P2PSIP
Internet-Draft
Intended status: Standards Track
Expires: May 22, 2008

M. Stiemerling
M. Brunner
NEC

November 19, 2007

Peer-to-Peer SIP Implementation Report
draft-stiemerling-p2psip-impl-02

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79. This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on May 22, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This memo is an implementation report about the peer-to-peer SIP system developed in the European IST Ambient Networks research project. This system replaces the traditional SIP proxy-registrar function with a distributed lookup mechanism, adds overlay functionality to the SIP signalling and to RTP traffic, takes care

about media/packet relay lookup and insertion into the SIP/RTP paths, plus automatic adaptation of the voice transmission according to changing network conditions. Standard, unmodified SIP user agents are used for communication. The presented system is work in progress and this memo is an attempt to gather IETF community feedback about the described approach.

Table of Contents

1. Introduction

This memo is an implementation report about the peer-to-peer SIP system developed in the European IST Ambient Networks research project. The presented system is work in progress and focusses on voice/video calls but does currently not consider presence and instant messaging.

The spread of Internet technology in many types of wireless networks, presenting different features and technologies, has been a revolution for the networking architecture, traditionally based on fixed networks. Nowadays, the interconnection of systems presenting heterogeneous network capabilities, access technologies and end-user devices is becoming a must for many applications and services.

The European Community's Sixth Framework Program Integrated Project called 'Ambient Networks' [AMBIENT] is trying to address these challenges, following generalisation and integration procedures. The main objective of the project is the definition of a complete innovative networking model, respondent to the development of mobile network solutions and the integration of different types of fixed and mobile networks.

One of the main interests in this kind of network organization is the provision of services to the end-users, in a way transparent to the network specific characteristics as much as possible. In this research area, Overlay Networks have been proved to be a powerful solution for abstracting services from the network connectivity and providing a decentralised network organization.

Ambient Networks is defining an overlay service with the concept of Service-aware Adaptive Transport Overlay (SATO). The main purpose is the design of a generalised structure which is able to realise overlay networks on demand, in order to carry multiple applications and to provide them useful functionalitiesrealised inside the network paths. The exchange of multimedia traffic like voice and video is one of the applications which can benefit of SATO. Peer-to-peer SIP has been realised on the SATO platform and this memo gives an overview of the current implementation.

The remainder of the document is structure in this way that Section 2 defines the goals of our peer-to-peer SIP system and introduces the system. Section 6 reports some findings.

You will find the P2PSIP terminology used within this memo being mainly aligned with the terminology in [I-D.willis-p2psip-concepts].

2. SATO System Overview

This section describes the overall system of the Service-Aware Transport Overlay (SATO) system used to built a peer-to-peer SIP system.

2.1. General concepts of SATO

The purpose of SATO in the Ambient Networks architecture is to provide a flexible and customisable transport service to the application layer by using overlay networks on top of the transport layer connectivity. Further service needed to run such a SATO are included, for instance a distributed lookup service (i.e., DHT but not limited to).

Service-aware Transport Overlay

Service-aware Adaptive Transport Overlays (SATOs) enable the flexible configuration of virtual networks consisting of Overlay Nodes (ONodes) on top of the underlying basic network connectivity. The Overlay Network topologies are responding to the application needs and can follow point-to-point, point-to-multipoint and multipoint-to-multipoint paradigms. Many SATOs can be created and deployed simultaneously.

Dynamic Inclusion of Network Elements

The novel SATO concept allows the transparent inclusion of network-side data processing elements (SATO Ports) in the end-to-end transport path (between a client and a server or Peer-to-Peer). These SATO Ports can provide value-added functions such as Overlay routing, smart caching, media adaptation, rate adaptation, synchronisation, filtering, metering, congestion control, etc.

On Demand Overlay Set Up and Tear Down

Service-aware Adaptive Transport Overlays are designed for accommodating the requests of the application services. As a consequence, they should be established on demand, based on particular requirements, and terminated when the service is not requested anymore (e.g., after the last user has disconnected).

