Addressing survivability and scalability of SIP networks by using Peer-to-Peer protocols

By Adrian Georgescu

September 5, 2005

Background

SIP (Session Initiation Protocol) is a protocol standardized by the Internet Engineering Task Force (IETF) for voice, video or other interactive sessions over the Internet. SIP is based on a client-server model where each SIP entity depending on the context behaves either like a client (UAC) or a server (UAS). Client-server models have to overcome one basic problem, the scalability combined with high availability, a problem that can be solved by using Peer-to-Peer techniques.

Pure Peer-to-Peer solutions, on the other hand, consist of protocol dependent closed islands. Peer-to-Peer implementations of real-time communications (like Skype) have proved themselves efficient in handling problems like NAT traversal or providing subscribers with an easy sign-in process (self-provisioning). Pure P2P model presents a number of disadvantages (see comparison table between the two models).

SIP is capable to enable new applications and is more effective in delivering of new services when compared with proprietary Peer-to-Peer applications that are limited in scope to their designated purpose. However Peer-to-Peer applications have proven to scale better for many users as well as resisting and recovering from disasters.

Combining SIP with Peer-to-Peer techniques offers advantages from both worlds, namely survivability and scalability.

SIP and Peer-to-Peer

There are two ways of merging SIP with Peer-to-Peer:

a) SIP User Agents can have an alternative lookup capability to DNS and can find other nodes using a distributed network lookup function for example by using a distributed hash table algorithm. This mechanism allows SIP User Agents to function without SIP Proxies but their ability to reach external networks or be reached from external networks is limited. Functions like NAT

traversal and accounting can be redistributed only randomly across un-trusted nodes. A pure P2P client approach to SIP requires modified User Agents.

b) SIP Proxy/Registrars in an overlay network can have the built-in capability to find other nodes and contribute to the proxying, registration and accounting requests for a large number of SIP clients they are responsible for. Access to ENUM and distributed NAT traversal capability guarantee a more reachable and a better quality service than un-trusted Peer-to-Peer User Agents. This approach requires no change in the existing SIP User Agents.

Analysis between P2P UA and P2P Proxy/Registrar

Feature	P2P Client	P2P Proxy/Registrar
Compatible with existing UAs	No	Yes
Detection of node failures	Yes	Yes
Keep-alive traffic rate	High	Low
Inter-vendor compatibility	Unlikely	Yes
Bootstrap server	Yes	Yes
Risk of Identity theft	Yes	No
Privacy control	Yes	Yes
Denial of Service possible	Yes	No
NAT traversal	Limited	Yes
Finding user (for N users)	log(N) steps	1 step
User aliases	No	Yes
Multiple devices	No	Yes
Geographic distribution	No	Yes
Business usage	No	Yes
Trace SIP messages	No	Yes
Public SIP URI	No	Yes
ENUM mappings	No	Yes
Disaster recovery	No	Yes
Trusted identities	No	Yes
Voicemail	Limited	Yes
Persistent identity	No	Yes
Call forwarding	No	Yes
Time of the day routing	No	Yes
Access to PSTN/Other islands	No	Yes
E.911 support	Improbable	Yes

Survivability

A survivable architecture can be implemented in the form of distributed P2P Proxy/Registrar. It acts like a mesh of super-nodes in pure P2P networks. Its architecture is scalable without the use of traditional load balancers and has no single point of failure.

A self-configurable network with zero maintenance is able to survive catastrophic failures like network connectivity loss or distributed Denial of Service attacks (DoS). A central provisioning system has full overview over the network topology and location of nodes and can decide to insert or remove nodes as necessary.

New P2P nodes join the network by installing and running network software. After joining the network, the node authorizes itself to the network and becomes available to serve SIP requests. On node failure, the SIP requests handled by that node are automatically distributed to surviving nodes without any manual intervention.

Scalability

The architecture can scale to handle millions of subscribers without depreciation of performance by simply adding new nodes into the network. As opposed to pure P2P client approach, the time required for user location and setup of an interactive session is almost instant (tens to hundreds of milliseconds).

Address space

The architecture is based on Universal Resource Identifiers (SIP URI), similar with the e-mail addresses. Standard DNS lookups are used for the name resolution. Parallel forking to multiple devices is possible. Privacy, when required, is realized by the SIP Proxy based on the subscriber profile.

SIP subscribers may be assigned a SIP address within predefined sets of public SIP domains for those who prefer total anonymity (typical private individuals) or may register an Internet domain (typical businesses). The domains are provisioned in the DNS and SIP Proxies providing a globally reachable identity.

Components

1. Peer-to-Peer overlay network
2. DNS and ENUM
3. Provisioning system
4. SIP Proxy/Registrar
5. CDR mediation accounting
6. NAT traversal using distributed media relays
7. User Agents

The DNS and the Provisioning system are the only centralized components. The DNS guarantees the integrity and availability of the name space to the Internet, while the central provisioning takes care that any change is done atomically and provides monitoring functions. All other components are distributed in the network to provide resilience and load sharing.

