------------------------------------------------------------------
```

Sections 5.2 through 5.4 describe in detail steps 2 through 4 above. Section 5.6 describes the conditions under which steps 2 and 3 may be skipped.

Note that the mechanism presented above requires infrequent KDC to KDC communication (as dictated by policy - this is discussed later). Without such an exchange, there are the following issues:

1) KDC_l would have to issue a ticket with the expectation that KDC_r will accept it.
2) In the message that the client sends to KDC_r, KDC_l would have to authenticate KDC_r with credentials that KDC_r trusts.
3) There is no way for KDC_r to convey policy information to KDC_l.
4) If, based on local policy, KDC_r does not accept a ticket from KDC_l, then the client gets stuck in the middle. To address such an issue would require modifications to standard client processing behaviour.

Therefore, the infrequent use of KDC to KDC communication assures that inter-realm KDC keys may be established in accordance with local policies and that clients may continue to operate without modification.

5.2. Local KDC's Request to Remote KDC

When the local KDC receives a request for cross-realm authentication, it first checks its ticket cache to see if it has a valid PKCROSS ticket, XTKT_(l,r). If it has a valid XTKT_(l,r), then it does not need to send a request to the remote KDC (see section 5.5).

If the local KDC does not have a valid XTKT_(l,r), it sends a request to the remote KDC in order to establish a cross realm key and obtain the XTKT_(l,r). This request is in fact a PKINIT request as described in the PKINIT specification; i.e. it consists of an AS-REQ with a PA-PK-AS-REQ included as a preauthentication field. Note, that the AS-REQ MUST have the PKCROSS flag (bit 9) set in the kdc_options field of the AS-REQ. Otherwise, this exchange exactly follows the description given in the PKINIT specification.

5.3. Remote KDC's Response to Local KDC

When the remote KDC receives the PKINIT/PKCROSS request from the local KDC, it sends back a PKINIT response as described in the PKINIT specification with the following exception: the encrypted part of the Kerberos ticket is not encrypted with the krbtgt key; instead, it is encrypted with the ticket granting server's PKCROSS key. This key, rather than the krbtgt key, is used because it encrypts a ticket used for verifying a cross realm request rather than for issuing an application service ticket. Note that, as a matter of policy, the session key for the XTKT_(l,r) MAY be of greater strength than that of a session key for a normal PKINIT reply, since the XTKT_(l,r) SHOULD be much longer lived than a normal application service ticket.

In addition, the remote KDC SHOULD include policy information in the XTKT_(l,r). This policy information would then be reflected in the cross-realm TGT, TGT_(c,r). Otherwise, the policy for TGT_(c,r) would be dictated by KDC_l rather than by KDC_r. The local KDC MAY enforce a more restrictive local policy when creating a cross-realm ticket, TGT_(c,r). For example, KDC_r may dictate a lifetime policy of eight hours, but KDC_l may create TKT_(c,r) with a lifetime of four hours, as dictated by local policy. Also, the remote KDC MAY include other information about itself along with the PKCROSS ticket. These items are further discussed in section 6 below.

5.4. Local KDC's Response to Client

Upon receipt of the PKINIT/CROSS response from the remote KDC, the local KDC formulates a response to the client. This reply is constructed exactly as in the Kerberos specification, except for the following:

A) The local KDC places XTKT_(l,r) in the TicketExtension field of the client's cross-realm, ticket, TGT_(c,r), for the remote realm.
   Where
      data-type equals 3 for TE-TYPE-PKCROSS-CLIENT
      data-value is ASN.1 encoding of XTKT_(l,r)

B) The local KDC adds the name of its CA to the transited field of TGT_(c,r).

5.5 Remote KDC's Processing of Client Request

When the remote KDC, KDC_r, receives a cross-realm ticket, TGT_(c,r), and it detects that the ticket contains a ticket extension of type TE-TYPE-PKCROSS-CLIENT, KDC_r must first decrypt the ticket, XTKT_(l,r) that is encoded in the ticket extension. KDC_r uses its PKCROSS key in order to decrypt XTKT_(l,r). KDC_r then uses the key obtained from XTKT_(l,r) in order to decrypt the cross-realm ticket, TGT_(c,r).

KDC_r MUST verify that the cross-realm ticket, TGT_(c,r) is in compliance with any policy information contained in XTKT_(l,r) (see section 6). If the TGT_(c,r) is not in compliance with policy, then the KDC_r responds to the client with a KRB-ERROR message of type KDC_ERR_POLICY.

