Network Working Group
Internet-Draft
Expires: December 22, 2011

S P. Romano
A. Amirante
University of Napoli
T. Castaldi
L. Miniero
Meetecho
A. Buono
Ansaldo Trasporti e Sistemi Ferroviari

June 20, 2011

A Framework for Distributed Conferencing
draft-romano-dcon-framework-09

Abstract

This document defines the framework for Distributed Conferencing (DCON). The framework draws inspiration from the work carried out in the XCON working group, which has defined a complete architecture for centralized conferencing. DCON is based on the idea that a distributed conference can be setup by appropriately orchestrating the operation of a number of XCON focus elements, each in charge of managing a certain number of participants. Interaction between each participant and the corresponding conference focus is based on the standard XCON framework, whereas inter-focus interaction is defined in this document.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

This document presents an architecture capable to move the XCON scenario towards a distributed framework. The requirements for DCON are presented in a separate document [I-D.romano-dcon-requirements]. In such an architecture a number of entities are used to manage conference setup in the presence of clients which are distributed across a geographical network. Each managing entity plays the role of a conference focus as defined in the XCON working group documents [RFC5239].

Indeed, an XCON focus is in charge of managing a certain number of clients falling under its own "realm". In order to move the XCON scope towards a distributed environment, we introduce inter-focus coordination, which is needed to effectively setup and manage conference instantiation and coordination. As in the centralized case, we define logical entities and naming conventions. An appropriate data model for distributed conferencing will be defined in a subsequent document and will extend, when needed, the XCON data model [I-D.ietf-xcon-common-data-model]. Furthermore, we propose the adoption of a suitable set of protocols which are complementary to the call signaling protocols and are needed to support advanced conferencing applications.

2. Conventions

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119] and indicate requirement levels for compliant implementations.

3. Terminology

Distributed conferencing uses, when appropriate, and expands on the terminology introduced in the both the SIPPING [RFC4353] and XCON conferencing frameworks. The following additional terms are defined for specific use within the distributed conferencing context.

Conferencing Cloud:

> A specific pair composed of a centralized focus entity (XCON) and its associated distributed focus (DCON). We will herein indifferently use both "cloud" and "island" to refer to a conferencing cloud.

DCON Focus:

A specific entity enabling communication of a centralized conferencing system with the outside world. A DCON focus allows for the construction of a distributed conferencing system as a federation of centralized conferencing components.

Focus Discovery:

The capability to detect the presence of new focus entities in a distributed conferencing framework.

Information Spreading:

The spreading of conference related information among the focus entities in a distributed environment.

Protocol Dispatching:

The capabilty to appropriately forward/distribute messages of a natively centralized protocol in order to let them spread across a distributed environment.

Label Swapping:

The opportune swap of labels assigned to a specific resource, needed to avoid conflicts in the assignment of labels across several point-to-point communications regarding the same resource.

4. Overview

In order to build distributed conferencing on top of the already available centralized conferencing framework, we basically need to introduce two major functions: (i) a coordination level among conference focus entities; (ii) a way to effectively distribute conference state information. As to the first point above, the coordination level is needed in order to manage a distributed conference along its entire life-cycle. For instance, once a user decides to create a new conference, the corresponding conference focus has to distribute conference information to all other foci, in such a way as to enable other potential participants to retrieve the needed data and possibly subscribe to the event. We herein assume that all the operations needed inside a single conference cloud are managed via the protocols and interfaces defined inside the XCON working group. Hence, each single cloud keeps on being based on a star-topology graph for all what concerns the call signaling part. The various available stars are then connected through an upper-layer mesh-based topology providing inter-focus communication. As depicted in Figure 1, the overall topology of the distributed conferencing scenario thus looks like an overlay network of focus entities, each managing an underlying "centralized" conferencing island. In the most general case, we envisage to exploit extended Instant Messaging (IM) protocols for inter-focus communication.

```
                   o      o     o
                    \     |    /
                     \    |   /
                      +------+
                  o---| XCON |---o
                      +---^--+
                          |
                      +---v--+
                      | DCON |
                      +------+
                    //        \\
                   //          \\
                  //            \\
                 //              \\
                //                \\
               //                  \\
              //                    \\
             //                      \\
            //                        \\
        +------+                   +------+
        | DCON |===================| DCON |
        +---^--+                   +---^--+
            |                          |
        +---v--+                   +---v--+
    o---| XCON |---o           o---| XCON |---o
        +------+                   +------+
       /    |   \                 /    |   \
      /     |    \               /     |    \
     o      o     o             o      o     o
```

