3GPP TSG SA WG3 Security — S3#28 **S3-030295**

6 – 9 May 2003, Berlin, Germany

3GPP TS ab.cde V0.32.0 (2003-054)

Technical Specification

3rd Generation Partnership Project; Technical Specification Group Service and System Aspects; Network Domain Security; Authentication Framework (Release x)

Keywords

<keyword[, keyword]>

3GPP

Postal address

3GPP support office address

650 Route des Lucioles - Sophia Antipolis
Valbonne - FRANCE
Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Internet

http://www.3gpp.org

Contents

Foreword

This Technical Specification has been produced by the 3rd Generation Partnership Project (3GPP).

The contents of the present document are subject to continuing work within the TSG and may change following formal TSG approval. Should the TSG modify the contents of the present document, it will be re-released by the TSG with an identifying change of release date and an increase in version number as follows:

Version x.y.z

where:

x the first digit:

1 presented to TSG for information;

2 presented to TSG for approval;

3 or greater indicates TSG approved document under change control.

y the second digit is incremented for all changes of substance, i.e. technical enhancements, corrections, updates, etc.

z the third digit is incremented when editorial only changes have been incorporated in the document.

Introduction

For 3GPP systems there is a need for truly scalable entity Authentication Framework (AF) since an increasing number of network elements and interfaces are covered by security mechanisms.

This specification provides a highly scalable entity authentication framework for 3GPP network nodes. This framework is developed in the context of the Network Domain Security work item, which effectively limits the scope to the control plane entities of the core network. Thus, *the Authentication Framework will provide entity authentication for the nodes that are using NDS/IP*.

Feasible trust models (i.e. how CA's are organized) and their effects are provided. Additionally, requirements are presented for the used protocols and certificate profiles, to make it possible for operator IPsec and PKI implementations to interoperate.

1 Scope

The scope of this Technical Specification is limited to authentication of network elements, which are using NDS/IP, and located in the inter-operator domain.

It means that this Specification concentrates on authentication of Security Gateways (SEG), and the corresponding Za-interfaces. Authentication of elements in the intra-operator domain is considered as an internal issue for the operators. This is quite much in line with [1] which states that only Za is mandatory, and that the security domain operator can decide if the Zb-interface is deployed or not, as the Zb-interface is optional for implementation.

However, NDS/AF can easily be adapted to intra-operator use. This is just a simplification of the inter-operator case as all NDS/IP NEs and the PKI infrastructure belong to the same operator. Validity of certificates may be restricted to the operator's domain.

NOTE: In case two SEG's interconnect separate network regions under a single administrative authority (e.g. owned by the same mobile operator) then the Za-interface is not subject to roaming agreements, but the decision on applying Za-interface is left to operators.

The NDS architecture for IP-based protocols is illustrated in figure 1.

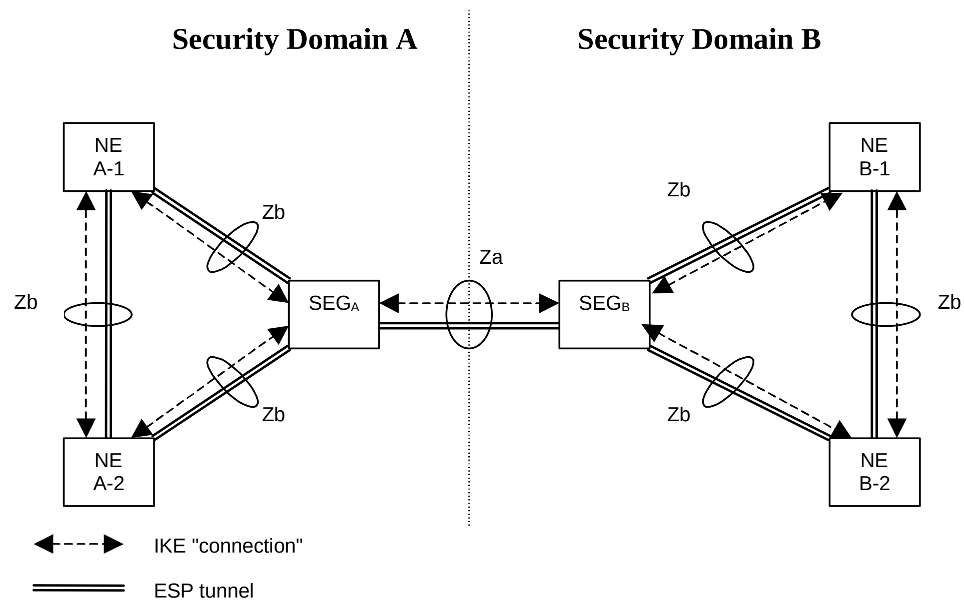

Figure 1: NDS architecture for IP-based protocols [1]

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, edition number, version number, etc.) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies. In the case of a reference to a 3GPP document (including a GSM document), a non-specific reference implicitly refers to the latest version of that document *in the same Release as the present document.*

[1] 3GPP TS 33.210: "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Network domain security; IP network layer security".

[2] IETF RFC 2986: "PKCS#10 Certification Request Syntax Specification Version 1.7"

[3] IETF RFC 3280: "Internet X.509 Public Key Infrastructure Certificate and CRL Profile "

[4] IETF Draft draft-ietf-pkix-rfc2510bis-08.txt: "Internet X.509 Public Key Infrastructure Certificate Management Protocol"

[5] IETF RFC 2252: "Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions"

[6] IETF RFC 1981: "Path MTU Discovery for IP version 6"

3 Definitions and abbreviations

3.1 Definitions

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

Local CRL: Repository that contains cross-certificate revocations

PSK: Pre-Shared Key. Method of authentication used by IKE between SEG in NDS/IP [1].

