3GPP TSG SA WG3 Security **S3-020542**

8th – 11th October, 2002

Munich, Germany

Agenda Item:	7.7
Source:	Nokia
Title:	Trust and PKI email discussion input paper
Document for:	Discussion/Decision

3GPP TSG SA WG3 Security

Email discussion input

Source:	**Nokia**
Title:	**Subscriber certification in cellular networks and the role of inter-operator PKI**
Document for:	**Discussion**

Subscriber certification in cellular networks and the role of inter-operator PKI

Source: Nokia 1
Title: Subscriber certification in cellular networks and the role of inter-operator PKI......................1
Document for: Discussion 1
Introduction...3
The requirements..3
Trust assumptions and security associations 3
1.1 Trust assumptions in the existing infrastructrure...3
1.2 Trust assumptions for the subscriber certificate feature 3
1.3 Summary of assumptions 4
Is it enough to have an inter-operator PKI?..4
1.4 Supporting subscriber certificates using inter-operator PKI 5
1.5 Implications of using standard PKI components for supporting subscriber certificates...6
1.6 Missing pieces...7
1.7 Summary..7
Subscriber certificates issued by the service operator domain...8
1.8 How is a subscriber certificate used? 10
1.9 How are disputes resolved?..11
Conclusions 11
References 11

Introduction

The purpose of this note is to first outline the goal of subscriber certification in cellular networks and the possible trust assumptions that are reasonable. In particular, this note is intended as input to the discussions of SA3's current tasks related to subscriber certification (see items 3 and 6 in [S3-020447]).

The requirements

The basic requirement is to find the means of providing a scalable authorization (and possibly charging) infrastructure to *third party service providers (SPs).* The claim is that this could be done quickly and effectively by bootstrapping support for subscriber certificates from the "existing" (i.e., in its currently planned form) 3GPP cellular business infrastructure (which includes, e.g.,3GPP authentication protocols, settlement procedures, inter-opreator roaming agreements etc.)

Trust assumptions and security associations

1.1 Trust assumptions in the existing infrastructrure

First, consider the implicit trust assumptions in the current 3GPP infrastructure used for authorizing access to cellular services. The subscriber has a *direct* business relationship with only one operator (the operator of the "home domain"). This relationship is represented by a security association embodied in the form of the USIM (or ISIM) in the subscriber's UICC. The home operator may revoke this relationship. The 3GPP Authentication and Key Agreement protocol provides the technical means for any operator to check if this relationship is in force.

The subscriber effectively trusts all cellular operators because any operator can generate a CDR for alleged use of cellular services by the subscriber. There is no technical way for the subscriber to demonstrate that he did not actually use those services. For this reason, operators are also required to trust each other. Operators do not trust the cellular subscriber. Instead, subscribers are expected to prove possession of their USIM/ISIM before they are granted access to cellular services.

1.2 Trust assumptions for the subscriber certificate feature

The goal of the subscriber certificate WID is to use this infrastructure to provide efficient authorization and settlement services to *third party service providers (SPs)*. Therefore it introduces a new player into the above model. We need to identify reasonable assumptions for trust relationships between existing players (subscribers and operators) and the SPs.

Ideally it would be nice to remove the restriction that a cellular operator must trust *all* other cellular operators with whom they have roaming agreements. However, using effectively the same trust relationships in the current 3GPP infrastructure will shorten the time it takes to deploy a bootstrapped infrastructure that can serve third party SPs. We recommend that in the short term, the architecture for supporting subscriber certificates should rely on the existing trust assumptions. However, we should also make sure that if the trust requirements among operators are made less stringent in the long term, this architecture could still be used.

One primary assumption, as explained in [S3-020378,S3-020365] is that SPs will typically have business relationships with only a small number of operators (presumably located in their geographic region). This implies that a SP who wants to authorize a visiting subscriber must be able to do so even if the SP and the subscriber do not have *direct* security associations with the same operator. This is also a requirement from S1 [S1-021685].