5.6. Short-Circuiting the KDC-to-KDC Exchange

As we described earlier, the KDC to KDC exchange is required only for establishing a symmetric, inter-realm key. Once this key is established (via the PKINIT exchange), no KDC to KDC communication is required until that key needs to be renewed. This section describes the circumstances under which the KDC to KDC exchange described in Sections 5.2 and 5.3 may be skipped.

The local KDC has a known lifetime for TGT_(c,r). This lifetime may be determined by policy information included in XTKT_(l,r), and/or it may be determined by local KDC policy. If the local KDC already has a ticket XTKT(l,r), and the start time plus the lifetime for TGT_(c,r) does not exceed the expiration time for XTGT_(l,r), then the local KDC may skip the exchange with the remote KDC, and issue a cross-realm ticket to the client as described in Section 5.4.

Since the remote KDC may change its PKCROSS key (referred to in Section 5.2) while there are PKCROSS tickets still active, it SHOULD cache the old PKCROSS keys until the last issued PKCROSS ticket expires. Otherwise, the remote KDC will respond to a client with a KRB-ERROR message of type KDC_ERR_TGT_REVOKED.

6. Extensions for the PKCROSS Ticket

As stated in section 5.3, the remote KDC SHOULD include policy information in XTKT_(l,r). This policy information is contained in a TicketExtension, as defined by the Kerberos specification, and the authorization data of the ticket will contain an authorization record of type AD-IN-Ticket-Extensions. The TicketExtension defined for use by PKCROSS is TE-TYPE-PKCROSS-KDC.

Where

data-type equals 2 for TE-TYPE-PKCROSS-KDC

data-value is ASN.1 encoding of CrossRealmTktData

CrossRealmTktData ::= SEQUENCE OF TypedData

CrossRealmTktData types and the corresponding data are interpreted as follows:

type	value	interpretation	ASN.1 data encoding
PLC_LIFETIME	1	lifetime (in seconds) for TGT_(c,r) - cross-realm tickets issued for clients by TGT_l	INTEGER

PLC_SET_TKT_FLAGS	2	TicketFlags that must be set - format defined by Kerberos specification	BITSTRING
PLC_NOSET_TKT_FLAGS	3	TicketFlags that must not be set - format defined by Kerberos specification	BITSTRING

Further types may be added to this table.

---

7. Usage of Certificates

In the cases of PKINIT and PKCROSS, the trust in a certification authority is equivalent to Kerberos cross realm trust. For this reason, an implementation MAY choose to use the same KDC certificate when the KDC is acting in any of the following three roles:

1) KDC is authenticating clients via PKINIT
2) KDC is authenticating another KDC for PKCROSS
3) KDC is the client in a PKCROSS exchange with another KDC

Note that per PKINIT, the KDC X.509 certificate (the server in a PKINIT exchange) MUST contain the principal name of the KDC in the subjectAltName field.

8. Transport Issues

Because the messages between the KDCs involve PKINIT exchanges, and PKINIT recommends TCP as a transport mechanism (due to the length of the messages and the likelihood that they will fragment), the same recommendation for TCP applies to PKCROSS as well.

9. Security Considerations

Since PKCROSS utilizes PKINIT, it is subject to the same security considerations as PKINIT. Administrators should assure adherence to security policy - for example, this affects the PKCROSS policies for cross realm key lifetime and for policy propagation from the PKCROSS ticket, issued from a remote KDC to a local KDC, to cross realm tickets that are issued by a local KDC to a client.

10. Bibliography

[KERBTLS] A. Medvinsky and M. Hur, "Addition of Kerberos Cipher Suites to Transport Layer Security (TLS)", RFC 2712, October 1999.

[KERB] J. Kohl and C. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.

[TLS] T. Dierks and C. Allen, "The TLS Protocol, Version 1.0", RFC 2246, January 1999.

[PKINIT] B. Tung, C. Neuman, M. Hur, A. Medvinsky, S. Medvinsky, J. Wray, J. Trostle. Public Key Cryptography for Initial Authentication in Kerberos. draft-ietf-cat-kerberos-pk-init-12.txt

[PKTAPP] A. Medvinsky, M. Hur, S. Medvinsky, C. Neuman. Public Key Utilizing Tickets for Application Servers (PKTAPP). draft-ietf-cat-kerberos-pk-tapp-03.txt

[X509] ITU-T (formerly CCITT) Information technology - Open Systems Interconnection - The Directory: Authentication Framework Recommendation X.509 ISO/IEC 9594-8

[NEUMAN] B.C. Neuman, "Proxy-Based Authorization and Accounting for Distributed Systems". Proceedings of the 13th International Conference on Distributed Computing Systems, May 1993