Overlay Adaptation

SATOs could adapt as a consequence of ONodes joining or leaving the virtual network or due to ONodes with changing context or as a result of adaptation requests from other Ambient Networks Functional Entities. This introduces the notion of Oadaptive overlaysO that dynamically re-configure in order to optimise the service delivery. Dynamic adaptation of a SATO can also take place in response to changes in other Functional Entities (FEs) or in the underlying network. SATO adaptation requires significant, time-consuming control space activities possibly involving interactions between many FEs. Highly dynamism or fast alterations within SATOs can therefore be addressed by means of internal SATO routing capabilities.

2.2. Elements of SATO

The core of the SATO P2PSIP systems are the SATO overlay nodes. These nodes are P2PSIP overlay peers participating in the peer-to-peer routing. Each overlay peer hosts a modified SIP proxy, RTP proxy, and the overlay manager. The overlay manager is in charge of controlling the overlay, providing a distributed lookup service for users and media relays, etc. Additionally, overlay peers can host media relays that are offered to other peers in order to assist in NAT/firewall traversal. Figure 1 shows the overlay peer building blocks.

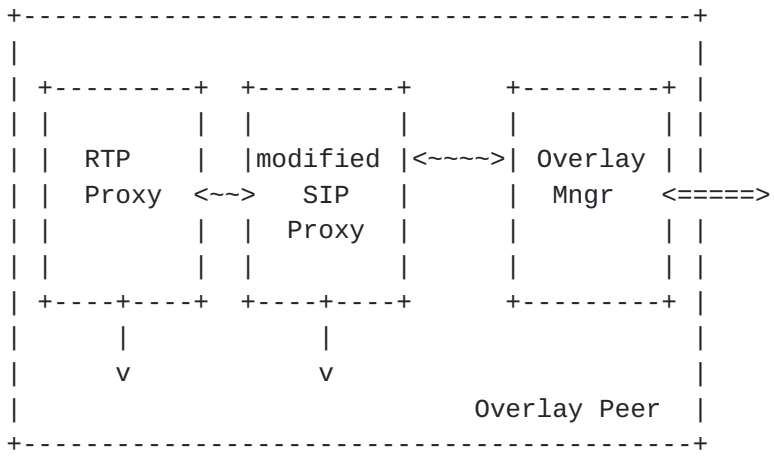

Figure 1: P2PSIP overlay peer

The overlay manager uses Web Services (chosen for rapid prototyping for the P2PSIP Overlay Peer Protocol) to communicate with other overlay managers on other peer nodes. The overlay peer internal communication is using Unix sockets. The SIP proxy is using an off the shelf SIP stack with modified routing, i.e., peer-to-peer routing, stores the user records in a DHT (see next paragraph), and a slightly changed interface to the IP sockets. The RTP proxy is off the shelf as well. The connections between two SIP proxies or two RTP proxies on different overlay peers can use any type of transport protocol or network layer setting, e.g., TCP running over IPsec. The used transport protocol, IP version, encryption, etc is negotiated via the overlay managers.

Not shown in Figure 1 is the distributed lookup service for storing P2PSIP overlay user records and location information for media relays. For both, the user records and the media relay location, a mapping of the respective information to the IP address of the overlay peer is stored. The bamboo DHT implementation [BAMBOO] is currently used for storing this information.

3. Running Peer-to-Peer SIP

The SATO P2PSIP system is built in such a way that traditional SIP UA can participate without any modification. The SIP UA use their standard way of SIP and RTP for signalling and media. The SIP UAs in the example below use SIP with UDP transport. The settings of the UAs need just to point for the SIP registrar and outbound proxy to their overlay peer (see also Figure 2. So, all outgoing SIP signalling messages will be forwarded to the respective P2PSIP overlay peer.

3.1. Registering and Deregistering a User

This section assumes an in advanced happening enrolment process for the user with the P2PSIP system. The enrolment process has not been defined and implemented.

