NSIS Working Group
Internet Draft
Document: draft-tschofenig-nsis-threats-00.txt
Expires: August 2002

Hannes Tschofenig
Siemens

May 2002

NSIS Threats
<draft-tschofenig-nsis-threats-00.txt>

Status of this Memo

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

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

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

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

Abstract

As the work in the NSIS working has begun to describe requirements and the framework people started thinking about possible security implication. This document should provide a starting point for the discussion at the NSIS interim meeting and at the NSIS working group mailing list regarding the security issues that have to be addressed. It does not describe threats for a particular published protocol. This memo is furthermore meant to create awareness for the security within the group. The threat scenarios in this document are matched against the security requirements described in [1].

1 Introduction

It is often argued that QoS signaling protocols are similar to other signaling protocols and one might re-use their security mechanisms for avoiding reengineering overhead. This is true up to some point: A QoS signaling protocol might borrow many security mechanisms from other protocols but different trust assumptions, and different protocol processing may demand different solutions or adaptations. This document tries to show security issues that need to be addressed by a QoS signaling protocol that claims to be secure. Although the base protocol might be sure, some extensions may cause problems when used in a particular environment. We think that it is necessary to investigate the kontext in which a QoS protocol is integrated and in which sequence protocols are executed (when combined together with other protocols). A particular focus of QoS signaling protocols should be given to the interaction with accounting and charging solutions: Without an appropriate integration of QoS and accounting protocols there is no good incentive for network operators to deploy them.

Independent of the threat scenarios described in Section 3 we indentify the following structural pieces, which require different security protection because of different trust relationships. The sub-parts are:_access network part, intra and inter-domain part, and the issues related to the end-to-end communication. These parts are briefly described. The threat scenarios in Section 3 can be assigned to the individual parts.

a) Access Network

This section addresses threats that arise when the QoS Inititiator (QI) is attached to access network and transmits and receives QoS signaling messages. There might not exist a pre-established trust relationship between a user and the access network, as in many mobility scenarios it is usually assumed.

Threat scenarios dealing with initial QoS security association setup, replay attacks, lack of confidentiality, denial of service, integrity violation, identity spoofing and fraud are applicable. From a security point of view this part of the network causes the most problems.

b) Intra-Domain

After receiving and verifying a QoS request at the access network the signaling messages traverse the network within the same administrative domain. Since the request has already been authenticated and authorized threats are different compared to those described in the previous section. To differentiate the user-to-access network interface with the intra-domain communication (i.e. communication within the core-network) we assume that no user hosts are attached to the core-network. (That is: the interface between any host and the first router is part of the access network). We furthermore assume that nodes within one administrative domain have a stronger trust relationship between each other.

c) Inter-Domain

The security considerations at the border between different administrative domains largely depends on how accounting is done. If one domain transmits forged QoS reservations (for example stating a higher QoS reservation than a aggregated number of user did) to next domain then it is likely that the originating network domain has also has to pay for the reservation. Hence in this case, there is no real benefit for the first network domain to forge a QoS reservation. But if the user is directly charged by intermediate domains too then this kind of attack may be reasonable. Security protection of messages transmitted between different administrative domains is still necessary to tackle attacks like spoofing, integrity violation, denial of service etc. The lower number of networks and higher trust relationship (compared in the access network case) cause fewer problems for a key management.

d) End-to-End

In our opinion end-to-end security for QoS signaling messages is rarely required if we assume that end-to-end issues like charging

and the selection which user has to pay for a reservation is already securely negotiated by preceding upper layer protocols (for example SIP). Information carried within a QoS signaling protocol for the purpose of charging is therefore assumed opaque to the QoS protocol itself and appropriately protected as part of the AAA interaction. For accounting data, the QoS signaling protocol is therefore only used as a transport mechanism. Note however that this assumption strongly depends on the chosen solution of a protocol interaction with AAA, QoS and application layer protocol. It is however possible to select a charging solution that requires end-to-end protection of information delivered within the QoS signaling protocol. The following example requires some sort of end-to-end protection: Alice wants Bob to pay for the QoS reservation. (reverse charging) Bob wants to be assured that the QoS signaling message he receives are transmitted by Alice because he is only willing to pay for particular users and not for everyone. Hence Bob requires Alice to authenticated the request.

2 Terminology

Some threat scenarios in this document use the entity user instead of the QoS Initiator (as introduced by [1]). This is mainly due to the fact that security protocols allow a differentiation between entities being hosts or users. Since the QoS Initiator as used in [1] also allows to act on behalf of various entities including a network it is reasonable to distinguish between these identities.

We use the term access network for a network to which a mobile node is attached. Other terms often used in this context are foreign or visited network. The missing direct trust relationship between the mobile node and the visited networks is characteristic for such an interface and complicates authentication and key agreement. Usually AAA protocols (like Radius or Diameter) are used for such a purpose. These protocols exploit the infrastructure and trust relationships between the access network and the home network of the user.

The term security association is used to describe established security-relevant data structure between two entities. This data structure consists of keys, algorithms including their parameters, values used for replay protection etc. Using this information two nodes are able to protect QoS signaling messages.

