ADTRAN

NetVanta Unified Communications Technical Note

The Purpose of a SIP-Aware Firewall/ALG

Introduction

This technical note will explore the purpose of a Session Initiation Protocol (SIP)-aware firewall/Application Layer Gateway (ALG) when having SIP trunks to a service provider. Provided in this document is a brief description of the SIP and Network Address Translation (NAT) technologies. In addition, this document will highlight problems that NAT causes for SIP, and how SIP firewalls/ALGs resolve the problem. There are detailed examples throughout the document to help understand the nature of the problem.

The NetVanta Unified Communication Server, in combination with a SIP firewall/ALG, will provide SIP Trunking connectivity to various carriers/service providers. The information in this document is intended to be a general description of the technology specific to the NetVanta Unified Communications Server product and SIP firewall/ALG as it applies to communication with Internet Telephony Service Providers (ITSPs), SIP Peers, and teleworker solutions.

Session Initiation Protocol (SIP)

SIP is an application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences.

SIP invitations are used to create sessions that carry session descriptions, which allow participants to agree on a set of compatible media types. SIP makes use of elements called proxy servers to help route requests to the user's current location, authenticate and authorize users for services, implement provider call-routing policies, and provide features for users. SIP also provides a registration function that allows users to upload their current locations for use by proxy servers. SIP runs on top of several different transport protocols.

The SIP requests and responses are written in plain text within the datagram of the IP Header. Contained in the SIP requests and responses are the addresses of the source and the destination of the participants. These addresses are SIP URI's, which have a UserInfo and Host Address, and this host address can either be an IP address or a domain name. Therefore, the routing of SIP is done using IPv4 addresses at the Application layer and does not route at the Transport or Network layer. Below is an example of the SIP request in the Application layer which contains the addresses of the source and the destination of the participants.

```
INVITE sip:3177@192.168.8.178 SIP/2.0
Via: SIP/2.0/UDP 192.168.8.44:5060
From: "Test Ext3" <sip:2003@192.168.8.44>
To: <sip:3177@192.168.8.178>
Contact: <sip:2003@192.168.8.44>
Call-ID: caf4e42e93a33579-270-B2BUA
CSeq: 1 INVITE
Supported: timer
Allow:
INVITE,ACK,CANCEL,BYE,OPTIONS,REFER,NOTIFY,UPDATE
Max-Forwards: 69
Content-Type: application/sdp
Session-Expires: 900;refresher=uac
Content-Length: 561

v=0
o=MxSIP 0 1324999754 IN IP4 192.168.8.57
s=SIP Call
c=IN IP4 192.168.8.57
t=0 0
m=audio 3000 RTP/AVP 0 18 101
a=rtpmap:0 PCMU/8000
a=rtpmap:18 G729/8000
a=rtpmap:101 telephone-event/8000
```

As the addressing and routing of SIP are done at the Application layer, the biggest problem the SIP protocol now has is the disconnect between the IPv4 addressing and routing at the Application layer versus the IPv4 addressing and routing at the Transport and Network layers. Network Address Translation (NAT) occurs at the Transport and Network layers, thus the challenge.

Network Address Translation (NAT)

The purpose of a NAT firewall with a business is to provide the translation between a single public IP address on the WAN and multiple private IP addresses for all of the workstations, Servers and other IP equipment within the LAN. The router running NAT should never advertise the LAN network addresses to the WAN network backbone. Only the networks with global addresses may be known outside the router. However, global information that NAT receives from the border router can be advertised in the LAN network the usual way. Typical or traditional firewalls apply NAT to the TCP/IP protocol at the Transport and Network layers.

NAT's basic operation is as follows: The network addresses inside a private domain can be reused by any other private domain. For instance, a single Class A address could be used by many private domains. At each exit point between a private domain and the public WAN backbone, NAT is installed. If there is more than one exit point it is of great importance that each NAT has the same translation table.

