

NSIS
Internet Draft

H. Tschofenig
Siemens
H. Schulzrinne
Columbia U.
R. Hancock
A. McDonald
Siemens/Roke Manor
X. Fu
Univ. Goettingen

Document: [draft-tschofenig-nsis-sid-00.txt](#)
Expires: December 2003

June 2003

Security Implications of the Session Identifier
<[draft-tschofenig-nsis-sid-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

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

Abstract

As one result of the analysis activities in the NSIS group it was realized that mobility and the ability to change the flow identifier causes problems with existing QoS reservations. To be able to associate a signaling message with existing state an identifier other than the flow identifier had to be used. Such an abstraction is achieved with the session identifier which allows identification of established state independently of the flow characteristics.

Security Implications of the Session Identifier June 2003

Although the introduction of a session identifier sounds simple and beneficial, it introduces a problem which is subsequently referred to as the session ownership problem.

This document describes the session ownership problem, the implications for an NSIS protocol and summarizes already discussed solutions.

Table of Contents

1.	Introduction.....	2
2.	Problem Description.....	3
2.1	Mobility.....	3
2.2	Local Repair Case.....	4
3.	Solution Discussion.....	5
3.1	Local solutions.....	6
3.1.1	Authorization Token.....	6
3.1.2	Context Transfer.....	6
3.1.3	Centralized Entity.....	7
3.2	Global Solutions.....	7
3.2.1	Purpose Built Key Based Approach.....	7
3.2.2	Hash Series Based Approach.....	8
3.2.3	Confidentiality protection of session identifier.....	9
4.	Pending Issues.....	10
5.	Summary.....	10
6.	Security Considerations.....	11
7.	Open Issues.....	11
8.	References.....	11
	Acknowledgments.....	12
	Author's Addresses.....	12

[1.](#) Introduction

As one result of the analysis activities in the NSIS group it was realized that mobility and the ability to change the flow identifier causes problems with existing QoS reservations. To be able to associate a signaling message with existing state an identifier other than the flow identifier had to be used. Such an abstraction is achieved with the session identifier which allows identification of established state independently of the flow characteristics.

Although the introduction of a session identifier sounds simple and

beneficial, it introduces a problem which is subsequently referred to as the session ownership problem. Although the problem is known for a very long time (discussion took place already at the 53rd IETF and even proposals for solving the problem have been mentioned) the topic still has some grey spots.

Security Implications of the Session Identifier June 2003

This document describes the session ownership problem, the implications for an NSIS protocol and summarizes already discussed solutions.

[2.](#) Problem Description

To allow signaling messages to refer to existing state some sort of identifier is required. In RSVP this identifier is based on the flow identifier.

To support mobility and to introduce the ability to change the flow identifier mid-session and mid-path an additional identifier is required. Throughout this text we call this additional identifier the session identifier. Section 4.5.2 of [\[NSIS-FW\]](#) provides a description of the different identifiers used in NSIS.

When a NSIS node receives a signaling message then it has to check whether state information already exists, or whether new state has to be established. The session identifier can quickly provide information whether state information is already available.

Some of the described problems are less problematic in non-mobile environments since the first NSIS-aware router (for example the edge or access router) can associate authentication state with the session identifier, and hence ownership can be verified. However, if we assume a mobility scenario then the movement of a node makes this verification step much more difficult since each NSIS-aware node along the path could possibly be forced to do this verification.

[2.1](#) Mobility

Figure 1 shows an NSIS Initiator which established state information at NSIS nodes along the path as part of the signaling procedure. As a result Access Router 1, Router 3 and Router 4 (and other nodes)

store state information including the Session Identifier SID-x.