Public CRL: Repository that contains revocations of SEG and CA certificates and can be accessed by other operators

Roaming CA: The CA that is responsible for issuing certificates for SEG that have interconnection with another operator

3.2 Abbreviations

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

AF	Authentication Framework
CA	Certification Authority
NDS	Network Domain Security
SEG	Security Gateway
Za	Interface between SEGs belonging to different networks/security domains (a Za interface may be an intra or an inter operator interface).
Zb	Interface between SEGs and NEs and interface between NEs within the same network/security domain

4 Introduction to Public Key Infrastructure (PKI)

[Editor's note: Serves as an introduction to PKI architecture and terminology. This should be kept relatively brief. Introduction of a certificate, certification authority, hierarchies, etc. Benefits of PKI: less secrets to be managed, n compared to n^2. Adding a new network element does not need configuration in other network elements. References: RFC 2459/3280]

4.1 Cross-certification

Cross-certification is a process that establishes a trust relationship between two authorities. When an authority A is cross-certified with authority B, the authority A has chosen to trust certificates issued by the authority B. Cross-certification process enables the users under both authorities to trust the other authority's certificates. Trust in this context equals to being able to authenticate.

4.1.1 Manual Cross-certification

Mutual cross certifications are done directly between the authorities and this approach is often called manual cross-certification. In this approach the authority does the decisions about the trust locally. When an authority A chooses to trust an authority B, the authority A signs the certificate of the authority B and distributes the new certificate (B's certificate signed by A) locally.

The down side of this approach is that it often results into scenarios where there needs to be lot of certificates available for the entities doing the trust decisions: There needs to be a certificate signed by the local authority for each security domain the local authority wishes to trust.

However, all the certificates can be configured locally and are locally signed, so the management of them is often flexible.

4.1.2 Cross-certification with a Bridge CA

The Bridge CA is a concept that reduces the amount of certificates that needs to be configured for the entity that does the certificate checking. The name "bridge" is descriptive; when two authorities are mutually cross-certified with the bridge, the authorities do not need to know about each other. Authorities can still trust each other because the trust in this model is transitive (A trusts bridge, bridge trusts B, thus A trusts B and vice versa). The Bridge CA acts like a bridge between the authorities. However, the two authorities shall also trust that the bridge does the right thing for them. All the decisions about the trust can be offloaded to the bridge, which is desirable in some use cases. If the bridge decides to cross certificate with an authority M, the previously cross-certified authorities start to trust the M automatically.

The bridge-CA style cross-certifications are useful in scenarios where all entities share a common authority that everybody believes to work correctly for them. If an authority needs to restrict the trust or access control derived from the bridge-CA, it additionally needs to implement those restrictions.

5 Use cases and profiling of the NDS/AF

[Editor's note: This section shall list the security requirements emerging from identified use cases.]

The roaming CA certificate of the owning operator shall be stored securely in the SEG. It defines who is the authority that the device trusts when connecting to the other devices. It is assumed that each operator domain could include 2 to 10 SEGs.

The NDS/AF is initially based on a simple trust model (see Annex A) that avoids introduction of transitive trust and/or additional authorisation information. The simple trust model implies manual cross-certification.

5.1 PKI architecture for NDS/AF

This chapter defines the PKI architecture for the NDS/AF. The goal is to define a flexible, yet simple architecture, which is easily interoperable with other implementations.

The architecture described below uses a simple access control method, i.e. every element which is authenticated is also provided service. More fine-grained access control may be implemented, but it is out of scope of this specification.

The architecture does not rely on bridge CAs, but instead uses direct cross certifications between the security domains. This enables easy policy configurations in the SEGs.

5.1.1 General architecture

Each security domain has at least one certification authority dedicated to it. The certification authority which the network elements use for inter-operator authentication is called roaming CA of the domain.

The roaming CA of the domain issues certificates to the SEG's in the domain. This specification describes the profile for the roaming CA and a profile for SEG. Also a method for creating the cross-certificates is described.

In general, all of the certificates should be based on the Internet X.509 certificate profile [3].

The roaming CA shall issue certificates for SEG's in the Za interface. When SEG of the security domain A establishes a secure connection with the SEG of the domain B, they shall be able to authenticate each other. The mutual authentication is checked using the certificates the roaming CAs issued for the SEGs. When a roaming agreement is established between the domains, roaming CAs cross-certify with each other. The created cross-certificates need only to be configured locally to each domain. The cross-certificate, which roaming CA of security domain A created for security domain B shall be available for the domain A SEG which provides Za interface towards domain B. Equally the corresponding certificate, which the roaming CA of the security domain B created for security domain A shall be available for the domain B SEG which provides Za interface towards domain A.