The SP is not likely to be trusted by the subscriber. It is also unlikely that the operators will trust every SP, even if the operator has a business relationship and security associations with the SP. The SP should not be required to trust the subscriber. The SP may trust the operators with whom he makes business relationships.

1.3 Summary of assumptions

In summary, we can assume direct security associations between a SP and a small number of cellular operators as well as a subscriber and his home operator. The cellular operators should not have to trust a SP or a subscriber. The SP may be required to trust the cellular operators. A subscriber trusts his home operator and possibly other operators. A subscriber should not have to trust a SP. Operators may be required to trust one another.

Is it enough to have an inter-operator PKI?

Suppose we have an inter-operator PKI. Is it enough to meet the above goal, subject to the listed assumptions? In particular, would we be able to build a system using off-the-shelf PKI tools? What additional specifications would be needed?

The 3GPP SA3 feasibility study report on the evolution of Network Domain Security [S3-020414] includes discussions on various options for building an inter-operator PKI. The focus of that study is limited to end entities in the PKI that are network elements. For subscriber certificates, the end entities also include subscribers. Please see [S3-020414] for more detailed explanations on PKI-related abbreviations, standards, and concepts.

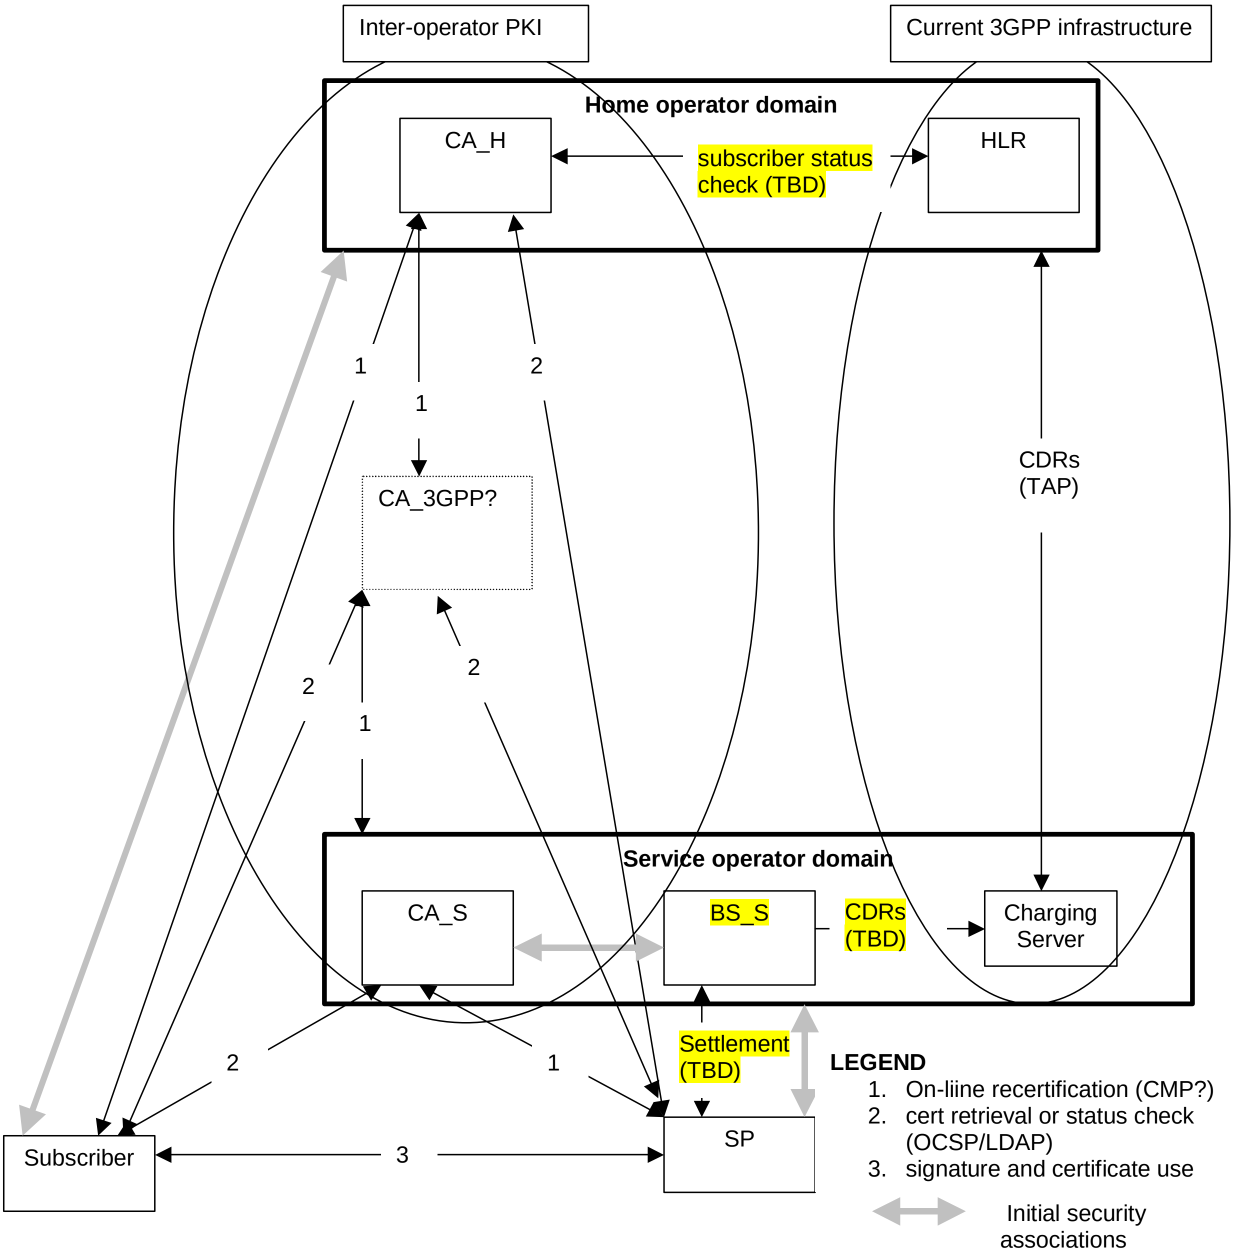

Inter-operator PKI
Current 3GPP infrastructure
Home operator domain
CA_H
subscriber status check (TBD)
HLR
1
2
1
CA_3GPP?
CDRs (TAP)
2
2
1
Service operator domain