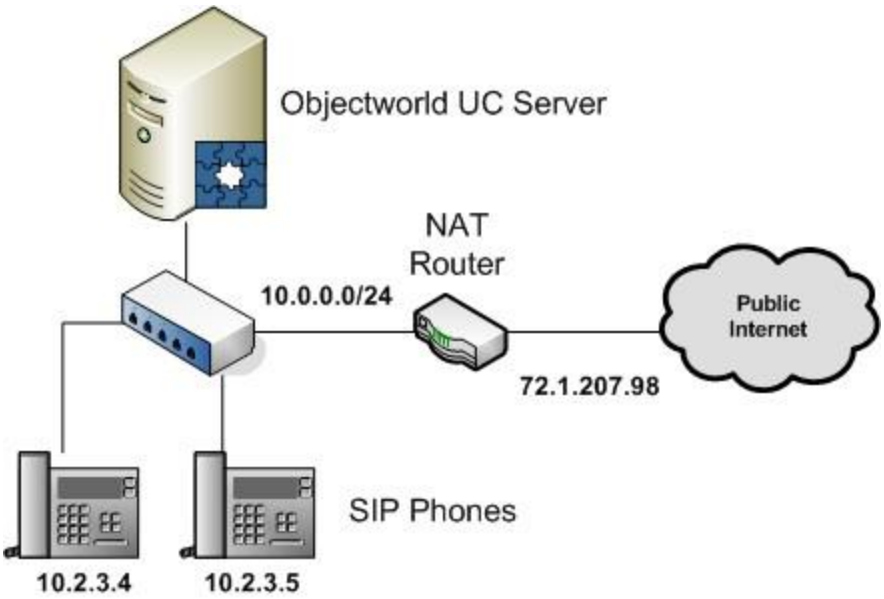

The following are two examples, one for incoming (before and after), the other for outgoing (before and after) of the behavior of NAT on the TCP/IP headers of an Ethernet packet through a NAT router/firewall, and then the same packet after NAT was applied.

Example 1. Incoming TCP/IP Packet before NAT:

```
⊞ Frame 40 (722 bytes on wire, 722 bytes captured)
⊞ Linux cooked capture
⊟ Internet Protocol, Src: 72.149.98.250 (72.149.98.250), Dst: 72.1.207.104 (72.1.207.104)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 706
    Identification: 0x65e7 (26087)
  ⊞ Flags: 0x00
    Fragment offset: 0
    Time to live: 239
    Protocol: UDP (0x11)
  ⊞ Header checksum: 0xa980 [correct]
    Source: 72.149.98.250 (72.149.98.250)
    Destination: 72.1.207.104 (72.1.207.104)
```

Incoming TCP/IP Packet after NAT:

```
⊞ Frame 41 (825 bytes on wire, 825 bytes captured)
⊞ Linux cooked capture
⊟ Internet Protocol, Src: 72.149.98.250 (72.149.98.250), Dst: 10.10.8.178 (10.10.8.178)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 809
    Identification: 0x0000 (0)
  ⊞ Flags: 0x04 (Don't Fragment)
    Fragment offset: 0
    Time to live: 64
    Protocol: UDP (0x11)
  ⊞ Header checksum: 0x5bae [correct]
    Source: 72.149.98.250 (72.149.98.250)
    Destination: 10.10.8.178 (10.10.8.178)
```

Notice how the NAT firewall has changed the Destination IP addresses of the IP Header to ensure that traffic will arrive at the correct destination.

Example 2. Outgoing TCP/IP packet before NAT:

```
⊞ Frame 42 (545 bytes on wire, 545 bytes captured)
⊞ Linux cooked capture
⊟ Internet Protocol, Src: 10.10.8.178 (10.10.8.178), Dst: 72.149.98.250 (72.149.98.250)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 529
    Identification: 0x7d56 (32086)
  ⊞ Flags: 0x00
    Fragment offset: 0
    Time to live: 127
    Protocol: UDP (0x11)
  ⊞ Header checksum: 0xe06f [correct]
    Source: 10.10.8.178 (10.10.8.178)
    Destination: 72.149.98.250 (72.149.98.250)
```

Outgoing TCP/IP packet after NAT:

```
⊞ Frame 43 (424 bytes on wire, 424 bytes captured)
⊞ Linux cooked capture
⊟ Internet Protocol, Src: 72.1.207.104 (72.1.207.104), Dst: 72.149.98.250 (72.149.98.250)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 408
    Identification: 0x0000 (0)
  ⊞ Flags: 0x04 (Don't Fragment)
    Fragment offset: 0
    Time to live: 64
    Protocol: UDP (0x11)
  ⊞ Header checksum: 0x7f92 [correct]
    Source: 72.1.207.104 (72.1.207.104)
    Destination: 72.149.98.250 (72.149.98.250)
```

Notice how the NAT firewall has protected the private LAN address and inserted the public address of the NAT firewall.

