SeaMoby Working Group L-N. Hamer
Internet Draft Peter Hazy
Nortel Networks
Document: draft-gopal-seamoby-ipsecctxt-issues-01.txt
Expires August 2002 Ram Gopal L
Govind Krishnamurthi
Senthil Sengodan
Nokia
February 2002

Issues in IPSec Context Transfer

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026 [1].

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

The distribution of this memo is unlimited. This memo is filed as <draft-gopal-seamoby-ipsecctxt-issues-01.txt>, and expires August 2002. Please send comments to the authors.

Copyright Notice

1. Abstract

The reasons and the motivation for transferring context have been described in [1]. The requirements for a context transfer protocol can be found in [2]. In this document, we describe issues that need to be considered for transferring security (specifically IPsec) related context between access routers.

There are a large number of IP access networks where one may wish to provide security for end user traffic, or secure the access network from unauthorized traffic. One protocol which may be used to

provide these services is IPSec, which requires a node to establish a security association with the access network in order to obtain these services. Traditionally, such an association is considered static, however, there are many situations in which the ability to move an IPSec security association (SA) from one security gateway (SG) to another within the access network may be beneficial. Examples of this include IPSec handover in mobile LANs and PANs, load-balancing between IPSec SGs, and fail-over applications where high-availability is required. Currently, in order to perform this transfer, it would be necessary to terminate the existing SA and re-negotiate a new SA at the new SG. However, this approach may be inappropriate in cases where high performance is required. Thus, in such cases, the ability to directly transfer an SA from one SG to another would be useful.

The intent of this draft is to describe the unique requirements for transfer of IPSec context and to detail the specific data which must be transferred in order to move an IPSec SA. In addition, a number of unique issues regarding IPSec context transfer will be addressed, and some potential solutions discussed.

2. Conventions used in this document

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

3. Introduction

There are a large number of IP access networks where one may wish to provide security for end user traffic, or secure the access network from unauthorized traffic. One protocol which may be used to provide these services is IPSec, which requires a node to establish a security association with the access network in order to obtain these services. Traditionally, such an association is considered static, however, there are many situations in which the ability to move an IPSec security association (SA) from one security gateway (SG) to another within the access network may be beneficial. Examples of this include IPSec handover in mobile LANs and PANs, load-balancing between IPSec SGs, and fail-over applications where high-availability is required. Currently, in order to perform this transfer, it would be necessary to terminate the existing SA and re-negotiate a new SA at the new SG. However, this approach may be inappropriate in cases where high performance is required. Thus, in such cases, the ability to directly transfer an SA from one SG to another would be useful.

Currently, the only approaches for re-establishing an IPSec session involve tearing down the old IPSec SAs and establishing new SAs with the aid of a key exchange protocol such as IKE [7] or KINK. Unfortunately there are many problems with this approach, the main ones being long latency and excessive signalling during handover. This document addresses another approach, where access routers exchange state information directly. This approach has many advantages, such as reduced latency during handover and minimal signalling from the mobile node.

Much of the work in establishing a generic context transfer framework has already begun [2][3]. These documents focus on the generic requirements and framework for context transfer. However, one must identify, for each feature to which context transfer is applicable, the data which must be transferred, and any unique requirements which are relevant. This document attempts to define the contents of the IPSec feature context and the specific requirements applicable to IPSec. In addition, a number of unique issues regarding the transfer of IPSec context will be discussed, along with potential solutions.

The organization of the document is as follows. Section 4 describes the terminology used in the draft. Section 5 describes several types of Security Associations (SA) and a model for illustrating where/how these SAs fit in. The coverage in this section is illustrative, and is not meant to be exhaustive. Irrespective of the type of SA (whether it is described in Section 5 or not), Section 6 describes the issues involved in IPSec security context for this SA. Section 7 describes the IPSec feature context itself.

4. Terminology

Much of the terminology used in this document is the same as that found in [2]. However, a number of additional definitions are provided below:

- MN - Mobile Node

- PR - Previous AR.

- NR - New AR.

- RN - Remote Node. The entity to which the message is intended.

- FP - Forwarding Path. The path traversed by packets when they go from their source to destination. FP may correspond to the control, data or management plane.