CA_S
BS_S
CDRs (TBD)
Charging Server
Settlement (TBD)
2
1
Subscriber
3
SP
LEGEND
1. On-liine recertification (CMP?)
2. cert retrieval or status check (OCSP/LDAP)
3. signature and certificate use
Initial security associations

Figure 1 Use of inter-operator PKI in subscriber certification (Note: “Service Operator” can be either the home, or the visited operator)

1.4 Supporting subscriber certificates using inter-operator PKI

In a typical mobile PKI, the initialisation consists of the following aspects (indicated by grey lines in Figure 1):

- each subscriber has one or more keypairs (e.g., in a WIM) and the home operator domain CA issues long-term certificates for the public keys. The subscriber also has an authentic copy of the home operator CA’s public key. This is stored as a root key on the subscriber’s UE.

- SPs and their chosen operators may also exchange public key certificates out-of-band. A SP will store the public keys of its operators as root keys.

We use the term “Service operator domain” to denote the domain of an operator with whom the SP has a business relationship. Note that this domain is not related to mobility; it may be the home operator domain or the visited domain. CA_S is the CA in the service operator domain.

The “home operator domain” is the domain of the home operator with whom the subscriber has a direct business relationship. The CA in the home operator domain is CA_H. When the service domain is the same as the home domain, CA_S is the same as CA_H. There may also be a 3GPP CA (CA_3GPP) which certifies CAs in 3GPP operator domains (alternatively, each operator CA may cross-certify other operator CAs, or, less likely, there can be deeper a hierarchy of CAs).