Application Layer Gateway (ALG)

An Application Layer Gateway, also known as ALG or Application-Level Gateway, consists of a security component that augments a firewall or NAT employed in a computer network. It allows legitimate application data to pass through the security checks of the firewall that would have otherwise restricted the traffic for not meeting its limited filter criteria.

Deep packet-inspection of all the packets handled by ALGs over a given network makes this functionality possible. For instance, an ALG can allow firewall traversal with SIP. If the firewall has its SIP traffic terminated on an ALG then the responsibility for permitting SIP sessions passes to the ALG instead of the firewall. An ALG can solve another major SIP headache: NAT traversal. Basically a NAT with a built-in ALG can re-write information within the SIP messages and can hold address-bindings until the session terminates.

In the context of SIP and SIP Trunking, an ALG may offer the following functions:

- Allows SIP clients, such as SIP phones, the UC server and SIP gateways to use dynamic UDP ports to communicate with the known ports used by the SIP Trunking service provider and other SIP client applications, even though a firewall configuration may allow only a limited number of known ports. In the absence of an ALG, either the ports would get blocked or the network administrator would need to explicitly open up a large number of ports in the firewall, rendering the network vulnerable to attacks on those ports.

- Converting the network layer address information found inside the SIP protocol payload between the addresses, acceptable by the hosts on either side of the firewall/NAT. This aspect introduces the term "gateway" for an ALG.

- Recognizing specific SIP protocol requests and responses and offering granular security controls over them.

- Synchronizing between multiple media streams and sessions between SIP clients. For example, a SIP Trunking phone call may use multiple and separate connections for a variety of purposes. During these phone calls the Voice RTP media is consistently in progress, while the SIP signaling control connection may remain idle. An ALG can prevent the control connection getting timed out by network devices before the voice call is complete.

Problems that NAT Causes for SIP

In the first section the SIP protocol was discussed, and we learned that the SIP protocol resides in the Application layer. The SIP addresses are formed as SIP URI's, with a Userinfo@host, where the host can be an IPv4 IP address or domain name. Thus SIP requests, responses and routing are controlled at the Application layer. In the next section NAT was discussed, and we learned that NAT provided an address translation between private LAN addresses and public WAN addresses. This translation was done at the Transport and Network layer within the Internet Protocol Header.

Typical firewalls do not apply NAT to the Application layer. As SIP is an Application layer protocol, the IPv4 addresses and domain resolution are not translated for Application layer routing. SIP traffic cannot effectively traverse these traditional enterprise firewalls and NAT devices, and as a result, the firewall/NAT device incorrectly routes all SIP traffic, which includes VoIP. Thus when a SIP phone call attempts to traverse a typical firewall, although the TCP/IP addressing NAT is correct, the IP addresses within the SIP protocol information are not corrected properly. As a result, when a far end WAN device receives a SIP request the SIP addresses are the private IP addresses of the SIP device behind the typical firewall. These private IP addresses are not routable back to the original source.

Here is an example of the before and after effects of a NAT firewall when a SIP Packet is translated incorrectly.

Below is the outgoing SIP Packet before the NAT firewall. Notice the Internet Protocol Source addresses, and in the SIP Protocol From: and Contact: headers, *all* are private LAN IP addresses.

