NSIS Working Group Internet-Draft Expires: January 17, 2005 C. Aoun Nortel Networks H. Tschofenig Siemens M. Stiemerling NEC July 19, 2004

NATFW NSLP Migration Considerations draft-aoun-nsis-nslp-natfw-migration-02

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with RFC 3668.

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/lid-abstracts.txt.

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

This Internet-Draft will expire on January 17, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This document discusses migration issues towards NSIS NAT/FW NSLP enabled NATs and Firewalls. In particular traversal of NSIS unaware NATs and NSIS proxy scenarios are addressed. This document will serve as input to the NSIS NATFW NSLP document.

Table of Contents

$\underline{1}$. Introduction			· ·	<u>3</u>
<u>2</u> . Terminology				<u>4</u>
3. Traversal of NSIS unaware NATs				<u>5</u>
<u>3.1</u> Abstract				<u>5</u>
<u>3.2</u> Introduction				<u>5</u>
<u>3.3</u> Class 1 NAT Handling				<u>6</u>
3.4 Class 2 NAT Handling				<u>6</u>
<u>3.5</u> Class 3 NAT Handling				7
3.6 Class 4 NAT Handling				9
3.7 Dealing with NSIS unaware NATs (Class 4 NAT	Handli	ng)		10
		0,		
<u>4</u> . NSIS Proxy Mode				<u>14</u>
<u>5</u> . NSIS unaware Firewall Traversal				<u>17</u>
<u>6</u> . NATFW NSLP NTLP requirements				<u>18</u>
<u>7</u> . Conclusion				19
<u>8</u> . Security Considerations				<u>20</u>
<u>9</u> . Contributors				<u>21</u>
<u>10</u> . Acknowledgements		• •	• •	<u>22</u>
11. References				23
11.1 Normative References				23
11.2 Informative References				23
			-	
Authors' Addresses				<u>24</u>
Intellectual Property and Copyright Statements				<u>25</u>

1. Introduction

This document discusses migration issues which are caused by the incremental deployment of NSIS NAT/Firewall NSLP nodes. As such, it is not only relevant for the NAT/FW NSLP but also for other NSLPs, such as the QoS NSLP.

- o The overall NSIS protocol suite (including the NAT/FW NSLP) is impacted by NSIS unaware NATs and Firewalls. This document covers the impacts as well as some suggestions to ease the deployments of the NSIS protocol suite until the installed base of NATs and Firewalls migrates to NSIS. <u>Section 3</u> addresses this issue.
- o The NAT/FW NSLP should allow an end host supporting NSIS to operate properly without the need for supporting true end-to-end NSIS signaling to its application correspondent. This is very practical during the initial phases of the NSIS migration and is applicable in simple network topologies. In the early phases of the NSIS NAT/FW NSLP migration, this situation will occur quite frequently, and hence this scenario must be supported. <u>Section 4</u> is addresses this scenario.

2. Terminology

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 [RFC2119].

The terminology used in this document is defined in [NSISNATFW].

3. Traversal of NSIS unaware NATs

3.1 Abstract

This section aims to describe different NAT handling classes in NSIS, and the relationship between different solutions provided in various NSIS documents. After a short introduction in <u>Section 3.2</u>, different NAT classes are discussed in <u>Section 3.2</u>. Various NAT scenarios are described in <u>Section 3.3</u>, in <u>Section 3.4</u>, in <u>Section 3.5</u>, and in <u>Section 3.6</u>. <u>Section 3.7</u> covers the details of addressing NSIS unaware NATs.

3.2 Introduction

In the past, mainly two approaches have been used for establishing NAT bindings:

- Static configuration This type of NAT bindings is typically used to allow inbound initiated communications. Ephemeral binds are often not established or required.
- Dynamic configuration: Dynamic NAT binding creation can be categorized into one of the following three categories:
 - * Implicit creation by outbound initiated communications. In this case, the translated address and port are selected from a configured address and port pool.
 - * Explicit creation by an Application Layer Gateway (ALG) either via an API call (if the NAT and the ALG are co-located) or otherwise via a separate protocol.
 - * Separate signaling protocols which request the creation of a NAT binding.

An alternative classification can be done by considering the trigger for the creation of a NAT binding. In many cases an outbound data packet itself is used to cause the allocation of a NAT binding. Alternatively, a signaling protocol can be used to establish the same goal by directly addressing the NAT itself. The Midcom and the NSIS working group are trying to develop protocols of the latter category.

