MMUSIC
Internet-Draft
Intended status: Standards Track
Expires: April 18, 2013

T. Reddy
Muthu A M. Perumal
Ram Mohan. Ravindranath
D. Wing
Cisco

October 15, 2012

STUN Extensions for Firewall Traversal
draft-reddy-mmusic-stun-auth-fw-traversal-00

Abstract

Some networks deploy firewalls configured to block UDP traffic. When SIP user agents or WebRTC endpoints are deployed behind such firewalls, media cannot be sent over UDP across the firewall, but must be sent using TCP (which causes a different user experience) or through a session border controller.

This draft describes an alternate model wherein extensions to ICE connectivity checks can be examined by the firewall to permit outgoing UDP flows across the firewall.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

To protect networks using real-time communications, firewalls or session border controllers are typically deployed.

Firewalls include Application Layer Gateway functionality, which intercepts and analyzes the session signaling traffic such as the Session Initiation Protocol (SIP) traffic and creates dynamic mapping to permit the media traffic. In particular, firewall extracts the media transport addresses, transport protocol and ports from the session description and creates dynamic mapping for media to flow through. This model has the following problems:

1. It does not work if the session signaling is end-to-end encrypted (say, using TLS).

2. It does not work if a non-standard session signaling is used that the firewall does not understand.

3. It does not work if the session signaling and media traverse different firewalls.

When an enterprise deploys WebRTC, the above problems are relevant because:

1. The session signaling between the WebRTC application running in the browser and the web server could be using TLS.

2. WebRTC does not enforce a particular session signaling protocol to be used. So, the firewall may not be able to understand it.

3. This session signaling and the peer-to-peer media may traverse different firewalls.

As a result the firewall may block ICE connectivity checks and media traffic.

Session Border Controllers (SBCs) are active participants with call signaling. Like firewalls, they also create dynamic mappings to permit media traffic. This forces call signaling and media through specific IP addresses, belonging to the SBC or an SBC-controlled media relay device.

TURN is also used as an alternative to permit media traffic, i.e. Use TCP transport between the client and TURN server because Firewalls are configured to block UDP entirely.

The use-case is explained in Section 4.2.4.1 of

[I-D.ietf-rtcweb-use-cases-and-requirements] refers to deploying a TURN server to audit all media sessions from inside the company premises to any external peer.

Using TURN for all such communication has the following problems:

- Single TURN server will result in single point of failure.

- TURN server could increase media latency and high-end TURN server would be needed to cater to all such calls.

- TURN server is just providing the 5-tuple details (source IP address, destination IP address, protocol number, source port number, and destination port number) but no other details of the WebRTC server using which the call is initiated

- Enterprise firewalls would typically have granular policies to permit call initiated using selected WebRTC servers (Dr.Good) it trusts and block others (Dr.Evil).

- It comes at a high cost to the provider of the TURN server, since the server typically needs a high-bandwidth connection to the Internet. As a consequence, it is best to use a TURN server only when a direct communication path cannot be found. When the client and a peer use ICE to determine the communication path, ICE will use hole punching techniques to search for a direct path first and only use a TURN server when a direct path cannot be found.

- The value of the Diffserv field may not be preserved.

- The Explicit Congestion Notification (ECN) field may be reset.

This draft has a solution where an authorized server (could be a Call Agent or a WebRTC server) generate a cryptographic token which is passed to the endpoints. The endpoint includes the token in its ICE connectivity checks. The firewall intercepts the ICE connectivity checks containing the token, validates it, and permits the ICE connectivity checks and the subsequent media flow through the firewall.

2. Notational Conventions

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].

This note uses terminology defined in [RFC5245].