Security Implications of the Session Identifier June 2003

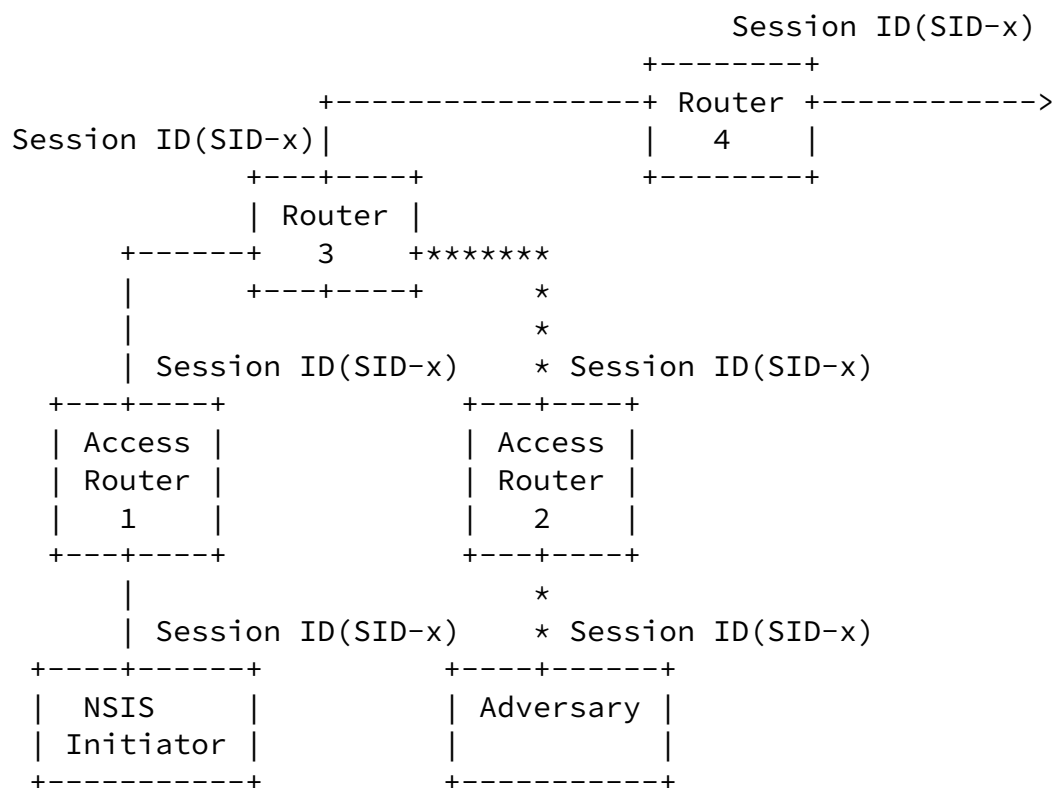


Figure 1: Mobility Scenario

The Session Identifier is included in signaling messages as a reference to the established state.

If an adversary were able to obtain the Session Identifier for

centralized entity) whereby the centralized entity needs to be queried to perform verification or some information is shipped around but allows local verification.

3.1.1 Authorization Token

The authorization token approach requires that an entity in the network produces a token during the initial signaling message exchange. This token is cryptographically protected and must be verifiable by entities in the local domain. The authorization token must be protected to prevent an adversary at the wireless link to intercept the token and to reuse it. The authorization token allows link subsequent actions to an initial authorization (e.g. by including the token into the signaling message after a handover). The end host only forwards the authorization token. Hence the token itself does not provide authentication.

This solution is referred as local since the token has to be granted and verified within the same domain (or within domains with a pre-defined trust relationship).

A more sophisticated version would use a concept similar to Kerberos tickets whereby the end host actively participates in the protocol by showing the knowledge of the session key. As authorization information the ticket could include the session identifier.

3.1.2 Context Transfer

Context Transfer is another approach the network to associate a new signaling message to a previous one. The Context Transfer protocol allows to move state information from one access router to another. The assumption thereby is that if state was create at one particular access router then the forwarded state would also allow the new access router to verify the incoming signaling message. Due to the transitive trust provided by hop-by-hop security protection in

intermediate router would trust that the signaling message have been correctly verified and only the authorized entity issued the signaling message.

A short discussion of context transfer in relationship with RSVP is provided in Section 1.2.6 of [[Tho02](#)]. The same considerations are also applicable to CT for NSIS.

[3.1.3](#) Centralized Entity

Using this approach a cross-over router where the new path hits the old path and where authorization is desired information inside the signaling message (e.g. a token which only points to state installed at the centralized entity or user credentials) could be used to perform this verification. For QoS signaling the Policy Decision Point would be a possible centralized entity.

Different to authorization token approach the token does not need to be verified at an individual node itself. For the authorization token approach it is necessary to provide the necessary information within the token itself. In case of a central entity only a reference to the stored information needs to be provided. The distinction between the authorization token and this approach is therefore, from an NSIS protocol point of view, marginal. A solution could possible support both mechanisms easily (e.g. [\[RFC3521\]](#) and [\[RFC3520\]](#)).