The SIP UA proceeds with the standard registration process as defined in [RFC3261] either UDP or TCP to the P2PSIP overlay peer, i.e., the P2PSIP Overlay Client Protocol is standard SIP. The P2PSIP registrar (part of the proxy here) uses in the SATO P2P system currently only the user part of the SIP URI for the registration process. The P2P SIP registrar stores the given user part of the URI in the DHT together with its (or one of its) IP addresses. The real location of the UA is only stored in the P2PSIP registrar.

The entry in the DHT and P2PSIP registrar is deleted, when the user

should be deregistered, i.e., either by timeout of the registration or explicitly request.

4. Running a P2P Call

Placing a call from one UA1 to UA2 (see Figure 2) is again following the possible call flows in [RFC3261] from UA to P2PSIP overlay peer. The SIP proxy at the overlay peer receiving the call request (SIP) is using the user part of the SIP URI for the DHT lookup. If the user is registered in the DHT, i.e., register at some overlay peer, the IP address of the serving overlay peer is obtained. Otherwise a 404 (Not Found) response is generated and send to the UA. This IP address is transfered from SIP1 to the local overlay manager OM1. OM1 uses the P2PSIP Overlay Peer Protocol to contact OM2 on the target overlay peer where UA2 is registered to. OM1 and OM2 negotiate the possible and by the SIP proxy desired transport connection between SIP1 and remote P2PSIP proxy SIP2. Once the transport connection between SIP1 and SIP2 is setup, both are notified about this. SIP1 forwards the SIP messages from UA1 (with modified values of some SIP headers to reflect the routing) to SIP2 via this transport connection (TCP in the example). SIP2 in turn forwards the message to UA2. Subsequent SIP messages are transfered via this transport connection.

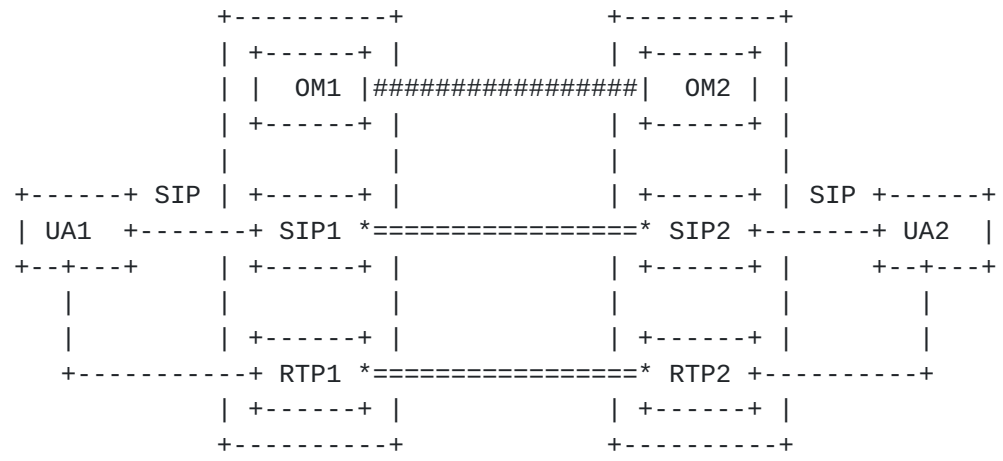

Figure 2: P2PSIP overlay network

In the process of the call signalling, both P2PSIP proxies (SIP1 and

SIP2) will replace the media IP address(es) and port(s) by an IP address and port of their local RTP proxy. The SDP part of UA1 will be replaced with IP addresses/ports of RTP2 and of UA2 with IP addresses/ports of RTP1. This behaviour can turned on or off per configuration in the respective P2PSIP proxies if the media should delivered directly between UA1 and UA2. OM1 and OM2 will take care of setting up the desired transport connection between RTP1 and RTP2. The example uses TCP but any other transport can used. Now the call can proceed and when the call ends all transport connections between SIP1/SIP2 and RTP1/RTP2 are torn down.