3. Problem Statement

In the below topology, an webRTC Server is deployed in the enterprise Data Center. Alice makes a webRTC call to Bob. For the two endpoints to successfully establish media sessions, firewalls FW1 and FW2 need to permit the ICE connectivity checks and media traffic. In such scenarios the mechanism described in this draft proposes a new comprehension-optional FW-FLOWDATA STUN attribute to be included in STUN Bind requests sent during ICE connectivity checks so that firewalls will permit media traffic between internal peers. This STUN attribute is created by the trusted WebRTC server and sent to the endpoints to be propagated by the respective ICE agents during ICE connectivity checks.

```
               =========================
               |  WebRTC Server        |
               =========================
                         |  Data Center
                         |
                         |
                 ==================
                 |    WAN         |-----+-+-------+---+----+-+-----+
                 ==================                        |
                           |              Branch office 2  |
  Branch office 1          |                               |
                           |                               |
                   +-------+-------+               +----+-------+
                   | Firewall 1    |               | Firewall 2 |
                   |               |               |            |
                   +-------+-------+               +----+-------+
                           |                               |
                           |                               |
                           |                               |
   ---+-+-----+-----------+-+-----+--------         -----+-+-----+------
                           |                               |
                        +-+------+                     +--------+
                        | Alice  |                     | Bob    |
                        +--------+                     +--------+
```

Figure 1: WebRTC service in enterprise - internal call

4. Solution Overview

This section gives an overview of the solution and the different components involved in the solution and the role of each component.

4.1. Different Components and the Trust model

Figure 1 above shows typical components involved in a WebRTC call scenario. As part of the call setup, the WebRTC endpoint would have to gather its candidates from a STUN/TURN server, send the candidates in the offer to the peer endpoint. On receiving the answer from the peer endpoint it starts the ICE connectivity checks. As discussed in the problem statement, firewalls would typically block these ICE connectivity checks and media flowing there after. To allow this traffic a firewall needs to authorize the flow.

o A new comprehensive optional STUN attribute called FW-FLOWDATA is defined as part of this draft. This is used by WebRTC endpoints requiring firewall traversal.

o This STUN attribute FW-FLOWDATA is generated by the WebRTC server in co-ordination with the WebRTC endpoint.

o Once the WebRTC session ends the firewall's dynamic mappings are closed after timeout.

o DISCUSSION: Could we could have a FW-FLOWDATA attribute sent in a STUN message to close the dynamic mappings in the firewalls?

5. Usage and Processing

An RTP endpoint which generates media can include the FW-FLOWDATA attribute in its STUN Binding requests used in ICE connectivity checks, to inform on-path firewalls to permit the flow.

5.1. **Generating** FW-FLOWDATA Attribute

The WebRTC server after processing the OFFER/ANSWER sends the FW-FLOWDATA STUN attribute to both the peers to be included in the ICE connectivity checks. The Authentication Tag field in the FW-FLOWDATA attribute contains the digest of the FW-FLOWDATA attribute for data origin authentication and integrity protection. The server first selects the candidate address info based on OFFER/ANSWER exchange and generates other fields of this attribute. The server then computes a digest for the FW-FLOWDATA attribute using HMAC-SHA1. The key for HMAC-SHA1 is provisioned using the technique in Section 7. The result of which is truncated to 96 bits (retaining the left most bits) to produce HMAC-SHA-1-96 and input into the Authentication Tag field. The mechanism to send FW-FLOWDATA attribute from the WebRTC server to the cient is outside the scope of this draft. But it is assumed that signalling protocols used for WebRTC call setup will be enhanced to deliver this new attribute to the WebRTC client. The

WebRTC server MUST provide a new FW-FLOWDATA to allow the media session to continue before Lifetime expires.

5.2. Sending FW-FLOWDATA Attribute in Binding Request

Once a WebRTC endpoint receives the FW-FLOWDATA, it is responsible for generating the STUN message and retransmitting the transactions per the STUN specification. The FW-FLOWDATA attribute should be placed before the FINGERPRINT attribute (if present) and after the MESSAGE-INTEGRITY attribute. The STUN length field is adjusted to point to the new end of the STUN message; that is, the STUN length field always accurately indicates the length of the STUN message (including the MESSAGE-INTEGRITY, FINGERPRINT, and FW-FLOWDATA attributes). This does not interfere with 3rd party receivers of the STUN message, as they will adjust the STUN length field to point to the end of the MESSAGE-INTEGRITY field. Receivers that do not understand the FW-FLOWDATA will ignore it.