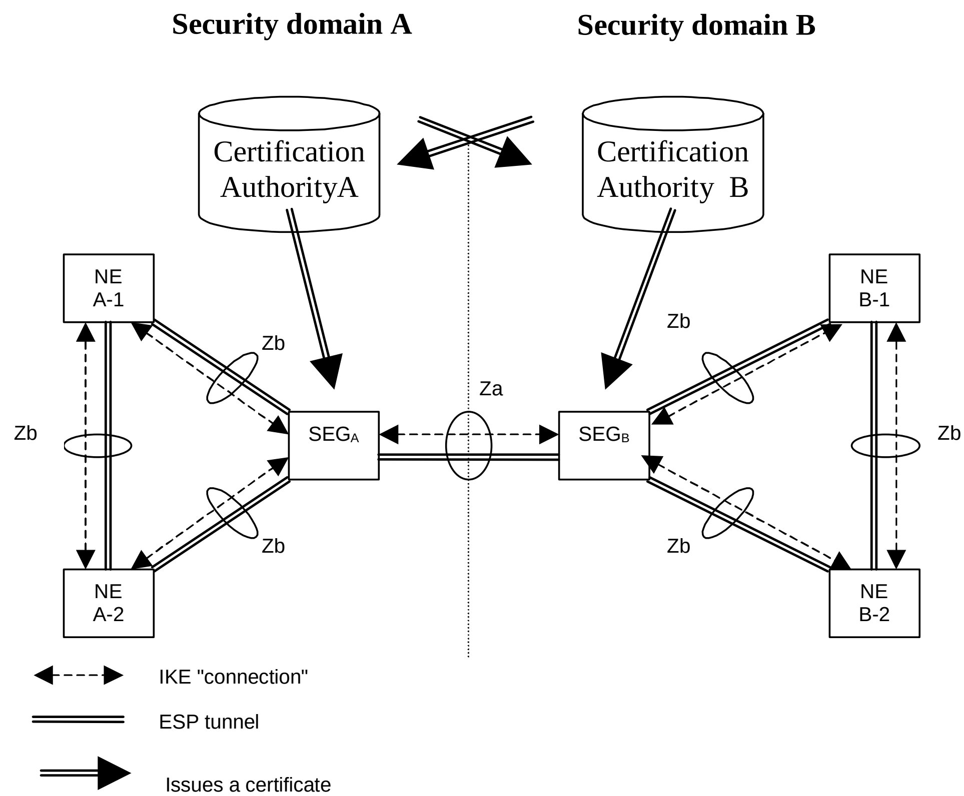

Figure 2: Trust validation path in context of NDS/IP

After cross-certification, the SEGa is able to verify the path: SEGb ->Authority B -> Authority A. Only the certificate of the roaming CA in domain A needs to be trusted by entities in security domain A.

Equally the SEGb is able to verify the path: SEGa -> Authority A -> Authority B. The path is verifiable in B domain, because the path terminates to a trusted certificate (roaming CA of the security domain B in this case).

The roaming CA signs the second certificate in the path. For example, in A domain, the certificate for roaming CA B is signed by roaming CA of the A domain when the cross-certification was done.

5.2 Use cases

5.2.1 Roaming agreement

Security gateways (SEG's) of two different security domains need to establish a secure tunnel, when the operators make a roaming agreement. The first technical step in creating the roaming agreement between domains is the cross-certification of the roaming CAs of the two domains.

Inter-operator cross-certification can be done using different protocols, but the certification authority shall support the PKCS#10 [2] method for certificate requests. Both roaming CAs create a PKCS#10 certificate request, and send it to the other operator. The method for transferring the PKCS#10 request is not specified, but the transfer method shall be secure. The PKCS#10 can be transferred e.g. in a floppy disk, or be send in a signed email. The PKCS#10 request contains the public key of the authority and the name of the authority. When roaming CA accepts the request, a new cross-certificate is created. The authority shall make that new certificate available to SEGs in his own domain, by storing the new cross-certificate into all SEGs that need to communicate with the other domain.

When creating the new cross-certificate, the roaming CA should use basic constraint extension (according to section 4.2.1.10 of [3]) and set the path length to zero. This inhibits the new cross-certificate to be used in signing new CA certificates. The validity of the certificate should be set sufficiently long. The cross-certification process needs to be done again when the validity of the cross-certificate is ending. The validity time could be e.g. 15 years. The start time of the validity should start e.g. a day before the actual roaming is set to start in order to avoid problems with different time zones. Problems in PKI are often due to the time differences.

When the new certificate is available for SEG, all that needs to be configured in SEG is the DNS name of the peering SEG gateway. The authentication can be done based on created cross-certificates.

When the cross-certification is implemented this way, the PKI architecture seems hierarchical to the network elements in the domain: At the very top of the hierarchy sits the roaming CA of the domain. At the second level, there are certificates directly issued by roaming CA for the SEGs together with the cross certificate issued for the peering domains. The certificates of the peer domains are located under the cross-certificates of the peer domains.

