STUN TURN and ICE Demystified

Redouane MEDDANE
Cisco Instructor CCSI 35458
3×CCNP Collaboration Security Enterprise

Before looking at the details of the different protocols and standards, it is important to understand why they are needed. ICE, STUN, and TURN are used to establish a MEDIA connection between two devices on different networks separated by firewalls and NAT servers. They do not apply to signaling between devices which means that if you cannot establish a signaling connection between the two devices then ICE, STUN and TURN are irrelevant. There must be a signaling connection established and only then will the devices attempt to establish a media connection.

The issues with NAT and collaboration media connections occurs during the signaling. The media IP address and port of each stream is shared inside the SDP messages during call setup. So even when a packet has the source addresses changed by NAT, the media information inside the SDP messages remain the same. When the called device receives the SIP message, it doesn't try streaming the media to the source address on the received SIP packet but to the IP media addresses inside the SDP message which is the internal address of the device and therefore the media connection cannot be established. To overcome this issue ICE, STUN and TURN are used.

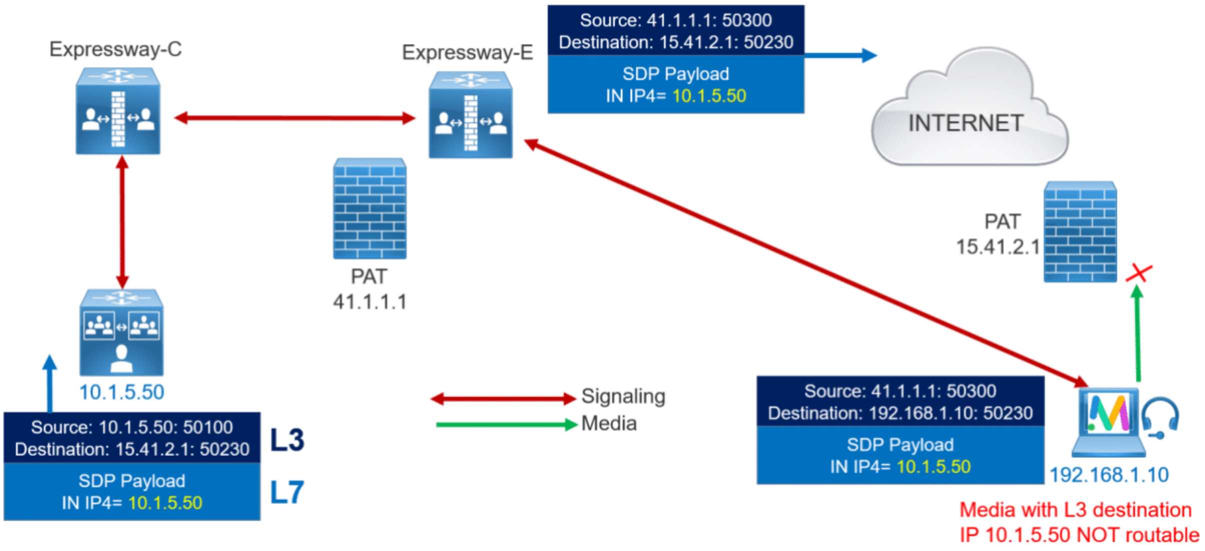

How they are used to traverse media through a Firewall and override the issue caused by Network Address Translation (NAT)? What is the role of each protocols? What are the differences? How the interaction occurs between STUN, TURN and ICE.

Let's start by looking at Session Traversal Utilities for NAT (STUN). The sole purpose of STUN is for a device behind a firewall to discover what its NAT'd address and port are when it sends UDP traffic outside. A STUN server is installed outside the firewall and the device inside sends a STUN request to the STUN server.

Before sending the INVITE, the calling endpoint jdoe starts a session with the STUN server to obtain its own client NATed IP 15.1.2.1 and port 19100.

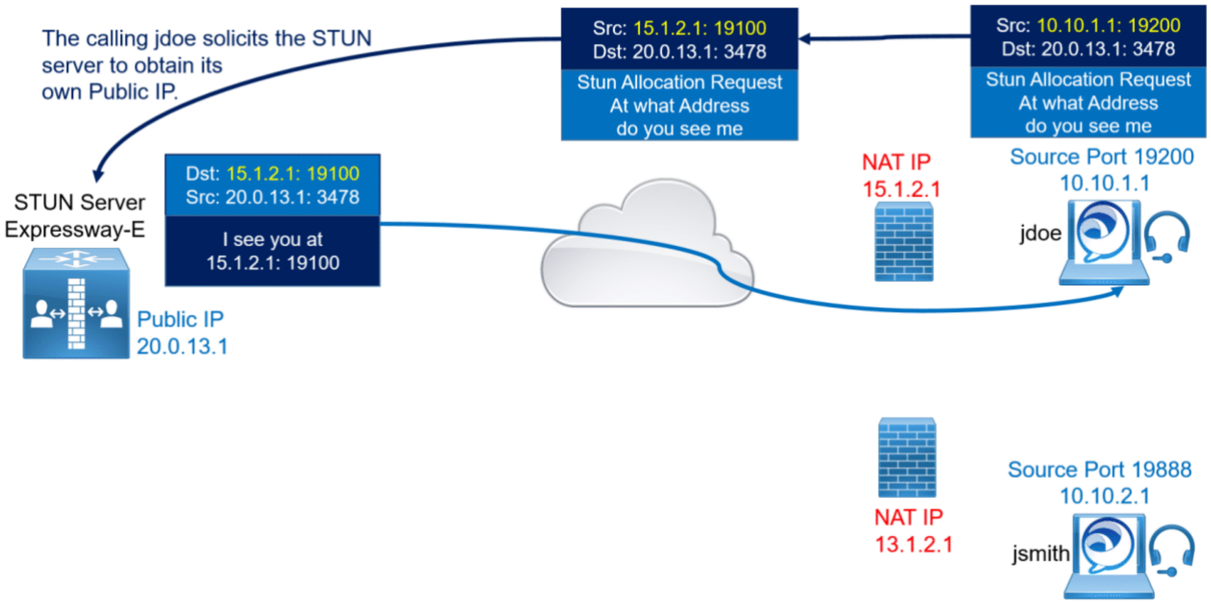

The client NATed IP 15.1.2.1 and port 19100 are sent in the SDP INVITE.
The INVITE is then forwarded by Expressway-E, Expressway-C, and Unified CM toward the called MRA device jsmith.