Figure 1: DCON architecture overview

As to the second point mentioned above, it looks clear that a way to propagate information about conferences is needed when switching the view from a centralized to a distributed perspective. Indeed, whenever a new conference is created (or an active conference changes its state) such an event has to be communicated to all interested (or active) participants. Given the intrinsic nature of the distributed framework (which actually expands the centralized one through the introduction of an overlay network of focus entities), the actual flow of information will always foresee the interaction among conference focus entities for both conference information exchanging and state changes notifications. The same obviously applies also to the involved natively centralized protocols defined in the XCON framework. A suitable mechanism has to be defined allowing for the dispatching of such centralized messages across the DCON network. The mechanism in question must be fully compliant with the existing operation of XCON clouds, which must keep their local participants totally unaware of the potential distributed nature of conferences.

Conference state propagation can take place in a number of alternative ways. For instance, each focus might flood the received information across the inter-focus communication mesh, thus guaranteeing that potential participants belonging to heterogeneous islands can be reached. In such case, focus entities are "stateful", i.e. each of them stores information about current sessions and forwards such information to all peering entities in order to get them up-to-date with respect to available conference sessions.

On the other hand, a distributed repository might be employed for the sake of storing conference information. Focus entities would access such repository, both to publish (either upon creation of a new conference, or to notify a change in the state of an active conference) and to retrieve information about active conferences (e.g. when a new participant wants to access the list of ongoing/ scheduled conference sessions he might be interested to join). In this last case, focus entities are "stateless".

Finally, a pure peer-to-peer approach can also be exploited for the purpose of conference state information spreading.

5. Towards Distributed Conferencing

In this section we first describe the overall architecture of a distributed conferencing framework, by highlighting both the involved entities and their interrelations. Then, we delve into the details of some key use cases which help understand the typical interaction paradigm of a decentralized environment.

As to the architecture, Figure 2 depicts how various XCON islands (just two in the picture to avoid confusion) can interact through the exchanging of synchronization messages between each pair of conferencing systems. Such messages are needed in order to circulate conference information among all involved entities. A dedicated protocol is obviously needed to take care of the communication between each pair: since its task is to synchronize the XCON and DCON pair, it will from now on be called XCON-DCON Synchronization Protocol (XDSP). The requirements for this protocol are further analyzed in a companion draft [I-D.romano-dcon-xdsp-reqs].

Inter-island coordination can be achieved via a number of available solutions (e.g. SIP/SIMPLE, XMPP). In this document we propose the exploitation of IM-based interaction. More precisely, we think that the Server-to-Server (S2S) module based on the XMPP protocol perfectly satisfies the requirements imposed by the new architecture.

Finally, media streams will directly flow between the XCON clouds once a distributed conference has been setup. Distributed mixing, however, will be only marginally discussed in this draft, in favour of the distribution of signaling and control messages.

```
  +-DCON-----------+                                 +-DCON-----------+
  | +------------+ |                                 | +------------+ |
  | | Dispatcher | <=== Inter-focus communication ===> | Dispatcher | |
  | +------------+ |                                 | +------------+ |
  +----------^-----+                                 +----^-----------+
     ^       |                                            |        ^
     |       | XDSP                                  XDSP |        |
     |       |                                            |        |
 +---|-------|----XCON-+                         +-XCON---|--------|---+
 |   |     +-v-------+ |                         | +------v--+     |   |
 |   |     | Gateway | |                         | | Gateway |     |   |
 |   |     +--^---^--+ |                         | +--^---^--+     |   |
 |   |    BFCP|   |CCP |                         | CCP|   |BFCP    |   |
 |   |        |   |    |                         |    |   |        |   |
 |   |     +--v---v--+ |                         | +--v---v--+     |   |
 |   +-----o Conf.   | |                         | | Conf.   o-----+   |
 |  Notify | Object  |<======= Media Flow ========>| Object  | Notify  |
 |         |         | |                         | |         |         |
 |         +---------+ |                         | +---------+         |
 +---------------------+                         +---------------------+

     XCON Cloud 1                                      XCON Cloud 2
```

