

Internet Engineering Task Force (IETF)  
Request for Comments: 7362  
Category: Informational  
ISSN: 2070-1721

E. Iovov  
Jitsi  
H. Kaplan  
Oracle  
D. Wing  
Cisco  
September 2014

Latching: Hosted NAT Traversal (HNT)  
for Media in Real-Time Communication

Abstract

This document describes the behavior of signaling intermediaries in Real-Time Communication (RTC) deployments, sometimes referred to as Session Border Controllers (SBCs), when performing Hosted NAT Traversal (HNT). HNT is a set of mechanisms, such as media relaying and latching, that such intermediaries use to enable other RTC devices behind NATs to communicate with each other.

This document is non-normative and is only written to explain HNT in order to provide a reference to the Internet community and an informative description to manufacturers and users.

Latching, which is one of the HNT components, has a number of security issues covered here. Because of those, and unless all security considerations explained here are taken into account and solved, the IETF advises against use of the latching mechanism over the Internet and recommends other solutions, such as the Interactive Connectivity Establishment (ICE) protocol.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7362>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction . . . . . 2
- 2. Background . . . . . 4
- 3. Impact on Signaling . . . . . 5
- 4. Media Behavior and Latching . . . . . 6
- 5. Security Considerations . . . . . 11
- 6. Acknowledgements . . . . . 13
- 7. References . . . . . 14
  - 7.1. Normative References . . . . . 14
  - 7.2. Informative References . . . . . 14

1. Introduction

Network Address Translators (NATs) are widely used in the Internet by consumers and organizations. Although specific NAT behaviors vary, this document uses the term "NAT" for devices that map any IPv4 or IPv6 address and transport port number to another IPv4 or IPv6 address and transport port number. This includes consumer NATs, firewall/NATs, IPv4-IPv6 NATs, Carrier-Grade NATs (CGNs) [RFC6888], etc.

The Session Initiation Protocol (SIP) [RFC3261], and others that try to use a more direct path for media than with signaling, are difficult to use across NATs. These protocols use IP addresses and transport port numbers encoded in bodies such as the Session Description Protocol (SDP) [RFC4566] and, in the case of SIP, various header fields. Such addresses and ports are unusable unless all peers in a session are located behind the same NAT.

Mechanisms such as Session Traversal Utilities for NAT (STUN) [RFC5389], Traversal Using Relays around NAT (TURN) [RFC5766], and Interactive Connectivity Establishment (ICE) [RFC5245] did not exist

when protocols like SIP began being deployed. Some mechanisms, such as the early versions of STUN [RFC3489], had started appearing, but they were unreliable and suffered a number of issues typical for UNilateral Self-Address Fixing (UNSAF), as described in [RFC3424]. For these and other reasons, Session Border Controllers (SBCs) that were already being used by SIP domains for other SIP and media-related purposes began to use proprietary mechanisms to enable SIP devices behind NATs to communicate across the NATs. These mechanisms are often transparent to endpoints and rely on a dynamic address and port discovery technique called "latching".

The term often used for this behavior is "Hosted NAT Traversal (HNT)"; a number of manufacturers sometimes use other names such as "Far-end NAT Traversal" or "NAT assist" instead. The systems that perform HNT are frequently SBCs as described in [RFC5853], although other systems such as media gateways and "media proxies" sometimes perform the same role. For the purposes of this document, all such systems are referred to as SBCs and the NAT traversal behavior is called HNT.

At the time of this document's publication, a vast majority of SIP domains use HNT to enable SIP devices to communicate across NATs despite the publication of ICE. There are many reasons for this, but those reasons are not relevant to this document's purpose and will not be discussed. It is, however, worth pointing out that the current deployment levels of HNT and NATs make the complete extinction of this practice highly unlikely in the foreseeable future.

The purpose of this document is to describe the mechanisms often used for HNT at the SDP and media layer in order to aid understanding the implications and limitations imposed by it. Although the mechanisms used in HNT are well known in the community, publication in an IETF document is useful as a means of providing common terminology and a reference for related documents.