FW-FLOWDATA attribute received by the WebRTC client is passed to the web browser's ICE agent (API to be added in in W3C WebRTC-API specification [I.D.w3c-webrtc]). The ICE agent includes the FW-FLOWDATA attribute with all ICE connectivity checks, so that on-path firewalls can validate and permit the ICE connectivity checks and forthcoming media. The token MUST included in the ICE binding indication packets (keepalive) (In case the lifetime expires)

For the FW-FLOWDATA attribute to be visible to the firewalls between the client and the TURN server, the FW-FLOWDATA should be included in the ALLOCATE request, channel bind or refresh messages going to the TURN server. This is to avoid firewalls having to look for STUN packets within STUN (TURN) packets.

5.3. Firewalls processing FW-FLOWDATA Attribute

Firewalls can reliably determine a UDP message is a STUN message because all STUN messages sent as ICE connectivity checks include the 32-bit STUN magic cookie and the FINGERPRINT attribute. STUN messages which are authenticated also include a MESSAGE-INTEGRITY attribute which authenticates the fields prior to the MESSAGE-INTEGRITY.

When the firewall receives a STUN binding request with FW-FLOWDATA attribute it stores the Authentication Tag in the FW-FLOWDATA attribute. The firewall then generates a digest for the FW-FLOWDATA attribute using HMAC-SHA1. The result of which is truncated to 96 bits (retaining the left most bits) to produce HMAC-SHA-1-96. If the value of the newly generated digest HMAC-SHA-1-96 is identical to the stored one, the firewall can ensure that the FW-FLOWDATA attribute

has not been tampered with. Otherwise the packet is discarded.

To facilitate timestamp checking for replay attacks, each firewall should perform the following check for each message:

When a message is received, the received timestamp, TSnew, is checked, and the packet is accepted if the timestamp is recent enough to the reception time of the packet, RDnew:

Lifetime + Delta > (RDnew - TSnew)

The recommended value for the allowed Delta is 30 seconds. If the timestamp is NOT within the boundaries then discard the STUN message.

The firewall also performs the following checks:

- Ensures that the source IP address and UDP port of the packet matches with one of the local CAI entries in the payload except for peer-reflexive cases.

- Ensures the destination IP address and UDP port of the packet matches with one of the local CAI entries in the packet payload except for peer-reflexive cases.

- Firewall if located after NAT(peer-reflexive cases) can skip CAI processing (It can be configurable option). For peer-reflexive case, destination CAI MUST match in case of outgoing STUN packet and source CAI MUST match incase of incoming STUN packet

If all the above checks pass then the firewall creates the 5-tuple dynamic mapping using the local candidate IP address, local candidate port, remote candidate IP address, remote candidate port, transport protocol. The session time of the dynamic mapping will be set to a short lifetime (default value of 60 seconds).

If the initial ICE connectivity check includes the ICE-CONTROLLING attribute but does not include USE-CANDIDATE, ICE connectivity check is successful and a subsequent ICE connectivity check includes both these attributes, the firewall can determine that the ICE agent is the controlling agent using regular nomination and this candidate pair is nominated for media flow. The firewall then sets the session time of the dynamic mapping equal to the Lifetime field in FW-FLOWDATA attribute.

If the initial ICE connectivity check includes the ICE-CONTROLLING attribute and the USE-CANDIDATE attribute, firewall can determine that the ICE agent is the controlling agent using aggressive using nomination. If the ICE connectivity check is successful It then

waits for the media traffic to flow before setting the session time of the dynamic mapping equal to Lifetime field in FW-FLOWDATA attribute.

DISCUSSION: If WebRTC implementations of RTP support multiplexing of multiple media sessions onto a single RTP session, FW-FLOWDATA attribute can be enhanced to carry a flag indicating the same so that firewall can immediately close the dynamic mapping created for other pairs in the ICE checklist once media starts flowing on one the candidate pairs. In case of multi-homing firewalls can track multiple host IP addresses using authentication supplicant or, for hosts lacking the supplicant, use address-based authentication method.

