H.323 FORUM

Open Community Specification

The H.323 NAT/FW Traversal Solution

January 2014

International Multimedia Communications Consortium

Summary

This document describes the NAT/FW traversal solution defined by the ITU-T in Recommendations H.460.17, H.460.18, H.460.19, H.460.23, H.460.24 and H.460.26. This document is intended to provide an overview of issues presented by network address translators and firewalls, concepts employed in order to overcome those obstacles, and a high-level overview of the signalling flows that comprise the solution. This document does not introduce any new or different signalling procedures from those already defined in the aforementioned ITU-T Recommendations.

Author(s)

Simon Horne, Spranto

Acknowledgements

Jan Willamowius, Founder GNU Gatekeeper Project.

Table of Contents

The H.323 NAT/FW Traversal Solution

1 What is a NAT?

Wikipedia defines Network Address Translation (NAT) as "the process of modifying network address information in datagram packet headers while in transit across a traffic routing device for the purpose of remapping a given address space into another." Most systems using NAT do so in order to enable multiple hosts on a private network, such as a home WiFi network or small enterprise LAN, to access the Internet using a single public Internet Protocol (IP) address.

In a typical configuration, a local network uses one of the designated "private" IP address subnets (192.168.x.x, 172.16.x.x through 172.31.x.x, or 10.x.x.x), and a router on that network has a private address (such as 192.168.0.1) in that address space. The router is also connected to the Internet with a single "public" address assigned by an Internet Service Provider (ISP). As traffic passes from the local network to the Internet, the source address in each packet is translated on the fly from the private addresses to the public address. The router tracks basic data about each active connection, particularly the destination address and port. When a reply packet returns to the router, it uses the stored connection tracking data to map the packet to an internal address and port. This mapping that allows packets from the outside to be passed into the internal LAN is often referred to as a "pinhole". To a system on the Internet, the router itself appears to be the only source/destination for this traffic.

2 Why is NAT a problem?

NAT breaks the fundamental model of end-to-end Internet connectivity, introducing significant complications in how peer devices initiate and maintain connectivity with each other.

Some of the main issues are:

1. Devices can communicate from behind the NAT by sending information to an endpoint that is located on an external network. However, endpoints on an external network cannot send information to a device on the internal network without having an open pinhole through the NAT already in place.
2. The NAT replaces the IP packet's actual internal endpoint address and port with a public address and port, making it difficult for devices to know precisely where communication packets are actually transmitted or received.
3. Pinholes are temporary, usually closing after only 30 seconds or so of inactivity.

3 NAT is here to stay

NAT has become the basic component of the Internet, largely due to the fact that IP address space is so scarce. However, even with IPv6 address space so plentiful (2^{128} possible addresses), NAT devices will likely continue to be employed since they provide an additional level of security and control.

The need to allow users on local networks to seamlessly access the public Internet has become a fact of file and NAT provides that service. The challenge for peer devices and networks is to adapt to an

increasingly fragmented environment to provision and provide simple, direct and scalable connectivity.

4 How are NAT different and why is that important?

While NAT issues are not unique to H.323 and, in fact, are issues shared by every communication protocol on the Internet, H.323 presents special challenges, since it is often the case that both communicating devices are behind two different NAT devices. Further, it is ideal to allow those devices to communicate directly, or to at least allow media to flow point-to-point between the devices in order to allow the H.323 multimedia systems to scale.

The basic premise of point-to-point media is to have the devices send media packets to each other directly, even if one or both communicating devices are behind a NAT. If only one device is behind a NAT and one knows which one that is, then one could instruct the device on the public Internet to wait for the media to be established from the device behind the NAT first. In so doing, the device behind the NAT creates pinholes in the NAT so the external device can send media to the internal device.

The problem becomes more challenging when both devices are behind a NAT and is complicated even further by the fact that there are a variety of different "NAT types". The key to solving the double NAT issue is to understand how different NAT devices work and to understand that the behaviour of some NAT devices is more accommodating to achieve direct media. By understanding NAT behaviour, it is then possible to exploit that behaviour in order to enable direct point-to-point media flows.