Peer-to-Peer overlay network

A set of nodes participates to serve SIP requests for the SIP domains managed by the operator. It is one administrative domain controlled by one provisioning engine.

The overlay network realizes the distribution of requests (REGISTER, INVITE, etc.) among multiple nodes. This overlay network is based on a protocol derived from the Chord protocol. This protocol realizes a deterministic distribution of requests among peers and automatic recovery and reconfiguration in case of nodes joining or leaving the network.

The overlay network correlates the network topology with the DNS and signals to the provisioning system the addition or removal of new nodes to the network.

DNS and ENUM

During SIP session setup or registration the DNS plays a central role for locating the SIP Proxy/Registrar for a given user domain or translating the E.164 numbers into SIP addresses by using the mechanisms specified in RFC 3263 and RFC 3761. DNS NAPTR and SRV records are used for these purposes.

The Peer-to-Peer algorithm linked into the DNS software achieves the dynamic load balancing based on the actual network topology.

Provisioning system

The role of the provisioning system is to handle the subscriptions, the atomic modification of subscriber data and for preserving the integrity of the name space.

Each node maintains only a copy of the local subscriber base; replication of full subscriber database is not necessary. Notifications of node join/leave are signaled to the provisioning system, which provides an overview of the network health.

The operator has full overview over the health of the P2P network and can trace messages and call flows across the mesh of nodes.

SIP Proxy/Registrar

Any SIP Proxy can be used providing it has an API for handling routing information. The Proxy handles session setup requests in behalf of user

agents. The SIP Proxy authorizes subscribers and grant access to external gateway or user location based of subscriber profile information. All proxies in the cloud may act as an outbound proxy for user agents and share the traffic load.

The registrar function is distributed across the network without the need of a central database. All instances of the same Address-of-Record are handled by the same SIP Registrar. For this reason, each REGISTER request is redirected to the serving SIP Registrar responsible for the SIP subscriber. The redirection is based on the Peer-to-Peer algorithm.

Profile database

The profile database contains information linked to each SIP subscriber with regards to aliases, ENUM numbers, call control, privacy and selective call accept/reject presence settings. It is available in the central provisioning system, which pushes all changes to the responsible P2P node.

In practice, the Peer-to-Peer overlay software, the SIP Proxy, SIP Registrar and profile database components run physically on the same node.

CDR mediation and accounting

CDR mediation and accounting fulfill different needs like accounting for PSTN termination, traces to network problems or quality control.

Accounting records are generated by the SIP Proxy for requests coming from User Agents and by the SIP Registrar that is responsible for user location for sessions originating from outside the P2P network.

The accounting records may also carry detailed information down to the SIP stack level that help trouble-shooting for issues reported by Subscribers. To overcome the load generated by heavy loaded SIP Nodes, a low level accounting feature can be enabled based on user profile or by remote control to the logging node.

Based on a real time monitoring of the rate between successful/failed-without-good-reason calls per node, certain nodes could be eliminated from the SIP cloud until the network problem is solved.

NAT traversal using distributed media relays

SIP clients should have ICE support and try to exchange RTP packets with the remote User Agents by using STUN techniques described by ICE methodology.

If ICE capability does not exist or does not work, a TURN solution will be used as follows. In contrast to the TURN protocol definition, the SIP Proxy and not the User agent does the session reservation for the media stream relay. This has the immediate advantage that the SIP UA does not have to have any TURN capability built in and secondary a database with user credentials does not need to be stored on both the TURN server and the client.

The TURN server has a trusted relationship with the SIP nodes. Another advantage that led to this approach is the fact that the SIP Proxy has always more clues about where is the best place to assign a media relay in between a SIP session than the UA itself. This allows per call allocation of a media relay session in an optimum place on the Internet and solves the load balancing and scalability of the media relay function.

MediaProxy is a NAT traversal solution based on this methodology.

User agents

The requirements for User Agents to function into the proposed architecture are mandatory symmetric signaling/media and optional ICE support.
To register to the service the User Agent should support DNS SRV record lookup for locating the Registrar

Basically any SIP compliant User Agents should be able to function in this P2P Proxy/Registrar architecture.

Implementation

SIP Thor™ is AG Projects solution for survivability and scalability of SIP communications. It is using Peer-to-Peer technology to realize scalability with no single point of failure.

SIP Thor is a much more scalable and flexible solution than pure Peer-to-Peer networks or stand-alone SIP Proxy/Registrars. SIP Thor combines the advantages of both models while eliminating the disadvantages presented by them.

For more information see:

http://www.ag-projects.com/SIPThor.html