

SIPPING WG	A. Hourì	
Internet-Draft	IBM	
Intended status: Informational	S. Parameswar	
Expires: March 6, 2010	Microsoft Corporation	
	E. Aoki	
	AOL LLC	
	V. Singh	
	H. Schulzrinne	
	Columbia U.	
	September 02, 2009	

[TOC](#)

## Scaling Requirements for Presence in SIP/SIMPLE draft-ietf-sipcore-presence-scaling-requirements-02.txt

### Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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 March 6, 2010.

### Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

## Abstract

The document lists requirements for optimizations of SIP/SIMPLE. These optimizations should reduce the load on the network and the presence servers due to inter-domain presence subscriptions. The need for the requirements is based on a separate document that provides scaling analysis for inter-domain presence over SIP/SIMPLE.

---

## Table of Contents

- [1.](#) Requirements notation
- [2.](#) Introduction
- [3.](#) Requirements
  - [3.1.](#) Backward Compatibility Requirements
  - [3.2.](#) Policy, Privacy, Permissions Requirements
  - [3.3.](#) Scalability Requirements
  - [3.4.](#) Topology Requirements
- [4.](#) Considerations for Possible Optimizations
- [5.](#) Security Considerations
- [6.](#) IANA Considerations
- [7.](#) Changes from Previous Versions
  - [7.1.](#) Changes in version 01
- [8.](#) Acknowledgments
- [9.](#) References
  - [9.1.](#) Normative References
  - [9.2.](#) Informational References
- [§](#) Authors' Addresses

---

### 1. Requirements notation

[TOC](#)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\] \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

---

### 2. Introduction

[TOC](#)

The document lists requirements for optimizations of SIP/SIMPLE. See [\[I-D.ietf-simple-simple\] \(Rosenberg, J., "SIMPLE made Simple: An Overview of the IETF Specifications for Instant Messaging and Presence](#)

[using the Session Initiation Protocol \(SIP\)," March 2009.](#)) for the list of RFCs and drafts that are considered as part of SIP/SIMPLE. These optimizations should reduce the load on the network and the presence servers due to inter-domain presence subscriptions. The need for the requirements is based on a separate scaling analysis document [\[I-D.ietf-simple-interdomain-scaling-analysis\] \(Hourì, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, "Presence Interdomain Scaling Analysis for SIP/SIMPLE," August 2009.\)](#).

The scaling analysis document shows that the following aspects of SIP/SIMPLE can be optimized: the number of bytes sent between two federating domains, the number of messages processed, and the amount of state managed by the presence server.

For example, for two peering networks that have a total of 20 million users, we calculated that approximately 19 billion messages per 8 hours work day are exchanged between the networks for the presence service. For very large session peering (150 million subscriptions), we calculated that the presence server needs to manage approximately a terabyte of state.

It may be that very large systems require the deployment of significant resources, but we should consider the following:

\*The scaling analysis document makes moderate assumptions about the number of presence status changes per hour and the the size of the presence document.

\*Even when applying all optimizations for presence as described by drafts and RFCs, we still calculated around 10 billion messages per 8 hours work day for a total of 20 million federating users. This is better than the base case, but not enough given the moderate assumptions and given that, when presence is deployed in a mass market, the number of federating users will be much larger than 20 million.

---

### 3. Requirements

[TOC](#)

This section lists the requirements for a solution that optimizes the inter-domain presence loads. The requirements are based on the presence scaling draft [\[I-D.ietf-simple-interdomain-scaling-analysis\] \(Hourì, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, "Presence Interdomain Scaling Analysis for SIP/SIMPLE," August 2009.\)](#).

---

[TOC](#)

### 3.1. Backward Compatibility Requirements

\*REQ-001: The solution SHOULD NOT deprecate existing protocol mechanisms defined in SIP/SIMPLE.

\*REQ-002: Existing SIP/SIMPLE clients SHOULD be able to communicate with clients and servers that implement new presence scaling features.

\*REQ-003: The solution SHOULD NOT constrain any existing RFC functional requirements for presence.

\*REQ-004: The solution MUST NOT constrain any existing RFC security requirements for presence.

\*REQ-005: Systems that do not use the new additions to the protocol SHOULD operate at the same level as they do today.

---

### 3.2. Policy, Privacy, Permissions Requirements

[TOC](#)

\*REQ-006: The solution SHOULD NOT limit the ability of presentities to present different views of presence to different watchers.

\*REQ-007: The solution SHOULD NOT restrict the ability of a presentity to obtain its list of watchers.

\*REQ-008: The solution MUST NOT create any new privacy holes or make any existing ones worse.