[KERB94] B.C. Neuman, Theodore Ts'o. Kerberos: An Authentication Service for Computer Networks, IEEE Communications, 32(9):33-38. September 1994.

[KERB-REV] C.Neuman, J. Kohl, T. Ts'o. The Kerberos Network Authentication Service (V5). draft-ietf-cat-kerberos-revisions-07.txt

11. Authors' Addresses

Matthew Hur
Cisco Systems
500 108th Ave. NE, Suite 500
Bellevue, WA 98004
Phone:
EMail: mhur@cisco.com

Brian Tung
Tatyana Ryutov
Clifford Neuman
USC/Information Sciences Institute
4676 Admiralty Way Suite 1001
Marina del Rey, CA 90292-6695
Phone: +1 310 822 1511
E-Mail: {brian, tryutov, bcn}@isi.edu

Ari Medvinsky
Keen.com
2480 Sand Hill Road, Suite 200
Menlo Park, CA 94025
Phone +1 650 289 3134
E-mail: ari@keen.com

Gene Tsudik

ICS Dept, 458 CS Building

Irvine CA 92697-3425

Phone: +1 310 448 9329

E-Mail: gts@ics.uci.edu

Bill Sommerfeld

Sun Microsystems

E-Mail: sommerfeld@east.sun.com

Annex E (normative): DNS locate specification

The DNS locate specification is currently still an IETF draft. The present document complies only with the version of the DNS locate draft that is listed in this annex. The IPCablecom security experts will continue to track progress of the DNS locate draft through the IETF and will advise the Study Group concerning how to best deal with the situation as the document becomes an RFC. The goal is to incorporate by reference and eliminate this annex.

INTERNET-DRAFT Ken Hornstein
<draft-ietf-cat-krb-dns-locate-02.txt> NRL
March 10, 2000 Jeffrey Altman
Expires: September 10, 2000 Columbia University

Distributing Kerberos KDC and Realm Information with DNS

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Distribution of this memo is unlimited. It is filed as <draft-ietf-cat-krb-dns-locate-02.txt>, and expires on September 10, 2000. Please send comments to the authors.

Abstract

Neither the Kerberos V5 protocol [RFC1510] nor the Kerberos V4 protocol [RFC????] describe any mechanism for clients to learn critical configuration information necessary for proper operation of the protocol. Such information includes the location of Kerberos key distribution centers or a mapping between DNS domains and Kerberos realms.

Current Kerberos implementations generally store such configuration information in a file on each client machine. Experience has shown this method of storing configuration information presents problems with out-of-date information and scaling problems, especially when

using cross-realm authentication.

This memo describes a method for using the Domain Name System [RFC1035] for storing such configuration information. Specifically, methods for storing KDC location and hostname/domain name to realm mapping information are discussed.

DNS vs. Kerberos - Case Sensitivity of Realm Names

In Kerberos, realm names are case sensitive. While it is strongly encouraged that all realm names be all upper case this recommendation has not been adopted by all sites. Some sites use all lower case names and other use mixed case. DNS on the other hand is case insensitive for queries but is case preserving for responses to TXT queries. Since "MYREALM", "myrealm", and "MyRealm" are all different it is necessary that the DNS entries be distinguishable.

Since the recommend realm names are all upper case, we will not require any quoting to be applied to upper case names. If the realm name contains lower case characters each character is to be quoted by a '=' character. So "MyRealm" would be represented as "M=yR=e=a=l=m" and "myrealm" as "=m=y=r=e=a=l=m". If the realm name contains the '=' character it will be represented as "==".

Overview - KDC location information

KDC location information is to be stored using the DNS SRV RR [RFC 2052]. The format of this RR is as follows:

Service.Proto.Realm TTL Class SRV Priority Weight Port Target

The Service name for Kerberos is always "_kerberos".

The Proto can be either "_udp" or "_tcp". If these records are to be used, a "_udp" record MUST be included. If the Kerberos implementation supports TCP transport, a "_tcp" record SHOULD be included.

The Realm is the Kerberos realm that this record corresponds to.

TTL, Class, SRV, Priority, Weight, Port, and Target have the standard meaning as defined in RFC 2052.

Example - KDC location information

These are DNS records for a Kerberos realm ASDF.COM. It has two Kerberos servers, kdc1.asdf.com and kdc2.asdf.com. Queries should be directed to kdc1.asdf.com first as per the specified priority.

Weights are not used in these records.