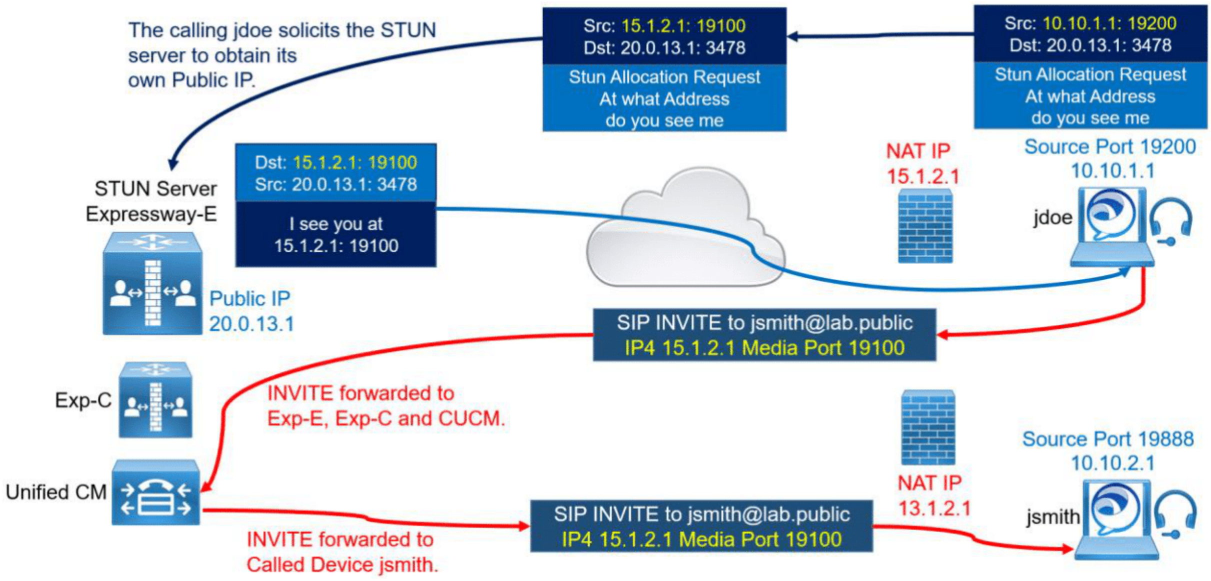
The calling jdoe solicits the STUN server to obtain its own Public IP.
Src: 15.1.2.1: 19100
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
Src: 10.10.1.1: 19200
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
STUN Server Expressway-E
Dst: 15.1.2.1: 19100
Src: 20.0.13.1: 3478
I see you at 15.1.2.1: 19100
Public IP 20.0.13.1
NAT IP 15.1.2.1
Source Port 19200
10.10.1.1
jdoe
SIP INVITE to jsmith@lab.public
IP4 15.1.2.1 Media Port 19100
Exp-C
INVITE forwarded to Exp-E, Exp-C and CUCM.
Unified CM
INVITE forwarded to Called Device jsmith.
SIP INVITE to jsmith@lab.public
IP4 15.1.2.1 Media Port 19100
NAT IP 13.1.2.1
Source Port 19888
10.10.2.1
jsmith

When the called device jsmith receives the SIP INVITE from jdoe, it starts communicating with the Expressway-E STUN server in order to obtain its own Client NATed IP 13.1.2.1 and port 19889.

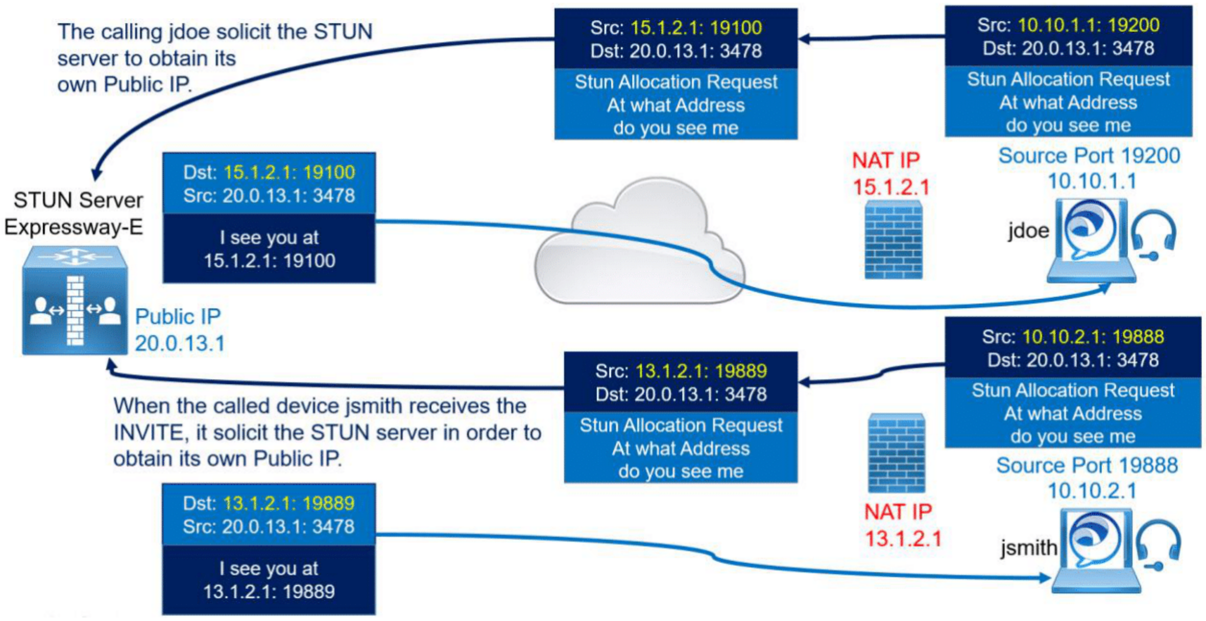
The calling jdoe solicit the STUN server to obtain its own Public IP.
Src: 15.1.2.1: 19100
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
Src: 10.10.1.1: 19200
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
STUN Server Expressway-E
Dst: 15.1.2.1: 19100
Src: 20.0.13.1: 3478
I see you at 15.1.2.1: 19100
NAT IP 15.1.2.1
Source Port 19200
10.10.1.1
jdoe
Public IP 20.0.13.1
Src: 13.1.2.1: 19889
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
Src: 10.10.2.1: 19888
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
When the called device jsmith receives the INVITE, it solicit the STUN server in order to obtain its own Public IP.
Source Port 19888
10.10.2.1
Dst: 13.1.2.1: 19889
Src: 20.0.13.1: 3478
I see you at 13.1.2.1: 19889
NAT IP 13.1.2.1
jsmith