---

### 3.3. Scalability Requirements

[TOC](#)

\*REQ-009: Presence systems (intra- or inter-domain) SHOULD scale in linear proportion to the number of watchers and presentities in the system.

\*REQ-010: The solution SHOULD NOT require a significant increase in the size of maintained state, compared to the current state size required by SIP/SIMPLE.

\*REQ-011: The solution MUST allow presence systems to scale. Note: we view scalability on the order of tens of millions of users in each peer domain.

\*REQ-012: The solution MUST support a high percentage of watcher/presentity intersections between the domains and it MUST support various intersection models such as either small or large percentage of users from each domain subscribing to users from the other domain.

\*REQ-013: Protocol changes MUST NOT prohibit optimizations in deployment models where there is a high number of inter-domain subscriptions.

\*REQ-014: New functionalities and extensions to the presence protocol SHOULD consider scalability with respect to the number of messages, state size, and management and processing load.

---

### 3.4. Topology Requirements

[TOC](#)

\*REQ-015: The solution SHOULD allow for arbitrary federation topologies including direct and indirect peering.

---

## 4. Considerations for Possible Optimizations

[TOC](#)

This section discusses the possible paths for optimization. One of the most important considerations is whether SIP, which was designed more as an end-to-end protocol, needs to be optimized for direct interactions between presence servers.

It is very possible that the issues described here are inherent to presence systems in general and not specific to SIP/SIMPLE.

Organizations need to be prepared to invest substantial resources in the form of networks and hardware in order to create sizable systems. However, it is apparent that additional protocol optimizations are possible and further IETF work is needed in order to provide better scalability of large presence systems.

We should remember that SIP was originally designed for end-to-end session creation and that the number and size of messages are of secondary importance for an end-to-end session negotiation protocol. For large scale and especially for very large scale presence, the number of messages and the size of each message are of extreme importance. Care must be taken to address scalability during the initial phase of protocol design; shoehorning scalability at a later phase will be doomed to failure.

We should also consider whether using the same protocol between clients and servers and between servers is a good choice. It may be that, between inter-domain servers or even intra-domain servers (such as between RLSes [\[RFC4662\] \(Roach, A., Campbell, B., and J. Rosenberg, "A Session Initiation Protocol \(SIP\) Event Notification Extension for Resource Lists," August 2006.\)](#) and presence servers), there needs to be a different protocol that is optimized for the load and that can make assumptions about the network (using only TCP, for example. In [\[I-D.ietf-simple-interdomain-scaling-analysis\] \(Hour, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, "Presence Interdomain Scaling Analysis for SIP/SIMPLE," August 2009.\)](#), see the section that calculates the number of bytes and messages for an imaginary, optimized SIP).

When a presence server connects to another server using current SIP/SIMPLE, there is an extreme number of redundant messages due to SIP's support of both TCP and UDP and due to privacy controls that cause the sending of multiple presence documents for the same presentity. A server-to-server protocol will have to address these issues. Some initial work to address these issues can be found in:

[\[I-D.ietf-simple-view-sharing\] \(Rosenberg, J., Donovan, S., and K. McMurtry, "Optimizing Federated Presence with View Sharing," November 2008.\)](#) and [\[I-D.ietf-simple-intradomain-federation\] \(Rosenberg, J., Hour, A., Smyth, C., and F. Audet, "Models for Intra-Domain Presence and Instant Messaging \(IM\) Bridging," July 2009.\)](#) and in other (still individual) drafts.

Another issue related to protocol design is whether NOTIFY messages should not be considered as media just like audio, video, and even text messaging. The SUBSCRIBE method may be extended to negotiate the route and other parameters of the NOTIFY messages, similar to the way the INVITE method negotiates media parameters. This way, the load can be shifted to specialized NOTIFY "relays" and taken off the control path of SIP. One of the possible ideas (Marc Willekens) is to use SIP for client/server NOTIFY but use a more optimized and controllable protocol for the server-to-server interface. Another possibility is to use the MSRP [\[RFC4975\] \(Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol \(MSRP\)," September 2007.\)](#), [\[RFC4976\] \(Jennings,](#)

[C., Mahy, R., and A. Roach, "Relay Extensions for the Message Sessions Relay Protocol \(MSRP\)," September 2007.](#)) protocol for the notifications.

---

## 5. Security Considerations

[TOC](#)

This document discusses the scalability requirements for the existing SIP/SIMPLE protocol and model. Many of the above-mentioned changes to the protocol will have security implications.

For example, a potential protocol change that may have security implications is the single sending of a presence document between domains in order to reduce the number of messages and network load. This possible optimization will delegate privacy protection from one domain to the other, and this delegated protection should be addressed during design.