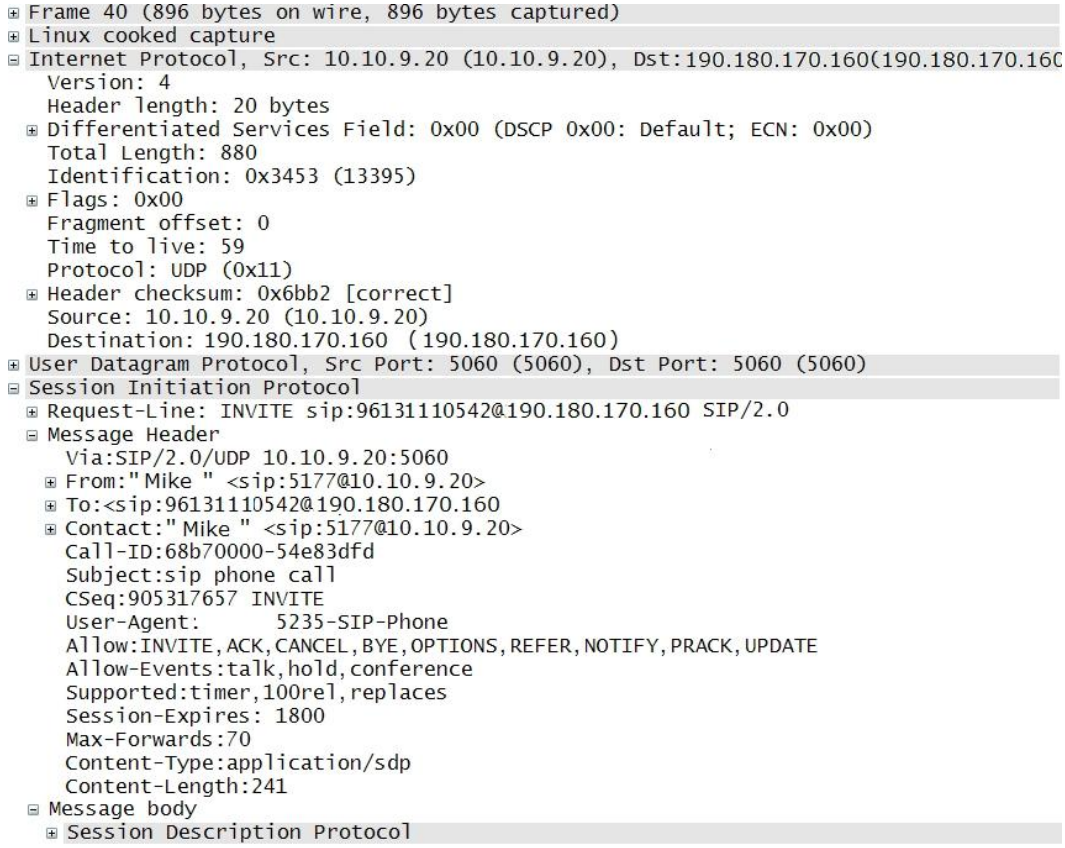

Below is the SIP Packet after it has been translated by the NAT firewall. Notice the Internet Protocol source addresses have been changed correctly, but nothing has been translated in the SIP addresses. This is an example of when NAT does not translate the Application layer of the SIP protocol, and the private addresses remain the same. The result is, when the SIP User Agent receives this SIP request, their SIP responses will be to the private LAN addresses listed within the SIP message. Thus the messages will not be directed to the public WAN IP address of the firewall. All of the SIP responses are effectively being sent to the wrong location.