This document does not attempt to make a case for HNT or present it as a solution that is somehow better than alternatives such as ICE. Due to the security issues presented in Section 5, the latching mechanism is considered inappropriate for general use on the Internet unless all security considerations are taken into account and solved. The IETF is instead advising for the use of the Interactive Connectivity Establishment (ICE) [RFC5245] and Traversal Using Relays around NAT (TURN) [RFC5766] protocols.

It is also worth mentioning that there are purely signaling-layer components of HNT as well. One such component is briefly described for SIP in [RFC5853], but that is not the focus of this document.

SIP uses numerous expressive primitives for message routing. As a result, the HNT component for SIP is typically more implementation-specific and deployment-specific than the SDP and media components. For the purposes of this document it is hence assumed that signaling intermediaries handle traffic in a way that allows protocols such as SIP to function correctly across the NATs.

The rest of this document focuses primarily on the use of HNT for SIP. However, the mechanisms described here are relatively generic and are often used with other protocols such as the Extensible Messaging and Presence Protocol (XMPP) [RFC6120], Media Gateway Control Protocol (MGCP) [RFC3435], Megaco/H.248 [RFC5125], and H.323 [H.323].

## 2. Background

The general problems with NAT traversal for protocols such as SIP are:

1. The addresses and port numbers encoded in SDP bodies (or their equivalents) by NATed User Agents (UAs) are not usable across the Internet because they represent the private network addressing information of the UA rather than the addresses/ports that will be mapped to/from by the NAT.
2. The policies inherent in NATs, and explicit in firewalls, are such that packets from outside the NAT cannot reach the UA until the UA sends packets out first.
3. Some NATs apply endpoint-dependent filtering on incoming packets, as described in [RFC4787]; thus, a UA may only be able to receive packets from the same remote peer IP:port as it sends packets out to.

In order to overcome these issues, signaling intermediaries such as SIP SBCs on the public side of the NATs perform HNT for both signaling and media. An example deployment model of HNT and SBCs is shown in Figure 1.

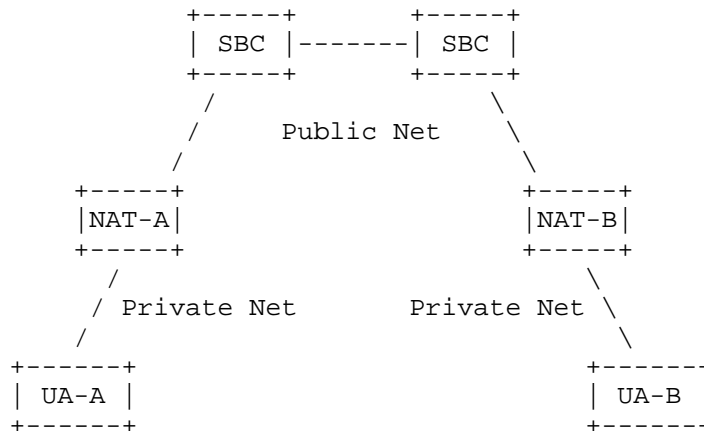


Figure 1: Signaling and Media Flows in a Common Deployment Scenario

3. Impact on Signaling

Along with codec and other media-layer information, session establishment signaling also conveys potentially private and non-globally routable addressing information. Signaling intermediaries would hence modify such information so that peer UAs are given the (public) addressing information of a media relay controlled by the intermediary.

In typical deployments, the media relay and signaling intermediary (i.e., the SBC) are co-located, thereby sharing the same IP address. Also, the address of the media relay would typically belong to the same IP address family as the one used for signaling (as it is known to work for that UA). In other words, signaling and media would both travel over either IPv4 or IPv6.

The port numbers introduced in the signaling by the intermediary are typically allocated dynamically. Allocation strategies are entirely implementation dependent and they often vary from one product to the next.

The offer/answer media negotiation model [RFC3264] is such that once an offer is sent, the generator of the offer needs to be prepared to receive media on the advertised address/ports. In practice, such media may or may not be received depending on the implementations participating in a given session, local policies, and the call scenario. For example, if a SIP SDP offer originally came from a UA behind a NAT, the SIP SBC cannot send media to it until an SDP answer is given to the UA and latching (Section 4) occurs. Another example is, when a SIP SBC sends an SDP offer in a SIP INVITE to a

residential customer's UA and receives back SDP in a 18x response, the SBC may decide, for policy reasons, not to send media to that customer UA until a SIP 200 response has been received (e.g., to prevent toll fraud).