3 Threat Scenarios

This section provides threat scenarios that are applicable to the quality of service signaling protocols.

Additionally, it might also be possible that the QoS initiator acts on behalf of an other user and must therefore interact with this node to be able to trigger the reservation setup. This issue however requires further investigation based on specific protocol proposals.

3.1 Man-in-the-Middle Attacks

This Section describes man-in-the-middle attacks of the following type: During the process of establishing a security association an adversary fools the QI with respect to the entity to which it has to authenticate. The man-in-the-middle adversary is able to modify signaling messages transmitted to the real network requesting different QoS parameters. The QI wrongly believes that it talks to the ôrealö network whereas it is actually attached to an adversary. Note that a solution for protecting QoS signaling messages does not necessarily need to establish a security association. In general it is however advisable to create one because of performance reasons.

For this attack to be successful, pre-conditions have to hold which are described with the two scenarios below:

a) No authentication

The first case considers the case that no authentication between the QI and other entity (access network, other networks, a single node)

takes places: Without authentication the QI is unable to detect an adversary.

b) Unilateral authentication

In case of only unilateral authentication (that is, a missing authentication of the access network to the QI) the QI is not able to discover the man-in-the-middle adversary. In the telecommunication world this type of attack is known as the false base-station attacks (if the unilateral authentication is executed between a user and the access network).

The two threats described above are a general problem of network access without appropriate authentication, not only for QoS. Still these issues need to be correctly addressed in a proposed protocol since the impacts may reach beyond the local network.

3.2 Missing real-time notifications of QoS reservation costs (cost control)

An other type of attack uses the fact that a user is not able to

authorize a particular network service provider (i.e. because of a large number of providers). A large number of service providers with complex roaming agreements create a non-transparent cost-structure. Using AAA protocols in a subscription-based scenario (i.e. user is registered with his home service provider) the user does not learn the identity of the network using a regular message exchange. The user is only authenticated to the home network (and possibly vice versa). The identity of the access network is possibly not revealed. Furthermore one service provider ôstealsö users from an other close-by service provider and because of a missing cost-notification the user is unable to refuse the more expensive service provider although he could route his traffic possible via both providers. The user is not able to select the ôcheapestö access router (in terms of QoS costs).

Although real-time notifications of quality of service reservation costs (cost control) to the user are outside the scope of a quality of service protocol itself there are still interactions with AAA and other protocols.

3.3 Eavesdropping and Traffic Analysis

This Section covers two threats: The first one is related to privacy concerns whereas the second addresses problems caused by weak authentication mechanisms and the increased risk of eavesdropping on the wireless link in absence of appropriate confidentiality protection.

The first threat case covers adversaries that are unable to actively participate in the QoS signaling (passive adversary) but eavesdrop messages. The collected signaling packets may serve for the purpose of traffic analysis or to later mount replay attacks as described in the next Section. By eavesdropping an adversary might violate a

userÆs privacy preference. Especially QoS signaling messages provide information that may be interesting for an adversary since the messages include user and/or application identities, policy information, information about the desired QoS reservation, etc. The information gathered by an adversary can be to learn usage patterns of users requesting resources and track QoS reservations.

The second threat case addresses weak authentication mechanisms whereby information transmitted within the QoS signaling protocol may leak passwords and may allow offline dictionary attacks. This threat is not specific to QoS signaling protocols by may also be applicable and countermeasures must be taken.

3.4 Adversary being able to replay signaling messages

This threat scenario covers the case where an adversary eavesdrops and collects signaling messages and replays them at a latter point in time (or at a different place, or uses parts of them at a different place or in a different way û e.g. cut and paste attacks). The adversary may use this technique in absence of appropriately protected messages to mount denial of service attacks. Furthermore also theft of service is possible.

A more difficult attack that may cause problems even in case of replay protection requires the adversary to crash a QoS aware node (router, broker, etc.) to lose synchronization and to be able to replay old QoS signaling messages.

3.5 Identity Spoofing

An adversary with the capability to spoof the identity may mount the following attacks:

Eve, acting as an adversary, claims to be the registered user Alice by spoofing the identity of Alice. Thereby Eve causes the network to charge Alice for the consumed network resources. Using unprotected messages Eve may experience no particular problems in succeeding.

In case that the signaling request is properly protected the situation becomes more difficult. This threat tries to address possible problems with network based QoS traffic classification based on some identifiers (IP address, ports, other header information etc.). The situation does not change when the data traffic is marked by the transmitting host (i.e. using DSCP).

After the network receives a properly protected reservation request, transmitted by the legitimate user Alice, traffic filters are installed at edge devices. These traffic filters allow data traffic originated from a given address to be assigned to a particular QoS class. The adversary Eve now spoofs the IP address of the Alice (or whatever identifier is used in the flow classification). Additionally AliceÆs host may be crashed by the adversary as a result of a denial of service attack or lost connectivity for a variety of other reasons. In any case Eve is now able to receive and

transmit data (for example RTP data traffic), that receives preferential QoS treatment, using AliceÆs IP address (or whatever identifier is used in the flow classification) until the next signaling message appears and forces Eve to respond with a protected signaling message. Again this issue is not only applicable to QoS traffic but the existence of QoS reservation causes more difficulties since this type of traffic is more expensive.