The called endpoint then includes this NATed IP 13.1.2.1 in the SDP response and port 19889.

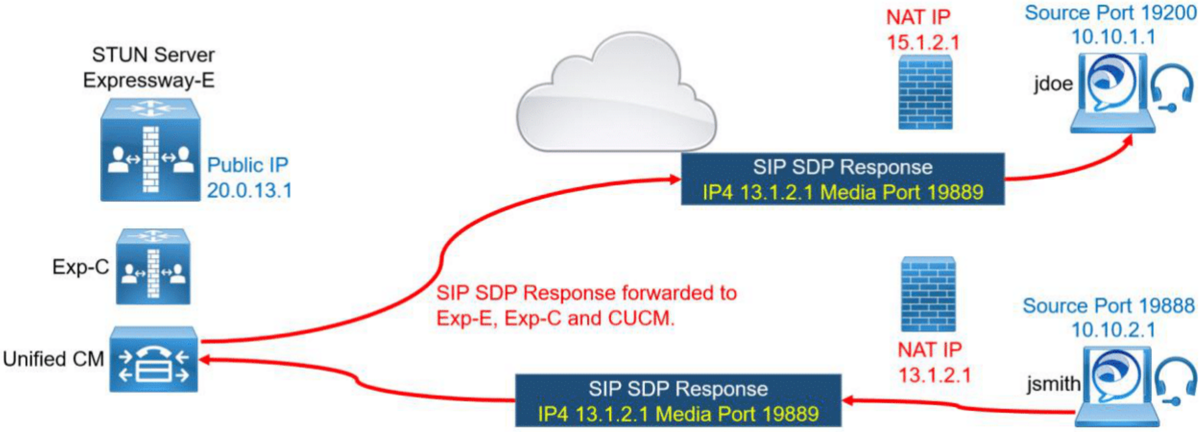

Once the signaling is completed, both endpoints have the respective Public IP Address that will be used to establish the RTP or media communication, and they can start the connectivity check phase.

The connectivity check is performed by using the STUN protocol. Each endpoint sends UDP traffic to the other endpoint's Public IP Address. If host-to-host and server reflexive to server reflexive binding requests fail, the media will then be sent through the TURN server.

Jsmith and jdoe test the connectivity using the public IP addresses 13.1.2.1 and 15.1.2.1 respectively.

Jdoe sends a UDP packet simulating the audio media traffic (STUN packet) to jsmith's Public IP Address, which is 13.1.2.1 with destination UDP port 19889. Jdoe's source transport address is 10.10.1.1 and UDP port 19200 (this address is translated by NAT before reaching jsmith's firewall external interface to 15.1.2.1: 19100). Once this packet reaches jsmith's firewall, it is dropped by the firewall because it is initiated from outside untrusted zone to inside trusted zone.

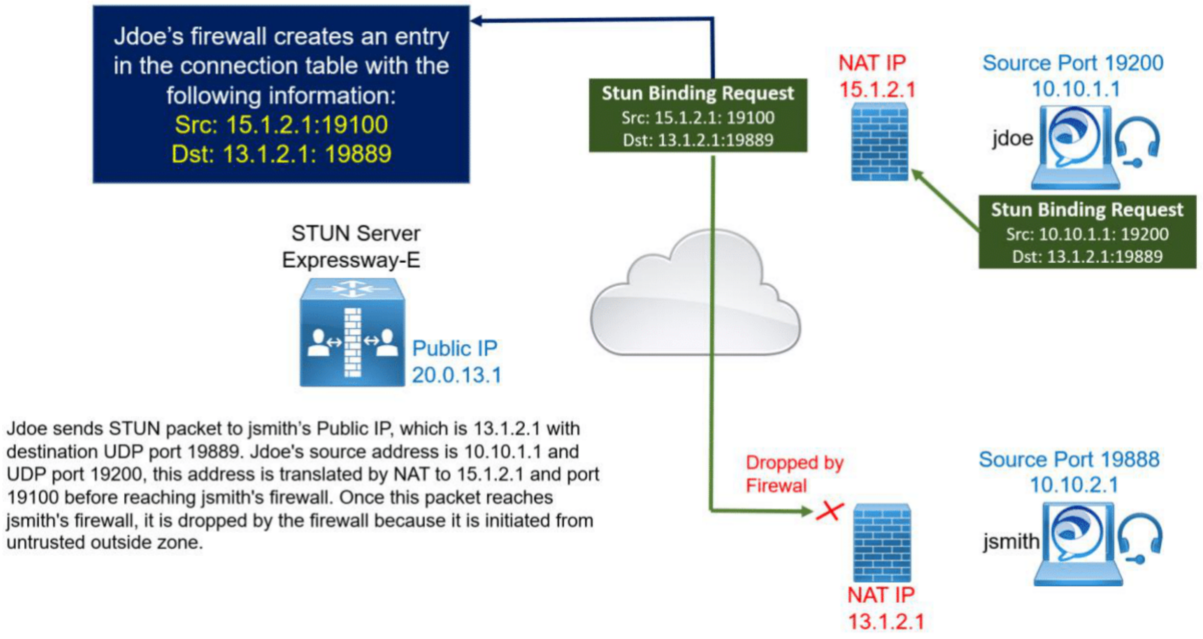

Jdoe sends STUN packet to jsmith's Public IP, which is 13.1.2.1 with destination UDP port 19889. Jdoe's source address is 10.10.1.1 and UDP port 19200, this address is translated by NAT to 15.1.2.1 and port 19100 before reaching jsmith's firewall. Once this packet reaches jsmith's firewall, it is dropped by the firewall because it is initiated from untrusted outside zone.