```
_kerberos._udp.ASDF.COM.        IN      SRV     0 0 88 kdc1.asdf.com.
_kerberos._udp.ASDF.COM.        IN      SRV     1 0 88 kdc2.asdf.com.
```

Overview - Kerberos password changing server location information

Kerberos password changing server [KERB-CHG] location is to be stored using the DNS SRV RR [RFC 2052]. The format of this RR is as follows:

Service.Proto.Realm TTL Class SRV Priority Weight Port Target

The Service name for the password server is always "_kpasswd".

The Proto MUST be "_udp".

The Realm is the Kerberos realm that this record corresponds to.

TTL, Class, SRV, Priority, Weight, Port, and Target have the standard meaning as defined in RFC 2052.

Overview - Kerberos admin server location information

Kerberos admin location information is to be stored using the DNS SRV RR [RFC 2052]. The format of this RR is as follows:

Service.Proto.Realm TTL Class SRV Priority Weight Port Target

The Service name for the admin server is always "_kerberos-adm".

The Proto can be either "_udp" or "_tcp". If these records are to be used, a "_tcp" record MUST be included. If the Kerberos admin implementation supports UDP transport, a "_udp" record SHOULD be included.

The Realm is the Kerberos realm that this record corresponds to.

TTL, Class, SRV, Priority, Weight, Port, and Target have the standard meaning as defined in RFC 2052.

Note that there is no formal definition of a Kerberos admin protocol, so the use of this record is optional and implementation-dependent.

Example - Kerberos administrative server location information

These are DNS records for a Kerberos realm ASDF.COM. It has one administrative server, kdc1.asdf.com.

_kerberos-adm._tcp.ASDF.COM. IN SRV 0 0 88 kdc1.asdf.com.

Overview - Hostname/domain name to Kerberos realm mapping

Information on the mapping of DNS hostnames and domain names to Kerberos realms is stored using DNS TXT records [RFC 1035]. These records have the following format.

Service.Name TTL Class TXT Realm

The Service field is always "_kerberos", and prefixes all entries of this type.

The Name is a DNS hostname or domain name. This is explained in greater detail below.

TTL, Class, and TXT have the standard DNS meaning as defined in RFC 1035.

The Realm is the data for the TXT RR, and consists simply of the Kerberos realm that corresponds to the Name specified.

When a Kerberos client wishes to utilize a host-specific service, it will perform a DNS TXT query, using the hostname in the Name field of the DNS query. If the record is not found, the first label of the name is stripped and the query is retried.

Compliant implementations MUST query the full hostname and the most specific domain name (the hostname with the first label removed). Compliant implementations SHOULD try stripping all subsequent labels until a match is found or the Name field is empty.

Example - Hostname/domain name to Kerberos realm mapping

For the previously mentioned ASDF.COM realm and domain, some sample records might be as follows:

```
_kerberos.asdf.com.              IN      TXT     "ASDF.COM"
_kerberos.mrkserver.asdf.com.    IN      TXT     "MARKETING.ASDF.COM"
_kerberos.salesserver.asdf.com. IN      TXT     "SALES.ASDF.COM"
```

Let us suppose that in this case, a Kerberos client wishes to use a Kerberized service on the host foo.asdf.com. It would first query:

```
_kerberos.foo.asdf.com. IN TXT
```

Finding no match, it would then query:

_kerberos.asdf.com. IN TXT

And find an answer of ASDF.COM. This would be the realm that foo.asdf.com resides in.

If another Kerberos client wishes to use a Kerberized service on the host salesserver.asdf.com, it would query:

_kerberos.salesserver.asdf.com IN TXT

And find an answer of SALES.ASDF.COM.

Security considerations

As DNS is deployed today, it is an unsecure service. Thus the information returned by it cannot be trusted.

Current practice for REALM to KDC mapping is to use hostnames to indicate KDC hosts (stored in some implementation-dependent location, but generally a local config file). These hostnames are vulnerable to the standard set of DNS attacks (denial of service, spoofed entries, etc). The design of the Kerberos protocol limits attacks of this sort to denial of service. However, the use of SRV records does not change this attack in any way. They have the same vulnerabilities that already exist in the common practice of using hostnames for KDC locations.

Current practice for HOSTNAME to REALM mapping is to provide a local configuration of mappings of hostname or domain name to realm which are then mapped to KDCs. But this again is vulnerable to spoofing via CNAME records that point to hosts in other domains. This has the same effect as when a TXT record is spoofed. In a realm with no cross-realm trusts this is a DoS attack. However, when cross-realm trusts are used it is possible to redirect a client to use a compromised realm.

This is not an exploit of the Kerberos protocol but of the Kerberos trust model. The same can be done to any application that must resolve the hostname in order to determine which domain a non-FQDN belongs to.

Implementations SHOULD provide a way of specifying this information locally without the use of DNS. However, to make this feature worthwhile a lack of any configuration information on a client should be interpreted as permission to use DNS.

Expiration

This Internet-Draft expires on September 10, 2000.

References

[RFC1510]
The Kerberos Network Authentication System; Kohl, Newman; September 1993.

[RFC1035]
Domain Names - Implementation and Specification; Mockapetris; November 1987

[RFC2782]
A DNS RR for specifying the location of services (DNS SRV); Gulbrandsen, Vixie; February 2000

[KERB-CHG]
Kerberos Change Password Protocol; Horowitz; ftp://ds.internic.net/internet-drafts/draft-ietf-cat-kerb-chg-password-02.txt

Authors' Addresses

Ken Hornstein
US Naval Research Laboratory
Bldg A-49, Room 2
4555 Overlook Avenue
Washington DC 20375 USA

Phone: +1 (202) 404-4765
EMail: kenh@cmf.nrl.navy.mil

Jeffrey Altman

The Kermit Project

Columbia University

612 West 115th Street #716

New York NY 10025-7799 USA

Phone: +1 (212) 854-1344

EMail: jaltman@columbia.edu

Annex F (informative): IPCablecom Admin guidelines & best practices

This annex 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.

F.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 a provisioned parameter that can be use in one the following ways:

In the case 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.

Annex G (informative): Example of MMH algorithm implementation

This annex 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 specification.

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
```