#### 4. Media Behavior and Latching

An UA that is behind a NAT would stream media from an address and a port number (an address:port tuple) that are only valid in its local network. Once packets cross the NAT, that address:port tuple will be mapped to a public one. The UA, however, is not typically aware of the public mapping and would often advertise the private address:port tuple in signaling. This way, while a session is still being set up, the signaling intermediary is not yet aware what addresses and ports the caller and the callee would end up using for media traffic: it has only seen them advertise the private addresses they use behind their respective NATs. Therefore, media relays used in HNT would often use a mechanism called "latching".

Historically, "latching" only referred to the process by which SBCs "latch" onto UDP packets from a given UA for security purposes, and "symmetric-latching" is when the latched address:port tuples are used to send media back to the UA. Today, most people talk about them both as "latching"; thus, this document does as well.

The latching mechanism works as follows:

1. After receiving an offer from Alice (User Agent Client (UAC) located behind a NAT), a signaling intermediary located on the public Internet would allocate a set of IP address:port tuples on a media relay. The set would then be advertised to Bob (User Agent Server (UAS)) so that he would use those media relay address:port tuples for all media he wished to send toward Alice (UAC).
2. Next, after receiving from Bob (UAS) an answer to its offer, the signaling server would allocate a second address:port set on the media relay. In its answer to Alice (UAC), the SBC will replace Bob's address:port with this second set. This way, Alice will send media to this media relay address:port.
3. The media relay receives the media packets on the allocated ports and uses their respective source address:ports as a destination for all media bound in the opposite direction. In other words, it "latches" or locks on these source address:port tuples.

4. This way, when Alice (UAC) streams media toward the media relay, it would be received on the second address:port tuple. The source address:port of her traffic would belong to the public interface of Alice's NAT, and anything that the relay sends back to that address:port would find its way to Alice.
5. Similarly, the source of the media packets that Bob (UAS) is sending would be latched upon and used for media going in that direction.
6. Latching is usually done only once per peer and not allowed to change or cause a re-latching until a new offer and answer get exchanged (e.g., in a subsequent call or after a SIP peer has gone on and off hold). The reasons for such restrictions are mostly related to security: once a session has started, a user agent is not expected to suddenly start streaming from a different port without sending a new offer first. A change may indicate an attempt to hijack the session. In some cases, however, a port change may be caused by a re-mapping in a NAT device standing between the SBC and the UA. More advanced SBCs may therefore allow some level of flexibility on the re-latching restrictions while carefully considering the potential security implications of doing so.

Figure 2 describes how latching occurs for SIP where HNT is provided by an SBC connected to two networks: 203.0.113/24 facing towards the UAC network and 198.51.100/24 facing towards the UAS network.

