P2PSIP
Internet-Draft
Intended status: Standards Track
Expires: August 28, 2007

C. Jennings
Cisco Systems
February 24, 2007

Security Mechanisms for Peer to Peer SIP
draft-jennings-p2psip-security-00

Status of this Memo

Copyright Notice

Abstract

This document describes an overview of some security mechanisms for P2P SIP. Specifically it discusses mechanisms that can be used to secure the stored data and the routing in the distributed storage.

This draft is an very early draft to outline the possible solution space and far more details would be needed. This work is being discussed on the p2psip@ietf.org mailing list.

1. Introduction

The P2P SIP work stores users registrations and possibly other data in a Distributed Hash table (DHT). This requires a solution to securing this data as well as securing, as best possible, the routing in the DHT. Each user of the system has a name, such as alice@dht.example.net. These names are unique and meant to be chosen and used by human much like an SIP Address of Record (AOR) or email address. When the user enrolls in the DHT and creates the name, they are also given an asymmetric key as an certificate that binds their name to that key in a way that can be validated by any user enrolled in this particular DHT. Note that since only users of this DHT need to validate a certificate, this usage does not require a global PKI.

The overview of the proposed approach is that the certificate and key can be used to sign any data stored in the DHT and any user retrieving the stored data can check that the data was not tampered with. In addition, when a peer goes to modify the routing data in the DHT, they can provide the information of which users they represent such that it is possible to know which user was associated with a change and possibly limit the number of peers that a single user can operates and position the peers in such a way to limit their ability to attack the routing. In addition, over longer periods of time, it may be possible to revoke that users credentials by allowing their certificate to expire.

The rest of this document is arranged into an abstract model of how the security work work that would apply to any protocol the working group might develop for the DHT. After the abstract model, a specific mapping of the model to SIP is described that would apply if the working group used SIP for the DHT protocol.

2. Conventions

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 RFC 2119 [1].

3. Data Protection Architecture

There are possibly several things a client may want to store in the DHT. The most obvious on is the registration information that indicates the IP address or route to where a given name or AOR can be found. There are other bits of information that could also be stored. Each chunk of information is stored in what will be referred to as a "record". The defined record types and what they stored

would be described in documents and registered with IANA. One of the record types would be the "registration" record where clients stored their registration information. Each user in the system would only have one registration record. The index in the DHT would be formed by taking using the concatenation of the AOR and the record type name.

When a client wants to store some information in a record, they sent a request that has: their AOR, the record type name, the time, the data to store in the record, and MUST include a signature over all that information. When a peer goes to store the information, it MUST check that the signature is correct. It SHOULD also check that the data looks appropriate for this type of record given by checking things like the size of the data is in an appropriate range. When a client retrieves data out of the DHT, it retrieves all the information that was signed and SHOULD verify the signature on the data.

Open Issue: how do we want to deal with checking time and also does the data have a Time To Live (TTL).

Open Issue: do we pass the certificate with the signature or do we provide some alternative scheme to get the certificates. I am leaning towards pass the certificate along with the signature. A problem with this is the message size. A possible problem with not doing it is that the signature are used to verify the constructions of the routing architecture and assuming that the routing architecture is in place before a signature can be checked may lead to problems.

4. Routing Protection Architecture

The goal of protecting the routing is stopping attacker from performing a DOS attack on they system by misrouting requests in the DHT. The data is already protected by the data protection scheme above so an attacker can't tamper with the data in a way the user can't detect but an attacker can make it look like no data is available. There are a few obvious observation to make about this. First, it is easy to ensure that attacker at least has to have an valid enrollment with this particular DHT. Second, this is a DOS attack and the value of successfully executing it is fairly low. Third, if a larger percentage of the peers on the DHT are controlled by the attacker, it is probably impossible to perfectly secure this.

When a peer sends a request that modifies the routing in the DHT, it MUST sign the request on behalf of a user that is currently responsible for the peer using that users certificate. A peer that

is changing the routing state based on this request to check the signature before performing the request.

To reduce attacks on routing, the design tries to limit the ability of an attacker to place peers at arbitrary locations in the DHT. Some possible ways to do this are:

L1: Limiting IP addresses: Other systems have done this by forcing the peer id to be a hash of a combination of the peers IP and port however this approach does not work with IPv6 where the users have an arbitrary number of IP addresses and the scheme is also difficult to make work with IPv4 and NATs.

L2: Limiting by AOR: The first step to doing this is limiting the number of AORs an attacker can enroll in the system. How to do this is out of scope. The next step would be forcing a peer ID to have the high order bits formed from an hash of the AOR and some low order bits chosen randomly or hashed from the IP address and port. Peers would check the Peer ID was appropriate for the given users that signed the request.

L3: Limited by assignment at enrollment: When enrolling, the user would be given a small set of peer IDs for their use. This is effectively equivalent to Limited by AOR but has the addition complexity of the certificates become more complex as a peer would need to sign with the appropriate peer id as well as the AOR.

Open Issue: how to do the limiting. At this point, the Limiting by AOR type approach looks most appealing.

5. Mapping to SIP

There are several ways this could be mapped to SIP.

M1: The simplest way from a specification point of view would probably be to put the information that needs to be signed in an Authenticated Identity Body (AIB)[RFC 3893] in the body of the SIP message and use S/MIME to sign it. It would also be possible to, instead of using the AIB, form a new body format for a particular record type and use S/MIME to sign it.

M2: An alternative proposal that does not use S/MIME would be to create a new way of computing a signature over the relevant data.

M3: The SIP Identity works provides certain sort of signatures but they are domain based instead of user based so it would be challenging to adapt them for use here. The problems revolves around certificates that can be used to sign for a one user in the DHT, would need to be limited such that the same certificate

could not be used to sign for a different user. Solutions to this are likely to end up being more or less the same as the proposal in the paragraph above this one.

All of these approaches would rely on the user enrollment providing an X.509 certificate that contained the users name in the SubjectAltName and signing the certificate with a root certificate that was also provided to all clients and peers as part of the enrollment.

Open Issue: Choose or design an envelope and signing scheme.

6. Security Considerations

TBD

7. IANA Considerations

This document does not require any actions from IANA.

8. Open Issues

Yes

9. Acknowledgments

Thanks to Eric Rescorla.

10. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] 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.

Author's Address

Cullen Jennings
Cisco Systems
170 West Tasman Drive
MS: SJC-21/2
San Jose, CA 95134
USA

Phone: +1 408 902-3341
Email: fluffy@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgment

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