Annex H (informative): Kerb error messages

Table H.1 shows the recommended handling of the Kerberos error situations as defined in security specification. The abbreviation EN is used for event notification. The abbreviation N/A stands for not applicable and indicates that IPCablecom does not support the indicated Kerberos error. A threshold designation on the sender event indicates that events should be throttled to prevent potential flooding at the SYSLOG server, at the trap/inform handler, or in the local log.

Table H.1. Kerberos error actions

Enumerated KDC_ERROR	Value	Explanation	Receiver| Action	Sender Event
KDC_ERR_NONE	0	No error	Valid response proceed	
KDC_ERR_NAME_EXP	1	Client's entry in database has expired. In the case of an MTA, the entry in the back office database expired.	This client will be unable to get any Kerberos tickets or establish any IPSec SAs. No automatic recovery is possible.	EN - Major
KDC_ERR_SERVICE_ EXP	2	Server's entry in database has expired	Indicate that security association is not valid and try getting a ticket for an alternate application server if possible	EN - Major
KDC_ERR_BAD_PVNO	3	Requested protocol version number not supported	Indicate that security association is not valid and use another KDC if possible	EN - Critical
KDC_ERR_C_OLD_ MAST_KVNO	4	Client's key encrypted in old master key	N/A	N/A
KDC_ERR_S_OLD_ MAST_KVNO	5	Server's key encrypted in old master key	N/A	N/A
KDC_ERR_C_PRINCIPAL _UNKNOWN	6	Client not found in Kerberos database. In the case of an MTA, the entry in the back office database is not found.	This client will be unable to get any Kerberos tickets or establish any IPSec SAs. No automatic recovery is possible.	EN - Major
KDC_ERR_S_PRINCIPAL _UNKNOWN	7	Server not found in Kerberos database	Indicate that security association is not valid and try getting a ticket for an alternate application server if possible	EN - Critical
KDC_ERR_PRINCIPAL_ NOT_UNIQUE	8	Multiple principal entries in database	Indicate that security association is not valid and use another KDC if possible	EN - Minor
KDC_ERR_NULL_KEY	9	The client has a null key	This error only applies to non-MTA clients (e.g. CMS getting a ticket for another CMS). This client will be unable to get any Kerberos tickets or establish any IPSec SAs. No automatic recovery is possible	EN - Major
KDC_ERR_NULL_KEY	9	The server has a null key	Indicate that security association is not valid and try getting a ticket for an alternate application server if possible	EN - Critical
KDC_ERR_CANNOT_ POSTDATE	10	Ticket not eligible for postdating	N/A - IPCablecom does not support postdating	N/A
KDC_ERR_NEVER_ VALID	11	Requested start time is later than end time	N/A -- IPCablecom clients are not allowed to specify start and end times.	N/A
KDC_ERR_POLICY	12	KDC policy rejects request	IKDC and client have incompatible software. No automatic recovery is possible.	EN - Major

Enumerated KDC_ERROR	Value	Explanation	Receiver| Action	Sender Event
KDC_ERR_BADOPTION	13	KDC cannot accommodate requested option	KDC and client have incompatible software. No automatic recovery is possible.	EN - Major
KDC_ERR_ETYPE_ NOSUPP	14	KDC has no support for encryption type	KDC and client have incompatible software. No automatic recovery is possible.	EN - Critical
KDC_ERR_SUMTYPE_ NOSUPP	15	KDC has no support for checksum type	KDC and client have incompatible software No automatic recovery is possible.	EN - Critical
KDC_ERR_PADATA_ TYPE_NOSUPP	16	KDC has no support for padata type	KDC and Client have incompatible software. No automatic recovery is possible.	EN - Major
KDC_ERR_TRTYPE_ NOSUPP	17	KDC has no support for transited type	N/A - IPCablecom does not support transit environments	N/A
KDC_ERR_CLIENT_ REVOKED	18	Client's credentials have been revoked	Fatal error - this client cannot get any more tickets.	EN - Major
KDC_ERR_SERVICE_ REVOKED	19	Credentials for server have been revoked	Stop using security association. Try getting a ticket for an alternate application server if possible.	EN - Critical
KDC_ERR_TGT_ REVOKED	20	TGT has been revoked	Send another AS Request to get another TGT.	EN - Major
KDC_ERR_CLIENT_ NOTYET	21	Client not yet valid try again later	This could occur if MTA was not yet valid in the back office subscriber database. (KDC gets KRB_MTAMAP_NOT_YET_VALID error when it queries the provisioning server for MTA FQDN.) The MTA should wait. Manual intervention may be required.	Thresholded EN - Minor
KDC_ERR_SERVICE_ NOTYET	22	Server not yet valid try again later	N/A	N/A
KDC_ERR_KEY_ EXPIRED	23	Password has expired change password to reset	N/A	N/A
KDC_ERR_PREAUTH_ FAILED	24	Pre authentication information was invalid	KDC and client software incompatibility. No automatic recovery is possible.	EN - Minor
KDC_ERR_PREAUTH_ REQUIRED	25	Additional pre authentication required [40]	KDC and client software incompatibility. No automatic recovery is possible.	EN - Minor
KDC_ERR_SERVER_NO MATCH	26	Requested server and ticket do not match	N/A - This error could occur in the user-to-user case as well as ticket renewal, forwarding, etc. None of these scenarios are supported by IPCablecom.	N/A
KDC_ERR_MUST_USE_U SER2USER	27	Server principal valid for user2user only	N/A - IPCablecom does not support user to user environments.	N/A
KDC_ERR_PATH_NOT_ ACCPETED	28	KDC Policy rejects transited path	N/A - IPCablecom does not support transited paths	N/A
KDC_ERR_SVC_ UNAVAILABLE	29	A service is not available	N/A - This error indicates that a KDC supports only AS Requests or only TGS Requests - not possible in IPCablecom.	N/A
KRB_AP_ERR_BAD_ INTEGRITY	31	Integrity check on decrypted field failed	This error code is never sent back to the client.	Threshold EN - Minor
KRB_AP_ERR_TKT_ EXPIRED	32	Ticket expired	Get a valid, non-expired ticket. If there is not a valid ticket in the local cache, get a fresh ticket with an AS Req or TGS Req. Then resend the reconstructed AP Req with the new ticket.	EN - Minor