To address the broad scope of NAT handling, this document tries to describe the design considerations for the work in the NSIS WG. An important impact for the design is caused by the introduction of the two layer architecture, by intermediaries, and by NSIS unaware NATS.

Four classes of NAT functionality can be distinguished in NSIS.

These are described in the following sub-sections.

<u>3.3</u> Class 1 NAT Handling

We refer to Class 1 NAT handling if NATs or Firewalls implement the NAT/Firewall NSLP.



Figure 1: Class 1 NAT Handling

The NSIS working group decided to work on two NSLP client layer applications: QoS and NAT/Firewall NSLP. The NAT/Firewall NSLP assumes that a NAT or a Firewall implements the signaling application (i.e., NTLP and NAT/Firewall NSLP implementation). [<u>NSISNATFW</u>] describes the signaling mechanisms in more detail.

<u>3.4</u> Class 2 NAT Handling

In Figure 2, a number of QoS NSLP nodes are shown, with one of the NSIS nodes being a NAT device. In this scenario, we assume that the NAT device does not contain a NAT/Firewall NSLP implementation. Incremental deployment can lead to such a configuration. A question raised by this scenario is whether the NSIS implementation (the NTLP for example) should offer a minimal NAT implementation which would allow it to request a NAT binding to update the flow identifier.



Figure 2: Class 2 NAT Handling

There is some desire to allow an NSLP signaling message exchange (e.g., QoS signaling) to work properly even in the presence of NAT. In this scenario, no NAT/Firewall NSLP implementation is available at the NAT, and the end host might not trigger an NAT/Firewall NSLP exchange. For some cases, the signaling message contains sufficient information to create a NAT binding based on the flow identifier in the NTLP layer. If additional security mechanisms have to be provided by the NAT/Firewall NSLP, then that approach will fail since, for example, a QoS NSLP will not be able to provide those mechanisms.

Another question of interest is whether it is possible to combine a NAT/Firewall signaling message with a QoS signaling message into a single protocol message (or to at least combine them using a shared session identifier). Error handling might be more complex because in addition to dealing with errors of the individual signaling applications, it will be necessary to deal with errors resulting from the combined applications.

3.5 Class 3 NAT Handling

We refer to Class 3 NAT handling if there is a NAT along the path which intercepts all NSIS signaling messages, but which does not contain the desired NSLP implementation. In Figure 3a, Host A signals for a QoS NSLP, but NAT 2 only offers an NTLP implementation. This NTLP could modify a flow identifier, if it is not integrity protected or encrypted. NAT 1 was already considered in <u>Section 3.4</u>.



Figure 3: Class 3a NAT Handling

Intermediaries make NSIS signaling message handling more complex. To avoid these problems it is possible to build functionality into the NTLP to make intermediate nodes to be invisible for NSIS signaling.

As a solution, it is suggested to make the discovery message (C-Mode) or the D-Mode in general cleverer to "discover" only those nodes which implement the desired NSLP functionality. (This was already discussed in the past.)

This requires indicating which NSLP functionality the signaling message is looking for along the path. A discovery message might, for example, want to know the next NSIS node along the path which supports a QoS NSLP implementation. It is for further study, whether a more fine-granular discovery is required (e.g., a QoS NSLP node which supports a certain QoS model). IANA registration would be required for the NSLPs as well as for QoS models.

Efficiency is an important issue here. An NSIS node which does not implement a certain NSLP application must be able to quickly distinguish whether it is interested in a message or not. Whether to encode the necessary information into a router alert option, into UDP port numbers, the Time-to-Live field, or into an IP protocol number was discussed in the past.

As a minor variation of the scenario described in Figure 3, it is possible that the end host does not contain a NAT/Firewall implementation, but the network itself provides a NAT/Firewall solution as shown in Figure 4.



Figure 4: Class 3b NAT Handling

The main advantage of the approach described in Figure 4 is the incremental deployment meaning that an administrative domain equips its NATs/Firewalls with NAT/Firewall NSLP implementations even though the end host might not support it. Note that the NAT/FW device might not offer the QoS NSLP implementation. The NAT/Firewall NSLP enabled Edge Router 1 might create a NAT binding or open a firewall pinhole based on an incoming QoS signaling message (even though the end host is not NAT/Firewall NSLP aware). It is also able to create NAT/ Bindings at the NAT/FW device independently of a signaling exchange. Such a signaling exchange might be necessary if the NAT/Firewall is not equipped with the NSLP application signaled by the end host - in our example this would mean that the NAT/FW would not run a QoS NSLP implementation.