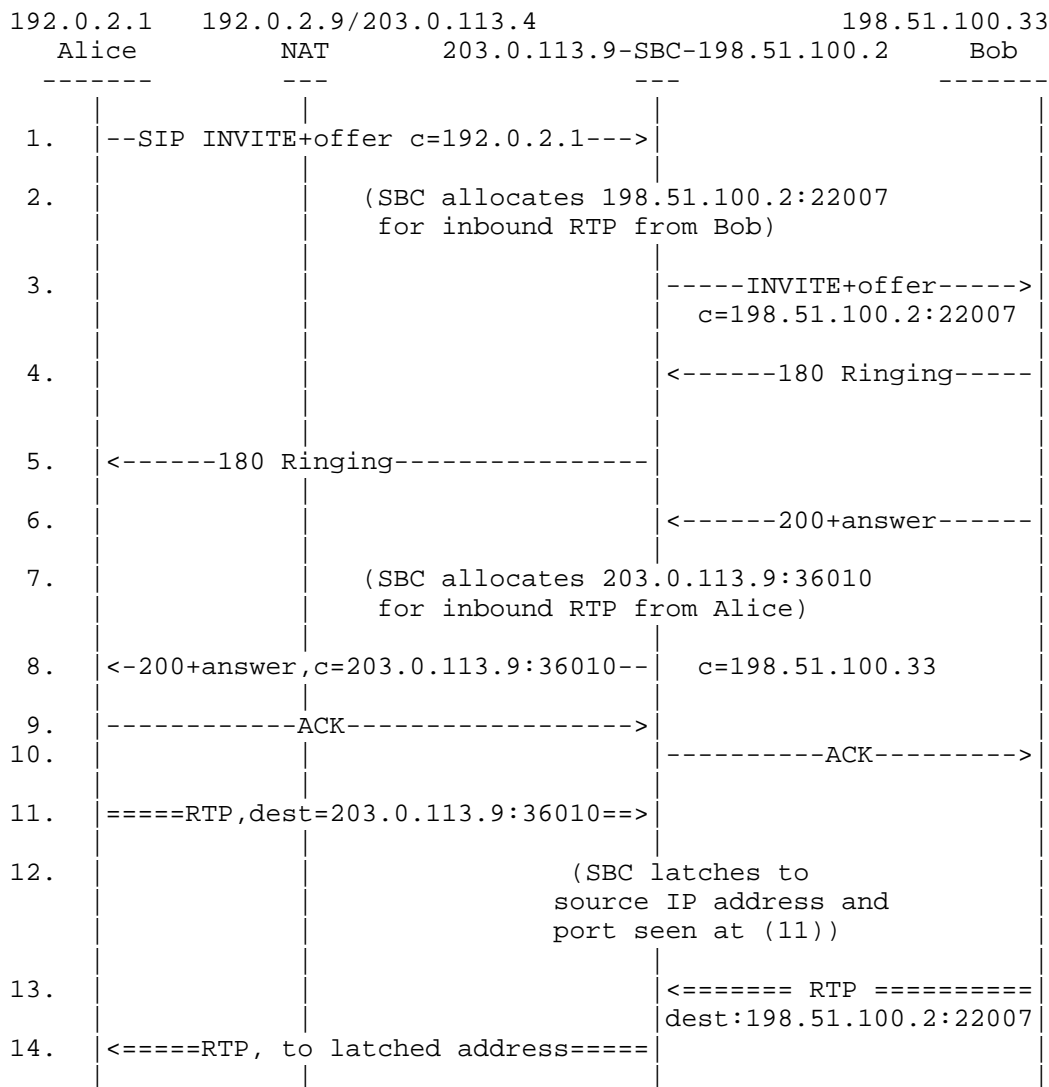


Figure 2: Latching by a SIP SBC across Two Interfaces

While XMPP implementations often rely on ICE to handle NAT traversal, there are some that also support a non-ICE transport called XMPP Jingle Raw UDP Transport Method [XEP-0177]. Figure 3 describes how latching occurs for one such XMPP implementation where HNT is provided by an XMPP server on the public Internet.



