

# The SIP Protocol and Firewall Traversal

## White paper by:

Lisa Hallingström and Janne Magnusson, 2001-10-11

**Ingate Systems AB**, [www.ingate.com](http://www.ingate.com), [info@ingate.com](mailto:info@ingate.com), tel: +46 8 7208930

Lars Berggren and Karl Erik Ståhl, 2001-12-02

**Intertex Data AB**, [www.intertex.se](http://www.intertex.se), [info@intertex.se](mailto:info@intertex.se), tel: +46 8 6282828

Revision date: 14 October 2005

**SIP (Session Initiation Protocol) is an Internet protocol for setting up sessions between users. It can be used for IP Telephony (voice and video), Presence, Instant Messaging, Conferencing and more. SIP is expected to have the same impact on Internet usage for IP Communication as SMTP had for email and HTTP had for the Web.**

**A powerful driving force for SIP is that Microsoft has announced that all future real time communication (RTC) will be based on the SIP standard. Windows Messenger, which can be downloaded at no charge, already has a SIP mode that provides the user with telephony, voice, video, presence, and instant messaging.**

**Protocols like SIP, setting up sessions between users present special problems when the users are residing on networks with a private address space (NAT) or when the users are protected by a firewall. Intertex (SOHO products) and Ingate (Enterprise products) have implemented unique SIP transparent Firewalls and NATs, by incorporating a SIP proxy and a SIP registrar dynamically controlling the firewall.**

## Why Use SIP?

The common usage of Internet has been to communicate between a host (a PC) and a server available on the public Internet. However, communication between endpoints has several highly usable applications. Telephony, fax, presence, video conferencing, instant messaging (chat), conferencing, file exchange are typical examples. A number of proprietary and other protocols have emanated for these types of services over the Internet.

SIP (Session Initiation Protocol) is a general standardized IETF protocol (RFC 3261) for these types of applications. It not only sets up sessions between endpoints; SIP also locates the persons and sets up person-to-person sessions. Unlike other protocols performing similar functions, SIP is a generalized protocol designed for scalability, for global interconnectivity and uses Internet services already available, e.g. DNS.

The other most known protocol for VoIP or IP Telephony is H.323. However, H.323 has become huge and complex, is limited in functionality, has demonstrated poor interoperability and scalability, and has therefore only been implemented in isolated islands, but most often without interoperability on the IP level between the different operators. The SIP standard is now taking over from the earlier H.323 and has grown rapidly during the last years.

## SIP and Firewalls

Initiating communication from the public Internet to a device on a private LAN is by definition a complex task for a firewall to handle. Supporting media streams (voice and video) transported over separate ports negotiated in the session setup, further adds to the complexity. Furthermore, the location of individual users on a private protected LAN also has to be

handled. Transparent SIP traversal through firewalls and NATs requires specific handling of these issues.

SIP initiates multimedia sessions carried over other protocols, for example RTP. This means that when a SIP session has been started, various other protocols are used as well, usually transmitted over TCP or UDP and using an arbitrary, dynamically assigned, port number. This is a problem for a firewall that generally opens up certain protocols and ports in advance. But since the ports are dynamically assigned, which ports are to be opened? Current firewalls do not support tunneling of dynamically allocated media streams.

Handling the SIP media streams through a firewall by opening a wide range of ports is, of course, unacceptable from a security point of view. Instead, a firewall that understands SIP can open up the ports for the right protocols just when the SIP traffic needs it.

A SIP message is destined for a specific user and therefore includes IP addresses of the session participants in header fields. This is a problem if a SIP session should be established through a firewall using NAT. The IP address on the hidden side (which appears in the SIP headers) will not be the same as the one that the clients on the outside should use.

This problem can be resolved in several ways. One is to make the SIP endpoints aware of NAT and let them discover what IP addresses and port numbers they can use. This has a several drawbacks. First, every SIP endpoint application needs to be modified – new software has to be written for them. Second, the NAT discovery is problematic since different NATs work in different ways. Thirdly, this solution works only for NAT since the firewall blocks IP packets not allowed to pass. New protocols has been developed to address this issue (STUN, TURN, ICE).

A more general way is to integrate a SIP proxy server that rewrites the SIP messages using proper global and local addresses on the respective public and private sides. The proxy server can also control the NATing firewall and open up for SIP initiated media streams.

A third problem is to find the specific user on a private LAN using private IP addresses. On such LANs, all SIP-signaling is normally addressed to port 5060 on the single global IP address of the Firewall. Thus, to handle incoming calls, the firewall must cooperate with a location function having a register of all SIP users inside the firewall. The location function can either be separate or, as in the Intertex and Ingate products, be integrated in the firewall in form of a SIP registrar.

Handling these three issues in a firewall is necessary to allow users to use SIP standard applications as they are designed today, no matter with whom they communicate.