Figure 2: Distributed conferencing framework

5.1. Setting up a distributed conferencing environment

In the following we are going to describe the typically required steps to setup a distributed conferencing environment. As described in the introductory sections, an overlay network of focus entities, each managing an underlying "centralized" conferencing island, will be needed, and the following points will help clarify how to effectively setup a distributed conferencing and manage it.

1. Overlay Creation and Management

   To enable effective operation of the distributed conferencing framework, an overlay network made of all interconnected conferencing clouds MUST be created. As an example, the mentioned overlay MAY be built by interconnecting all focus entities (with each such entity being the root of a local centralized conferencing cloud) through a full-meshed topology. Once the overlay is created, appropriate management of its structure SHOULD be envisaged; this includes, for example, dynamic updating of the topology information at the occurrence of relevant events (grafting/pruning of new centralized conferencing islands, etc.).

2. Focus Discovery

   An appropriate mechanism for the discovery of peering focus entities SHOULD be provided. Given the sensitive nature of the shared information, an appropriate authentication mechanism SHOULD be adopted. The trigger of the discovery process MAY be related to the concept of "presence"; in such case, an Instant Messaging (IM) based paradigm is RECOMMENDED. Alternatively, a logically centralized, physically distributed repository (e.g. UDDI) MAY be employed as a single reference point for the discovery of peering entities. A pure peer-to-peer approach can also be exploited for the same purpose.

3. Self-configuration

   At the occurrence of events like the grafting of a new cloud onto the overlay distributed conferencing network, the needed configuration steps SHOULD be performed in an automated fashion. This entails that all the news are appropriately exchanged across the overlay and, if needed, notified to the underlying centralized clouds as well.

4. Information Sharing

   The core of the operation of a distributed conferencing framework resides in the possibility to exchange information among all involved entities. The information sharing process SHOULD be made as effective as possible, e.g. by limiting the information that is forwarded outside a single centralized conferencing cloud to the data that are strictly necessary in order to guarantee that the overall state of the overaly is consistent, yet not redundant. Information sharing MAY be achieved either by exploiting a request/response paradigm, or through the adoption of asynchronous notification messages. A combined use of both the aforementioned paradigms is RECOMMENDED.

5. Dynamic Update

   All the clouds participating in the distributed overlay MUST keep the peers updated with respect to worth-noting events happening in their local realm. This MAY be achieved either by exploiting a request/response paradigm, or through the adoption of asynchronous notification messages. A combined use of both the aforementioned paradigms is RECOMMENDED. A pure peer-to-peer approach can also be exploited for the same purpose.

6. Distributed Conference Management

   In order to allow users' access to remotely created conferences, appropriate mechanisms MUST be provided by the framework. Such mechanisms SHOULD enable transparent management of both locally- and remotely-created conference instances. A pure peer-to-peer approach can be exploited for the same purpose.

7. Centralized Protocols Routing and Dispatching

   Focus entities MUST forward any centralized protocol message to their related peer in the distributed overlay whenever the message is directed to a receiver who does not belong to the local centralized system. Natively centralized protocol messages include, but are not limited to, any protocol defined and specified in the XCON framework (e.g. conference control management and floor control) as well as DTMF messages propagation. An example could be BFCP [RFC4582] messages the local floor control server might need to send to a user who is remotely participating in the conference (because she/he does not belong to the current XCON cloud). Another example concerns BFCP messages a local user might send to the remote floor control server handling the remote distributed conference she/he is involved in. Any message sent by local entities to local entities has to be treated in the usual centralized way according to the relative protocol specifications (i.e. dispatching shall not be involved).

8. Distributed Mixing

As soon as two or more centralized conferencing islands get connected in order to provide for a distributed conferencing scenario, the need arises to cope with the issue of mixing media flows generated by the conference participants. This challenging issue is currently considered out-of-scope in this document, which mainly focuses on the distribution of conference signalling/control information rather than addressing media management.