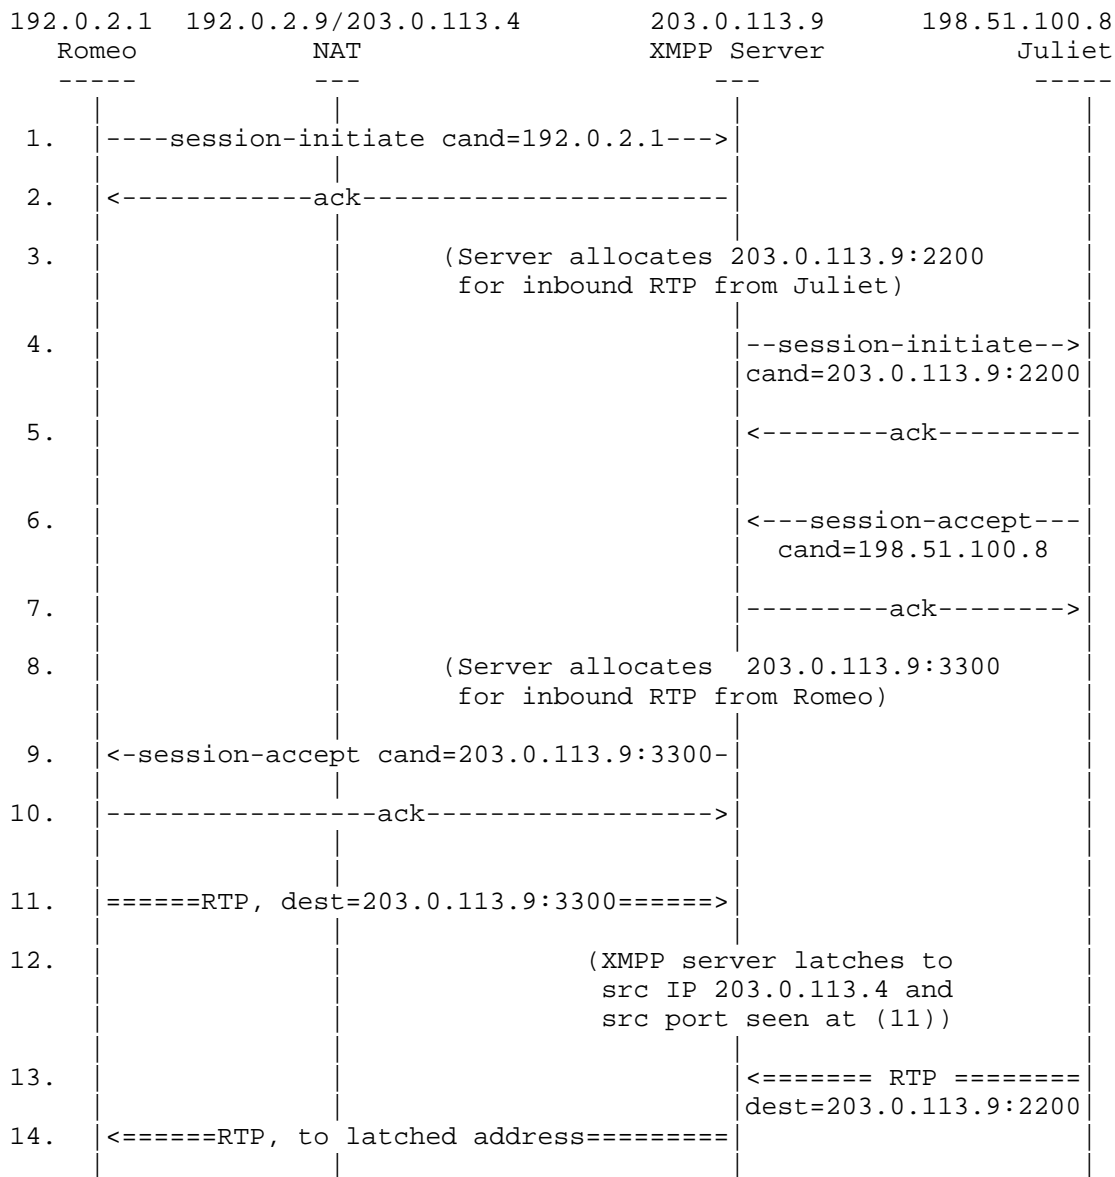


Figure 3: Latching by an XMPP Server across Two Interfaces

The above is a general description, and some details vary between implementations or configuration settings. For example, some intermediaries perform additional logic before latching on received

packet source information to prevent malicious attacks or latching erroneously to previous media senders -- often called "rogue-rtp" in the industry.

It is worth pointing out that latching is not exclusively a "server affair", and some clients may also use it in cases where they are configured with a public IP address and are contacted by a NATed client with no other NAT traversal means.

In order for latching to function correctly, the UA behind the NAT needs to support symmetric RTP. That is, it needs to use the same ports for sending data as the ones it listens on for inbound packets. Today, this is the case with almost all SIP and XMPP clients. Also, UAs need to make sure they can begin sending media packets independently without waiting for packets to arrive first. In theory, it is possible that some UAs would not send packets out first, for example, if a SIP session begins in 'inactive' or 'recvonly' SDP mode from the UA behind the NAT. In practice, however, SIP sessions from regular UAs (the kind that one could find behind a NAT) virtually never begin in 'inactive' or 'recvonly' mode, for obvious reasons. The media direction would also be problematic if the SBC side indicated 'inactive' or 'sendonly' modes when it sent SDP to the UA. However, SBCs providing HNT would always be configured to avoid this.