## **SIP Within the Intertex and Ingate Firewalls**

The Intertex and Ingate firewalls relay SIP traffic and keep track of which ports should be used for NAT, enabling machines on different sides of a firewall to send and receive media streams just as if there was no firewall at all. They implement the SIP protocol as described in RFC 3261 including:

- SIP user registration
- SIP header rewriting for NAT addresses
- SIP request relaying
- SIP user authentication

### ***SIP User Registration - the SIP Registrar***

For outgoing SIP requests, only a SIP proxy is needed. Incoming SIP requests however, need some device that keeps track of the local users so that the request can be relayed to the right machine and user. This is particularly important when NAT is used, since no SIP registrar on the outside will know the IP addresses on the internal networks.

The Intertex and Ingate firewalls manage user registrations, allowing the SIP module to keep track of where to send incoming session requests. It is also possible to make restrictions on which users are allowed to register and/or from where they can register. You can also monitor which users are currently registered.

The integrated registrar can be the main registrar or only be a passive registrar, monitoring and storing information from registration done at an outside registrar. In both cases, the registrar keeps the required information to locate users inside the firewall.

Each registration has a timeout after which it is removed unless the client extends it.

### ***SIP Header Rewriting - Hiding Information***

The SIP proxy server in the firewalls handles the SIP-NAT combination by rewriting the SIP headers to give them the right IP addresses. This can be done, as it is the firewall itself that provides the NAT addresses.

### ***SIP Request Forwarding - the SIP Proxy***

The Intertex and Ingate firewalls relay SIP requests for a user through the firewall to the device (computer, telephone, etc) from which the user has registered. In this way the user can be contacted by other SIP users behind other interfaces of the firewall. If no registration exists for a user, the firewall returns a SIP error message when someone tries to contact him.

The firewall rules are temporarily changed to let the media streams through. The user can monitor which sessions are currently active.

A SIP session media stream can consist of many different MIME types. The clients agree on what MIME types they both understand and can handle. On top of this, the user can choose what MIME types the firewall should forward. Common MIME types like text/plain and text/html should probably be forwarded, but you can block types that you don't want to allow through. The user can also restrict the number of concurrent media streams for a session.

Users can choose whether to process SIP requests in the firewall or forward them to an external, outbound SIP proxy. This can be useful if you want the firewall to keep a registry of local users only, and forward (and NAT if needed) all requests for external users to the external proxy (which in turn probably forwards requests to other proxies).

### ***SIP User Authentication - Who Are You?***

SIP user authentication can be performed by Digest Access Authentication like in HTTP. This is an authentication method that uses checksums, which means that the required Shared Secret is never sent in the clear. The Digest method used is *auth*, which allows for NAT as it does not use the IP addresses in the headers of the message as part of the checksum.

## More About the Packet Handling

### Incoming SIP Requests

This is what a NATing Intertex or Ingate firewall will do for incoming SIP requests:

- Catch any packet on port 5060 going through the firewall or to the firewall itself.
- Inspect the headers and the body of the packet.

The body will probably contain SDP information about which media streams the sender wishes to receive, and which IP address and ports it wants to use for this. Some of this information will be stored in the firewall, depending on the characteristics of the SIP request (stateful transaction).

- Forward the packet.

The packet will be forwarded to the device that the user has registered from.

- Allocate a port on one of the firewall's outside IP addresses. The firewall has a set of ports dedicated for SIP traffic.

- Replace all IP information in the headers and body with the new IP address and port.

Also replace the Via header of the client with one containing information about the firewall.

- Set up a firewall rule to let the media stream through.
- Intercept all following packets for this media stream and rewrite their IP headers.
- Remove the firewall rule when the session is terminated.

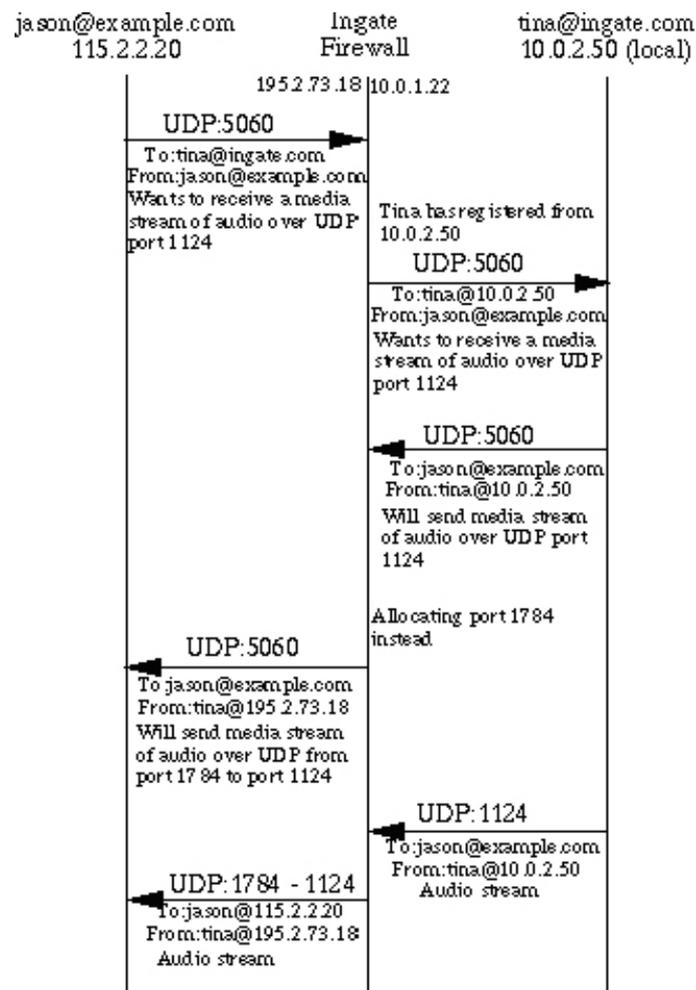


Fig. 1. Incoming SIP request.