These NAT types are classified according to behaviour and were first documented in Section 5 of IETF RFC 3489, referred to as the "Classic STUN" specification. RFC 3489 classified NAT implementations into two broad categories, Cone and Symmetric. The Cone NAT category is divided further into three subcategories: Full, Restricted and Port Restricted. In effect, there are four distinct NAT types one must consider and are discussed more thoroughly in the subsequent sections.

4.1 Full Cone NAT

Once an internal address is mapped to an external address, *any* external host can send packets to the mapped external address and port to reach the internal address.

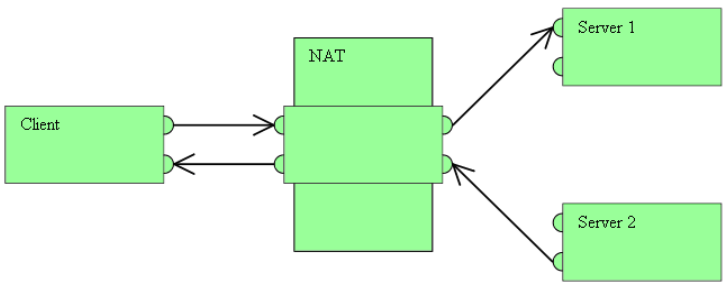

Figure 1 - Full Cone NAT

4.2 Restricted Cone NAT

Once an internal address is mapped to an external address, any other external host can send packets to the mapped external address and port to reach the internal address *only if* the internal address has first sent a packet to the same address using the same mapping (port is irrelevant).

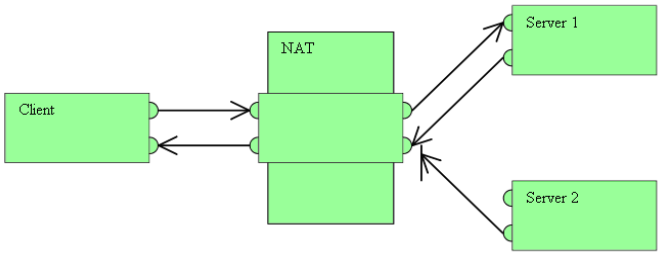

Figure 2 - Restricted Cone NAT

4.3 Port restricted cone NAT

Once an internal address is mapped to an external address, any other external host can send packets to the mapped external address and port to reach the internal address *only if* the internal address has first sent a packet to the same address and port using the same mapping.

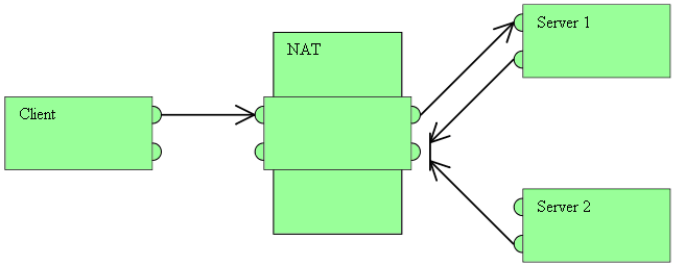

Figure 3 - Port Restricted NAT

4.4 Symmetric NAT

Each request from the same internal IP address and port to a specific external IP address and port is mapped to a unique external source IP address and port. If the same internal host sends a packet even with the same source address and port but to a different destination, a different mapping is used. Only an external host that receives a packet from an internal host can send a packet back.

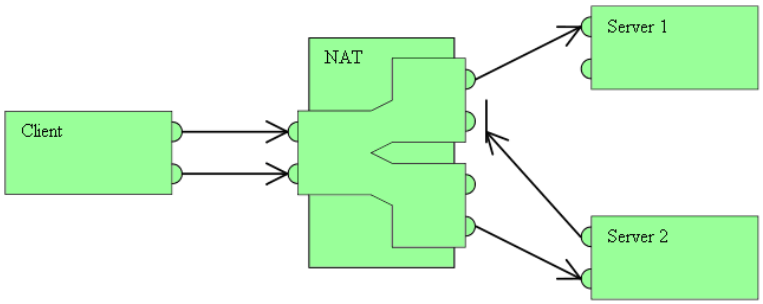

Figure 4 - Symmetric NAT

Classic STUN (RFC 3489) will work with Full Cone NAT but not with Symmetric NAT. That said, devices implementing the Classic NAT procedures may see mixed results with Restricted and Port Restricted NAT. For this reason STUN should not be considered a complete solution when dealing with traversing media across NAT.