- PFP - Previous FP. The FP when the access router that the MN is connected to is PR.

- o NFP - New FP. The FP when the access router that the MN is connected to is NR. The destination itself may be different along the NFP compared to that along the PFP.

- o Data FP - The FP for packets in the data plane.

- o Control FP - The FP for packets in the control plane.

- o Management FP - The FP for packets in the management plane.

- o G - Gateway. This is an entity from which point onwards the FP to the RN remains unchanged.

- o SA - Security Association. Defined in [4] as a "simplex connection that affords security services to the traffic carried by it." This may be either a Data SA, a Control SA or a Management SA.

- o SG - Security gateway. A network entity with which a node may establish one or more security associations, either from the node to the gateway or vice versa.

- o Data SA - SA that provides security services to packets in the Data FP. Such an SA may be between any two entities along the Data FP.

- o Control SA - SA that provides security services to packets in the Control FP. Such an SA may be between any two entities along the Control FP.

- o Management SA - SA that provides security services to packets in the Management FP. Such an SA may be between any two entities along the Management FP.

Note: An IPSec SA itself is agnostic to the particular type of message that it protects. The definition of Data SA, Control SA and Management SA is for illustration purposes only, but from an IPSec perspective, they are identical. Hence, from an IPSec context transfer perspective, these different SAs are treated identically.

5. Illustrating Various Types of SAs

While discussing IPSec context transfer between ARs, in order to describe several types of IPSec SAs for which the IPSec context needs to be transferred, the following model is illustrative:

```
             +- Access Network -+  +--------------+
+--------+   |     +-----+      |  |       +----+ |   +--------------+
| Mobile |===|=====| PR  |======|==|=======| G  |=|===|       |
| Node   |   |     +-----+      |  |       +----+ |   |   Remote     |
| (MN)   |   |                  |  | Internet *   |   |    Node      |
|        |   |                  |  |          *   |   |    (RN)      |
|        |   |     +-----+      |  |          *   |   |              |
|        |***|*****| NR  |******|**|***********   |   |              |
|        |   |     +-----+      |  |              |   |              |
+--------+   +------------------+  +--------------+   +--------------+
```

Figure 1: Model for Problem Statement

In Figure 1, the line from an MN to RN denoted using "=" represents the previous forwarding path (PFP), while the line from the MN to RN denoted using "*" represents the new forwarding path (NFP). "G" denotes the network entity beyond which the forwarding path to RN remains unchanged. The case of RN representing different physical nodes along PFP and NFP is also subsumed by the above model.

It may be noted that when the AR changes from PR to NR, the change in the forwarding path (FP) may either be marginal or significant. At one extreme (marginal case), G is collocated with PR, while at the other extreme (significant case), G is collocated with RN. The FP may correspond to Data, Control or Management FP. When the Data FP is being considered, RN corresponds to the correspondent node (CN). When the Control FP is being considered, RN corresponds to an entity with which MN exchanges control messages. Similarly, when the Management FP is being considered, RN corresponds to an entity with which MN exchanges messages in the Management plane.

Note:

1. The model has been generalized to accommodate data, control and management plane messages between the MN and a remote node. While it may be the case that management plane messages between MN-RN may either not exist or may not be of interest, the framework covers such aspects as well should they be deemed necessary in certain scenarios.
2. The model may be extended to cover the case of data, control and management plane messages between any two entities, and not restricting itself to the case where one of the entities is MN. It may be reiterated that the model is not meant to be exhaustive, but is merely intended to illustrate many types of SAs whose context may be transferred.

5.1 Types of SAs

For the scenario depicted in Figure 1, several types of SAs are possible:

(1) Type 1 SA: Each endpoint of such an SA lies between G (inclusive) and RN (inclusive). There is no change in such SAs when the FP changes from PFP to NFP.
(2) Type 2 SA: One endpoint of such an SA is at MN, while the other endpoint is between G (inclusive) and RN (inclusive). There is no change in such an SA when the FP changes from PFP to NFP.
(3) Type 3 SA: One endpoint of such an SA is between PR (inclusive) and G (exclusive), while the other endpoint is between G (inclusive) and RN (inclusive).
(4) Type 4 SA: Each endpoint of such an SA lies between PR (inclusive) and G (exclusive). When the FP changes from PFP to NFP, then both endpoints of such SAs change.
(5) Type 5 SA: One endpoint of such an SA is at MN, while the other lies between PR (inclusive) and G (exclusive). When the FP changes from PFP to NFP, then one endpoint of such SAs change.