5.2. Use-case scenarios and examples

In this subsection we provide some examples of the operation of the distributed conferencing framework.

5.2.1. Creating a new distributed conference

Figure 3 illustrates how a distributed conference can be created and managed in a distributed environment. A participant contacts its corresponding focus entity in order to request the creation of a new conference instance. With respect to the centralized scenario, upon conference instantiation, in this case the focus has to publish conference information by notifying its related DCON focus. This is done in order to allow other remote focus entities to get up-to-date information about available conferencing sessions.

```
   Client                     XCON                   DCON
     |                          |                       |
     |  Request creation of     |                       |
     |  a new XCON conference   |                       |
     |------------------------->|                       |
     |                          |--+ Schedule           |
     |                          |  | new XCON           |
     |                          |<-+ conference         |
     |                          |                       |
     |                          |  Notify scheduling    |
     |                          |  of new conference    |
     |                          |---------------------->|
     |                          |                       |--+ Manage
     |                          |                       |  | new XCON
     |                          |                       |<-+ conference
     |                          |                       |
     |                          |                       | Spread new
     |                          |                       | information
     |                          |                       |-------->  ~~~
     .                          .                       .
     .                          .                       .
```

Figure 3: Creating a new conference

5.2.2. Retrieving information about available conferences

Figure 4 illustrates how information about available centralized and distributed conferences can be retrieved. A participant contacts its corresponding focus entity in order to request the above information. With respect to the centralized scenario, upon reception of a participant's request, the XCON focus has to forward the request to the related DCON focus. It will be up to the distributed focus entity to provide such information, which will include the list of both centralized (local) and distributed (remote) conferences. This way, a participant will be able to transparently keep on contacting the XCON focus to get all the information she/he needs in both cases.

```
   Client                    XCON                     DCON
     |                         |                        |
     | Request list of         |                        |
     | available conferences   |                        |
     |------------------------>|                        |
     |                         |--+ Process             |
     |                         |  | client's            |
     |                         |<-+ request             |
     |                         |                        |
     |                         | Forward request        |
     |                         | to DCON focus          |
     |                         |----------------------->|
     |                         |                        |--+ Process
     |                         |                        |  | forwarded
     |                         |                        |<-+ request
     |                         |  Send conferences list |
     |   Send conferences list |<-----------------------|
     |<------------------------|                        |
     .                         .                        .
     .                         .                        .
```

Figure 4: Retrieving information about available conferences

5.2.3. Joining a conference hosted by a foreign island

Figure 5 illustrates how a participant can join a conference which is managed by a focus entity belonging to a foreign centralized island. The whole sequence diagram has been split in three parts to better help understanding all the required steps. A participant contacts its corresponding focus entity in order to send the join request. With respect to the centralized scenario, upon reception of the participant's request, the local focus has to forward join information to the focus entity belonging to the island in which the conference in question was created.

The following steps are performed in sequence:

1. once the client has locally joined the distributed conference by placing a SIP call to the focus she/he belongs to (XCON (A)), the focus chooses a new label for the client (A) which will be needed to opportunely dispatch all the messages related to her/him;

2. XCON (A) at this point forwards the join request to its related DCON focus entity (DCON (A)); in this example this is done by sending, through the XDSP protocol, a message called AddUser

containing the newly assigned client's label A;

3. DCON (A) receives the join request; since it regards a new client, the DCON focus entity chooses a new label (e.g. XYZ) and associates it with the just received label A; depending on the distributed conference the client wants to join, it associates the label (XYZ) with the DCON focus entity managing the XCON focus physically hosting the conference (DCON (B)) and forwards the join request to it;

4. DCON (B) receives the forwarded message through the XMPP-based S2S channel; since it regards a new client, DCON (B) chooses a new label (e.g. B) and associates it with the just received label XYZ; since the conference the remote client (A) wants to join is physically hosted by XCON (B), the join request is forwarded there using the XDSP protocol, with an AddUser message containing the newly assigned label B which identifies the remote client;

5. XCON (B) receives the request, and thus associates the received label B with the remote Client (A); all the operations needed to add the new user to the conference are performed, and the new information is sent back to the client through the same path. All the involved labels (B, XYZ, A) will be opportunely swapped to route all the XCON protocols messages between the two entities.

