SIPPING WG	A. Houri	
Internet-Draft	IBM	
Intended status: Informational	S. Parameswar	
Expires: July 31, 2009	Microsoft Corporation	
	E. Aoki	
	AOL LLC	
	V. Singh	
	H. Schulzrinne	
	Columbia U.	
	January 27, 2009	

TOC

Scaling Requirements for Presence in SIP/SIMPLE
draft-ietf-sipping-presence-scaling-requirements-03.txt

Status of this Memo

Copyright Notice

Abstract

The document provides a set of requirements for enabling interdomain scaling in presence for SIP/SIMPLE.

Table of Contents

---

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 the SIP/SIMPLE protocol. 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 the SIP/SIMPLE protocol. These optimizations should reduce the load on the network and the presence servers in interdomain presence subscriptions. The need for the requirements is based on a separate scaling analysis

document [I-D.ietf-simple-interdomain-scaling-analysis] (Houri, 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 have shown that there is much room for optimizations in the SIP/SIMPLE protocol. The need for optimizations is in the number of bytes that are sent between two federating domains, the number of messages that need to be processed and the amount of state that needs to be managed by the presence servers.
For example, for two peering networks that have total of 20 million users, we got around 19 billion messages per 8 hours work day that needs to be exchanged between the networks only for supporting the presence service.
For very large session peering (150 million subscriptions) we got a state close to a tera byte that needs to be managed by the server in order to manage presence.
It may be that when deploying a very large systems big resources need to be allocated but we should take into the considration the following:

* The assumptions that have been used in the scaling analysis document are very moderate from the aspect of number of presence status changes per hour and the the size of the presence document that was assumed.

* Even when applying all current drafted and/or RFCd optimizations for presence we still got around 10 billion messages per 8 hours work day for a total of 20 million fedearting users. This is good but not enough given the moderate assumptions that we have used and given that when presence will be deployed to a mass market the number of federating users will be much more then 20 million federating users.

---

3. Requirements

TOC

This section lists requirements for a solution that will optimize the interdomain presence loads. The requirements are based on the presence scaling draft [I-D.ietf-simple-interdomain-scaling-analysis] (Houri, 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 are not using 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 for 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 or make worse any existing privacy holes.

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 state to maintain, 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: There may be various usage patterns when users of one domain subscribe to users from another domain. It may be that only small percentage of users from each domain will subscribe to users from the other domain, it may be that most watchers will be from the other domain while there will be few watchers from the same domain. The solution MUST support high percentage of watcher/presentity intersections between the domains and it MUST support various intersection models.

*REQ-013: Protocol changes MUST NOT prohibit optimizations in deployment models where there is a high level of cross subscriptions between the domains.

*REQ-014: New functionalities and extensions to the presence protocol SHOULD take into account 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

The document provides an initial list of requirements for a solution of scalability of interdomain presence systems using the SIP/SIMPLE

protocol. The issue of scalability was shown in a separate document [I-D.ietf-simple-interdomain-scaling-analysis] (Houri, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, "Presence Interdomain Scaling Analysis for SIP/SIMPLE," August 2009.).

The following is a discussion of the various possible paths for optimizations. One of the most important considerations is whether there is a need to adapt SIP that was designed more as an end to end protocol to be much more optimized when presence server interacts directly with another presence server.

It is very possible that the issues that are described in this document are inherent to presence systems in general and not specific to the SIP/SIMPLE protocol. 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 work is needed in the IETF 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 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 that are needed and the size of each message are of extreme importance. Adequate care must be taken in addressing scalability as part of the initial protocol design phase; Trying to to shoehorn 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 in interdomain or even between servers in the same domain (as between RLSs [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 is a need to have a different protocol that will be very optimized for the load and can assume some assumptions about the network (for example do not use unreliable protocol as UDP but only TCP, see the section that calculates the number of bytes and messages for imaginary very optimized SIP).

When a presence server connects to another server using the current SIP/SIMPLE protocol, there will be an extreme number of redundant messages due to the overhead in the SIP protocol of supporting both TCP and UDP and due to the need to send multiple presence documents for the same watched user because of privacy issues. 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. McMurry, "Optimizing Federated Presence with View Sharing," November 2008.) and [I-D.ietf-simple-intradomain-federation] (Rosenberg, J., Houri, 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 that is more 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, in a similar way that the INVITE method negotiates media parameters. This way the load can be offloaded to a specialized NOTIFY "relays" thus not loading the control path of SIP. One of the possible ideas (Marc Willekens) is to use the SIP protocol for client/server NOTIFY but make use of 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 scalability requirements for the existing SIP/ SIMPLE protocol and model. Many of the changes to the protocol will have security implications as mentioned in some of the requirements above.
One example of possible protocol changes that may have security implications is sending a presence document only once between domains in order to optimize the number of messages and network load. This possible optimization will delegate privacy protection from one domain to another domain and should be addressed when designing protocol optimizations
Important part of work on the requirements and optimizations will be to make sure that all the security aspects are covered.

---

6. IANA Considerations

TOC

This document has no IANA actions.

---

7. Acknowledgments

TOC

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

8. References

TOC

8.1. Normative References

TOC

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).

8.2. Informational References

TOC

[I-D.ietf-simple-interdomain-scaling-analysis]	Houri, A., Aoki, E., Parameswar, S., Rang, T., Singh, V., and H. Schulzrinne, "Presence Interdomain Scaling Analysis for SIP/SIMPLE," draft-ietf-simple-interdomain-scaling-analysis-08 (work in progress), August 2009 (TXT).
[I-D.ietf-simple-intradomain-federation]	Rosenberg, J., Houri, A., Smyth, C., and F. Audet, "Models for Intra-Domain Presence and Instant Messaging (IM) Bridging," draft-ietf-simple-intradomain-federation-04 (work in progress), July 2009 (TXT).
[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)," draft-ietf-simple-simple-05 (work in progress), March 2009 (TXT).
[I-D.ietf-simple-view-sharing]	Rosenberg, J., Donovan, S., and K. McMurry, "Optimizing Federated Presence with View Sharing," draft-ietf-simple-view-sharing-02 (work in progress), November 2008 (TXT).
[RFC3857]	Rosenberg, J., "A Watcher Information Event Template-Package for the Session Initiation Protocol (SIP)," RFC 3857, August 2004 (TXT).
[RFC3858]	Rosenberg, J., "An Extensible Markup Language (XML) Based Format for Watcher Information," RFC 3858, August 2004 (TXT).
[RFC4662]	Roach, A., Campbell, B., and J. Rosenberg, "A Session Initiation Protocol (SIP) Event Notification Extension for Resource Lists," RFC 4662, August 2006 (TXT).
[RFC4975]	Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)," RFC 4975, September 2007 (TXT).

[RFC4976]	Jennings, C., Mahy, R., and A. Roach, “Relay Extensions for the Message Sessions Relay Protocol (MSRP),” RFC 4976, September 2007 (TXT).

Authors' Addresses

TOC

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