jsmith's endpoint, however, also sends a STUN packet Jsmith sends a STUN packet to jdoe 's Public IP which is 15.1.2.1.
The packet has 10.10.2.1:19888 as the source transport address and 15.1.2.1:19100 as the destination address. This packet is also translated by NAT by jsmith's firewall, and it reaches jdoe's firewall with the source address of 13.1.2.1:19889. With UDP stateful filtering inspection, the packet is allowed by jdoe's firewall because it'a part of an existing connection created by the Stun packet initiated by jdoe and the connectivity check will be successful. The media is established using the Public IP Addresses jdoe and jsmith.

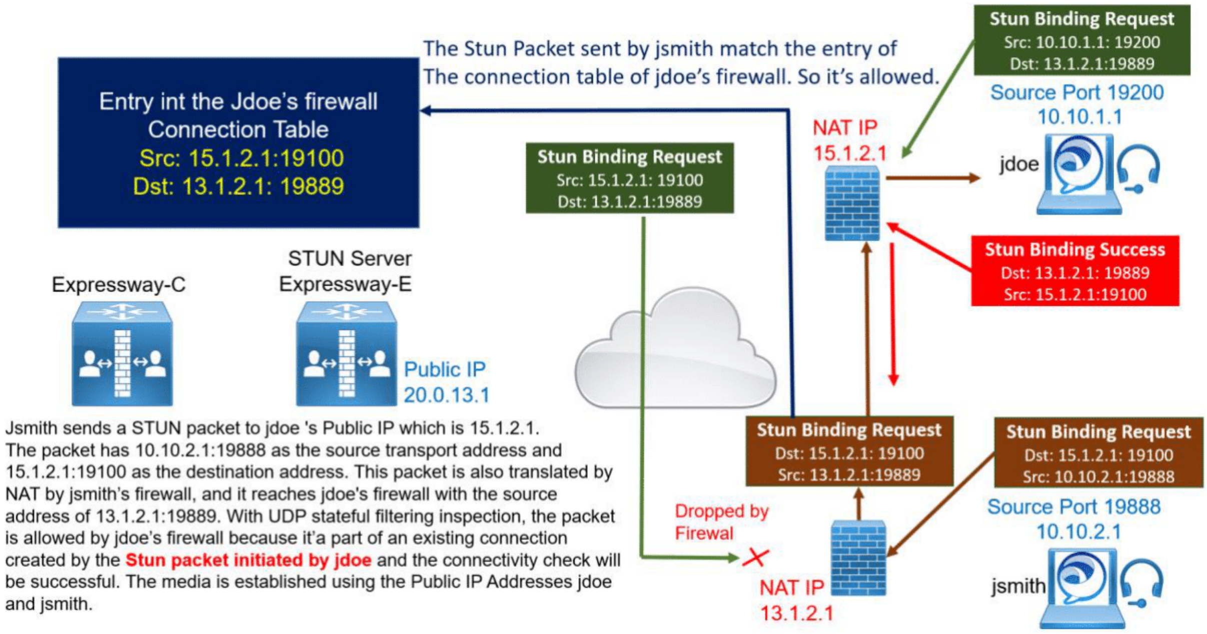

Jsmith sends a STUN packet to jdoe 's Public IP which is 15.1.2.1. The packet has 10.10.2.1:19888 as the source transport address and 15.1.2.1:19100 as the destination address. This packet is also translated by NAT by jsmith's firewall, and it reaches jdoe's firewall with the source address of 13.1.2.1:19889. With UDP stateful filtering inspection, the packet is allowed by jdoe's firewall because it'a part of an existing connection created by the **Stun packet initiated by jdoe** and the connectivity check will be successful. The media is established using the Public IP Addresses jdoe and jsmith.

STUN only works with less-secure NATs, so-called "full-cone" NATs. STUN cannot be used with symmetric NATs. This protocol may be a disadvantage in many situations as most enterprise firewalls use symmetric NAT. For some people who traditionally work with different vendors like Cisco, F5, Fortinet, different terms of NAT are used, the most common NAT Types used in many vendors are: PAT, Static NAT SNAT, Dynamic NAT.

The question is there a differences with Full Cone NAT and Symmetric NAT used as a references to explain Stun and Turn protocols?

Section 5 : NAT Variations

```
   It is assumed that the reader is familiar with NATs.  It has been
   observed that NAT treatment of UDP varies among implementations.  The
   four treatments observed in implementations are:

   Full Cone: A full cone NAT is one where all requests from the
      same internal IP address and port are mapped to the same external
      IP address and port.  Furthermore, any external host can send a
      packet to the internal host, by sending a packet to the mapped
      external address.

   Symmetric: A symmetric NAT is one where all requests from the
      same internal IP address and port, to a specific destination IP
      address and port, are mapped to the same external IP address and
      port.  If the same host sends a packet with the same source
      address and port, but to a different destination, a different
      mapping is used.  Furthermore, only the external host that
      receives a packet can send a UDP packet back to the internal host.
```

If we look at RFC 3489 in section: 5. Variations, the Full Cone NAT means that the private IP and port of an internal host are always mapped or translated to the same public IP and port when going to internet, this means that from this point of view, this internal host is reachable from internet using its own public IP and Port. In vendor's language, this type of NAT behavior is called Static NAT.