Along any data FP, control FP or management FP, there may be zero, one or more SAs of any given SA type. While the SAs discussed here are with respect to the PFP, it may be noted that an SA type may only be known after the NFP is determined and any SGs along this NFP are discovered. A reason for this is that "G" is not known until the NFP is known.

5.2 Illustrating SA types

The notation used in the illustrations below is as follows. S1 represents one endpoint of the SA, while S2 represents the other. It is immaterial if the SA itself is directed from S1 towards S2, or vice versa. When an endpoint of the SA (either S1 or S2) may not be collocated within a certain entity (i.e., exclusive case) then this is indicated by lines with "%". For example, in Figure 4, since S1 or S1' may not be collocated with G, two sides of the box have a "%"

Type 1 SA: As illustrated in Figure 2, the two endpoints of the SA (S1 and S2) lie between G (inclusive) and RN (inclusive). Action: No security context needs to be transferred for such SAs during context transfer.

```
   +--------+  +-----+    +----+   +--+      +--+      +--+
   | Mobile |==| PR |=====| G |====|S1|======|S2|======|RN|
   | Node   |  +-----+    +----+   +--+      +--+      +--+
   | (MN)   |                *
   |        |                *
   |        |  +-----+      *
   |        |**| NR  |*******
   |        |  +-----+
   +--------+
```

Figure 2: Type 1 SA between S1 and S2

Type 2 SA: As illustrated in Figure 3, S1 is collocated with MN, while S2 lies between G (inclusive) and RN (inclusive).
Action: No security context needs to be transferred for such SAs during context transfer.

```
   +--------+  +----+     +---+          +--+      +--+
   | Mobile |==| PR |=====| G |==========|S2|======|RN|
   | Node   |  +----+     +---+          +--+      +--+
   | (MN)   |               *
   |        |               *
   |        |  +-----+      *
   |  S1    |**| NR  |*******
   |        |  +-----+
   +--------+
```

Figure 3: Type 2 SA between S1 and S2

Type 3 SA: As illustrated in Figure 4, S1 lies between PR (inclusive) and G (exclusive), while S2 lies between G (inclusive) and RN (inclusive). This type is relevant in environments where, for example, a network entity will provide IPSec security to flows from a specific user, as opposed to the Mobile Node equipment.
Action: In this case, S1' may represent the new entity to which context at S1 is moved. Note that one special case of Type 3 SAs is the case where S1 is collocated with PR, while S1' is collocated with NR.

```
   +--------+  +----+    +--+      %---+     +--+      +--+
   | Mobile |==| PR |====|S1|======% G |=====|S2|======|RN|
   | Node   |  +----+    +--+      %%%%%     +--+      +--+
   | (MN)   |                        *
   |        |                        *
   |        |  +----+    +---+       *
   |        |**| NR |****|S1'|********
   |        |  +----+    +---+
   +--------+
```

Figure 4: Type 3 SA between S1 and S2

Type 4 SA: As illustrated in Figure 5, both S1 and S2 lie

between PR (inclusive) and G (exclusive). This type is relevant since, while some network based IPSec tunnels are static and non-user specific, other SA may be user specific. In the latter case, the IPSec SA information MUST be transfered.
Action: In this case, S1' may represent the new entity to which context at S1 is moved, while S2' may represent the new entity to which context at S2 is moved. Note that one special case of Type 4 SAs is the case where S1 is collocated with PR, while S1' is collocated with NR.

```
   +--------+  +----+    +--+    +--+     %---+          +--+
   | Mobile |==| PR |====|S1|====|S2|=====% G |==========|RN|
   | Node   |  +----+    +--+    +--+     %%%%%          +--+
   | (MN)   |                              *
   |        |                              *
   |        |  +----+    +---+   +---+     *
   |        |**| NR |****|S1'|***|S2'|*******
   |        |  +----+    +---+   +---+
   +--------+
```