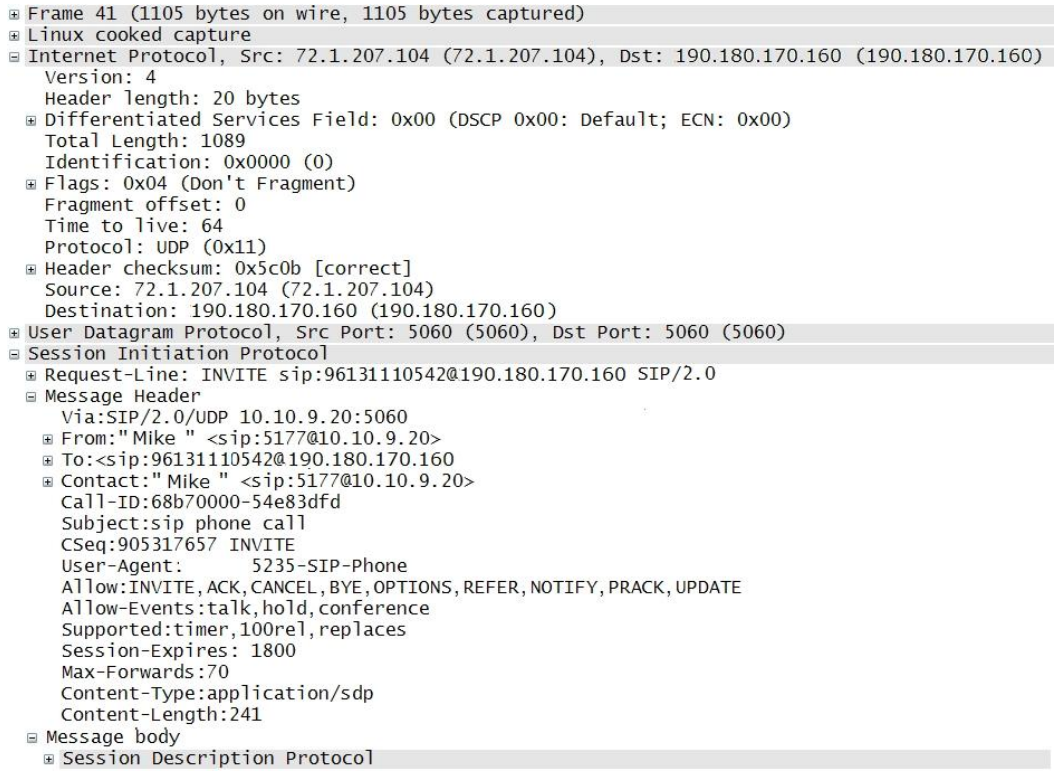

Why a SIP-Aware Firewall/ALG is needed

“SIP-aware” firewalls contain a SIP ALG above the typical firewall NAT capabilities. This SIP ALG component allows the additional functionality of the traversal of the IP addresses within the SIP protocol. This ALG capability within the firewall will inspect each SIP packet, and substitute the private IP addresses with the public IP address of the firewall/ALG.

The following is an example of the before and after effects of a SIP–aware firewall/ALG with NAT, when a SIP Packet is translated correctly.

First, the outgoing SIP Packet before the NAT firewall; notice the Internet Protocol Source addresses and in the SIP Protocol From: and Contact: headers, *all* are private LAN IP addresses.