<u>3.6</u> Class 4 NAT Handling

We refer to a scenario as Class 4 NAT handling scenario if a NAT is within the path which does not understand NSIS at all.



Figure 5: Class 4 NAT Handling

To allow NSIS signaling messages to traverse an NSIS unaware NAT, it is required that they are sent in non-raw IP mode. This is necessary to allow NAPT to perform modification of the transport protocol port numbers. For an IPsec protected signaling message, UDP encapsulation MUST be used. If IKE or IKEv2 [I-D.ietf-ipsec-ikev2] is used, then NAT traversal functionality is necessary to dynamically detect the presence of a NAT. The relevant work in this area can be found in [I-D.ietf-ipsec-nat-t-ike], in [I-D.ietf-ipsec-ikev2] and in [IPSECNAT].

3.7 Dealing with NSIS unaware NATs (Class 4 NAT Handling)

This section describes NSIS signaling in a Class 4 NAT handling scenario. It is in general not possible to reuse the NAT binding created with the NSIS signaling also for the data traffic. An exception is a NAT which maps the source IP address of all outgoing IP packets to the same external public IP address (without modifying the port number). In order to update the flow identifier, the NSIS NSLP daemon has to interact with a NAT in a non-NSIS fashion (such as STUN [STUN] or a MIDCOM protocol), or to reuse an "NSIS-STUN-alike" mechanism. We will describe the "NSIS-STUN-alike" mechanism in this section.

In many cases it might be sufficient to detect the presence of an NSIS unaware NAT. This might be useful for those cases where the NSLP would break in such cases. The NSIS unaware NAT discovery functionality could be a built-in feature of the NTLP [NTLP], allowing its usage on any NE regardless of the supported NSLPs. In order to discover a NAT the following procedure can be applied to D-mode messages (which includes also discovery messages).

The initiator of the discovery message includes the source IP address

Internet-Draft

and the source port of the transmitted message into the signaling message payload.

An NSIS unaware NAT then modifies the source IP address (and possibly the source port) of the NSIS signaling message. This procedure represents typical NAT handling. The responder of the discovery message will notice the modification by comparing information in the IP header with the content of the discovery message. If both are equal then no NAT was present. If the responder sees a deviation, then an NSIS unaware NAT was located along the path. The responder returns the source IP address and port number as a payload in the discovery reply message. Unfortunately, the information found does not help to update the flow identifier for the data traffic.

Traversing an NSIS unaware NAT (from inside to outside) dynamically creates a NAT binding. Please note that only a NAT binding for the signaling traffic is created. More complexity is introduced by creating NAT bindings for data traffic. It seems to be reasonable that neighboring NSIS nodes control the NAT and also the Firewall (of the same administrative domain) via MIDCOM protocols, such as SNMPv3. The usage of SNMPv3 for this purpose is simple but requires the NSIS unaware NAT to implement this protocol.

Another option is to include an "NSIS-STUN-alike" mechanism into NSIS which has the property of an in-band signaling mechanism. Ideally NSIS signaling messages should look like regular data traffic to experience the same treatment as data traffic. The discovery mechanisms is thereby an important part which we will investigate in more detail. Discovery messages are addressed to the same IP address as the data traffic. The source IP address can, however, be the source IP address of the data traffic, or the source IP address of the signaling message, in which case it would be equal to the IP address of the NSIS node transmitting it. The source port can be either equal to the source port number of the transmitted data traffic, or to the source port of the transmitting NSIS daemon. Setting the source port to the port number of the application traffic makes it very difficult for the NSIS daemon to intercept the response to a discovery message. Further investigations are required to verify its practicability. But this step would be very important since responses need to be addressed to the port number which was modified by the NAT (see symmetric NATs). It is suggested to set the destination port number to the port number of the data traffic destination. The Router Alert Option will allow an NSIS node to intercept the message and to distinguish it from a regular data packet. But note that this is only true for D-mode messages. For C-Mode messages, an additional problem is created if the transport layer protocol of the data traffic does not match the transport protocol of the signaling traffic. Furthermore, it seems to be very

NATFW NSLP Migration

difficult for an end host to distinguish the data traffic from the signaling traffic. If we can assume that the discovery message exchange is (for most parts) indistinguishable from data traffic, then this exchange can be used by NSIS signaling messages and data traffic to traverse an NSIS unaware NAT. This, however, additionally assumes that only flows are signaled, rather than aggregates.