6. STUN Attribute Format

```
        0                   1                   2                   3
        0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                          Lifetime                             |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                          Nonce                                |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                          Timestamp                            |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       | LCA  Count    |  RCA  Count   |    Reserved                   |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                                                               |
       |                 Candidate Address Info                        |
       |                                                               |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                                                               |
       |                 Authentication Tag                            |
       |                                                               |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: FW-FLOWDATA Attribute

Lifetime: 32-bit unsigned integer. The length of time in seconds that the STUN attribute is valid for the purpose of firewall creating dynamic mapping. The lifetime of the firewall dynamic mapping is set to this value. After the lifetime expires the mapping is deleted, unless the lifetime is extended using a another FW-FLOWDATA attribute.

Nonce: 96-bit unsigned integer. Random value chosen by the WebRTC Server that uniquely identifies the STUN attribute.

Timestamp: 64-bit unsigned integer field containing a timestamp. The value indicates the number of seconds since January 1, 1970, 00:00 UTC, by using a fixed point format. In this format, the integer number of seconds is contained in the first 48 bits of the field, and the remaining 16 bits indicate the number of 1/64K fractions of a second.

LCA Count: 8-bit unsigned integer. Number of local candidate addresses.

RCA Count: 8-bit unsigned integer. Number of remote candidate addresses.

Reserved: 16-bit unsigned integer. An unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Candidate Address Info:

```
      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |Family         | Protocol      |         Port                  |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                                                               |
     |                 Address (32 bits or 128 bits)                 |
     |                                                               |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

It consists of an 8-bit address family, L4 protocol (for example 17 for UDP, 6 for TCP) and a 16-bit port, followed by a fixed-length value representing the IP address. If the 16-bit port value is 0 it indicates "all ports". The address family can take on the following values: 0x01:IPv4 and 0x02:IPv6. An endpoint may send Zero or more CAI in its FLOWDATA

Authentication Tag: A 96-bit field that carries the Message Authentication Code for the FW-FLOWDATA STUN attribute.

7. Key Provisioning

Static keys are preconfigured, either manually or through a network management system. The simplest way to implement FW-FLOWDATA validation is to use static keys. The provisioning of static keys requires either manual operator intervention on the WebRTC Server and each firewall in the enterprise or a network management system

performing the same task.

Alternatively using Dynamic Group Key Distribution, group keys are dynamically distributed among the WebRTC server and enterprise firewalls using GDOI [RFC6407]. In this way, each firewall requests a group key from a key server as part of an encrypted and integrity-protected key agreement protocol. Once the key server has authenticated and authorized the firewalls, it distributes a group key to the group member. The authentication in this model can be based on public key mechanisms, thereby avoiding the need for static key provisioning.

8. Security Considerations

Hosts using WebRTC calls will see lot of FW-FLOWDATA attributes. They determine the key by trying a number of candidate keys and seeing if one of them is correct. The attack works when the keys have low entropy, such as a word from the dictionary. This attack can be mitigated by using strong keys with large entropy. In situations where even stronger mitigation is required, the keys can be dynamically changed using GDOI. The WebRTC server controls how long a firewall session is kept open via the Lifetime value and WebRTC server could use different Lifetime values depending on the anticipated level of trust of the device (e.g. company provided laptop might be trusted more than a Bring Your Own Device (BYOD)); the device with more trust need to obtain its authentication attribute less often). Firewalls in addition to timestamp checking can also maintain a cache of used Nonces, IP source addresses associated with used Nonces as an effective countermeasure against replay attacks.

All the security considerations applicable to STUN [RFC5389] and ICE [RFC5245] are applicable to this document as well.

9. IANA Considerations

Allocate new STUN attribute value for FW-FLOWDATA from the [STUN-ATTR] registry.

10. Acknowledgements

The authors would like to thank Prashanth Patil and Ramesh Nethi for review comments.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[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.

[RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, October 2011.

11.2. Informative References

[I-D.ietf-rtcweb-use-cases-and-requirements] Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use-cases and Requirements", draft-ietf-rtcweb-use-cases-and-requirements-09 (work in progress), June 2012.

[STUN-ATTR] IANA, "IANA: STUN Attributes", December 2011, <http://www.iana.org/assignments/stun-parameters/stun-parameters.xml#stun-parameters-3>.

Authors' Addresses

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com

Muthu Arul Mozhi Perumal
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: mperumal@cisco.com

Ram Mohan Ravindranath
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: rmohanr@cisco.com

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

Email: dwing@cisco.com