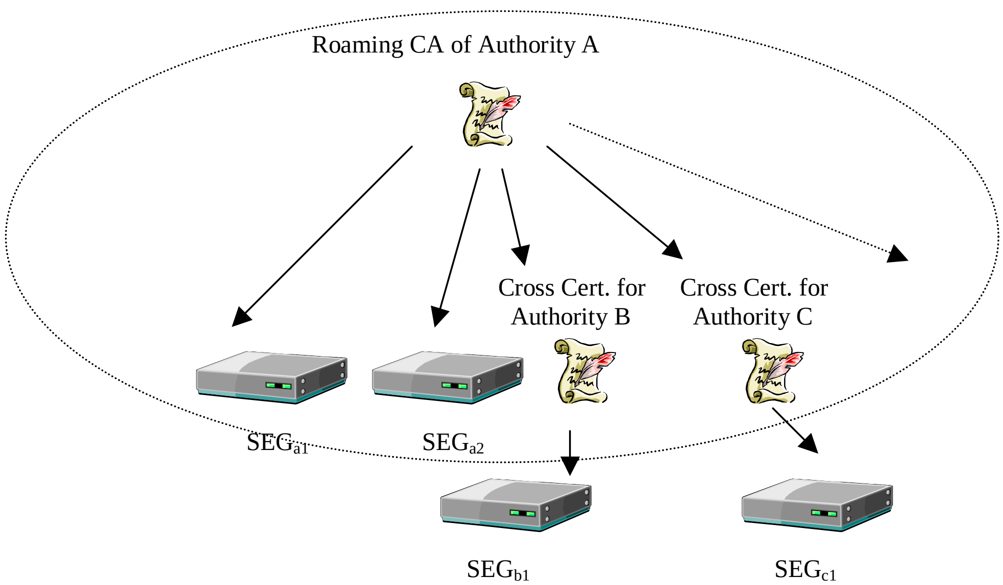

Figure 3: Security domain A illustrated. The PKI is hierarchical inside the domain.

5.2.2 VPN tunnel establishment

After establishing a roaming agreement and finishing required preliminary certificate management operations as specified in the previous section, the operators configure their SEGs for SEG-SEG connection, and the SAs are established as specified by NDS/IP [1].

In each connection configuration, the remote SEG DNS name is specified. Only local roaming CA is configured as the trusted CA. Because of the cross-certification, any operator whose roaming CA has been cross-certified, can get access using this VPN connection configuration. If access to a certain local subnet is allowed for only certain operators, the VPN connection configuration shall include limitations for certificate issuer name.

[Editor’s note: These limitations for certificate issuer name are ffs.]

Following is the flow of connection negotiation from the point of view of Operator A’s SEG (initiator). Operator B SEG (responder) shall behave in a similar fashion.

- During connection initiation, the initiating Operator A's SEG A provides its own SEG-certificate and the corresponding digital signature in Main Mode message 3
- SEG A receives the remote SEG B certificate and signature;
- SEG A validates the remote SEG B signature;
- SEG A verifies the validity of the SEG B certificate by a CRL check to both the Operator A and B CRL databases. IKE Phase-1 SA is established, and the Phase-2 SA negotiation proceeds as described with NDS/IP [1] with PSK authentication.

NOTE: This specification provides authentication of SEGs in an "end-to-end" fashion as regards to roaming traffic (operator to operator). If NDS/AF (IKE) authentication were to be used for both access to the transport network (e.g. GRX) and for the end-to-end roaming traffic, IPsec mechanisms and policies such as iterated tunnels or hop-by-hop security would need to be used. However, it is highlighted that the authentication framework specified is independent of the underlying IP transport network.

5.2.3 Operator deregistration

When a roaming agreement is terminated or due to an urgent service termination need, all concerned peers shall remove the SAs using device-specific management methods. Each concerned operator shall also list the cross-certificate created for the roaming CA of the terminated operator in his own local CRL.

5.2.4 SEG deregistration

If a SEG is removed from the network, the SAs shall be removed as above. The operator of the SEG shall have the certificate of the SEG listed in his CRL.

[Editor's note:

Two new paragraphs needed to describe the involved actions for revocation and check our model !?

Roaming CA certificate revocation ?

A) of the own roaming CA

B) of a partner roaming CA

SEG revocation

A) own SEG

B) SEG of a roaming partner]

5.3 Profiling

[Editor's note: "Motivation" statements marked with italic in chapters 5.3.1 and 5.3.2 are included in the drafting stage of the TS, but will be removed before submission for approval to TSG SA.]

5.3.1 Certificate profiles

[Editor's note: A more detailed check on using RFC3280 and draft-ietf-ipsec-pki-profile-02.txt *as the main profiling base is needed. It needs to be assessed why and how we want to deviate from these papers]*

5.3.1.1 Common rules to all certificates

- Version 3 certificate

    Motivation: This is the current state of the art [3].

- Hash algorithm for use before signing certificate: Sha-1 mandatory to support, MD-5 shall not be used.

Motivation: SHA-1, is state of the art, MD-5 shall not be used anymore as it is considered weaker

- Subject and issuer name format. Note that C is optional element. : (C=<country>), O=<Organization Name>, CN=<Some distinguishing name>. Organization and CN shall be in UTF8 format.

  Motivation: RFC3280 states in clause 4.1.2.4 Issuer that The UTF8String encoding in RFC 2279 is the preferred encoding, and all certificates issued after December 31, 2003 MUST use the UTF8String encoding of DirectoryString (except in some migration cases).

- CRLv2 support with LDAPv3 [5] retrieval shall be supported as the primary method of certificate revocation verification.

~~5.3.1.1~~5.3.1.2 CA Certificate profile

In addition to clause 5.3.1.1, following requirements apply:

- The RSA key length shall be at least 2048-bit

  Motivation: "RSA Laboratories currently recommends key sizes of 1024 bits for corporate use and 2048 bits for extremely valuable keys like the root key pair used by a certifying authority "

  see http://www.rsasecurity.com/rsalabs/faq/3-1-5.html

- Extensions:
  - Optionally non critical authority key identifier
  - Optionally non critical subject key identifier
  - Mandatory critical key usage: At least keyCertSign and CRL Sign should be asserted
  - Mandatory critical basic constraints: CA=True, path length unlimited or at least 2.