## ***Outbound SIP Requests***

This is what a NATing Intertex or Ingate firewall will do for outbound SIP requests:

- Catch any packet on port 5060 bound to or through the firewall.
- Inspect the headers and the body of the packet.

Many headers, such as the Via header and the routing headers, will contain the IP address of the sender. Since the firewall uses NAT, this information will have to be rewritten as the packet is passing through the firewall.

The body will probably contain SDP information about which media streams the sender wishes to receive, and which IP address and ports it wants to use for this. Some of this information will be stored in the firewall, depending on the characteristics of the SIP request (stateful transaction).

- Allocate ports on one of the firewall's outside IP addresses. The firewall has a set of ports dedicated for SIP traffic. Replace all IP information in the headers and body with this IP address and port.

It will also replace the Via header of the client with one containing information about the firewall itself.

- Forward the packet.

If an outbound SIP proxy is defined, the packet will be sent to this address.

If no outbound SIP proxy is defined, the Intertex Firewall will try to locate a SIP server for the receiver by asking in DNS for SRV and/or A records for the SIP server. The packet will then be sent to this SIP server.

If this fails, the firewall will try to resolve the domain itself in DNS and send the packet there. If the domain cannot be resolved, an error is reported to the client.

- Set up a firewall rule to let the media stream through.
- Intercept all following packets for this media stream and rewrite the IP headers.
- Remove the firewall rule when the session is terminated.

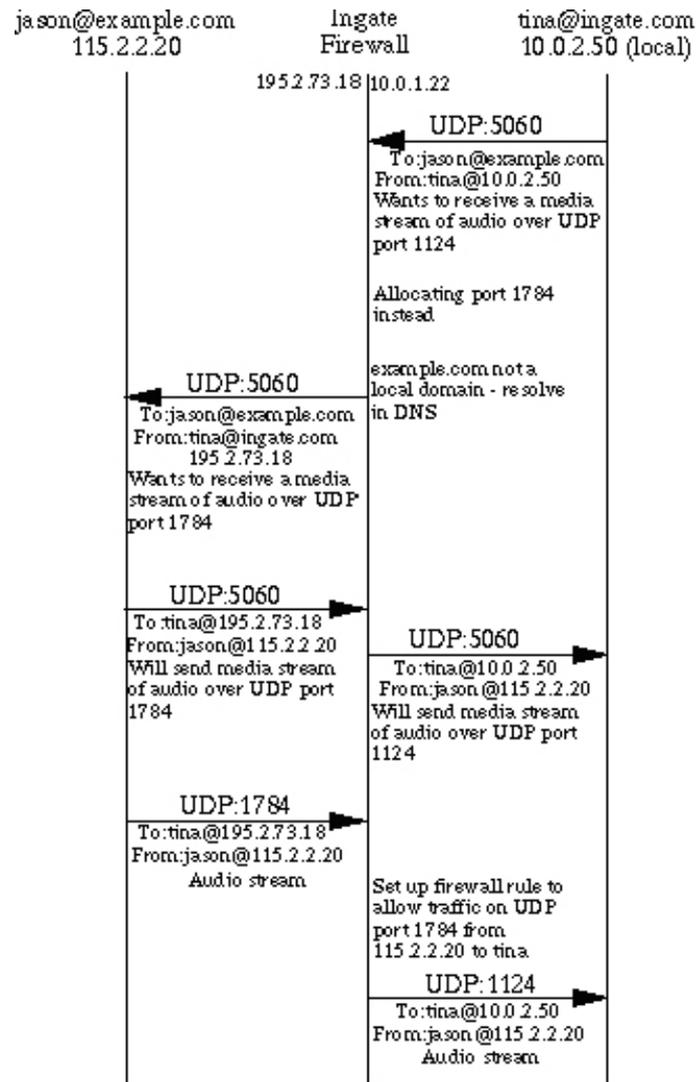


Fig. 2. Outbound SIP request.

## Implemented Parts of the SIP Protocol

- Registrar server for acceptance of REGISTER requests
- Proxy server, performing address mapping on any type of SIP request
- Stateful or stateless transaction, depending on the properties of the SIP request
- Parallel forking of a request to multiple destinations
- Digest authentication
- Session timer
- Record routing
- Via hiding
- DNS SRV record support
- DNS NAPTR record support
- Support for Transport Layer Security (TLS)
- ENUM