1.5 Implications of using standard PKI components for supporting subscriber certificates

In a typical PKI, when a subscriber wants to open an encrypted connection (e.g., TLS connection) to a peer, he needs to first get the peer’s public key, and enough certificates so that a certificate chain from a trusted root key to the peer’s public key may be formed.

Similarly, when a subscriber wants to send a signed message to the peer, the peer needs to acquire sufficient certificates so that a suitable certificate chain may be formed in order to verify the signature.

These certificates have expiry dates. If the lifetime period is long (e.g., several years), it is possible that a certificate may be revoked long before its expiry date is reached. So a relying party must also be sure that a certificate in a chain is not revoked. This is achieved by CRLs or protocols like OCSP (See [S3-020414] for references).

In the case of certificates issued to network elements, use of CRLs to implement revocation is quite suitable [S3-020414]. But in the case of subscriber certificates CRLs are unrealistic because (a) the number of certificates involved can be very high, and (b) a revocation even could happen at a very small granularity (within minutes). On-line status checking (OCSP) is a more likely candidate to manage revocation, but it is questionable that the current OCSP servers are able to support the kind of load that is likely on a cellular network: the typical load on OCSP server products intended for the Internet market is small compared to what an operator OCSP server is likely to be faced with if subscriber certificates are used widely.

Typically mobile devices are relieved from performing revocation checks by requiring that servers provide short-lived certificates. This is the approach taken for WTLS server certificates in WAP 2.0. The same approach could also relieve service providers from having to always perform revocation checks. But short-lived certificates imply that there must be a way for on-line recertification. It is necessary to select or specify an appropriate certificate lifecycle management protocol so that certificate can be periodically refreshed. Also, if the chain is short (e.g., if there is a CA_3GPP then chains would consist of 2 certificates only), the task of the relying party is less heavy. Ideally, both (short chains and short-lived certificates) may be used in combination.

So all the CAs involved (CA_H, CA_S, and possibly CA_3GPP) should support:

- On-line recertification interface for refreshing certificates (e.g., using CMP). This interface is indicated by arrows marked “1” in Figure 1.

- On-line certificate retrieval and validation (e.g., using LDAP and OCSP). This interface is indicated by arrows marked “2” in Figure 1.

1.6 Missing pieces

In the preceding section, we described the implications of using off-the-shelf PKI components in order to support subscriber certificates. However, there are still some missing pieces. These are highlighted (in yellow) in Figure 1.

First, how does CA_H decide whether a subscriber’s request for certificate validation must be granted (or, alternatively, whether a subscriber certificate is to be revoked)? In the current 3GPP infrastructure, HLR is the element that knows whether a subscriber’s account has been revoked. So a “subscriber status check” interface between HLR and CA_H must be specified. Since CA_H and HLR already trust each other, a simple database query would suffice; nevertheless it will be a new interface that has to be specified. However, note that there is already a well-defined means for SGSNs/MSCs/IMS elements of directly interacting with the HLR and verifying the validity of the UE: the 3GPP AKA protocol. This protocol requires the involvement of the subscriber. So, it cannot be used as the simple subscriber status check protocol as shown in Figure 1.

Also, when settlement is necessary, how does SP do it? SP has collected signatures from the subscriber. Therefore, there must be a new function in the service operator domain which can receive and verify signatures. In Figure 1 we call it a “billing server” (BS_S). BS_S should also generate CDRs or equivalent from these subscriber signatures.

Finally, one common concern with long-term certificates is privacy: for example, transactions that can be verified using the same signature verification key are linkable by the parties that see those signatures (in our case, these are the third party SPs). The simple way to avoid this is to change keypairs when unlinkability is desired. But this implies getting a new certificate.

1.7 Summary

In summary, an inter-operator PKI, even if it exists is not enough for providing a scalable authorization and charging infrastructure for third party SPs. If off-the-shelf PKI components are used, then

1. preferably, certificates should have a **short lifetime***;* this minimizes the scope for revocation, and makes it possible for the subscriber to use multiple public keys, should it become necessary.

2. CAs must export a **recertification interface**; certificates with short lifetimes imply the need for an easy-to-use recertification procedure.