~~5.3.1.2~~5.3.1.3 SEG Certificate profile

In addition to clause 5.3.1.1, following requirements apply:

- The RSA key length shall be at least 1024-bit

  Motivation: "RSA Laboratories currently recommends key sizes of 1024 bits for corporate use and 2048 bits for extremely valuable keys like the root key pair used by a certifying authority "

  see http://www.rsasecurity.com/rsalabs/faq/3-1-5.html

- Issuer name is the same as the subject name in the Domain authority cert.
- Extensions:
  - Optionally non critical authority key identifier
  - Optionally non critical subject key identifier
  - Mandatory critical key usage: At least digitalSignature shall be set.
  - Optional critical enhanced key usage: If present, at least server authentication and IKE intermediate shall be set
  - Mandatory non critical Distribution points: CRL distribution point

~~5.3.1.3~~5.3.1.4 Cross Certificate profile~~ion between domains~~

In addition to clause 5.3.1.1, following requirements apply:

- Subject name is the same, which the authority of the other domain uses in it's certificates
- Issuer Name is the same as used for signing our entities

- Extensions:
  - o Optionally non critical authority key identifier
  - o Optionally non critical subject key identifier
  - o Mandatory critical key usage: At least keyCertSign and CRL Sign, should be asserted
  - o Mandatory critical basic constraints: CA=True, path length 0.

5.3.2 IKE negotiation and profiling

[Editor's note: A more detailed check on using draft-ietf-ipsec-pki-profile-02.txt *as the main profiling base is needed. It needs to be assessed why and how we want to deviate from these papers]*

5.3.2.1 IKE Phase-1 profiling

The Internet Key Exchange protocol shall be used for negotiation of IPsec SAs. The following requirements on IKE in addition to those specified in NDS/IP [1] are made mandatory for inter-security domain SA negotiations over the Za-interface.

For IKE phase-1 (ISAKMP SA):

- The use of RSA signatures for authentication shall be supported.
- Initiating/responding SEG are required to send certificate requests in the IKE messages

  Motivation: suggested by draft-ietf-ipsec-pki-profile-02.txt to avoid interoperability problems

- Cross-certificates shall not be send by the peer SEG as they are pre-configured in the SEG.

  Motivation: avoiding known problems (see clause 5.3.5.2)

- The SEG shall always send its own certificate in the certificate payload of the last (third) Main Mode message

  Motivation: avoids the need to cache Peer SEG certificates.

- The certificates in the certificate payload shall be encoded as type 4 (X.509 Certificate – Signature).
- The lifetime of the Phase-1 IKE SA shall be limited to at most the remaining validity time of the peer SEG certificate.

5.3.2.2 Potential interoperability issues

Some PKI-capable VPN gateways do not support fragmentation of IKE packets, which becomes an issue when more than one certificate is sent in the certificate payloads, forcing IKE packet fragmentation. This means that direct cross-certification or manually importing the peer CA certificate to the local SEG and trusting it is preferable to bridge CA systems. When IKE is run over pure IPv6 the typical MTU sizes do not increase and long packets still have to be fragmented (allowed for end UDP hosts even for IPv6, see Path MTU Discovery for IPv6 – [6]), so this is a potential interoperability issue.

Certificate encoding with PKCS#7 is supported by some PKI-capable VPN gateways, but it shall not be used.

5.3.3 Path validation

5.3.3.1 Path validation profiling

- Validity of certificates received from the peer SEG shall be verified by CRLs retrieved with LDAP, based on the CRL Distribution Point in the certificates.
- A SEG shall not validate received certificates from the peer SEG whose validity time has expired, but end the path validation with a negative result.

- A SEG shall not validate received certificates from the peer SEG whose CRL distribution point field is empty, but end the path validation with a negative result.
- Certificate validity calculation results shall not be cached for longer than the resulting IKE phase-1 lifetime.

5.3.4 Services utilising inter-domain PKI

[Editor's note: Subscriber certificates are feasible to implement without Authentication Framework (AF), but AF could help as inter-domain PKI provides the validation path for certificate usage.]

6 Security features

[Editor's note: Subsections may have to be moved to suitable places.]

6.1 Repositories

During VPN tunnel establishment, each SEG has to verify the validity of it's peer SEG's certificate according to section 5.2.2. Any certificate could be invalid because it was revoked (and replaced by a new one) or a SEG or operator has been deregistered.

SEG_B has to verify that

a) the cross-certificate of CA_A is still valid

b) the certificate of SEG_A is still valid

SEG_A performs according checks from its own perspective.

Check a) can be performed by querying the local CRL. For check b), a CRL of the peering CA shall be queried. At this point of time, the VPN tunnel is not yet available, therefore the public CRL of the peering CA shall be accessible for a SEG without utilising Za interface.