Enumerated KDC_ERROR	Value	Explanation	Receiver| Action	Sender Event
KRB_AP_ERR_TKT_ NYV	33	Ticket not yet valid	This is due to a clock skew between the KDC and the application server - if they are not properly synchronized. The client can switch to an alternate application server, if possible.	EN- Major
KRB_AP_ERR_REPEAT	34	Request is a replay	This error code is never sent back to the client.	Threshold EN - Minor
KRB_AP_ERR_NOT_US	35	The ticket is not for us	This error code is never sent back to the client.	Threshold EN - Major
KRB_AP_ERR_ BADMATCH	36	Ticket and authenticator do not match	This error code is never sent back to the client.	Threshold EN - Minor
KRB_AP_ERR_SKEW	37	Clock skew too great	Adjust skew based on the information in the error message, save the clock skew for future requests and send another AS Request, TGS Request (see clause 6.4.3.1.2.1.1), or AP Request (see clause 6.5.5.3).	Threshold EN - Minor
KRB_AP_ERR_ BADADDR	38	Incorrect net address	N/A. In IPCablecom, IP address is not included in the tickets.	N/A
KRB_AP_ERR_ BADVERSION	39	Protocol version mismatch	This error code is never sent back to the client.	EN - Critical
KRB_AP_ERR_MSG_ TYPE	40	Invalid message type	This error code is never sent back to the client.	Threshold EN - Minor
KRB_AP_ERR_ MODIFIED	41	Message stream modified	N/A - IPCablecom does not support message streams	N/A
KRB_AP_ERR_ BADORDER	42	Message out of order	N/A - IPCablecom does not support sequence numbers	N/A
KRB_AP_ERR_BADKEYV ER	44	Specified version of key is not available	Get a new ticket with AS Req or TGS Req and resend the reconstructed AP Req with the new ticket (see clause 6.4.6.3.1).	EN - Minor
KRB_AP_ERR_NOKEY	45	Service key not available	This error code is never sent back to the client.	EN - Major
KRB_AP_ERR_MUT_ FAIL	46	Mutual authentication failed	N/A	N/A
KRB_AP_ERR_ BADDIRECTION	47	Incorrect message direction	N/A	N/A
KRB_AP_ERR_METHOD	48	Alternative authentication method required	N/A	N/A
KRB_AP_ERR_BADSEQ	49	Incorrect sequence number in message	N/A - IPCablecom does not support sequence numbers	N/A
KRB_AP_ERR_INAPP_ CKSUM	50	Inappropriate type of checksum in message	This error code is never sent back to the client.	Threshold EN - Minor
KRB_AP_PATH_NOT_ ACCEPTED	51	Policy rejects transited path	N/A - IPCablecom does not support transit environments	N/A
KRB_ERR_RESPONSE_ TOO_BIG	52	Response too big for UDP	N/A - Only UDP is supported. Long frames may use fragmentation	N/A
KRB_ERR_GENERIC	60	Generic error	An application server can send this error. The actual error code is SNMPv3 or IPSec-specific and is indicated in the e-data field as specified by IPCablecom.	Threshold EN - Minor
KRB_ERR_FIELD_ TOO LONG	61	Field is too long for this implementation	This error code is never sent back to the client.	Threshold EN - Minor