If no means of controling the NAT are available, then the STUN protocol [<u>STUN</u>] can be used. The usage of STUN and other protocols (such as TURN [<u>I-D.rosenberg-midcom-turn</u>]) should be investigated in future versions of this document.

In Figure 6, we consider a scenario where an NSIS aware initiator also hosts a STUN implementation. Note that more complex topologies are possible but not investigated in detail in this section.

+	+	
		++
++	I	STUN
Apps	I	Server
++	++	++
STUN	NAT	
CLIENT		
	++	
ANY_NSLP	I	
	I	
++	I	
Host A	I	
+	+	

Figure 6: STUN usage for NSIS unaware NATs

Within Host A, shown in Figure 6, the NSIS API could invoke the services of the STUN client upon determination that an NSIS unaware NAT is on the path. NSLPs, such as a QoS NSLP, would use the STUN returned global scoped address for the flow identifier of the NSIS signaling message. If some NSLPs are between Host A and the NSIS unaware NAT, then a wrong flow identifier would be communicated to these devices. This might be problematic for a QoS NSLP and would not really provide a solution. Without learning a globally routable IP address via STUN, the correct flow identifier (i.e. the private IP address) would be installed between Host A and the NSIS unaware NAT, but a wrong flow identifier between the NAT and the destination host, since the private IP address used as flow identifier is not converted to a public IP address.

The consequences might be different if STUN is used by an entity

along the path and not the end host. Figure 4 shows such an example. This impact needs to be studied in more detail in a future version of this document.

4. NSIS Proxy Mode

When NSIS NAT/FW signaling will start to be deployed, it is quite possible that an NI sends an NSIS message without having an NR to respond to it. The NATFW NSLP should be able to handle this type of deployments. NSIS NATFW NSLP signaling for a data receiver behind a NAT already has just a local scope (the REA message is not forwarded beyond the edge NAT, see [NSISNATFW]). This mechanism only works for data receivers behind a NAT, but not for data receivers behind a firewall.

Since the purpose of this section is to discuss how end to end signaled messages are handled when no NRs are available on the end-host, only Firewalls (the NFs) are discussed within the example networks.

The local trust domain (from an NI perspective) has at least one NSIS aware Firewall, there is no NR on the far end, nor an NSIS aware NAT or Firewall. Goal of this exchange is to keep NSIS signaling local within network A. The solution of this approach is similar to [lrsvp], but the NSIS messages do not include any scoping information.

Figure 7 shows this scenario graphically.





To terminate the NSIS signaling exchange within Network A, two approaches are feasible: explicit and implicit scoping. With explicit scoping, Host A has to indicate that the NSIS signaling message should terminate locally. With implicit scoping, the NI simply sends its firewall policy rule creation message. The message traverses the first NF (its own firewall), but there is no NR to respond back. As a consequence a timer will expire since no response message is received. The last NF will respond back to the NI with a notification that NSIS signaling had to terminate somewhere along the path without reaching the NR. Using the network deployment shown in Figure 7, the message exchange in Figure 8 takes place.

Host A		Host B	
NI F	-W++	Expected	NR
1-NSIS Init msg			
>			
	2-NSIS Init msg		
I	+>		
	NATFW NSLP ON		
	Timeout		
3-NSIS Init msg Ack	(V		
No NR			
<			

Figure 8: Detecting the last NSIS peer

Figure 9 provides the message sequences when more than one NSIS aware NAT or Firewall is deployed within the same trust domain. Upon determination of a previous NSIS hop, an NSIS aware node will notify the previous NSIS hop of its existence to avoid launching the timer that triggers sending of an NSIS message back to the NI. The current NTLP message association establishment procedures supports this behavior. The last NF on the path will launch the timer since no valid downstream NSIS neighbor responded back.