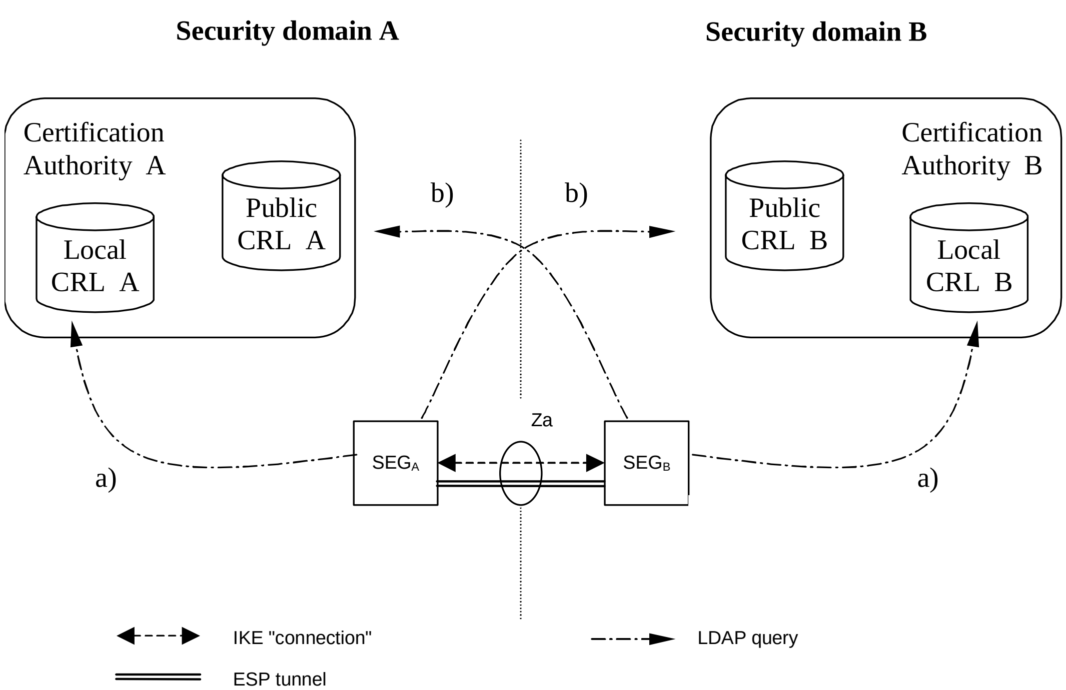

Figure 4: CRL Repositories

The public and local CRL repositories of a CA may be implemented as two separate databases or as a single database which is accessible via two different interfaces. Access to the "public" CRL is public with respect to the interconnecting transport network (e.g. GRX). The public CRL should be adequately protected (e.g by a firewall) and the owner of the public CRL may limit access to it according to his roaming agreements.

SEGs shall use LDAP to access the CRL repositories.

[Editor's note: Further specification of public CRL interface and its relation to Za is ffs.]

6.2 Life cycle management

Certificate management protocol v2 (CMPv2, [4]) shall be the supported protocol to provide certificate lifecycle management capabilities. All SEGs and Roaming CAs shall support initial enrolment by SEG from CA via CMPv2, i.e. receiving a certificate from the Roaming CA, and updating the key of the certificate via CMPv2 before the certificate expires.

[Editor's note: CMPv2 is still at draft status, but is already widely supported (see 'CMP Interop Project': http://www.ietf.org/proceedings/00dec/slides/PKIX-4/), and expected to move to Draft Standard status in the near future. Thus it is expected that CMPv2 receives a RFC status before the NDS/AF specification is completed. Additionally, CMPv2 is preferred to CMPv1(RFC2510), because of the interoperability issues with CMPv1.]

7 Security mechanisms

[Editor's note: This section shall describe the security mechanisms that are provided for inter-domain authentication, i.e. the actual description of what the Authentication framework consists of.]

7.1 Authentication

8 Evolution path

[Editor's note: This chapter describes the evolution path from using NDS/IP towards optional PKI structure.]

8.1 Backward compatibility

Annex A (informative):
~~<Informative annex title>~~

Decision for the simple trust model

A.1 Introduction

In order to document the decision for the "simple trust model", which requires manual cross-certification, this section discusses technical advantages and disadvantages of two basic approaches to providing inter-operator trust for purposes of roaming traffic protection, namely **cross-certification** and a **Bridge CA**. The Bridge CA is an extension of the cross-certification approach, and identified as one of the recommendable solutions for providing inter-operator trust in NDS/AF feasibility study (TR33.810). Taking into account the current state of PKI software and the general need for simple solutions when there is a choice, there is pressure to make the cross-certification without a Bridge CA as the working assumption for the NDS/AF TS. This document discusses the background motivation for such direction.

The direct cross-certification without Bridge CA model is associated strongly with the current practice in the Internet IPsec world, where each IPsec connection is configured with a list of trusted CAs, and anyone with a certificate that has a trust path that can be followed up to such trusted CA (trust anchor) is allowed access. In this model, cross-certification is done at the time the roaming agreement is made. We call this the "**simple trust model**."

The Bridge CA model assumes that all operators wishing to establish a roaming agreement with other operators will first get certified by the Bridge CA for purposes of identification by other operators. This is a necessary preliminary step. Next, when the roaming agreement is done, the operators will configure their IPsec tunnels, with information about which one of the identifiable operators (who have a certificate issued by the Bridge CA) can use that tunnel. This is called the "**extended trust model**", or "separated trust and access control."

This Annex does not discuss the benefits of certificates vs. Pre-Shared Keys. The benefit of cross-certification vs. the explicit listing of roaming peer CAs includes the easier evolution path to a possible eventual Bridge CA model.

A.2 Requirements for trust model in NDS/AF

The following is a list of requirements for the trust model for NDS/AF:

A. *Simplicity and ease of deployment.* PKI brings many benefits when a large number of operators need to tunnel traffic in a mesh configuration, but its adoption should not be hindered by an unnecessarily complex technical solution. The required technical and legal operations necessary for exchanging traffic with another operator should be as easy and straightforward as possible.

B. *Compatibility with existing standards.* Unless there are explicit requirements why existing PKI standards should be extended to accommodate 3GPP environment, the 3GPP specifications should be accommodated to the existing standards. This allows best choice of equipment for operators and allows interoperability with non-3GPP environments.

C. *Usable by both GRX and non-GRX operators.* Both operators making use of GRX providers and those without (using leased lines or even the public Internet), should be able to make use of NDS/AF measures to exchange traffic securely.