5 Problems with Classic STUN (RFC 3489)

The success of any mechanism to allow for a direct media path is only as good as the reliability of the input information. Significant work has been conducted into verifying STUN classifications and has shown that under certain conditions such as heavy load or mapping arrangements, certain NAT devices may exhibit different behaviour, thus resulting in a change their classification. For this reason the provisions were removed from of the revised standard RFC 5389. This does not mean the results from the Classic STUN test are not useful. To the contrary, well-behaved NAT devices work perfectly following the procedures defined in Classic STUN. Further, by implementing mechanisms such as random port assignments and periodic STUN testing and reporting, the probability of enabling direct media flows with ill-behaved NAT devices is greatly improved. If a NAT device is ill-behaved and port mapping behaviour is unpredictable, then the NAT must be treated as if it is a Symmetric NAT device which may result in media having to be proxied.

6 The ALG Dilemma

ALG (or application layer gateway) is a component increasingly being deployed with NAT. It is a component that augments (or assists) NAT to provide application specific address and port translation. It is designed to help H.323 (and other) devices to traverse NAT by inspecting and altering information contained in existing H.323 messages as they pass through the NAT. It may alter addresses and ports information in registration, call signalling and automatically opening pinholes in the NAT to allow media flow. For existing non-NAT aware H.323 devices, it can be a valuable tool in assisting to seamlessly work with NAT. This advantage may also provide a hindrance in getting standard based H.323 NAT traversal mechanisms to work.

7 Advantages of the H.323 Direct Media solution

The H.323 direct media solution has several important advantages over other solutions.

1. **Compatible**. The solution utilises existing valid signalling messages and does not interfere in the RTP media stream making it suitable to use in existing deployed H.323 networks. Devices that do not support Direct Media are still valid inputs in the direct media calculation and pathways are formulated with existing devices in mind.
2. **Controllable**. Media pathways can be predefined. Decisions on where media or signalling is routed can be controlled allowing IT departments to suitably provision equipment, border elements and apply strict firewall rules in accordance with network security policy.
3. **Reliable**. The media pathway is formulated prior to the establishment of the call, so the Gatekeeper will know even before a call is made whether media connectivity is possible and since the pathway is not simply the first detected, media is established in a predictable manner.

4. **Efficient**. Since the media pathway is known when placing a call there are no media delays and the endpoints don't need to spend valuable time and resources at call establishment collecting information, waiting and probing for a media pathway.
5. **Interoperable**. Since Direct Media detects the presence of ALG etc. it will work with ALG and is not susceptible to call failure due to these types of devices.
6. **Extendable**. Support may be added for port mapping protocols like UPnP.

8 The six basic steps to establish Direct Media with NAT

There are 6 basic steps needed to achieve direct media through NAT.

1. **Detect.** When registering the H.323 Gatekeeper compares the supplied source address with the datagram packet source address of the endpoint. If they are different then there may be a NAT. The Gatekeeper tells the endpoint to perform a test to detect and if present, determine the type of NAT device employed.
2. **Interrogate.** The endpoint conducts the prescribed tests of NAT device type determination. This usually involves a classic STUN test.
3. **Report.** The endpoint reports the results of the NAT test to the Gatekeeper.
4. **Collect.** When initiating a call, the NAT information for both the calling and called parties is collected.
5. **Compare.** The Gatekeeper then compares the NAT information for both parties; a media transmission solution is formulated, which may be one that enables direct media flow or indirect media flow.
6. **Execute.** The parties are instructed to initiate connectivity in the prescribed manner. This usually involves the parties initiating the media flows in a particular sequence to achieve point-to-point media flows.