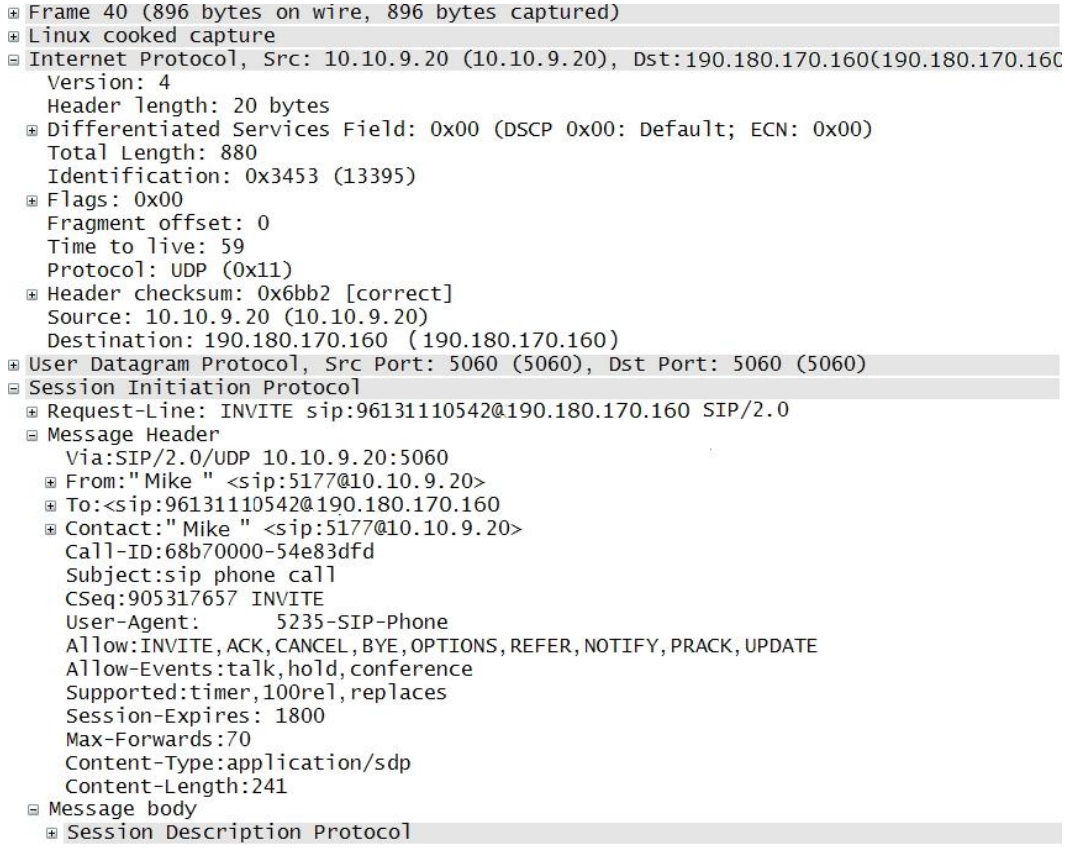

Below is the SIP Packet after it has been translated by the NAT firewall. Notice the Internet Protocol source addresses have been changed correctly, and in addition, the From: and Contact: headers have been translated in the SIP addresses. This is an example of when the ALG translates the Application layer of the SIP protocol and the private addresses are changed to the public WAN IP address. The result is, when the SIP Peer device receives this SIP request, their SIP responses will be to the public WAN IP address listed within the SIP message. Thus the response messages will be directed back to the public WAN IP address of the firewall. All of the SIP responses are being sent to the correct location.

```
⊞ Frame 41 (1105 bytes on wire, 1105 bytes captured)
⊞ Linux cooked capture
⊟ Internet Protocol, Src: 72.1.207.104 (72.1.207.104), Dst: 190.180.170.160 (190.180.170.160)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 1089
    Identification: 0x0000 (0)
  ⊞ Flags: 0x04 (Don't Fragment)
    Fragment offset: 0
    Time to live: 64
    Protocol: UDP (0x11)
  ⊞ Header checksum: 0x5c0b [correct]
    Source: 72.1.207.104 (72.1.207.104)
    Destination: 190.180.170.160 (190.180.170.160)
⊞ User Datagram Protocol, Src Port: 5060 (5060), Dst Port: 5060 (5060)
⊟ Session Initiation Protocol
  ⊞ Request-Line: INVITE sip:96131110542@190.180.170.160  SIP/2.0
  ⊟ Message Header
      Via: SIP/2.0/UDP 72.1.207.104:5060
      Via: SIP/2.0/UDP 72.1.207.104:5060
    ⊞ From:" Mike " <sip:5177@72.1.207.104 >
    ⊞ To:<sip:96131110542@190.180.170.160
    ⊞ Contact: <sip:ec_jX56w64LzCMfqhNSVcf0lyQw1ZLeBMiQk00SmEABw.@72.1.207.104>
      Call-ID:68b70000-54e83dfd
      Subject:sip phone call
      CSeq:905317657 INVITE
      User-Agent:      5235-SIP-Phone
      Allow:INVITE,ACK,CANCEL,BYE,OPTIONS,REFER,NOTIFY,PRACK,UPDATE
      Allow-Events:talk,hold,conference
      Supported:timer,100rel,replaces
      Session-Expires: 1800
      Max-Forwards: 69
      Content-Type:application/sdp
      Content-Length: 245
      Record-Route: <sip:4a7d34ac501e33ba@72.1.207.104;lr>
  ⊟ Message body
    ⊞ Session Description Protocol
```

The SIP-aware firewall/ALG will allow the network traversal of SIP Trunking calls to various carriers/service providers from the UC server. The firewall controls both incoming and outgoing SIP communications and routes the SIP communication to the intended users and devices. The advantage of the SIP-aware firewall/ALG is that it will allow all voice traffic as well as data traffic to traverse the enterprise firewall/NAT/ALG. NAT firewall with ALGs enables enterprises to utilize SIP trunks to ITSPs while continuing to manage data traffic.

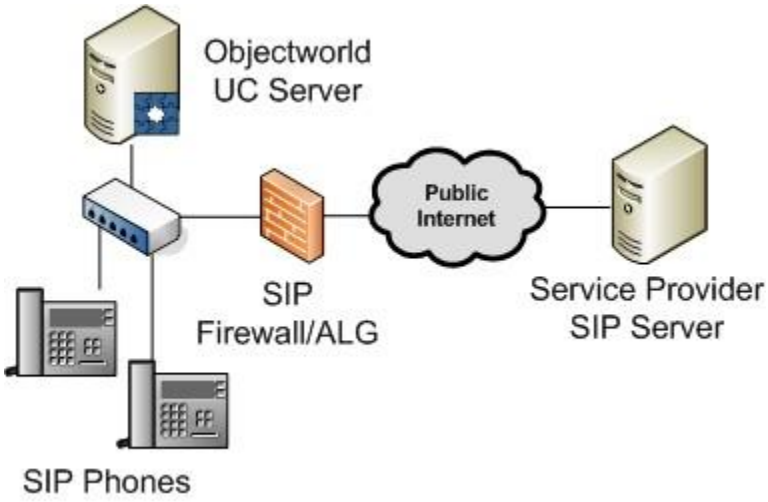

The general rule is NAT breaks SIP, unless you have a SIP ALG.