3. Each operator PKI must export an **interface for certificate downloading and verification**.

4. a **new “subscriber status check” protocol** between the CA and the HLR must be defined

5. a **new “Billing server” element** which verifies signatures and translates them into CDRs (or equivalent) needs to be defined. (it should also store signatures as potential future evidence; See Section 1.8)

The concerns regarding the use of off-the-shelf PKI components for cellular subscriber certificates are:

1. **scalability:** loads on PKI servers will be too high if they are used in the standard way.

2. **need for new interfaces**: handling revocation of subscribers imply the need for new interfaces.

3. **privacy concerns**: with long-term subscriber certificates, there is a potential concern with linkability of a subscriber's activities by *third parties*.

Based on the above discussion, we can conclude that even if an inter-operator PKI is available, using off-the-shelf PKI elements in the traditional way is not enough to support the subscriber certificates capability.

But, there is a way to leverage existing features of the 3GPP architecture (such as the AKA protocol) to design support for subscriber certificates so that it works *even without* assuming an inter-operator PKI. Furthermore, when an inter-operator PKI is available the same design is still usable and would provide greater security guarantees. Such an approach is described next.

Subscriber certificates issued by the service operator domain

Subscriber certificates issued by the service operator domain constitutes an alternative that addresses the issues discussed above.

The primary differences are the following:

- CA_S issues a short-term subscriber certificate to the subscriber; unlike in traditional PKIs, the certificate issuing process is secured by AKA authentication. I.e., the service operator domain will be able to link the subscriber's AKA identity (IMSI or IMPI) with the granted certificate. "Authenticator" is the entity in the sevice operator domain which is responsible for ensuring the authentication of the UE. The authenticator functionality may be co-located with the CA or with an existing element in the core network. Depending on where the authenticator is located, it may be necessary to specify an interface between CA_S and the authenticator.

- CA_S also issues certificates to SPs with whom it has a business relationship.

- AKA authentication could also be used to deliver an authentic copy of the public key of CA_S to subscribers.

Figure 2 shows this architecture. This implies that

- Even if there is no inter-operator PKI this solution will work; if there is eventually an inter-operator PKI, this solution can be used along with the inter-operator PKI providing better security (e.g., relaxing the trust assumption regarding service domain operators).

- Usually, neither the SP nor the subscriber needs to process any certificate chains (CA_S issues all certificates).

- Certificates have short lifetimes. The need for revocation is greatly reduced.

- Optionally CA_S (or more generally, the PKI in the service operator domain) may provide an on-line certificate retrieval and status check interface.

- If necessary, it is possible for the subscriber to obtain certificates on different public keys, even newly generated ones.

The specification of the BS_S element and associated interfaces is common for both cases. But note that this interface and element are internal to the domain of an operator. It does not involve traversing domain boundaries.