9 Overview of the H.323 NAT/FW traversal protocol specifications

The diagram below depicts the relationships between the various NAT/FW traversal protocols as they relate to previously-existing H.323 protocols specifications.

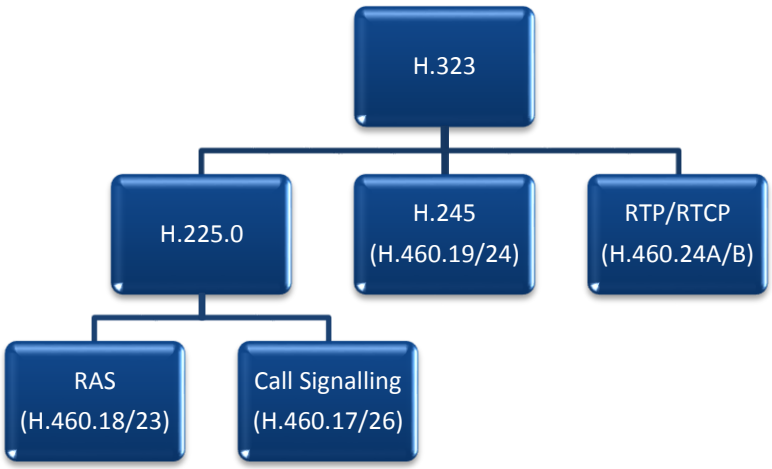

Figure 5 - H.323 Protocol Relationships

The standards H.460.17/.18/.19/.19M/.23/.24/.24A/B work together to provide a complete Direct Media solution for H.323.

In summary:

H.460.17 provides a mechanism to register and maintain a call listener on a single TCP socket.

H.460.18 provides a mechanism for the traversal of Call signalling across NAT.

H.460.19 provides a mechanism for the traversal of media across NAT by proxying media through an intermediary.

H.460.19 Multiplexing: All media is transported over the same known UDP media channel. This allowing strict firewalls to be applied.

H.460.19 Traversal Zone: A gatekeeper resides on the inside of the NAT and one on the outside forming a partnership to allow signalling and media to be proxied/tunnelled between them to traverse the NAT.

H.460.23 provides a mechanism to test the characteristics of any intermediary NAT.

H.460.24 provides a strategy using the information collected in H.460.23 to where possible avoid the need to proxy media via H.460.19 and have the media travel directly between the devices.

H.460.24 Annex A/B provide mechanisms to probe to realise direct connectivity where H.460.24 is irresolvable such as where the two devices are behind the same NAT or are behind 2 distinct difficult NAT.

H.460.26 Where no open UDP ports are available, media can be tunnelled over a single TCP call signalling connection. Used in conjunction with H.460.17 all registrations, signalling and media can be conducted over the same single persistent TCP connection on commonly available open TCP ports such as 80 or 443.

Important to note: The H.323 NAT Traversal solution is compatible with previously existing H.323 specification and will interwork with existing H.323 devices.

9.1 H.460.17

Registrations are made over a regular TCP call signalling channel to a gatekeeper on the outside of the network. This TCP connection is a permanent kept alive connection which is used to receive calls. Media may be carried over H.460.19 below or tunnelled over the same connection as defined in H.460.tunnel.

9.2 H.460.18

Upon registration a pinhole is opened for the Registration, Admission, and Status (RAS) messages. RAS messages are transmitted in the usual standard H.323 manner, but they contain extra signalling that indicates the usage of the H.460.18 standard. The endpoint maintains an open pinhole by transmitting "keep-alive" messages on a defined interval. When an incoming call arrives at the Gatekeeper, meaning the Gatekeeper received an H.225.0 Setup message intended for the internal endpoint, a series of messages are exchanged between the Gatekeeper and the internally registered endpoint, utilizing the existing pinhole as shown below.