Given that, in order for latching to work properly, media relays need to begin receiving media before they start sending, it is possible for deadlocks to occur. This can happen when the UAC and the UAS in a session are connected to different signaling intermediaries that both provide HNT. In this case, the media relays controlled by the signaling servers could end up each waiting upon the other to initiate the streaming. To prevent this, relays would often attempt to start streaming toward the address:port tuples provided in the offer/answer even before receiving any inbound traffic. If the entity they are streaming to is another HNT performing server, it would have provided its relay's public address and ports, and the early stream would find its target.

Although many SBCs only support UDP-based media latching (in particular, RTP/RTCP), many SBCs support TCP-based media latching as well. TCP-based latching is more complicated; it involves forcing the UA behind the NAT to be the TCP client and sending the initial SYN-flagged TCP packet to the SBC (i.e., be the 'active' mode side of a TCP-based media session). If both UAs of a TCP-based media session are behind NATs, then SBCs typically force both UAs to be the TCP clients, and the SBC splices the TCP connections together. TCP splicing is a well-known technique, as described in [TCP-SPLICING].

HNT and latching, in particular, are generally found to work reliably, but they do have obvious caveats. The first one usually raised by IETF participants is that UAs are not aware of it occurring. This makes it impossible for the mechanism to be used with protocols such as ICE that try various traversal techniques in an effort to choose the one that best suits a particular situation. Overwriting address information in offers and answers may actually completely prevent UAs from using ICE because of the ice-mismatch rules described in [RFC5245].

The second issue raised by IETF participants is that it causes media to go through a relay instead of directly over the IP-routed path between the two participating UAs. While this adds obvious drawbacks such as reduced scalability and increased latency, it is also considered a benefit by SBC administrators: if a customer pays for "phone" service, for example, the media is what is truly being paid for, and the administrators usually like to be able to detect that the media is flowing correctly, evaluate its quality, know if and why it failed, etc. Also, in some cases, routing media through operator controlled relays may route media over paths explicitly optimized for media and hence offer better performance than regular Internet routing.

## 5. Security Considerations

A common concern is that an SBC (or an XMPP server -- all security considerations apply to both) that implements HNT may latch to incorrect and possibly malicious sources. The ICE [RFC5245] protocol, for example, provides authentication tokens (conveyed in the ice-ufraq and ice-pwd attributes) that allow the identity of a peer to be confirmed before engaging in media exchange with her. Without such authentication, a malicious source could attempt a resource exhaustion attack by flooding all possible media-latching UDP ports on the SBC in order to prevent calls from succeeding. SBCs have various mechanisms to prevent this from happening or to alert an administrator when it does. Still, a sufficiently sophisticated attacker may be able to bypass them for some time. The most common example is typically referred to as "restricted-latching", whereby the SBC will not latch to any packets from a source public IP address other than the one the SIP UA uses for SIP signaling. This way, the SBC simply ignores and does not latch onto packets coming from the attacker. In some cases, the limitation may be loosened to allow media from a range of IP addresses belonging to the same network in order to allow for use cases such as decomposed UAs and various forms of third-party call control. However, since relaxing the restrictions in such a way may provide attackers with a larger attack

surface, such configurations are generally performed only on a case-by-case basis so that the specifics of individual deployments can be taken into account.

All of the above problems would still arise if the attacker knows the public source IP of the UA that is actually making the call. This would allow attackers to still flood all of the SBC's public IP addresses and ports with packets spoofing that SIP UA's public source IP address. However, this would only impact media from that IP (or range of IP addresses) rather than all calls that the SBC is servicing.