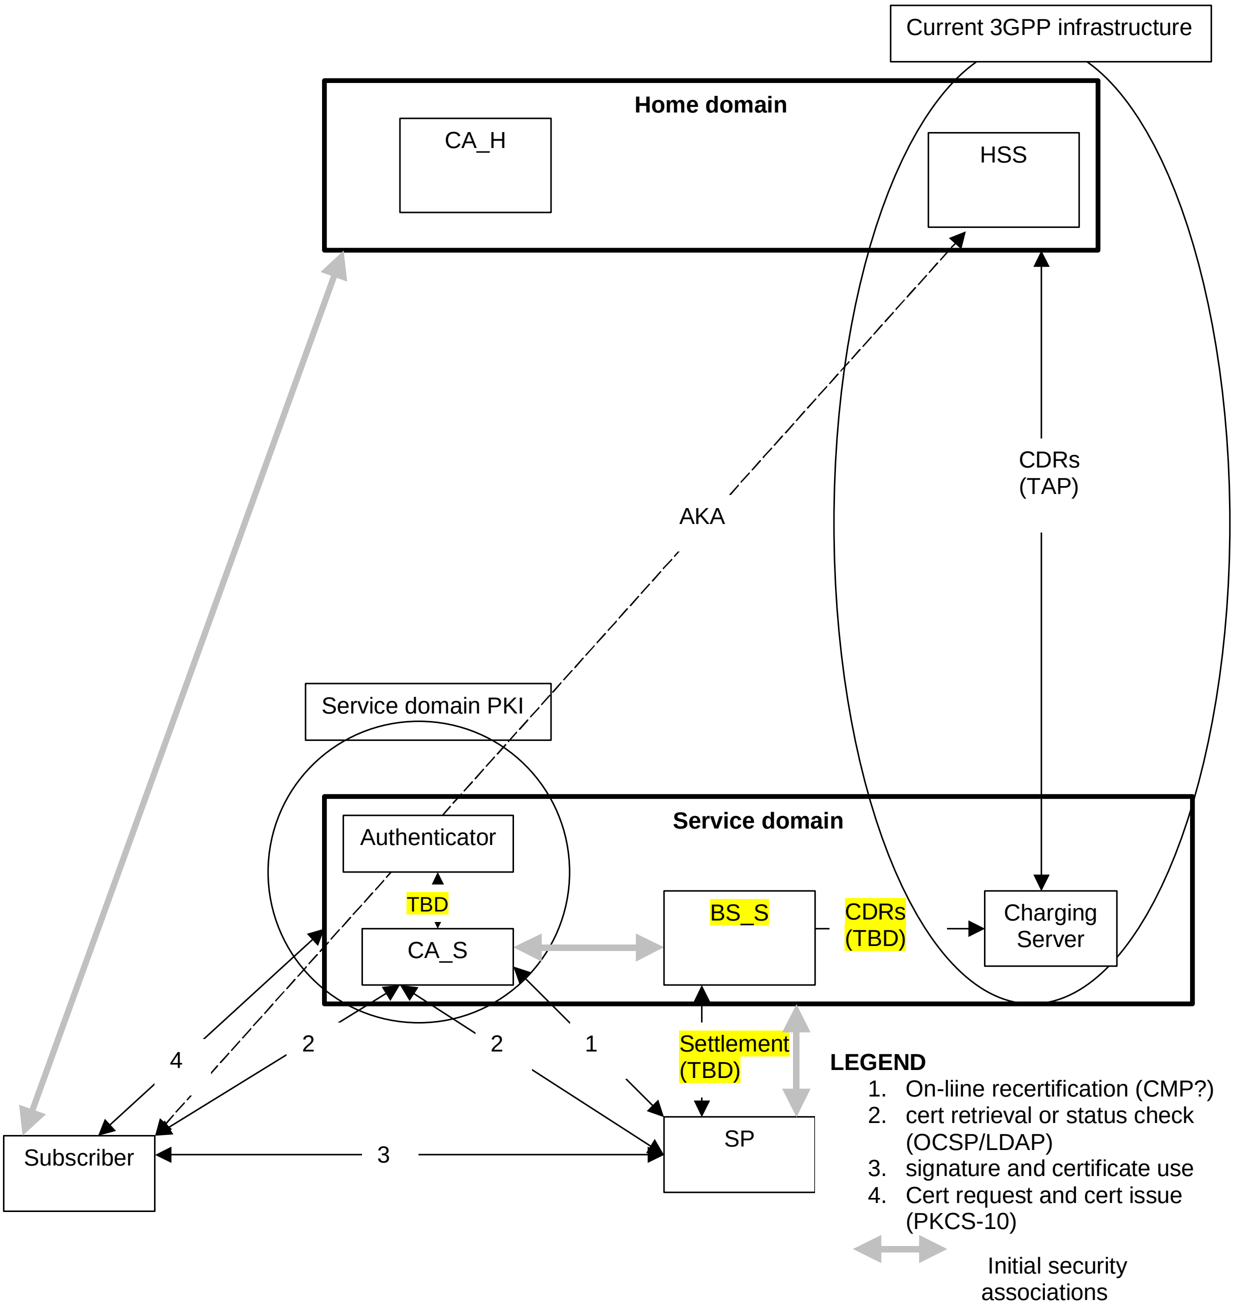
Current 3GPP infrastructure
Home domain
CA_H
HSS
CDRs
(TAP)
AKA
Service domain PKI
Service domain
Authenticator
TBD
CA_S
BS_S
CDRs
(TBD)
Charging
Server
Settlement
(TBD)
SP
Subscriber
1
2
2
3
4
LEGEND
1. On-liine recertification (CMP?)
2. cert retrieval or status check (OCSP/LDAP)
3. signature and certificate use
4. Cert request and cert issue (PKCS-10)
Initial security associations