Once XCON (A) receives the confirmation that the user has been successfully added to the remote conference, together with the needed information, the client (A) is updated through a SIP REINVITE containing the BFCP information she/he will need to communicate with the Floor Control Server. All BFCP messages sent from now on by the client to the Floor Control Server will be intercepted by the gateway, and then forwarded to the Floor Control Server of XCON (B). This case will be furtherly presented and discussed in the next section.

```
 Client(A)              XCON(A)           DCON(A)                DCON(B)
    |                      |                  |                       |
    |  SIP INVITE          |                  |                       |
    |--------------------->|                  |                       |
    |         SIP Trying   |                  |                       |
    |<---------------------|                  |                       |
    |             SIP OK   |                  |                       |
    |<---------------------|                  |                       |
```

```
    |  SIP ACK            |                 |                      |
    |-------------------->|                 |                      |
    |                     |--+ Choose a     |                      |
    |                     |  | Label (A)    |                      |
    |                     |<-+ for new user |                      |
    |                     |                 |                      |
    |                     |  AddUser (A)    |                      |
    |                     |---------------->|                      |
    |                     |                 |--+ Choose a Label    |
    |                     |                 |  | (XYZ) and find    |
    |                     |                 |  | destination       |
    |                     |                 |<-+ (DCON (B))        |
    |                     |                 |                      |
    |                     |                 |--+ Label             |
    |                     |                 |  | Swap              |
    |                     |                 |<-+ (A=>XYZ)          |
    |                     |                 |                      |
    |                     |                 |  XMPP (AddUser)      |
    |                     |                 |----  ~~(S2S)~~  ---->|
    |                     |                 |                      |
    .                     .                 .                      .
    .                     .                 .                      .


            DCON(A)              DCON(B)             XCON(B)
              .                    .                    .
              .                    .                    .
              |                    |                    |
              |--+ Label           |                    |
              |  | Swap            |                    |
              |<-+ (A=>XYZ)        |                    |
              |                    |                    |
              |  XMPP (AddUser)    |                    |
              |---  ~~(S2S)~~  --->|                    |
              |                    |--+ Choose a        |
              |                    |  | Label (B)       |
              |                    |<-+ for new user    |
              |                    |                    |
              |                    |  AddUser           |
              |                    |------------------->|
              |                    |                    |--+ Assign new
              |                    |                    |  | User ID
              |                    |                    |  | to remote
              |                    |                    |<-+ participant
              |                    |      UserAdded (B) |
              |                    |<-------------------|
              |           Label +--|                    |
```

```
              |            Swap |  |                    |
              |        (B=>XYZ) +->|                    |
              |                    |                    |
              |  XMPP (UserAdded)  |                    |
              |<---  ~~(S2S)~~  ---|                    |
     Label +--|                    |                    |
      Swap |  |                    |                    |
  (XYZ=>A) +->|                    |                    |
              |                    |                    |
              .                    .                    .
              .                    .                    .


  Client(A)             XCON(A)          DCON(A)                DCON(B)
    .                     .                 .                      .
    .                     .                 .                      .
    |                     |                 |                      |
    |                     |                 |     XMPP (UserAdded) |
    |                     |                 |<----  ~~(S2S)~~  ----|
    |                     |        Label +--|                      |
    |                     |         Swap |  |                      |
    |                     |     (XYZ=>A) +->|                      |
    |                     |                 |                      |
    |                     |   UserAdded (A) |                      |
    |                     |<----------------|                      |
    |  Manage received +--|                 |                      |
    | info on new user |  |                 |                      |
    |       (e.g. IDs) +->|                 |                      |
    |                     |                 |                      |
    |                     |                 |                      |
    |SIP REINVITE (+BFCP) |                 |                      |
    |<--------------------|                 |                      |
    |  SIP Trying         |                 |                      |
    |-------------------->|                 |                      |
    |  SIP OK             |                 |                      |
    |-------------------->|                 |                      |
    |            SIP ACK  |                 |                      |
    |<--------------------|                 |                      |
    |                     |                 |                      |
    .                     .                 .                      .
    .                     .                 .                      .
    .                     .                 .                      .
```

Figure 5: Joining a foreign conference

5.2.4. Dispatching XCON protocols in DCON

Figure 6 illustrates how natively centralized XCON protocols (BFCP, in the figure) can be opportunely dispatched in order to let them spread across a distributed environment. Such mechanism would allow users participating in distributed conferences to avoid knowing the transport addresses needed to communicate with remote focus entities, and to keep transparently referring to the local focus entity instead.