Figure 5: Type 4 SA between S1 and S2

Type 5 SA: As illustrated in Figure 6, S1 is collocated with MN, while S2 lies between PR (inclusive) and G (exclusive).
Action: In this case, S2' may represent the new entity to which context at S2 is moved.

```
    +--------+  +----+    +--+      %---+           +--+
    | Mobile |==| PR |====|S2|======% G |===========|RN|
    | Node   |  +----+    +--+      %%%%%           +--+
    | (MN)   |                        *
    |        |                        *
    | (S1)   |  +----+    +---+       *
    |        |**| NR |****|S2'|********
    |        |  +----+    +---+
    +--------+
```

Figure 6: Type 5 SA between S1 and S2

5.3 Grouping SA types: Multi-lateral and Multi-level SAs

The five SA types, described in this document,
may be grouped in a variety of fashions within the actual topology.
These are illustrated below :

- o SAs may be grouped in a multi-lateral fashion, such that one or more of the same or different types of SAs are present in a sequential fashion. Figure 7 illustrates this scenario, where S1-S2 represents one SA, while Sa-Sb represents another.

```
+--------+  +-----+   +--+    +--+     %---+          +--+
| Mobile |==| PR |====|Sa|====|Sb|=====% G |==========|RN|
| Node   |  |(S2)|    +--+    +--+     %%%%%          +--+
| (MN)   |  +-----+                      *
|        |                               *
|        |  +------+  +---+   +---+      *
| (S1)   |**| NR   |***|Sa'|***|Sb'|*******
|        |  |(S2')|   +---+   +---+
+--------+  +------+
```

Figure 7: Illustrating multi-lateral SAs

o SAs may be grouped in a multi-level fashion, such that one or more of the same or different types of SAs are stacked one upon the other. Security associations may be combined into bundles in two ways: transport adjacency and iterated tunneling, as described in section 4.3 of [4]. Figure 8 illustrates this scenario, where S1-S2 represents one SA, while Sa-Sb represents another.

```
+--------+  +-----+   +--+    +--+     %----+          +--+
| Mobile |==| PR |====|Sa|====|Sb|=====% G  |==========|RN|
| Node   |  |    |    +--+    +--+     %(S2)|          +--+
| (MN)   |  +-----+                    %%%%%%
|        |                                *
|        |  +------+  +---+   +---+       *
| (S1)   |**| NR   |***|Sa'|***|Sb'|********
|        |  |     |   +---+   +---+
+--------+  +------+
```

Figure 8: Illustrating multi-level SAs

5.4 Some Specific Scenarios

Some typical scenarios where security context needs to be transferred are now discussed.

5.4.1 MN-PR control FP SA

Figure 9 illustrates the case where a control SA that exists between

MN-PR is replaced by a control SA that exists between MN-NR.

```
+--------+  +-----+
| Mobile |==| PR |
| Node   |  |(S2)|
| (MN)   |  +-----+
|        |
|        |  +------+
| (S1)   |**| NR  |
|        |  |(S2')|
+--------+  +------+
```

Figure 9: Illustrating MN-PR Control FP SA

5.4.2 MN-SG data FP SA

Figure 10 illustrates the case where a data SA that exists between MN-SG1 is replaced by a data SA that exists between MN-SG2.

```
+--------+  +-----+    +------+         %---+          +--+
| Mobile |==| PR |====| SG1 |=========% G |==========|RN|
| Node   |  |(S2)|    |(S2) |         %---+          +--+
| (MN)   |  +-----+    +------+           *
|        |                               *
|        |  +------+   +------+           *
| (S1)   |**| NR  |***| SG2 |************
|        |  |(S2')|   |(S2')|
+--------+  +------+   +------+
```

Figure 10: Illustrating MN-SG Data FP SA

6. Context Transfer Requirements

All of the specific requirements defined in [3] are applicable to the transfer of IPSec context. In addition, one of the primary requirements of IPSec is that the identity of the security gateway MUST be known to the mobile node, in order to properly encapsulate packets for transmission.