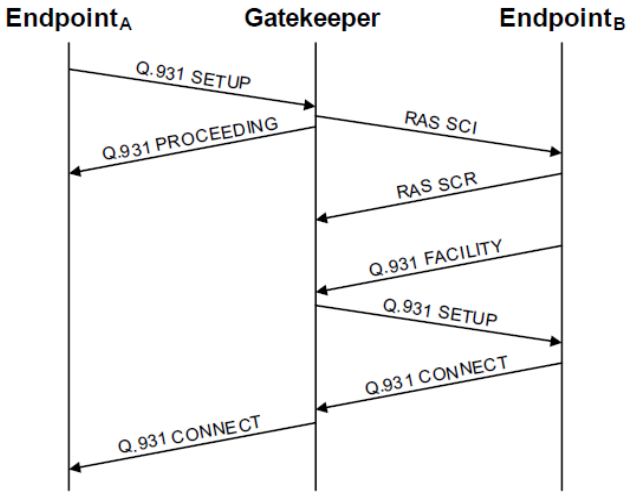

Figure 6 - H.460.18 Call Signalling Messages

When the Gatekeeper receives a Setup message, a RAS Service Control Indication (SCI) message is sent via the open RAS pinhole between the Gatekeeper and the internal Endpoint B. This message indicates to the internal endpoint that an incoming call is being attempted.

The Gatekeeper will then receive a Service Control Response (SCR) message back from the internal Endpoint B. The internal Endpoint B then sends a special H.225.0 Facility message, which opens a pinhole through the firewall. By recognizing the content of the message, the Gatekeeper can link this newly opened connection to the incoming call and forward the H.225.0 Setup message to the endpoint. The internal Endpoint B is then able to exchange any number of H.225.0 messages with the calling Endpoint A via this newly established communication channel.

9.3 H.460.19

As stated earlier, H.460.19 defines the procedures to allow media to be transported between a private network and a public network by utilizing a media proxy device. To allow media to traverse the NAT, empty UDP packets (keep-alive messages) are sent from the endpoint behind the NAT to the address of an intermediary (called an H.460.19 Server) on the outside of the NAT, thereby opening and maintaining a pinhole through the NAT. The pinhole IP and port become the destination the intermediary will use to forward media to reach the endpoint. The intermediary waits to receive these keep-alive messages and will then proxy media to and from the endpoint behind the NAT.

9.4 H.460.19 Multiplexing

This is an extension on H.460.19 which allows media (audio/video/content/camera controls) to transported over a single defined UDP port allowing strict firewall rules to be applied for video

conferencing. It can be used to proxy media between gatekeepers inside and outside the network or used as part of H.460.24 below to send media directly point to point without requiring a media proxy.

9.5 H.460.19 Traversal Zones

While H.460.19 can be used between an endpoint and a gatekeeper to get calls through firewalls, it is also possible to use gatekeepers, one resides on the outside of the NAT and the other on the inside. This has advantages in allowing existing non-H.460 endpoints inside the NAT to traverse and communicate with the public internet. Firewalls can be configured to completely deny all traffic to other internal IPs other than the internal gatekeeper.

In this scenario, a gatekeeper on the inside of the organization and another gatekeeper on the outside of the NAT form want are called a "traversal zone".

Usually the gatekeepers in the traversal zone use H.460.18/.19, where the internal gatekeeper acts as client to the external server. Features like H.460.19 RTP Multiplexing may be used as well to reduce the number of RTP ports. For inbound or outbound calls, the gatekeepers query each other much like a traditional H.323 neighbors do and the media and signalling is tunnel between the two gatekeepers to traverse the NAT/FW.

9.6 H.460.23

H.460.23 utilises H.460.18 for call signalling and then applies methodology to formulate media pathways to avoid, where possible, the need for proxy via an H.460.19 Server. In H.460.23, the Gatekeeper, on receipt of the RRQ message, detects whether there may be a NAT between the endpoint and itself by comparing the received RAS address and the apparent source of the RRQ message. It then instructs the endpoint to perform a STUN test to assist in the classification of the NAT. This information is then transmitted back to the gatekeeper (via another RRQ) to be used as input in the media pathway calculation of described H.460.24. Where a gatekeeper has not detected a NAT being present, the Gatekeeper supplies back to the endpoint a RAS address supplied to it in the RRQ for which the endpoint can use to compare what was sent to what the gatekeeper received. If there is a discrepancy then a H.323 intermediary (or ALG) may be acting on behalf of the endpoint. The detection of an ALG indicates that there may be an inconsistent behaviour in how the H.323 NAT Traversal mechanism functions and the Gatekeeper is advised to disable all NAT support for this endpoint and allow the ALG to perform the required NAT traversal.