Aoun, et al. Expires January 17, 2005 [Page 15]

Internet-Draft

Trust domain A Trust domain B <.....> <-----> Host B Host A NI FW++ FW++ Expected NR | NSIS Init msg | | -----> | NSIS Init msg | | -----> | NSIS Init msg | NATFW NSLP ON |----->| | | with Token | | Valid . | NATFW NSLP ON | | NSIS Neighbor | | |<----| | |---->| | Timeout Ack | V | <....+ | NSIS Init msg Ack| 1 | NSIS Init msg Ack | No NR NO NR | <....|

Figure 9: Detecting the last NSIS peer (multiple FWs)

Aoun, et al. Expires January 17, 2005 [Page 16]

5. NSIS unaware Firewall Traversal

In case an NSIS unaware firewall is traversed by NSIS messages, NSIS messages should be allowed to go through it, as well as the exchanged data flows between the user applications. This is not necessarily an obvious task to perform in case the NSIS messages cannot be identified by the NSIS unaware firewall. The same applies to the user application data flows.

NSIS message identification should be supported by existing firewalls.

Currently firewalls support flow identification by using the 5 tuple or a sub-set of it. The authors are still expecting feedback from firewall vendors to see if we can assume that existing firewalls will not drop packets including the Router Alert Option (RAO) [<u>RFC2113</u>]. In case existing firewalls drop packets with the router alert option set, then the RAO should not be the only element used to identify packets to be dropped.

User application data flow identification should be deterministic at a specific address and port range level. This means that the application uses a combination of an address and specific transport port range.This combination should be configured on the firewall.

In case a NAT is deployed on the path and it is NSIS-NATFW, the assigned bind should be consistent with policy rules configured in the NSIS unaware firewall.

Even though the deployed Firewall is not NSIS aware, the application data would still be forwarded if existing interim solutions were used, such as a mix of stateless policy rules and flow based states with initial packets sent in the outbound direction (from inside a trust domain to outside the trust domain).

Aoun, et al. Expires January 17, 2005 [Page 17]

6. NATFW NSLP NTLP requirements

In this section we list two requirements for the NTLP raised by this document.

- o When NSIS signaling is used in the presence of NSIS unware NATs, then raw IP MUST NOT be used. Network address and port translation requires transport layer identifiers as means to direct inbound traffic to the right recipient.
- o For the traveral of NSIS unaware NATs, UDP is more likely to be supported than DCCP or SCTP.
- o If IPsec is used to secure NSIS signaling messages, then UDP encapsulation for IPsec protected packets (see [<u>IPSECNAT</u>]) MUST be used to ensure that IPsec does not break. IKE with extensions or IKEv2 is able to detect the presence of a NAT along the path.

Aoun, et al. Expires January 17, 2005 [Page 18]

7. Conclusion

To handle NAT devices properly it is necessary to address the different NAT handling scenarios individually:

The impact of intermediaries causes complexity for signaling message handling. It is therefore recommended to avoid Class 2 and Class 3 NAT handling scenarios by incorporating additional knowledge into the discovery message.

Class 4 NAT handling requires some interaction with other protocols such as MIDCOM or STUN. The ability to reuse the NTLP discovery mechanisms to create NAT bindings for the signaling and the data traffic is briefly outlined but requires more investigations.

To deal with firewalls it is also necessary to

- o allow the NSIS signaling message to pass, and
- o also to create pinholes for subsequent data traffic.

This is mainly an authorization problem and requires depends on the environment where NSIS is used.

It is important to keep in mind to differentiate NAT bindings for the signaling traffic and those for the data traffic. This separation is necessary since the NSIS signaling message and the subsequent data traffic are different in terms of the flow identifier observed by the NAT. The same is true for firewall pinholes.

Aoun, et al. Expires January 17, 2005 [Page 19]

8. Security Considerations

This document discusses various security issues for NAT/Firewall signaling in migration scenarios.