For Symmetric NAT, RFC 3489 explains this kind of NAT as follow: internal host's IP/Port are translated to different public/port when going to different destination on the internet. And it add an interesting sentence: only the external host that receives a packet can send packet back to the internal host, this means that this internal host is not reachable directly from internet like the Full Cone NAT but after initiating a traffic then the response is sent back to this host.

From the vendors's perspective such as Cisco, the Dynamic NAT and PAT are similar to Symmetric as explained by RFC 3489.

So finally Full Cone NAT is our Static NAT we usually configure on our routers and firewalls while the Symmetric NAT is finally the Dynamic NAT or PAT we commonly use on our infrastructure.

With Symmetric NAT (IP address and TCP/UDP ports are translated differently) then STUN technology will not work so media peer to peer is not possible. Challenge is to find a fallback mechanism.
Solution the media will be handled by a central server using a protocol called Traversal Using Relays around NAT (TURN). This central server is the Turn Server.

Jabber endpoint or phone devices addresses can be translated dynamically by NAT. However, if one of the two clients is behind a router or firewall performing symmetric NAT, direct media will always fail, and the media will flow through the TURN server.

Symmetric NAT happens when the source transport address is translated in different ways based on the destination address.

The calling jdoe solicits the Stun server to obtain its own Public IP and port 15.1.2.1: 19222. The client NATed IP 15.1.2.1 and port 19222 is sent in the SDP INVITE.
When the called device jsmith receives the INVITE, it solicits the Stun server in order to obtain its own Public IP and port 15.1.2.1: 19999. The called endpoint then includes this NATed IP 13.1.2.1 and port 19999 in the SDP response.

Once the signaling is completed, both endpoints have the respective Public IP Address and port that will be used to establish the RTP or media communication, and they can start the connectivity check phase.

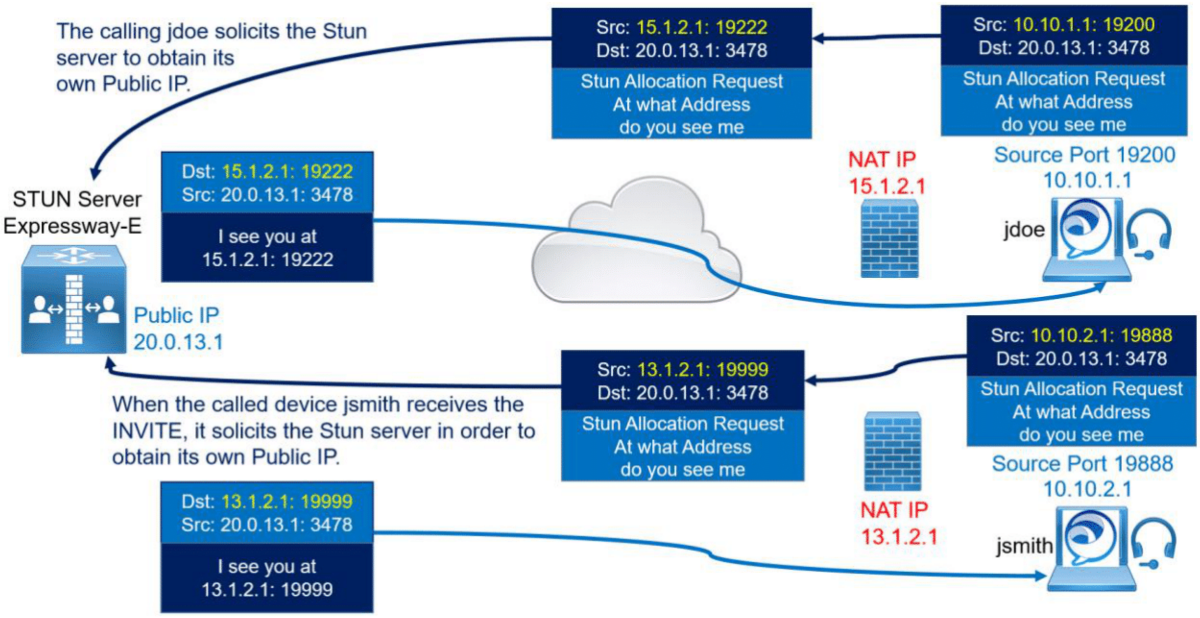

The calling jdoe solicits the Stun server to obtain its own Public IP.
Src: 15.1.2.1: 19222
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
Src: 10.10.1.1: 19200
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
NAT IP
15.1.2.1
Source Port 19200
10.10.1.1
jdoe
STUN Server
Expressway-E
Dst: 15.1.2.1: 19222
Src: 20.0.13.1: 3478
I see you at 15.1.2.1: 19222
Public IP
20.0.13.1
Src: 10.10.2.1: 19888
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
Src: 13.1.2.1: 19999
Dst: 20.0.13.1: 3478
Stun Allocation Request At what Address do you see me
When the called device jsmith receives the INVITE, it solicits the Stun server in order to obtain its own Public IP.
Source Port 19888
10.10.2.1
NAT IP
13.1.2.1
jsmith
Dst: 13.1.2.1: 19999
Src: 20.0.13.1: 3478
I see you at 13.1.2.1: 19999

Jdoe sends a UDP packet simulating the audio media traffic (STUN packet) to jsmith's Public IP Address, which is 13.1.2.1 with destination UDP port 19999. jdoe's source transport address is 10.10.1.1 and UDP port 19200 (this address is translated by NAT before reaching jsmith's firewall external interface to 15.1.2.1: 20200 (because the Symmetric NAT, another source port is used). Once this packet reaches jsmith's firewall, it is dropped by the firewall because it is initiated from outside untrusted zone to inside trusted zone.

Jsmith’s endpoint, also sends a STUN packet. Jsmith sends a STUN packet to jdoe 's Public IP which is 15.1.2.1.
The packet has 10.10.2.1:19888 as the source transport address and 15.1.2.1:19222 as the destination address. This packet is also translated by NAT by jsmith’s firewall, and it reaches jdoe's firewall with the source address of 13.1.2.1:19777. Due to Symmetric NAT, jdoe can receive traffic on port 20200, but jsmith is sending Stun Packet to port 19222.
Connectivity check will fail because jdoe's firewall is expecting to receive a packet with L3/L4 informations Src: 15.1.2.1:20200 and Dst: 13.1.2.1: 19999 that matches the entry's connection table.