In order to understand who the actual receiver of a message shall be, all messages are intercepted by a logical entity, called Gateway, belonging to the XCON focus. The Gateway will understand whether a message is directed to a local entity (e.g. a user belonging to the XCON focus, or the local Floor Control Server) or to a remote entity belonging to another focus (e.g. a remotely participating user, or a remote Floor Control Server).

```
   +----------------------------------------------------------+
   | Client 1: Label A (Server Leg: FCS)                      |
   | Client 2: Label B (Server Leg: Remote FCS-->Dispatcher)  |
   | Client 3: Label C (Server Leg: FCS)                      |
   +----------------------------------------------------------+

                   +--DCON-------------+
                   |                   |
                   |    +------------+ |
                   |    | Dispatcher |<=== (BFCP: Label Swap) ===...>
                   |    +------------+ |
                   +-------------^-----+
                                 |
                                 |  XDSP Message: Label B
                                 | (BFCP encoded in Base64)
                                 |
                           +-----|-------------------XCON--+
                           |     |                         |
                           |     +--- Notify (B) ---+      |
                           |     |                  |      |
   +----------+            |     v                  v      |
   | Client 1 |<---+       | +---------+       +---------+ |
   +----------+    |       | |  Floor  |<--A-->|  Floor  | |
                   +-A------>| Control |       | Control | |
   +~~~~~~~~~~+            | |  Server |<--C-->|  Server | |
   i Client 2 i<------B----->| Gateway |       +---------+ |
   +~~~~~~~~~~+            | +---------+                   |
                           +---^---------------------------+
   +----------+                |
   | Client 3 |<-------C-------+
   +----------+
```

Figure 6: Centralized protocols dispatching

To make the whole thing clearer, the example in figure Figure 7 will be used. As in the previous case, the whole sequence diagram has been split in three parts to better help understand all the required steps. In this example, a user (Client (A)) belonging to XCON (A) is remotely participating to a distributed conference hosted by XCON (B). Since XCON (B) is physically hosting the conference, floor control will be entirely managed by its Floor Control Server. To allow Client (A) to communicate with Floor Control Server (B) and viceversa, appropriate dispatching of BFCP messages between the two peers will be needed. We have already seen how labels are assigned and swapped: the same labels will be used for dispatching.

The flow of a typical message exchange can be seen as follows:

1. The Client (A) sends a BFCP message to the Floor Control Server; the message is intercepted by XCON (A)'s gateway; the label assigned to client (A) is retrieved, and used to forward the BFCP message to the DCON (A) Dispatcher; of course, since BFCP messages are binary, an opportune treatment (e.g. through Base64 encoding) should be done to encapsulate the message in a text-based protocol message (as XDSP will probably be);

2. Once DCON (A) receives the encapsulated BFCP message, the labels are opportunely swapped (in this case, A=>XYZ) and the message is routed to the right destination (DCON (B));

3. DCON (B) will receive the message and swap labels again (XYZ=>B); at this point, the encapsulated message will be forwarded to the underlying XCON (B) Gateway to be further processed there;

4. The XCON (B) Gateway will receive the encapsulated (and probably Base64-encoded) BFCP message; after decoding it (if needed), the Gateway will analyze the label marked in the message (B, in this case), and will understand it is a message sent by a remote user (Client (A)) to the local Floor Control Server. It will forward the (now 'natively' binary) message there, where it will be processed;