5. Using HIP in P2PSIP

This section briefly describes how the P2PSIP system is using Host Identities and the Host Identity Protocol (HIP) (see [RFC4423] and [I-D.ietf-hip-base]) for establishing a secure communication between different nodes and to securely identify nodes.

Each node participating in the Ambient Networks Peer-to-Peer SIP network has assigned a Host Identity and publishes his Host Identity Tag (HIT) with all of its information in the distributed database. Typically, the HIT is used to setup the communication between both hosts and to maintain it even when the IP address is changing (mobility or just assignment of another IP address through DHCP). The P2PSIP implementation relies on HIP's IPsec usage (preferably the BEET mode [I-D.nikander-esp-beet-mode]) for providing authenticated and encrypted data transmission between two nodes. The HIT is also used to ensure that the contacted host is rely the desired host, as stored in the distributed database.

The P2PSIP implementation assumes that HIP is able to find the mapping between a HIT and the locator set, i.e., resolving the HIT used by P2PSIP to an IPv4 or IPv6 address usable by HIP itself. It should be noted that this step is not yet fully solved in HIP, other than doing some add-ons to the DNS.

HITs are not revealed to the SIP user agents, as they just see and use IPv4 addresses and host names in the regular SIP signalling. However, within the P2PSIP overlay, HITs are used within the SIP signalling and also used to route messages. HITs simply appear as IPv6 addresses in the SIP signalling and can so easily be handled by existing SIP stacks.

The current prototype implementation goes even a step beyond HIP, as it is using the Node ID architecture. This architecture is described in [ref.node-id].

6. Conclusions

This memo sketches the Ambient Networks SATO-based peer-to-peer SIP system, which is one possible way to realise such a P2PSIP system. It uses a generalised overlay deployment system, a well-known distributed lookup service and unmodified SIP stacks in proxies and user agents. Not yet described but also implemented is the lookup of media/packet relays and their inclusion into the overlay network. This enables support for NAT/firewall traversal of signalling and data. The exact changes in the SIP routing and the SIP header values are not yet listed in this memo, but will be added to the next revision.

The intention of this document is raise further discussions and not to give a final recommendation or similar.

7. Security Considerations

To be done.

8. Acknowledgements

Martin Stiemerling and Marcus Brunner are partly funded by Ambient Networks, a research project supported by the European Commission under its Sixth Framework Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Ambient Networks project or the European Commission.

9. References

9.1. Normative References

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

9.2. Informative References

[AMBIENT] "IST project 507134 Ambient Networks (AN)", Web Site http://www.ambient-networks.org, October 2006.

[BAMBOO] "BAMBOO DHT", Web Site http://bamboo-dht.org/,

October 2006.

[I-D.ietf-hip-base]
Moskowitz, R., "Host Identity Protocol", draft-ietf-hip-base-06 (work in progress), June 2006.

[I-D.nikander-esp-beet-mode]
Melen, J. and P. Nikander, "A Bound End-to-End Tunnel (BEET) mode for ESP", draft-nikander-esp-beet-mode-06 (work in progress), August 2006.

[I-D.willis-p2psip-concepts]
Willis, D., "Concepts and Terminology for Peer to Peer SIP", draft-willis-p2psip-concepts-03 (work in progress), October 2006.

[RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.

[ref.node-id]
Ahlgren, B., Arkko, J., Eggert, L., and J. Rajahalme, "A Node Identity Internetworking Architecture", April 2006.

Authors' Addresses

Martin Stiemerling
NEC Network Laboratories Europe/University of Goettingen
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 113
Fax: +49 6221 4342 155
Email: stiemerling@netlab.nec.de
URI: http://www.netlab.nec.de/

Marcus Brunner
NEC Network Laboratories Europe
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 129
Fax: +49 6221 4342 155
Email: marcus.brunner@netlab.nec.de
URI: http://www.netlab.nec.de/

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).