A malicious source could send media packets to an SBC media-latching UDP port in the hopes of being latched to for the purpose of receiving media for a given SIP session. SBCs have various mechanisms to prevent this as well. Restricted latching, for example, would also help in this case because the attacker can't make the SBC send media packets back to themselves since the SBC will not latch onto the attacker's media packets, not having seen the corresponding signaling packets first. There could still be an issue if the attacker happens to be either (1) in the IP routing path where it can thus spoof the same IP as the real UA and get the media coming back, in which case the attacker hardly needs to attack at all to begin with, or (2) behind the same NAT as the legitimate SIP UA, in which case the attacker's packets will be latched to by the SBC and the SBC will send media back to the attacker. In the latter case, which may be of particular concern with Carrier-Grade NATs, the legitimate SIP UA will likely end the call anyway when a human user who does not hear anything hangs up. In the case of a non-human call participant, such as an answering machine, this may not happen (although many such automated UAs would also hang up when they do not receive any media). The attacker could also redirect all media to the real SIP UA after receiving it, in which case the attack would likely remain undetected and succeed. Again, this would be of particular concern with larger-scale NATs serving many different endpoints, such as Carrier-Grade NATs. The larger the number of devices fronted by a NAT is, the more use cases would vary, and the more the number of possible attack vectors would grow.

Naturally, Secure RTP (SRTP) [RFC3711] would help mitigate such threats and, if used with the appropriate key negotiation mechanisms, would protect the media from monitoring while in transit. It should therefore be used independently of HNT. Section 26 of [RFC3261] provides an overview of additional threats and solutions on monitoring and session interception.

With SRTP, if the SBC that performs the latching is actually participating in the SRTP key exchange, then it would simply refuse to latch onto a source unless it can authenticate it. Failing to implement and use SRTP would represent a serious threat to users connecting from behind Carrier-Grade NATs [RFC6888] and is considered a harmful practice.

For SIP clients, HNT is usually transparent in the sense that the SIP UA does not know it occurs. In certain cases, it may be detectable, such as when ICE is supported by the SIP UA and the SBC modifies the default connection address and media port numbers in SDP, thereby disabling ICE due to the mismatch condition. Even in that case, however, the SIP UA only knows that a middlebox is relaying media but not necessarily that it is performing latching/HNT.

In order to perform HNT, the SBC has to modify SDP to and from the SIP UA behind a NAT; thus, the SIP UA cannot use S/MIME [RFC5751], and it cannot sign a sending request, or verify a received request using the SIP Identity mechanism [RFC4474] unless the SBC re-signs the request. However, neither S/MIME nor SIP Identity are widely deployed; thus, not being able to sign/verify requests appears not to be a concern at this time.

From a privacy perspective, media relaying is sometimes seen as a way of protecting one's IP address and not revealing it to the remote party. That kind of IP address masking is often perceived as important. However, this is no longer an exclusive advantage of HNT since it can also be accomplished by client-controlled relaying mechanisms such as TURN [RFC5766] if the client explicitly wishes to do so.

## 6. Acknowledgements

The authors would like to thank Flemming Andreassen, Miguel A. Garcia, Alissa Cooper, Vijay K. Gurbani, Ari Keranen, and Paul Kyzivat for their reviews and suggestions on improving this document.

## 7. References

### 7.1. Normative References

- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.
- [RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, March 2011.
- [XEP-0177] Beda, J., Saint-Andre, P., Hildebrand, J., and S. Egan, "XEP-0177: Jingle Raw UDP Transport Method", XSF XEP 0177, December 2009.

### 7.2. Informative References

- [H.323] International Telecommunication Union, "Packet-based multimedia communication systems", ITU-T Recommendation H.323, December 2009.
- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.

- [RFC3435] Andreassen, F. and B. Foster, "Media Gateway Control Protocol (MGCP) Version 1.0", RFC 3435, January 2003.
- [RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.
- [RFC5125] Taylor, T., "Reclassification of RFC 3525 to Historic", RFC 5125, February 2008.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.
- [RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, April 2013.
- [TCP-SPLICING] Maltz, D. and P. Bhagwat, "TCP Splice for application layer proxy performance", Journal of High Speed Networks Vol. 8, No. 3, 1999, pp. 225-240, March 1999.

## Authors' Addresses

Emil Ivov  
Jitsi  
Strasbourg 67000  
France

EEmail: [emcho@jitsi.org](mailto:emcho@jitsi.org)

Hadriel Kaplan  
Oracle  
100 Crosby Drive  
Bedford, MA 01730  
USA

EEmail: [hadrielk@yahoo.com](mailto:hadrielk@yahoo.com)

Dan Wing  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134  
USA

EEmail: [dwing@cisco.com](mailto:dwing@cisco.com)