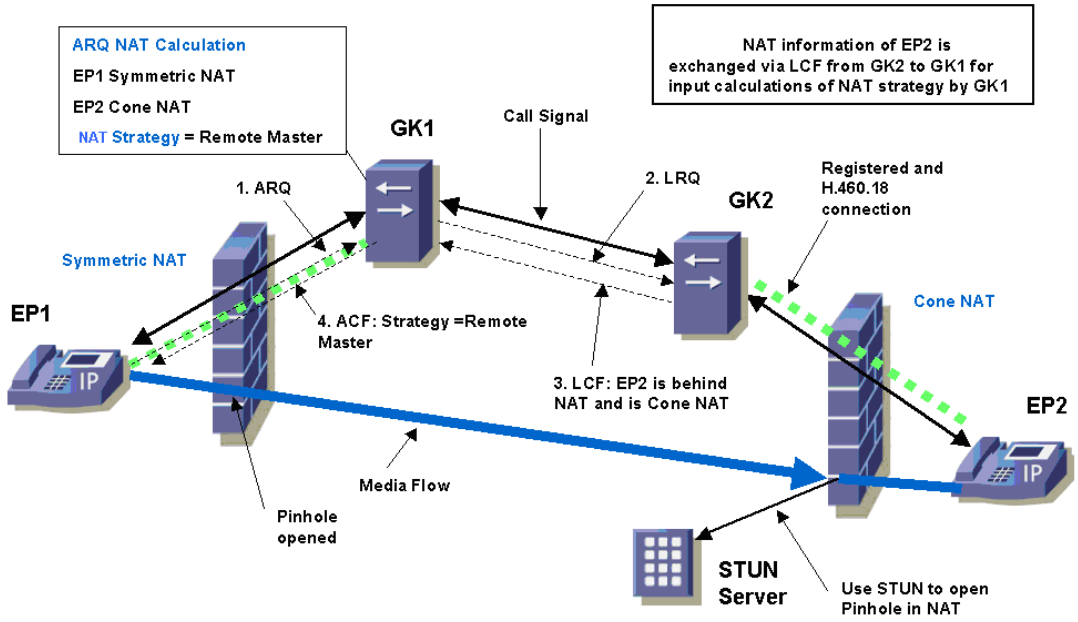

Figure 7 - Direct Media with Double NAT

9.7 H.460.24

When establishing a call the local Gatekeeper compares the NAT information of the calling endpoint with that of the called endpoint. This information includes the detected call signalling address (external NAT IP address), whether the endpoint is behind NAT, the NAT classification, whether the call must be proxied (in the case of a border Gatekeeper), and whether the endpoint supports direct media or Annex A. As the information is compared, an optimized media pathway is formulated. The results of the computation allow the Gatekeeper to instruct the endpoints to employ “Classic STUN” to open a pin holes or to initiate media before the other endpoint, as example. One or both endpoints may utilize an H.460.19 server in the case where a direct media pathway cannot be resolved, such as when both devices are behind a Symmetric NAT. If both endpoints are behind the same NAT, they may be instructed to use an H.460.19 server initially and then use H.460.24 Annex A to probe for a direct pathway. The calling endpoint is notified of the strategy ID to use (via an ACF) , configures itself to use that strategy and includes the strategy ID in the Setup message to the remote endpoint. The remote endpoint then configures itself according to the specified strategy, this enabling media establishment in the prescribed way. The computation also supports existing H.323 devices by designing the media pathway in such a way as to allow these devices to communicate with endpoints behind a NAT.