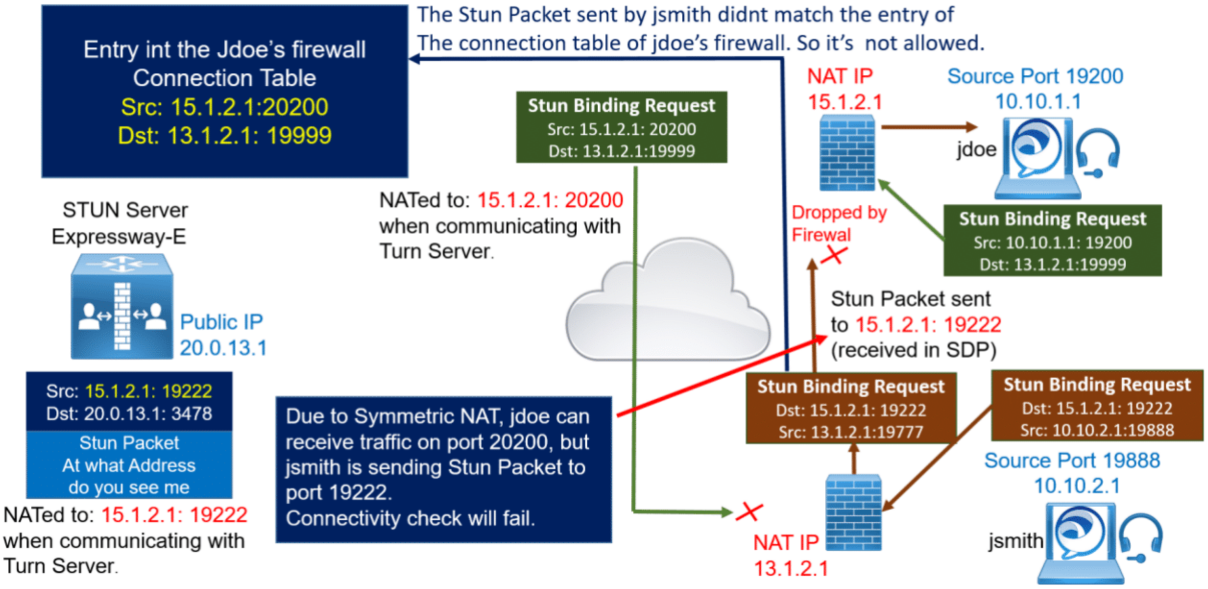
Entry int the Jdoe's firewall Connection Table
Src: 15.1.2.1:20200
Dst: 13.1.2.1: 19999
The Stun Packet sent by jsmith didnt match the entry of The connection table of jdoe's firewall. So it's not allowed.
Stun Binding Request
Src: 15.1.2.1: 20200
Dst: 13.1.2.1:19999
NAT IP
15.1.2.1
Source Port 19200
10.10.1.1
jdoe
STUN Server
Expressway-E
NATed to: 15.1.2.1: 20200 when communicating with Turn Server.
Dropped by Firewall
Stun Binding Request
Src: 10.10.1.1: 19200
Dst: 13.1.2.1:19999
Public IP
20.0.13.1
Stun Packet sent to 15.1.2.1: 19222 (received in SDP)
Src: 15.1.2.1: 19222
Dst: 20.0.13.1: 3478
Stun Packet
At what Address do you see me
Due to Symmetric NAT, jdoe can receive traffic on port 20200, but jsmith is sending Stun Packet to port 19222.
Connectivity check will fail.
Stun Binding Request
Dst: 15.1.2.1: 19222
Src: 13.1.2.1:19777
Stun Binding Request
Dst: 15.1.2.1: 19222
Src: 10.10.2.1:19888
Source Port 19888
10.10.2.1
NATed to: 15.1.2.1: 19222 when communicating with Turn Server.
NAT IP
13.1.2.1
jsmith

With Symmetric NAT (IP address and TCP/UDP ports are translated differently) then STUN technology will not work so media peer to peer is not possible. Challenge is to find a fallback mechanism.
Solution the media will be handled by a central server using a protocol called Traversal Using Relays around NAT (TURN). This central server is the Turn Server.

Interactive Connectivity Establishment (ICE), defined in RFC 8445, is a framework that combines STUN and TURN.

Using ICE, devices can determine:

If there is direct connectivity between them and will then apply the STUN Protocol.

If direct media connectivity cannot be achieved, the endpoints will fall back to the TURN server and will send their UDP traffic centrally instead of going peer-to-peer.

Stun to discover the Endpoint's NATed IP and Port.
Turn to discover the Relayed Public IP of Turn Server (Ex. Expressway-E)
ICE to discover optimal RTP path either direct RTP flow between endpoints if they use dedicated Public IP, or through the Relayed Turn Server (Ex. Expressway-E), or through the legacy RTP flow (Ex. Through Expressway-C for MRA Endpoints).

In the initial phase, a client communicates with the TURN server to obtain the other peer's transport addresses. A transport address is defined as the combination of an IP address and UDP port number. There are three different transport addresses that an endpoint might use. The native transport IP address of the client is called the host address (Address A in the figure); the IP address and UDP port as seen by the TURN server after NAT has been applied is called the server reflexive address (Address B in the figure), and the endpoint-mapped address on the TURN server is called the relayed address (address C).

The TURN server sends the endpoint the server reflexive address and the relayed address during this initial phase. Those addresses (A, B, and C) populate the SIP SDP payload and offered as ICE candidates, and after the signaling has gone through, both endpoints will have the remote party ICE candidate addresses. It is at that point that the endpoints do a connectivity check by sending STUN messages to one another in an attempt to punch transport holes in the firewalls in order to establish media connectivity between peers.