An important part of work on the requirements and optimizations will be to ensure that all the security aspects are covered.

---

## 6. IANA Considerations

[TOC](#)

This document has no IANA actions.

---

## 7. Changes from Previous Versions

[TOC](#)

### 7.1. Changes in version 01

[TOC](#)

Editorial language changes.

---

## 8. Acknowledgments

[TOC](#)

We would like to thank Jonathan Rosenberg, Ben Campbell, Markus Isomaki, Piotr Boni, David Viamonte, Aki Niemi, Marc Willekens, and Gonzalo Camarillo for their ideas and input. Special thanks to Jean-Francois Mule, Vijay K. Gurbani and Hisham Khartabil for their a detailed review. Very sprecial thanks A. Jean Mahoney for reviewing this draft.

---

## 9. References

[TOC](#)

---

### 9.1. Normative References

[TOC](#)

- |           |  |
|-----------|--|
| [RFC2119] | <a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ). |
|-----------|--|

---

### 9.2. Informational References

[TOC](#)

[I-D.ietf-simple-interdomain-scaling-analysis]	Houri, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, " <a href="#">Presence Interdomain Scaling Analysis for SIP/SIMPLE</a> ," draft-ietf-simple-interdomain-scaling-analysis-08 (work in progress), August 2009 ( <a href="#">TXT</a> ).
[I-D.ietf-simple-intradomain-federation]	Rosenberg, J., Hour, A., Smyth, C., and F. Audet, " <a href="#">Models for Intra-Domain Presence and Instant Messaging (IM) Bridging</a> ," draft-ietf-simple-intradomain-federation-04 (work in progress), July 2009 ( <a href="#">TXT</a> ).
[I-D.ietf-simple-simple]	Rosenberg, J., " <a href="#">SIMPLE made Simple: An Overview of the IETF Specifications for Instant Messaging and Presence using the Session Initiation Protocol (SIP)</a> ," draft-ietf-simple-simple-05 (work in progress), March 2009 ( <a href="#">TXT</a> ).
[I-D.ietf-simple-view-sharing]	Rosenberg, J., Donovan, S., and K. McMurtry, " <a href="#">Optimizing Federated Presence with View Sharing</a> ," draft-ietf-simple-view-sharing-02 (work in progress), November 2008 ( <a href="#">TXT</a> ).
[RFC4662]	Roach, A., Campbell, B., and J. Rosenberg, " <a href="#">A Session Initiation Protocol (SIP) Event Notification Extension for Resource Lists</a> ," RFC 4662, August 2006 ( <a href="#">TXT</a> ).
[RFC4975]	Campbell, B., Mahy, R., and C. Jennings, " <a href="#">The Message Session Relay Protocol (MSRP)</a> ," RFC 4975, September 2007 ( <a href="#">TXT</a> ).
[RFC4976]	Jennings, C., Mahy, R., and A. Roach, " <a href="#">Relay Extensions for the Message Sessions Relay Protocol (MSRP)</a> ," RFC 4976, September 2007 ( <a href="#">TXT</a> ).



---

## Authors' Addresses

[TOC](#)

	Avshalom Hour i
	IBM
	3 Pekris Street, Science Park
	Rehovot,
	Israel
Email:	<a href="mailto:avshalom@il.ibm.com">avshalom@il.ibm.com</a>
	Sriram Parameswar
	Microsoft Corporation
	One Microsoft Way
	Redmond, WA 98052
	USA
Email:	<a href="mailto:Sriram.Parameswar@microsoft.com">Sriram.Parameswar@microsoft.com</a>
	Edwin Aoki
	AOL LLC
	401 Ellis St.
	Mountain View, CA 94043
	USA
Email:	<a href="mailto:aoki@aol.net">aoki@aol.net</a>
	Vishal Singh
	Columbia University
	Department of Computer Science
	450 Computer Science Building
	New York, NY 10027
	US
Email:	<a href="mailto:vs2140@cs.columbia.edu">vs2140@cs.columbia.edu</a>
URI:	<a href="http://www.cs.columbia.edu/~vs2140">http://www.cs.columbia.edu/~vs2140</a>
	Henning Schulzrinne
	Columbia University
	Department of Computer Science
	450 Computer Science Building
	New York, NY 10027
	US
Phone:	+1 212 939 7004
Email:	<a href="mailto:hgs+ecrit@cs.columbia.edu">hgs+ecrit@cs.columbia.edu</a>
URI:	<a href="http://www.cs.columbia.edu/~hgs">http://www.cs.columbia.edu/~hgs</a>