6.1 Entities involved in context transfer

It was seen that no security context needs to be moved for Type 1 or

Type 2 SAs, but security context may need to be moved for Type 3, 4 or 5 SAs. Irrespective of the type of SA (Type 3, 4 or 5) whose context needs to be transferred or the number of such SAs for which context needs to be transferred, the context is always transferred from the PR to the NR. This has the following implications:

- All IPSec security context for any SA in the PFP may need to be made available to the PR.

- After the IPSec security context for one or more SAs in the PFP has been transferred from PR to NR, a mechanism is needed to move these contexts to the appropriate SA endpoint in the NFP.

As an illustration of this approach, consider a Type 4 SA in Figure 5, where context at S1 needs to be moved to S1', and context at S2 needs to be moved to S2'. In this case, the contexts at S1 and S2 are made available at PR, which then transfers them to NR, and NR subsequently moves these contexts to S1' and S2' respectively.

For the case where the security context already exists at PR and where the new location of the security context is NR (for instance, the case of Section 4.4.1), no context needs to be moved either from a different entity to PR along PFP or from NR to a different entity along NFP.

6.2 Discovery and Update of SG Identity at MN

Due to the peer to peer nature of the IPSec architecture, it is necessary for the nodes participating in an SA to know the identity of each other. Under normal circumstances, this is not an issue. However, in the case of context transfer, the identity of the SG may change on-the-fly. As a result, there must be some way to ensure that the mobile node transmits packets to the correct SG, even after a context transfer. There are two primary solutions to this issue:

Direct SG Communication:

Direct SG communication requires that the node be able to discover the address of the initial SG through some means (DHCP, DNS, etc). In addition, during handover to a new SG, the node must be notified of the new SG address through some form of signaling so that the local SAD in the MN may be updated to reflect the address of the new SG.

Indirect SG Communication:

This form of communication requires some form of tunnel to be set up from the node to a virtual SG. This is achieved with the use of a virtual address for the SG. The node must somehow retrieve (via DHCP, DNS, etc) the virtual address necessary to communicate to the SG currently serving the MN. Then, during handover, the network takes care of correctly re-directing traffic destined to the new SG, making the process transparent to the mobile node.

Each method has its pros and cons. Direct communication reduces the complexity in the network but requires additional signaling, and thus added latency. The indirect form requires added complexity at the network, but is transparent to the node. One must weigh these factors when considering an appropriate mechanism for solving this problem.

Moreover, from a system perspective, when the FP changes from PFP to NFP, one may also need to be concerned about whether the NFP and/or SGs along the NFP are known prior to or after the security context transfer takes place:

o One or more security contexts associated with SAs along the PFP, may be transferred from PR to NR "prior" to discovering either the NFP and/or SGs along the NFP.

o One or more security contexts associated with SAs along the PFP, are transferred from PR to NR "only after" the NFP is known and the SGs along the NFP have been discovered.

There are trade-offs to each of these two approaches. In the former case, the knowledge of the PFP, SGs and their associated security context along the PFP at the NR may help in establishing the NFP. The NFP, of course, would still have to be in conformance with the user's static profile. In the latter case, knowing the NFP, the NR may selectively request security context transfer (or may even not have to request any security context transfer) from PR.

Such discovery of SGs along any FP is not within the scope of this document nor is within the scope of any security context transfer mechanism defined. It is assumed that protocol(s) for discovery of SGs will be standardized by other WGs in the IETF.

6.3 PMTU Rediscovery

The IPSec architecture, as defined in [4], requires that the nodes participating in an SA be aware of the Path MTU over which the treated packets are traveling. However, if a context transfer occurs, and the new SG is in another location in the network, it is possible for the PMTU of the underlying network to change. This is not a problem if the PMTU of the new path is greater than that of the old path. However, if the PMTU decreases, this may cause problems for applications which rely on knowledge of the PMTU. Possible solutions to this problem may include PMTU rediscovery, or even network engineering to avoid the problem entirely. However, this issue is really one of general context transfer, and thus will not be discussed here.

6.4 SA Conflict Resolution

During a context transfer, it is possible that the new SG which has been targeted as the candidate for context transfer may not be able to support the SAs being transferred (i.e., unavailability of ciphering algorithms, etc). The method for dealing with this situation is beyond the scope of this document, but may include selection of a new candidate SG, or the termination of the IPSec SAs, forcing the mobile node to establish a new SA pair with the new SG, allowing for re-negotiation of the SA parameters.