[3.2](#) Global Solutions

[3.2.1](#) Purpose Built Key Based Approach

This approach makes use of a cryptographic session identifier and follows the idea described in [\[PBK\]](#). The identifier is created as:

Session ID = PRF(Public Key)

The output of the PRF function needs to be truncated if necessary to fit the length requirements of the Session ID. As a PRF function MD5 or SHA-1 could be used.

The signaling initiator would therefore create a public / private key pair before starting NSIS signaling for a specific session.

Every time the end host roams from one location to another the following information is added to the NSIS signaling message:

- Session ID
- Public Key

- Digital Signature of some signaling message payloads including a timestamp

A receiving entity (e.g. cross-over router) can verify that

- the Session ID matches the hash of the public key
- the private key, which was used to compute the digital signaling, matches to the public key (by verifying the digital signature)
- the authorization indication is fresh (verifying the timestamp)

Note that this approach does not require a public key infrastructure. It only makes use of the inability of an adversary to later compute a key pair whereby the hash of public matches the session identifier. Since the session identifier is stored at the individual routers along the path it is not possible adversary to masquerade the owner of the session id.

As described above replay protection is provided with the help of timestamps.

The disadvantages of this approach are:

- Performance for the public key operation
- Size of the messages (the public key has to be included in addition to other fields)
- Only the creator of the key pair is able to authorize actions (if we ignore delegation approaches). This seems to be very restrictive.

[WC02] follows a similar approach by distributing a public key along the path. Whenever an update is required then the message containing the previous session identifier, the new session identifier, the public key and a sequence number. The sequence number field is not only a short random number; instead it is a digital signature computed over the care-of address and the sequence number. A successful verification at the cross-over router stores the new values along the entire path.

[3.2.2](#) Hash Series Based Approach

The Hash Series approach provides better performance than the PBK-based approach. To set up the protocol a random T_0 number is created and hashed n times as shown below. The length of the hash series is chosen by the creator with the value n :

$$\begin{aligned}T_1 &= \text{hash}(T_0) \\ T_2 &= \text{hash}(T_1) \\ &\dots \\ T_n &= \text{hash}(T_{n-1})\end{aligned}$$

The session identifier is chosen in such a way that it equals T_n .

Security Implications of the Session Identifier June 2003

The hash values are then released in the reverse order. Every hash value is used only once and the number of the latest hash value has to be stored. Since the total number of hash values has to be set at protocol startup it is necessary to "change" the hash series after all values are exhausted. Since the hash series is associated to the session identifier it is also necessary to change the session identifier or to restart the protocol with a new session identifier.

A signaling message would therefore include:

- Session ID (=Tn)
- Total number of hash values (n)
- Current hash value (Ti)
- Index of current hash value (i)

A verifying entity would therefore recompute the hash chain to verify that the session id is valid. Furthermore it is necessary to compare the received hash value with the lasted one received (if no hash value got lost) T_{x+1} would have to equal $\text{hash}(T_x)$.

To prevent reuse of a hash value by an adversary it is necessary that all nodes along the path store the latest valid value.

[3.2.3](#) Confidentiality protection of session identifier

This approach is very simple and follows the following arguments:

- An adversary can only mount an attack if it knows the Session ID
- The end host has to trust the intermediate nodes and networks to perform according to the expected behavior. Due to the flexible protocol operation it is necessary for intermediate nodes to act on behalf of the end host.

Providing confidentiality protection to protect the Session ID makes it more very difficult for an adversary to eavesdrop the session identifier and to reuse it for a subsequent attack.

Confidentiality protection of the session identifier therefore addresses attacks from outsiders (entities which do not actively participate in the NSIS signaling protocol). Hence it must be assured that the session identifier is never transmitted in clear between two signaling entities (e.g. in clear over the wireless link). Adversaries along the path (i.e. an NSIS node which was intercepted by an adversary) are not addressed by this approach.

It is obvious that the session identifier must be chosen in a way which does not allow an adversary to guess it. One possibility is to choose the value for the Session Identifier randomly with each session. It must be ensured that the identifier is sufficient large (e.g. 128 bits).

Security Implications of the Session Identifier June 2003

This approach was selected for CASP [[CASP](#)].

[4.](#) Pending Issues

- Replay protection

Replay protection for the solutions described in [Section 3](#) is hard. Assuming globally synchronized clocks for timestamp-based replay protection is possibly hard to justify.

- Authorizing entity

To keep the NSIS protocol flexible it seems that it is undesirable to restrict certain actions only to a single entity (e.g. to the signaling initiator). Some solutions discussed in [Section 3](#) tend to force such an approach. The question therefore is: "Which entity is allowed to authorize which actions?"