Enumerated KDC_ERROR	Value	Explanation	Receiver| Action	Sender Event
KDC_ERR_CLIENT_ NOT_TRUSTED	62	Client not trusted	No automatic recovery is possible. KDC database must be modified to correct problem.	EN - Major
KDC_ERR_KDC_NOT_ TRUSTED	63	KDC not trusted	N/A - This error could occur if the KDC did not have a certificate signed by a Certification Authority that is trusted by this client. This is not possible in IPCablecom.	N/A
KDC_ERR_INVALID_ SIG	64	Invalid Signature	The KDC should never reply with such an error code. An attacker could use it to experiment with different keys until the signature matches.	Thres-hold EN - Minor
KDC_ERR_KEY_TOO_ WEAK	65	Key chosen is very week	This could occur if the PKINIT client chose Diffie-Hellman Oakely group #1 (768-bit) and the KDC wants the client to use group #2 (1024-bit). The client must retry with another AS Request using DH Oakley group #2.	EN - Minor
KDC_ERR_ CERTIFICATE_ MISMATCH	66	Certificate mismatch	N/A. This could happen if a PKINIT Request specifies which certificate KDC should use in the reply (and KDC does not have that certificate). A IPCablecom client is not allowed to specify a KDC certificate.	N/A
KDC_ERR_CANT_ VERIFY_CERTIFICATE	70	Not possible to parse the manufacturer certificate	The MTA is out of service until it gets an updated manufacturer certificate from manufacturer.	EN - Critical
KDC_ERR_CANT_ VERIFY_CERTIFICATE	70	Not possible to parse the device certificate	The MTA is out of service until it gets an updated device certificate from manufacturer.	EN - Major
KDC_ERR_INVALID_ CERTIFICATE	71	Invalid manufacturer certificate	The MTA is out of service until it gets an updated manufacturer certificate from manufacturer.	EN - Critical f
KDC_ERR_INVALID_ CERTIFICATE	71	Invalid device certificate	The MTA is out of service until it gets an updated device certificate from manufacturer.	EN - Major
KDC_ERR_REVOKED_ CERTIFICATE	72	Manufacturer certificate made void	The MTA is out of service until it gets an updated manufacturer certificate from manufacturer	EN - Critical
KDC_ERR_REVOKED_ CERTIFICATE	72	Device certificate made void	The MTA is out of service until it gets an updated device certificate from manufacturer	EN - Major
KDC_ERR_ REVOCATION_STATUS_ UNKNOWN	73	Manufacturer revocation status unknown	The MTA is out of service until it gets an updated manufacturer certificate from manufacturer	EN - Critical
KDC_ERR_ REVOCATION_STATUS_ UNKNOWN	73	Device revocation status unknown	The MTA is out of service until it gets an updated device certificate from manufacturer	EN - Major
KDC_ERR_ REVOCATION_STATUS_ UNAVAILABLE	74	Revocation status not available	All MTAs are out of service until the KDC can re-establish connection with the CRL publisher.	EN - Critical

Enumerated KDC_ERROR	Value	Explanation	Receiver| Action	Sender Event
KDC_ERR_CLIENT_ NAME_MISMATCH	75	Client name mismatch	For an AS Request, this indicates that the specified MTA MAC address and FQDN do not match - could be a synchronization problem between the DHCP server and back office subscriber database. For a TGS Request, this could mean that the client FQDN changed since the time a TGT was created. The client should get another TGT.	EN - Minor
KDC_ERR_KDC_NAME_ MISMATCH	76	KDC name mismatch	N/A. Obsolete error code from the PKINIT draft.	N/A

Annex I (informative): Bibliography

- Planned ETSI Technical Specification TS 101 909-16: " Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 16:Signalling for Call Management Server".
- B. Schneier, "Applied Cryptography," John Wiley & Sons Inc, second edition, 1996.
- Randomness Recommendations for Security, IETF Draft (Donald Eastlake, Stephen Crocker and Jeff Schiller) Internet Proposed Standard, RFC 1750, 1994.
- How to Protect DES Against Exhaustive Key Search, J. Killian, P. Rogaway, (Edited version presented at Proceedings of Crypto '96), July 1997.
- S. Halevi and H. Krawczyk, "MMH: Software Message Authentication in Gbit/sec Rates," Proceedings of the 4th Workshop on Fast Software Encryption, (1997) vol. 1267 Springer-Verloag, pp. 172-189
- FIPS 140-1 (1994): "Federal Information Processing Standards, Security Requirements for Cryptographic Modules".
- IETF RFC 2205 (1997): "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification Proposed Standard".
- Draft ETSI Technical Specification TS 101 909-20: " Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 20: Lawful Interception".

History

Document history		
V1.1.1	July 2001	Publication
V1.2.1	July 2002	Publication