A.3 Cross-certification approaches

A.3.1 Manual Cross-certification

The trust model of manual cross certification is characterized by the clause: "Trust nobody unless explicitly allowed". Issuing a certificate for the authority we wish to trust creates the allowances. The manual cross certification is easy to understand. Also the security of this depends only on the decisions done locally.

A.3.2 Cross-certification with a Bridge CA

The trust model of bridge-CA can be characterized by the clauses:

"Trust everybody that the Bridge-CA trusts unless explicitly denied". Explicit denials are handled by writing the restrictions (in the form of name constraints) to the certificate issued to the bridge.

"Trust everybody listed in the certificate which I issued to the bridge". Explicit allowances are listed in the certificate issued to the bridge (in the form of name constraints).

Name constraint is a rarely used extension for X.509 certificates. In essence it is a clause that says who to trust or who not to trust based on names on certificates. The fact that they are relative rarely used and the fact that there is so little official documentation about them is a risk. Name constraints also require that there is some organization doing registration of names in order to avoid name collisions.

A.4 Issues with the Bridge CA approach

A.4.1 Need for nameConstraint support in certificates or strong legal bindings and auditing

If no precautions are taken, it is possible that an operator (M) whose Roaming CA has been signed by the Bridge CA (= certified by the Bridge), creates certificates that resemble another operator's (A) certificates, letting M access to operator (B)'s network, even without authorization.

Let's say operator B has the following configuration for access to her subnetwork reserved for handling roaming traffic:

Local-Subnetwork = some ipv6 subnetwork address

TrustedCA's = BridgeCA

AllowedCertificateSubject = O=Operator A or O=Operator C or O=Operator D

Note: The IP addresses of the remote SEGs are not limited, as authentication is done based on certificates, and all trusted operators are allowed similar access. If different foreign operators would require to access different subnetworks, there would be several configuration blocks like the above, with the IP addresses appropriately specified.

Such "AllowedCertificateSubject" feature (the term name is imaginary) is widely supported by PKI-capable IPSec devices.

If Operator M used certificates of the following form for her certificates, she would not be allowed in:

Subject: CN=SEG 1, O=Operator M

Signer: CN=Roaming CA, O=Operator M

However, she can fabricate certificates of the following form:

Subject: CN=SEG 1, O=Operator A

Signer: CN=Roaming CA, O=Operator M

Using such certificates would allow full but illegitimate access to Operator B's network revealed for use by Operator A.

Now, there are the following possibilities to circumvent the problem:

1. Checking also the Signer name when authenticating foreign operators, either by a) a proprietary "AllowedCertificateSigner" property or b) support for nameConstraints in the Bridge CA certificate issued to operator M.
2. Establishing strong legal bindings and auditing that would discourage Operator M from such illegitimate fabrication of Operator A certificates.

The problem with solution 1.a is that such "AllowedCertificateSigner" is not commonly supported by current PKI end-entity products, being in conflict with requirement B.

The problem with solution 1.b is that such "nameConstraints" attribute in certificates is not commonly supported by current PKI CA or end-entity products, being in conflict with requirement B.

The problem with solution 2 is that first of all, an organization willing to run a Bridge CA has to be found before any pair of operators can exchange roaming traffic with NDS/AF mechanisms. Next, there shall be established paperwork and auditing procedures to make sure that the exploit described here can be detected. This is in conflict with requirement A. Also, the illegitimate act described could not be technically prevented beforehand.

If name constraints are used, every time a new roaming agreement is made, each operator shall update the certificate they issue for the Bridge, adding the new roaming partner's name into the certificate. From the point of view of one operator, the number of new certificate signing operations is the same whether a Bridge CA or a direct cross certification model is in use.

A.4.2 Preventing name collisions

If name constraints are used to prevent the additional "bureaucracy" involved with the Bridge CA, the names written into the certificate need to be registered with a third party to prevent two operators accidentally or on purpose using the same name in their certificates. This is in conflict with requirement B.

A.4.3 Two redundant steps required for establishing trust

As described in the introduction, with the "extended trust model", each operator shall first be certified by the bridge (authentication), and then as the second step, enumerate the trusted operators when configuring the IPSec tunnel (access control).

For the Bridge CA model to work, there is a need for organization that all the other parties involved can trust - and the trust shall be transitive! If you trust the bridge, you shall also trust the other organizations joining to the bridge via the cross certification. If Operator A and the Bridge CA cross certify with each other, Operator A will automatically trust every other certified operator to obey the rules. And this trust is not related to the roaming traffic tunnel; the tunnel has to be configured independently of the PKI.

So even if we avoid configuring new certificates in the SEG's when we use cross certification, we shall configure and maintain the roaming information in the SEG some other way. And the hard part: How do we combine the trust provided by the PKI and the roaming agreements, because clearly in this case PKI provided trust is not the same as roaming agreements.

We would need two steps:

1. building "trust" through Bridge CA => authenticating the peer SEG
2. specify in the tunnel configuration which peering SEGs we can trust

If the cross-certification is done without a Bridge CA, the steps can be combined into one. What is the additional value of the PKI provided trust (step 1), if the peering SEGs have to be restricted in any case?