3.6 Adversary being able to inject/modify messages

The next type of threat is caused by an integrity violation: An adversary modifies signaling messages (e.g. by acting as a man-in-the-middle) to achieve an unexpected network behavior with the bogus request. Possible actions are reordering, delaying, dropping, injecting and modifying.

Using a different identity the adversary may forward a modified a QoS signaling message requesting a large amount of resources (using a different identity). If granted it causes other user's resource-request not to be successful and a different user to pay for the reservation. This attack is only useful in absence of user authentication or if the adversary is able to spoof someoneÆs identity since the attack is useless if the adversary itself is charged for the huge resource reservation.

3.7 Missing Non-Repudiation Property

Repudiation in this context refers to a problem where one party later denies to have made a reservation. This issue comes in two flavors:

From a service provider point-of-view the following threat may be worth an investigation because a user may deny to have issued reservation requests for which he was charged. A service provider may then like to prove that a particular user issued the reservation request.

The same threat can be interpreted from the users point-of-view. A service provider claims to have received a number of reservation requests. The user in question thinks that he never issued those requests and wants to have a proof for correct service usage for a given set of QoS parameters.

3.8 Malicious Edge-Router

Network elements within a domain (intra-domain) experience a different trust relationship with regard to the security protection of signaling messages compared to edge routers. Assuming that edge routers have the responsibility to perform cryptographic processing (authentication, integrity and replay protection, authorization and accounting). If however an adversary manages to take over an edge router then the security of the entire network is affected. An adversary can then launch a number of attacks including denial of service, integrity violation, replay attacks etc. Note that this

problem is not only restricted to the QoS protocols. In such a case even the chain-of-trust principle does not prevent the network from being vulnerable: If we assume that the adversary, with access to the edge router, is able to access the keys used to secure messages to other nodes.

Thus the edge router is a critical component that requires strong security protection. This does not necessarily imply that all routers within the core network do not need to cryptographically verify signaling messages and that these routers cannot have any security effect if they act maliciously. If the (hop-by-hop) chain-of-trust principle is deployed then the security of the path (in this case within the network of a single administrative domain) is as strong as the weakest link. In our case the edge router is the most critical component of this network that may also act as a security gateway/firewall for incoming/outgoing traffic. For outgoing traffic this device has to act according to the security policy of the local domain to apply the appropriate security protection.

3.9 Denial of Service in a two phase reservation

This threat tries to address potential denial of service attacks when the reservation setup is split into two phases i.e. path and reservation. For this example we assume that the node transmitting the path message is not charged for this message and is able to issue a high number of reservation request (possibly in a distributed fashion). The reservations are however never intended to be successful because of various reasons: for example the destination node cannot be reached or is not responding node or rejects the reservation. An adversary can benefit from the fact that resources are already consumed along the path for various processing tasks including path pinning.

3.10 Denial of Service with a bogus reservation request

With a resource reservation request received at a network element (for example by the first QoS aware router) processing is required for authentication and authorization (processing by other nodes including policy server, LDAP server, etc. is also possible depending on the network architecture). The verification of the provided credentials requires computations and resources to be allocated memory for state maintenance, setting timers, additional messages transmitted to other nodes, cryptographic computations). If an adversary is able to transmit a large number of reservation request (flooding) with bogus credentials and assuming that the verification is expensive in terms of resource consumption then the verifying node may not be able to process further reservation messages by legitimate user.

3.11 Disclosing the networking structure

In some architectures a network provider does not want to reveal its internal network structure to the outside world. An adversary might

be able to use NSIS messages for network mapping (e.g. discovering which nodes existed, which used NSIS, what version etc.). This requirement might conflict with a protocol solution that provides a mean to automatically discover NSIS aware nodes and their identity (the identity required for security protection).

3.12 Modification of subsequent reservation request

An adversary might be able to modify an existing reservation which had already been established within the network as a result of a previous QoS signaling message. This means that a QoS signaling messages that modifies established state must be subject to security protection comparable to the original signaling message setting up the reservation. Furthermore it might be necessary to provide (possibly cryptographic) information to assure a correct binding to a specific state/session.

3.13 Faked Error/Response messages

An adversary may be able to use false error/response messages as part of a denial of service attack. This could be either at the reservation level or at the protocol level.

4 Security Considerations

This entire memo discusses security issues.

5 References

[1] Brunner, M., "Requirements for QoS Signaling Protocols", draft-ietf-nsis-req-02.txt, Work In Progress, May 2002.

6 Acknowledgments

I would like to thank (in alphabetical order) Marcus Brunner, Jorge Cuellar, Xiaoming Fu and Robert Hancock for their comments to this draft. Jorge and Robert gave me an extensive list of comments for this memo and provided more information on additional threats that should be added.

7 Author's Addresses

Hannes Tschofenig
Siemens AG

Otto-Hahn-Ring 6
81739 Munchen
Germany
Email: Hannes.Tschofenig@mchp.siemens.de