9.8 H.460.24 Annex A

H.460.24 Annex A may be used to probe for a direct media pathway when two devices exist behind the same NAT device. In the case of the endpoints residing behind the same NAT, it is impossible for the Gatekeeper to know whether the two endpoints are on the same network. One or both endpoints may reside on different segments separated by nested NAT devices and unable to transmit media directly. The endpoints are instructed to utilize an H.460.19 server initially, but once the media pathway has been established the endpoints send out RTCP probe packets to the other endpoints RTCP address to see if direct connectivity is possible. Once an endpoint receives a probe

packet, it responds and direct connectivity has been established. The endpoints will then initiate the media pathway in the same order as the verified connectivity path thereby traversing any internal NAT.

9.9 H.460.24 Annex B

H.460.24 Annex B may be used to probe for a direct media pathway when the two endpoints exist behind 2 distinct NAT devices. In cases where media cannot be established directly due to the NAT combinations used, media is first proxy via H.460.19 and Annex B is used to probe for and to open pathways through the NATs between the endpoints and redirect the media to flow directly.

The media server (proxy) using the detected IP address of the two devices instructs the endpoints to send RTCP probe packets to the detected IP address of the other endpoint in such a sequence to expose a direct connectivity pathway through the NATs. Once the connectivity pathway has been verified, the endpoints initiate the media pathway in the same order as the verified connectivity path thereby creating a pathway directly between the endpoints. Figure 8 describes the call flows of Annex B.

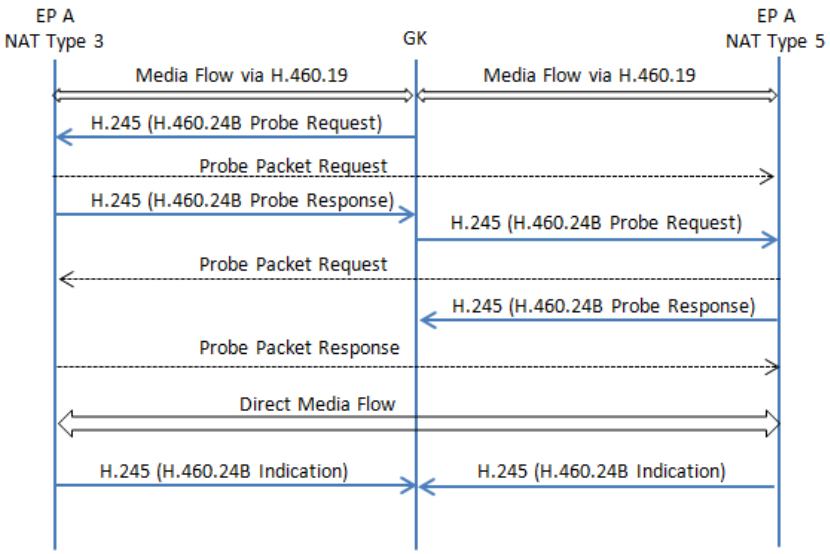

Figure 8 – H.460.24 Annex B Call Flows

9.10 H.460.26

H.460.26 allows media to be transported over an H.460.17 TCP connection thereby allowing everything (registration, signalling & media) to be carried over a single persistent call signalling connection. This allows traversal of the most secure NAT/FW by utilizing commonly open TCP port 80 and 443. Media packets are binary loaded into information messages and transmitted via the TCP connection.

10 Integrating Port Mapping systems like UPnP

UPnP IGD is a port mapping protocol commonly used in most home/small office NAT devices. It allows devices on the internal network to request ports mappings through the NAT. Ports are

opened on the external address and traffic is forwarded to the internal address. UPnP (and other port mapping systems) behave not dislike FULL Cone NAT for the purpose of Direct Media. When instructed these systems may be used as a replacement for STUN to open port mappings to establish direct media pathways.

11 Devices that support Direct Media

H.460.17/.18/.19/.19M/.23/.24/.24AB/.26 is supported in

H323plus www.h323plus.org as of version 1.24

GnuGk www.gnugk.org as of version 3.4

Spranto Softphone www.spranto.com