A.4.4 Long certificate chains connected with IKE implementation issues

If Bridge CA is used, a Roaming CA certificate has to be sent in the certificate payload in addition to the local end entity (SEG) certificate. This leads in Ethernet environments to the fragmentation of the IKE packet, which some current IKE implementations do not support. It is a problem in the implementation, not the protocol. Even in IPv6, the IKE UDP packets need to be fragmented, posing a potential interoperability problem. Clearly it is not a solution to use a different protocol, but instead the current implementations should be fixed. Still, taking into account requirement B, it is safer to avoid the problem altogether by not forcing the fragmentation of IKE packets by not using a Bridge CA.

A.4.5 Lack of existing relevant Bridge CA experiences

The Federal PKI in the USA is an example deployment where a Bridge CA is used to connect together CAs of the various federal agencies. It seems to be however the only documented one of its kind, and is connected with very heavy policy documentation and obviously heavy auditing practices, even within one organization, the federal government. The bridge approach is warranted in the case, because they want to automatically check whether some entity has legal rights to sign some document. The number of entities doing cross-domain PKI validation can be several millions, and it is impossible for one validating entity to keep count of individual signers.

In 3G roaming, the situation is in many ways different. When a new operator is born, the other ones do not automatically want to exchange roaming traffic with the new one, but a legal agreement with that operator and a technical tunnel establishment shall be done. In Federal PKI, the situation is the opposite: nothing should need to be done and still be able to trust the other.

In the Federal PKI, the paperwork and processes make name constraints in certificates unnecessary, and IKE is supposedly not used together with the Bridge CA.

A.5 Feasibility of the direct cross certification approach

This chapter discusses the direct cross certification, i.e. manual cross certification approach, where operators are doing the cross certification operation only when agreeing to set up a tunnel with another operator. This tunnel setup is a legal and technical operation in any case, so it is feasible to do also the cross-certification at this time, removing the need for the initial step to cross-certify with the Bridge CA.

There is no technical difference regarding the feasibility of direct cross certification or Bridge CA in the context of GRX or non-GRX environment. GRX might be one possible choice for providing the Bridge CA services.

A.5.1 Benefits of direct cross certification

The benefits of the direct cross certification is that as a mechanism it is well known, supported widely by current PKI products and there even exists an evolution path to a Bridge CA solution if the products come to support it adequately, a Bridge CA is established, and the number of operators becomes so large to warrant the use of the Bridge CA technology. Bridge CA uses the cross certification mechanisms in any case.

The tunnel configuration would look like the following:

Local-Subnetwork = some ipv6 subnetwork address

TrustedCA's = LocalCA

The information of which operator is allowed access is implicit in the direct cross certifications that have been done by the LocalCA, thus authentication and access control are tightly connected. If different foreign operators need to access different subnetworks, there would be separate tunnel configurations with SEG IP address for each, including an "AllowedCertificateSubject" limitation. The "AllowedCertificateSigner" limitation is not needed as necessary in this model (compared to the bridge CA model), since the set of operators who we are able to authenticate are only the ones, we have previously agreed to trust when doing the direct cross certification. In the bridge CA case, the set of operators we are able to authenticate includes all operators who have joined to the bridge.

A.5.2 Memory and processing power requirements

In case of direct cross certification, each operator shall store the certificates issued for the other operators locally. They could be stored in the SEG devices, or then in a common repository.

If an operator makes roaming agreements with 500 other operators, this would require roughly 1000 kilobytes of memory, if the operator signs the certificates herself, and one certificate takes 1 kilobyte of memory. This should be quite feasible taken into account the high-end nature of SEG hardware.

Processing power benchmark for validating certificates:

Hardware: 800 MHz Pentium III, 256 MB of memory.

200 x 1024-bit RSA certificates, 1 Root CA (operator's own CA), 200 Sub CAs (other operator CAs) and 200 end entity (SEG) certificates. Also CRLs were verified. Both certificates and CRLs were loaded from disk during the test. The whole test took 3.5 seconds, with probably disk I/O taking most of the time.

In this test 200 certificate chains were validate up to the trusted root.

A.5.3 Shortcomings

As discussed in the previous section, the Bridge CA approach saves memory or storage space in SEGs, because all the other operators Roaming CA certificates do not need to be stored with other operators. Just the Bridge CA certificate would be stored, and other certificates retrieved during IKE negotiation.

A.5.4 Possible evolution path to a Bridge CA

If needed, it is possible to take the Bridge CA into use gradually, given that the support by PKI products becomes reality. From one operator's point of view, the bridge CA would be like any other operator so far, and a cross-certification would be made, but additionally the name constraints in the certificate issued for the Bridge CA should be updated every time a new roaming agreement is made.

Annex B (informative): Decision for the CRL repository access protocol

In order to document the decision for the protocol to access CRL repositories, this section summarises technical advantages and disadvantages of the two candidates.

LDAP

+ implemented by all PKI products (unless purely manual)

+ scalability

+ flexibility (integration possibility to other systems, automatic public key retrieval possibility)

- complexity

HTTP

+ simple

- not supported by all PKI products (although widely supported)

LDAP was chosen as the more future-proof protocol. Although more complex than HTTP, LDAP is well established amongst PKI vendors and operators.

Annex <C~~B~~> (normative): <Normative annex title>

Annex <X> (informative): Change history

Change history							
Date	**TSG #**	**TSG Doc.**	**CR**	**Rev**	**Subject/Comment**	**Old**	**New**
02-2003					TOC proposal for SA3#27		0.0.1
02-2003					Content of SA3#27 approved TDoc S3-030083 added and meeting comments incorporated	0.0.1	0.1.0
04-2003					Editorial changes and corrections	0.1.0	0.2.0
05-2003					Updated according to SA3#28 decisions	0.2.0	0.3.0