6.5 IPv4 and IPv6 address support

This requirement fits within the framework of general context transfer requirements. Context transfer should be possible irrespective of whether one or both ARs have either an IPv4 or an IPv6 address.

6.6 Private Addressing support

This requirement fits within the framework of general context transfer requirements. Context transfer should be possible irrespective of whether one or both ARs have a private IPv4 address. One mechanism that may be used to handle the case of private addresses is as described in [9].

7. IPSec Feature Context

When determining the contents of the IPSec feature context, one must examine all the state, which is maintained at the SG. The actual data, which is stored in the gateway is collected in the Security Policy Database (SPD) and the Security Association Database (SAD) [5].

The IPSec feature context itself is comprised of several components:

- Part or whole of the static profile associated with the 'user'.
- Selector fields of an SA
- SPI value
- The static attributes of an SA
- The dynamic attributes of an SA
- Replay window parameters

Each of these is discussed in greater detail below.

7.1 User's static profile

Security requirements pertaining to a user are stored within the user's static profile. Among other things, such a profile may specify the range of security mechanisms, security algorithms, key lengths etc. that may be used for security associations pertaining to the user. The various options provided in the profile may also be in a certain preferred order, so that a certain option is only chosen if an option of higher preference is not available. For instance, a decreasing order of preference for encryption algorithms could be "AES, 3-DES, DES", implying that 3-DES is to be used only if AES is cannot be used, and that DES is to be used only if both AES and 3-DES cannot be used. Similarly, a decreasing order of preference for key lengths used within AES could be "256, 192, 128", implying that a 192-bit key is to be used only when a 256-bit key cannot be used, and that a 128-bit key is to be used only when both 256-bit and 192-bit keys cannot be used.

The reason that the user's static profile may have to be known at the NR is that along the NFP, it may be possible to use a certain option in the static profile that has a higher preference than which is used within an SA along the PFP. For instance, if we consider a Type 5 SA (as shown in Figure 6), it may be possible that 3-DES is used between S1-S2 because S2 does not support AES, but that S2' supports AES. In this case, although not always, it may be desirable to support the more preferred option along the NFP. If the user's static profile is not sent, then the NR would not have this choice.
It should be noted that the changing of attributes for a security association may be achieved by updating an existing SA or establishing a new SA. In both cases messaging may be required to the SA endpoints. For instance, changing the encryption algorithm for a type 5 SA would require messages sent to the mobile node.

It may be noted that the security policy database contains some static entries, containing general policies, which are established by the operator of the access network. These are user independent and should not be transferred. Generally, these SPD entries are the same on all SGs within the operator domain.

7.2 Selector fields of an SA

The SPD also contains selector parameters used to support SA management to facilitate control of SA granularity. In fact, an SA may be fine-grained or coarse-grained, depending on the selectors used to define the set of traffic for the SA. Selector fields are used to determine which packets are provided the services of a particular SA. While some selector fields are always sent, others are optional. Selectors that are always sent are the source IP address and the destination IP address. Optional selectors are the source port, destination port, TOS byte etc.

The selectors used to define the SA MUST be context transferred.

Selector Fields:

Source and Destination IP Address
Source and Destination Port
Transport Layer Protocol
Name
Sensitivity Level

These fields are used by the gateway to identify packets for inbound and outbound processing. Note, fields not used to match packets against this SA MAY be omitted. Therefore, if, for example, only the source and destination IP addresses are used as a selector, the other fields MAY be excluded.

For inbound processing, the following packet fields are used to look up the SA in the SAD:

Outer Header's Destination IP address.
IPSec Protocol (AH or ESP)
SPI: the 32-bit value used to distinguish among different SAs terminating at the same destination and using the same IPSec protocol.

These fields are used by a gateway to look up the SA in the SAD. Therefore, they MUST be context transferred.

A major problem that may occur during the context transfer of an IPSec SA is when the SPI value of an IPSec SA to be transferred is already in use at the new SG. In this scenario, three possible solutions are to be considered:

-Deny the context transfer.
-Accept the context transfer but force re-negotiation of the IPSec SA.
-Assign a new SPI entry unused at the new SG and signal this information back to the mobile node. This operation MUST be secure since it may open up some security holes.