Two important security threats are worth being highlighted:

- o The proxy mode of operation, described in <u>Section 4</u>, demands the property that NSIS signaling messages terminate somewhere along the path. This functionality should allow NSIS to capture additional scenarios not envisioned by RSVP. As a consequence it makes it very hard to allow end-to-end security mechanisms to be applied. These end-to-end security mechanisms have been proposed to enable delayed authorization by both end hosts, and to tie the NSIS end-to-end signaling together with application layer signaling. The same is true if NSIS signaling is triggered by a node other than the end host.
- o Providing mechanisms to traverse NSIS unaware NATs also has security implications. The impact can be related to an NSIS signaling message, or even to the data traffic (based on signaling of the flow identifier). <u>Section 3</u> provides the details of traversal of NSIS unaware NATs. An adversary along the path (a non-NSIS node) is able to redirect NSIS signaling to another NSIS node to cause denial of service attacks. If an adversary is additionally able to modify the flow identifier, then it is possible to cause NSIS to create an arbitrary policy rule which would allow the adversary to inject traffic from an arbitrary location. Note that an adversary along the path is always able to cause denial of service attacks by dropping or delaying signaling messages. Furthermore, it is also able to inject data packets, but only with a flow identifier chosen by the signaling initiator, and possibly modified by an NSIS aware NAT along the path.

Further security considerations can be found in [<u>NSISNATFW</u>] and [<u>I-D.fessi-nsis-natfw-threats</u>].

Aoun, et al. Expires January 17, 2005 [Page 20]

9. Contributors

We would like to thank Marcus Brunner and Miquel Martin for their contribution to this draft.

<u>10</u>. Acknowledgements

We would like to thank Joachim Kross for this comments to this draft.

Internet-Draft

<u>11</u>. References

<u>11.1</u> Normative References

[NSISNATFW]

Stiemerling, M., Martin, M., Tschofenig, H. and C. Aoun, "A NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", DRAFT <u>draft-ietf-nsis-nslp-natfw-03.txt</u>, July 2004.

[NTLP] Schulzrinne, H. and R. Hancock, "GIMPS: General Internet Messaging Protocol for Signaling", <u>draft-draft-ietf-nsis-ntlp-00</u> (work in progress), May 2004.

<u>11.2</u> Informative References

[I-D.fessi-nsis-natfw-threats]

Martin, A., Stiemerling, M., Thiruvengadam, S., Tschofenig, H. and C. Aoun, "Security Threats for the NATFW NSLP", DRAFT <u>draft-fessi-nsis-natfw-threats-01.txt</u>, July 2004.

[I-D.ietf-ipsec-ikev2]

Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", draft-ietf-ipsec-ikev2-14 (work in progress), May 2004, <reference.I-D.ietf-ipsec-ikev2.xml>.

[I-D.ietf-ipsec-nat-t-ike]

Kivinen, T., "Negotiation of NAT-Traversal in the IKE", draft-ietf-ipsec-nat-t-ike-08 (work in progress), February 2004, <reference.I-D.ietf-ipsec-nat-t-ike.xml>.

[I-D.rosenberg-midcom-turn]

Rosenberg, J., "Traversal Using Relay NAT (TURN)", <u>draft-rosenberg-midcom-turn-01</u> (work in progress), March 2003.

[IPSECNAT]

A. Huttunen et all, A., "UDP Encapsulation of IPsec Packets", DRAFT <u>draft-ietf-ipsec-udp-encaps-07.txt</u>, Jan 2003.

[RFC2113] Katz, D., "IP Router Alert Option", <u>RFC 2113</u>, February 1997.

- [RFC3304] Swale, R., Mart, P., Sijben, P., Brim, S. and M. Shore, "Middlebox Communications (midcom) Protocol Requirements", <u>RFC 3304</u>, August 2002.
- [STUN] Rosenberg, J., Weinberger, J., Huitema, C. and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", <u>RFC 3489</u>, March 2003.
- [lrsvp] Manner, J., "Localized RSVP", draft-manner-lrsvp-03 (work in progress), January 2004.

Authors' Addresses

Cedric Aoun Nortel Networks

France

EMail: cedric.aoun@nortelnetworks.com

Hannes Tschofenig Siemens AG Otto-Hahn-Ring 6 Munich 81739 Germany

Phone: EMail: Hannes.Tschofenig@siemens.com URI:

Martin Stiemerling Network Laboratories, NEC Europe Ltd. Kurfuersten-Anlage 36 Heidelberg 69115 Germany

Phone: +49 (0) 6221 905 11 13 EMail: stiemerling@ccrle.nec.de URI:

Aoun, et al. Expires January 17, 2005 [Page 24]

Internet-Draft

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in <u>BCP 78</u> and <u>BCP 79</u>.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in <u>BCP 78</u>, and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.