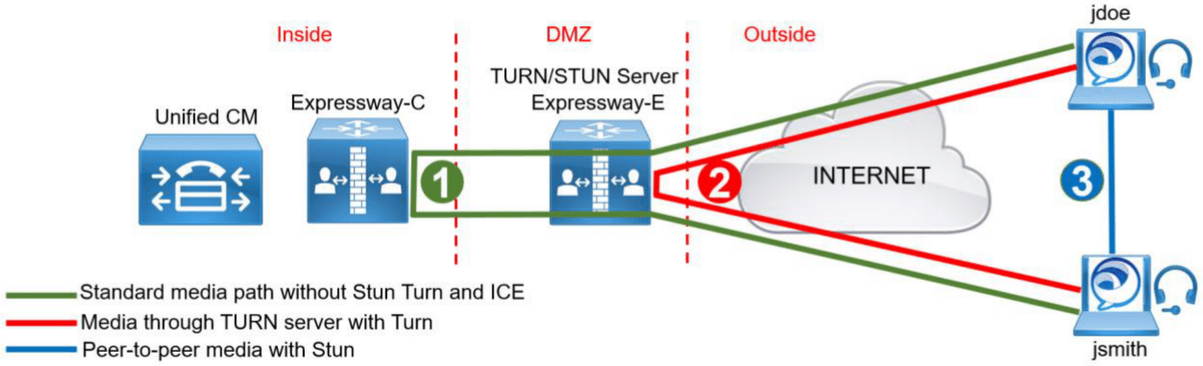
Inside
DMZ
Outside
jdoe
TURN/STUN Server
Expressway-E
Expressway-C
Unified CM
1
2
3
INTERNET
jsmith
Standard media path without Stun Turn and ICE
Media through TURN server with Turn
Peer-to-peer media with Stun

If the direct connection via the host address or server reflexive address cannot be achieved, then connectivity through the TURN server should be granted because the prerequisite for ICE to be set up successfully is for the TURN server to be reachable.

If there is an ICE path (via the host, server reflexive, or relayed address) the Turn Server will send a re-INVITE to the other endpoint, this time specifying only the chosen candidate. This re-INVITE will go through Expressway-E, Expressway-C, and Unified CM. But since Expressway-C and Expressway-E understand that the host-to-host connectivity check or the server reflexive to server reflexive connectivity check is successful, they will not put themselves into the media path, so the media will flow directly between the two endpoints.

If the only successful checks are those against the TURN server, the media will flow through Expressway-E using the TURN server ports.

If one of the two endpoints does not support ICE or does not satisfy the ICE prerequisites, the call will follow the legacy media path through Expressway-E and Expressway-C.

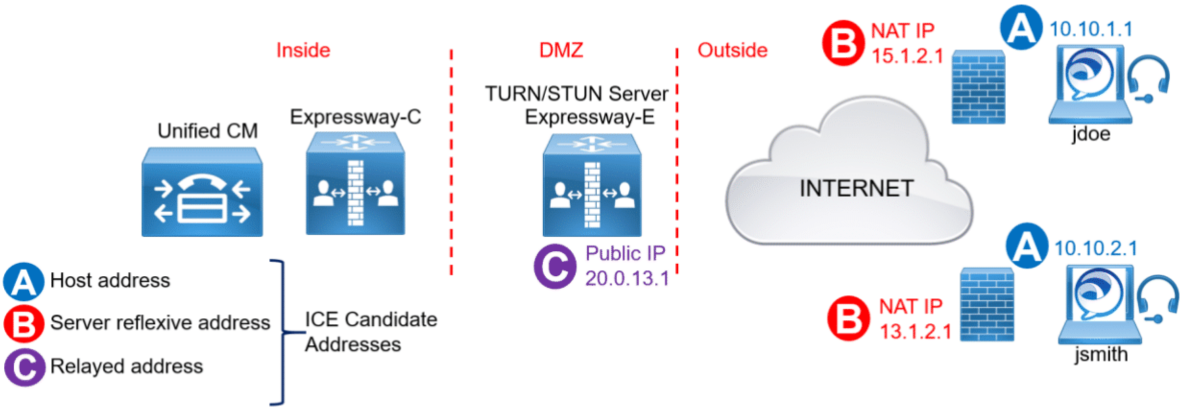

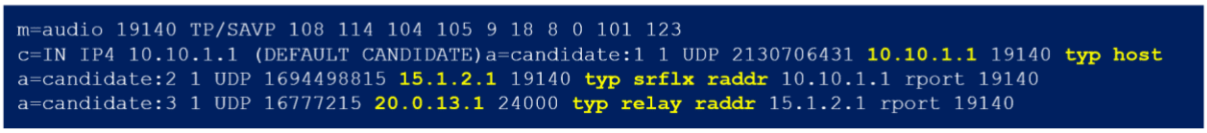

```
m=audio 19140 TP/SAVP 108 114 104 105 9 18 8 0 101 123
c=IN IP4 10.10.1.1 (DEFAULT CANDIDATE)a=candidate:1 1 UDP 2130706431 10.10.1.1 19140 typ host
a=candidate:2 1 UDP 1694498815 15.1.2.1 19140 typ srflx raddr 10.10.1.1 rport 19140
a=candidate:3 1 UDP 16777215 20.0.13.1 24000 typ relay raddr 15.1.2.1 rport 19140
```

When a user joins a space, before sending any communications to the Web Bridge.

1. The Cisco Meeting Server web app will first send a STUN request to the STUN server.
2. The STUN server replies with its Server Reflexive Candidate.
3. The Cisco Meeting Server web app also sends a TURN Allocation Request to the TURN server to request that the TURN server allocate it a dedicated port on the

TURN server. This address is referred to as the Cisco Meeting Server web app TURN Relay Candidate.

When the Call Bridge receives the call request, before replying, it will perform the same STUN and TURN Allocation Request to identify its Server Reflexive Candidate and Turn Relay Candidate.