Figure 2 Use of local PKIs with AKA-authenticated subscriber certification (Note: "Service Operator" can be either the home, or the visited operator)

1.8 How is a subscriber certificate used?

In either model above, the service provider verifies a signature. The SP must also store the signatures as potential evidence in future disputes. In the second model, the service operator domain (e.g., BS_S) should also verify signatures during the settlement phase (if there is one) and store them as evidence. In the first model, the service operator may not be required to verify or store signatures *if* the signatures are forwarded to the home operator domain. This would require either changing the existing CDR transfer protocol (TAP) or defining a new protocol for transferring evidence to the home operator domain. However, in practice, the service operator domain will also verify, and possibly store signatures in order to reduce the likelihood and extent of impact of potential fraud.

1.9 How are disputes resolved?

The subscriber has a direct business relationship with his home operator only. So that is where he will go if there is a dispute about an entry in the subscriber's bill. In the approach described in Section 0, the home operator does not have the signature that gave rise to the CDR. But the service operator domain does have this. So a means to either transfer the signature to the home operator, or to request a dispute check to the service operator domain will be needed. Which one is sensible depends on the refined trust assumptions. The CA_H and the subscriber are required to trust other cellular operators (as is the case for current cellular services).

Without an inter-operator PKI the use of signatures provides non-repudiation by subscribers with respect to service providers but it does **not** provide non-repudiation with respect to the service operator domains! If service operator domain is not trusted by the home operator, then in the latter model above (Figure 2) we need to add:

- an inter-operator PKI with long-term subscriber certificates issued by CA_H
    - the certificate request from the subscriber to CA_S will additionally contain the above long-term subscriber certificates as well.
- alternately, if the subscriber has a long-term public key, it could be added to the subscriber profile.

Although full non-repudiation service is not possible without an inter-operator PKI, the use of subscriber certificates for signature verification keys as described in Section 0 already improves the situation in one respect: if a subscriber obtains a certificate from CA_S and uses it in multiple transactions, he must try to repudiate *all* of them. He cannot selectively repudiate a subset. Linkability of signatures plays a useful role here.

Conclusions

We described what could be reasonable assumptions about trust and security associations for augmenting the 3GPP security architecture with the subscriber certificate capability. We considered the use of off-the-shelf PKI components and an inter-operator PKI in designing this capability. Based on this consideration, we concluded that the "traditional" approach to building a PKI is not suitable in this case. We also showed that the approach based on short-term certificates issued by operators achieves the objectives even in the absence of an inter-operator PKI. Furthermore, when an inter-operator PKI is available, the same approach can still be used, but with an improved level of security guarantees. In this sense, the proposed approach can be seen as a step in the eventual migration towards an inter-operator PKI.

References

[S3-020365] Siemens, "Analysis of Subscriber Certificates Concept", July 2002

[S3-020378] Nokia, “Security and other requirements for subscriber certificates”, July 2002

[S3-020077] Nokia, Usage scenarios for subscriber certificates , February 2002

[S3-020447] SA3 Liasion Statement to SA1 and SA2, “LS on architecture and requirements for subscriber certificates”, July 2002.

[S1-021685] SA1 Liaision Statement to SA2, SA3 and T2, “Liaison Statement on subscriber certificates”, August 2002.

[S3-020414] SA3 technical report, “Network Domain Security / Authentication Framework (NDS/AF); Feasibility Study to support NDS/IP evolution”, July 2002