7.3 Treatment Fields:

Sequence Number
Sequence Number Overflow Flag
Antireplay Window
AH Algorithm, keys, etc
ESP Encryption Algorithm, keys, IV Mode, IV, etc
ESP Authentication Algorithm, keys, etc
Lifetime
Protocol Mode
Path MTU

Treatment fields are used by the IPSec stack in actually processing the packets, once it has been determined that they must be treated. Again, fields which are not applicable to this SA MAY be omitted. For example, depending on the protocol mode, either the ESP or AH fields need not be transferred. Others may not need to be transferred depending on the IPSec implementation (for example, some IPSec stacks allow disabling of sequence number checking, thus these fields may not need to be transferred).

7.3.1 Static attributes of an SA

Static attributes of an SA refer to those attributes that are instantiated at the time of SA establishment, and which do not change with time. Examples of such attributes include:

- Authentication/encryption algorithm
- Key length
- Block size
- Algorithm mode

7.3.2 Dynamic attributes of an SA

Dynamic attributes of an SA refer to those attributes that change with time. Examples of such attributes include:

- SA duration (in terms of seconds or bytes transmitted) before key refresh

- Replay window parameters (discussed below)

A receiver of an IPSec SA may decide to activate anti-replay protection. Parameters relevant to anti-replay protection are:

- window size
- highest sequence number of an authenticated packet
- indication of whether packets within window have been successfully received or not

When the IPSec security context is transparently sent from PR to NR, the replay window parameters are also sent. Using these parameters, the entity that receives this security context from the NR in the NFP, may transparently start providing anti-replay protection.

These fields may initially cause concern, as they must be updated in real time, and should reflect the current state of the IPSec SA. The concern is that these fields may not be entirely accurate after context transfer because of the loss of some user packets. Careful considerations reveal this is not a problem. In fact, IPSec anti-replay functionalities were designed to accommodate minor packet loss [4].

8. Security Considerations

Careful consideration needs to be taken to ensure that the context transfer of an IPSec SA is secure, especially when transferring SA information such as keys. In fact, as defined in [3], context transfer MUST be secure. How to secure the context transfer is dependent on the network environment. In a trusted environment, no additional security mechanism is needed. But in an un-trusted environment, a security mechanism MUST be utilized.

In order to keep the context transfer protocol simple, re-use of existing security technologies is recommended. All security requirements MAY be provided at the network layer with IPSec or at the transport layer with TLS.

9. References

[1] S. Bradner, "keywords for use in RFCs to Indicate Requirement Levels", RFC2119 (BCP), IETF, March 1997.

[2] The seamoby CT design team, "Context transfer: problem statement", draft-ietf-seamoby-context-transfer-problem-stat-00.txt.

[3] The seamoby CT design team, "General Requirements for a Context Transfer Framework", draft-ietf-seamoby-ct-reqs-00.txt.

[4] S. Kent et. Al., "Security Architecture for the Internet Protocol" RFC-2401, November 1998

[5] Kent, S., and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

[6] Kent, S., and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[7] Harkins, D., and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[8] D. Trossen et. al., "Protocol for Anticipated Candidate Access Router Discovery for Seamless IP-level Handovers," draft-trossen-seamoby-CARdiscovery-anticipated-00.txt, Aug 2001.

10. Acknowledgments

The authors would like to thank Brett Kosinski who participated in the first version of this draft. We would also like to thank the following people for their useful comments and suggestions related

to this draft: Hamid Syed, Gary Kenward, Jerry Chow and Bill Gage.

11. Author's Addresses

Louis-Nicolas Hamer
Peter Hazy

Nortel Networks
Ottawa, ON
CANADA
Email: {nhamer,hazyp}@nortelnetworks.com

Ram Gopal.L
Govind Krishnamurthi
Senthil Sengodan

Nokia Research Center
5 Wayside Road
Burlington, MA 01803
USA
Email: {ram.gopal,govind.krishnamurthi,senthil.sengodan}@nokia.com

Full Copyright Statement

Expiration Date

This memo is filed as <draft-gopal-seamoby-ipsecctxt-issues-01.txt>, and expires August 2002.