- Certain solutions discussed in [Section 3](#) require the distribution (and storage) of information along the path.

- Signaling message behavior

NSIS signaling messages do not always travel end-to-end. Instead, in mobility scenarios it is sometimes desirable to start or to terminate the signaling message exchange somewhere along the path. Is this still valid or do all message have to travel end-to-end?

- Local or global solution

It is obviously much simpler to provide a solution which works locally. Since the ownership problem could possibly require verification at any node along the path it seems to be that a global solution should be targeted.

[5.](#) Summary

To provide proper security for the session ownership problem a solution has to face many challenges including performance, state maintenance, replay protection and most important - the flexibility of the NSIS protocol itself.

The above-described problem of authorization is not restricted to communication at the edge as described above. The problem basically occurs anywhere in the network whenever an old path becomes invalid and a reservation along a new path has to be established. The merge point (or cross-over router from the above mobility scenario) has to make sure that only the legitimate owner of the reservation issued this request.

Introducing a session identifier for the purpose of more efficient mobility handling needs to be carefully compared to the additionally introduced complexity (for example by the corresponding security mechanism). The benefits gained by this new concept can easily be destroyed by heavy-weight security mechanisms or by introducing new security vulnerabilities.

[6.](#) Security Considerations

This document addresses security issues of the Session ID. To provide a more detailed threat analysis it is necessary to resolve the pending issues listed in [Section 4](#).

This threat is also briefly described in [\[THREATS\]](#).

The solutions described in this document do not aim to provide protection for signaling messages itself.

[7.](#) Open Issues

Adding multicast handling to an NSIS adds a number of further open issues. In case of multicast it is possible that nodes join and leave the multicast group. If sensitive information is transmitted to the active signaling entities then a previously joined node can later perform some actions even after leaving the multicast group.

Sooner or later it is necessary to come up with a definition of the problem we are aiming to solve. Such a definition might look like:

"An NSIS message is authentic if it originates from the initiator for a session, or from an NSIS node that has been authorized to act on behalf of the initiator by virtue of taking part in the NSIS signaling session."

8. References

[THREATS] H. Tschofenig and D. Kroeselberg: "Security Threats for NSIS", Internet Draft, Internet Engineering Task Force, March 2003. Work in progress.

[NSIS-FW] R. Hancock, I. Freytsis, G. Karagiannis, J. Loughney and S. Van den Bosch: "Next Steps in Signaling: Framework", Internet Draft, Internet Engineering Task Force, March 2003. Work in progress.

[PBK] S. Bradner, A. Mankin, and J. Schiller, "A framework for purpose built keys (PBK)", Internet Draft, Internet Engineering Task Force, November 2002. Work in progress.

Security Implications of the Session Identifier June 2003

[WC02] C. Westphal and H. Chaskar: "QoS Signaling Framework for Mobile IP", Internet Draft, Internet Engineering Task Force, June 2002. Expired.

[CASP] H. Schulzrinne, H. Tschofenig, X. Fu, and A. McDonald, "CASP - Cross-Application Signaling Protocol", Internet Draft, Internet Engineering Task Force, 2003. Work in progress.

[RFC3521] L. Hamer, B. Gage, and H. Shieh, "Framework for session set-up with media authorization," [RFC 3521](#), Internet Engineering Task Force, April 2003.

[RFC3520] L. Hamer, B. Gage, B. Kosinski, and H. Shieh, "Session Authorization Policy Element", [RFC 3520](#), Internet Engineering Task Force, April 2003.

[Tho02] M. Thomas: "Analysis of Mobile IP and RSVP Interactions", Internet Draft Internet Engineering Task Force, (work in progress), October 2002.

Acknowledgments

We would like to thank Rainer Falk for his comments to this draft.

Author's Addresses

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Hannes.Tschofenig@siemens.com

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
EMail: schulzrinne@cs.columbia.edu

Robert Hancock
Roke Manor Research
Old Salisbury Lane
Romsey
Hampshire
SO51 0ZN
United Kingdom
Email: robert.hancock@roke.co.uk

Andrew McDonald
Roke Manor Research
Old Salisbury Lane
Romsey, Hampshire
UK
EMail: andrew.mcdonald@roke.co.uk

Xiaoming Fu
Institute for Informatics
University of Goettingen
Lotzestrasse 16-18
37083 Goettinge
Germany EMail: fu@cs.uni-goettingen.de

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of

Security Implications of the Session Identifier June 2003

developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.