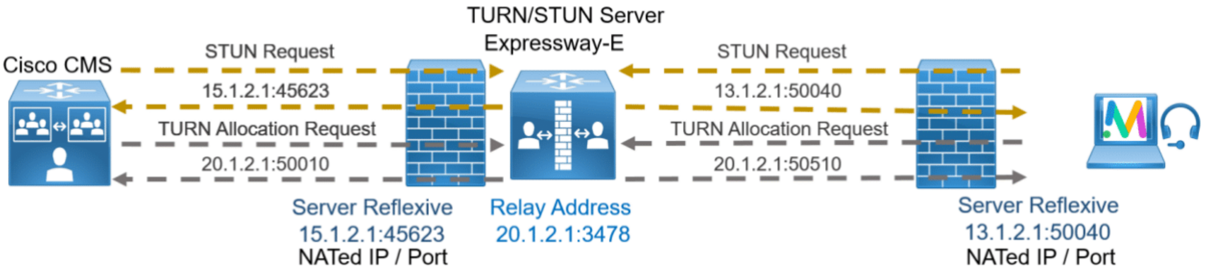

ICE Prioritizes media connection as follow:

1. Use the host candidate (host address if it succeeds.
2. Use the server-reflexive candidate if it succeeds.
3. Use the TURN server relayed candidate if all else fails.

ICE compliant devices will exchange three addresses with each other during call setup. Their local host address and port, their Server Reflexive candidate, and their TURN Relay candidates. These three addresses are the addresses that the other device can use to try to establish a media connection. ICE is the protocol that determines how devices should attempt to establish a media connection.

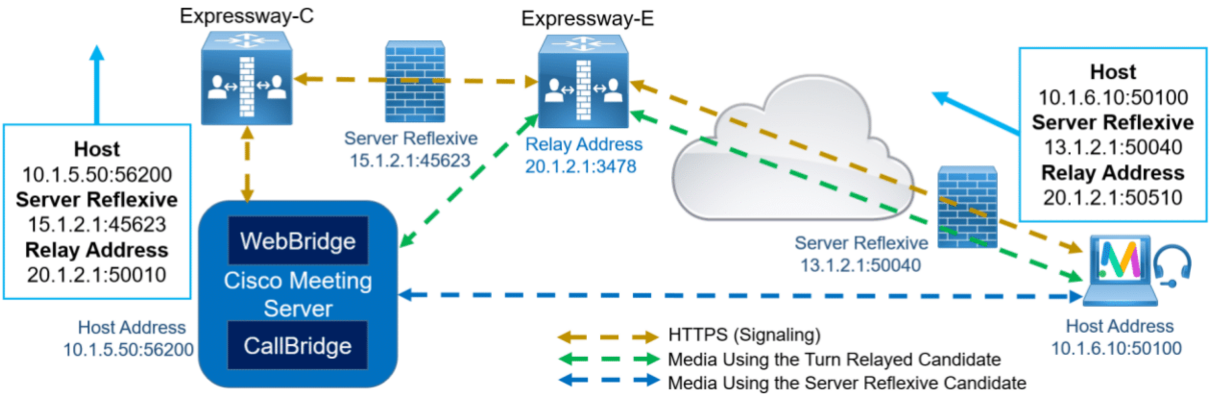

For users or guests to use the Cisco Meeting Server web app when outside the corporate network, access through the firewalls needs to be provided using the Cisco Expressways to proxy the WebRTC traffic from outside to in. The Cisco Expressway-E will proxy the HTTPS traffic back through an SSH tunnels to the Cisco Expressway-C. The Cisco Expressway-C will then route the traffic to an internal Web Bridge. For the Cisco Meeting Server web app to be able to access Web Bridge3.0 the Expressways must be running version X12.6 or later and if the Cisco Meeting Server web app is to use the Cisco Expressway-E TURN server then the TURN option key must be installed.

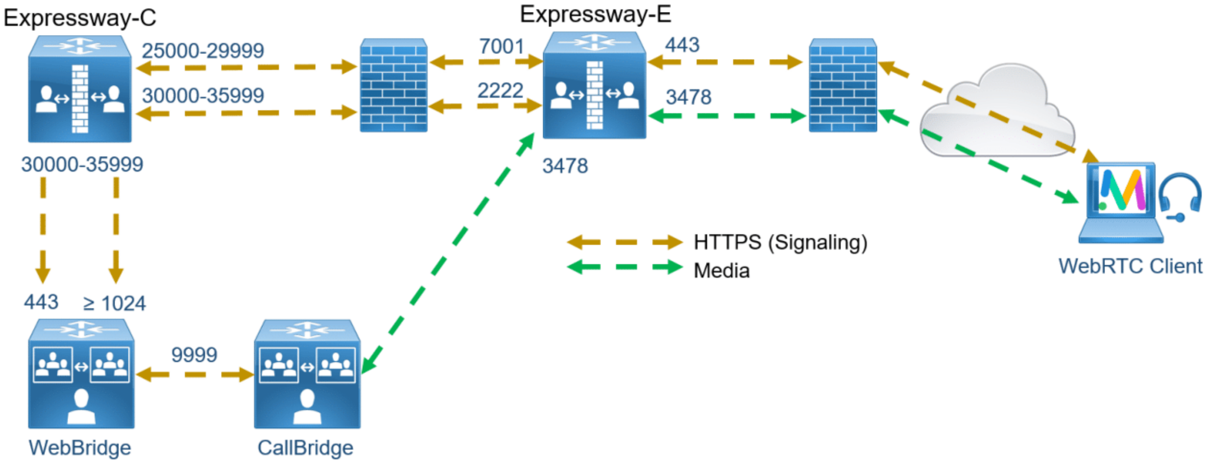

To enable the WebRTC traffic to be proxied, firewall ports will need to be opened from inside to out on the internal firewall to enable a connection to be established between the two Expressways through which the SSH tunnel is negotiated. The internal firewall must also allow the Call Bridge to connect to the TURN server media ports, as a minimum 3478 but preferably the higher TURN Relay Candidate ports as well.

The external firewall will need to allow HTTPS and media UDP traffic to connect to the Cisco Expressway-E. The HTTPS traffic connects to the Expressway on port 443 and media as a minimum will require access to the TURN Relay Address on the Expressway-E (3478) but as with the Call Bridge preferably the Cisco Meeting Server web app should be able to connect to the higher TURN media ports as well.