5. In case the FCS (B) needs to send a message to Client (A), exactly the same operations will be performed, and the same path will be followed through the needed label swaps among the involved peers. FCS (A), while not actually managing the floors related to the remote conference Client (A) is participanting in, will however be notified upon each floor status change, so to opportunely update the local media mixes when needed (e.g. to mute Client (A) excluding her/him from XCON (A)'s local mix if FCS (B) has decided so).

```
 Client(A)           XCON(A)               DCON(A)               DCON(B)
   |               (Gateway)                |                     |
   |                    |                   |                     |
   |  BFCP Message      |                   |                     |
   |------------------->|                   |                     |
   |                    |--+ Get Label (A)  |                     |
   |                    |  | assigned to    |                     |
   |                    |<-+ client/FCS     |                     |
```

```
   |                    |                    |                     |
   |                    |  BFCP encoded      |                     |
   |                    |  in Base64         |                     |
   |                    |--(Label A)--------->|                     |
   |                    |                    |--+ Label            |
   |                    |                    |  | Swap             |
   |                    |                    |<-+ (A=>XYZ)         |
   |                    |                    |                     |
   |                    |                    |--+ Get destination  |
   |                    |                    |  | from label XYZ   |
   |                    |                    |<-+ (DCON B)         |
   |                    |                    |                     |
   |                    |                    | XMPP                |
   |                    |                    | (BFCP in Base64)    |
   |                    |                    |----  ~~(S2S)~~  --->|
   |                    |                    |                     |
   .                    .                    .                     .
   .                    .                    .                     .


 DCON(A)               DCON(B)            XCON(B)             XCON(B)
   |                     |               (Gateway)             (FCS)
   .                     .                   .                   .
   .                     .                   .                   .
   | XMPP                |                   |                   |
   | (BFCP in Base64)    |                   |                   |
   |---  ~~(S2S)~~  ---->|                   |                   |
   |                     |                   |                   |
   |                     |--+ Label          |                   |
   |                     |  | Swap           |                   |
   |                     |<-+ (XYZ=>B)       |                   |
   |                     |                   |                   |
   |                     |  BFCP encoded     |                   |
   |                     |  in Base64        |                   |
   |                     |-----(Label B)---->|                   |
   |                     |                   |--+ Check Label (B)|
   |                     |                   |  | assigned to    |
   |                     |                   |<-+ FCS/client     |
   |                     |                   |                   |
   |                     |                   |  BFCP Message     |
   |                     |                   |------------------>|
   .                     .                   .                   .
   .                     .                   .                   .
   |                     |                   |     BFCP Message  |
   |                     |                   |<------------------|
   |                     |  Get Label (B) +--|                   |
   |                     |    assigned to |  |                   |
```

```
   |                     |     FCS/client +->|                   |
   |                     |                   |                   |
   |                     |      BFCP encoded |                   |
   |                     |         in Base64 |                   |
   |                     |<-----(Label B)----|                   |
   |            Label +--|                   |                   |
   |             Swap |  |                   |                   |
   |         (B=>XYZ) +->|                   |                   |
   |                     |                   |                   |
   |  Get destination +--|                   |                   |
   |   from label XYZ |  |                   |                   |
   |         (DCON A) +->|                   |                   |
   |                     |                   |                   |
   |                XMPP |                   |                   |
   |    (BFCP in Base64) |                   |                   |
   |<---  ~~(S2S)~~  ----|                   |                   |
   |                     |                   |                   |
   .                     .                   .                   .
   .                     .                   .                   .


 Client(A)           XCON(A)               DCON(A)               DCON(B)
   |                (Gateway)                |                      |
   |                    |                    |                      |
   |                    |                    |                XMPP  |
   |                    |                    |    (BFCP in Base64)  |
   |                    |                    |<----  ~~(S2S)~~  --- |
   |                    |           Label +--|                      |
   |                    |            Swap |  |                      |
   |                    |        (XYZ=>A) +->|                      |
   |                    |                    |                      |
   |                    |       BFCP encoded |                      |
   |                    |          in Base64 |                      |
   |                    |<---------(Label A)-|                      |
   | Check Label (A) +--|                    |                      |
   |     assigned to |  |                    |                      |
   |      client/FCS +->|                    |                      |
   |                    |                    |                      |
   |      BFCP Message  |                    |                      |
   |<-------------------|                    |                      |
   |                    |                    |                      |
   .                    .                    .                      .
   .                    .                    .                      .
   .                    .                    .                      .
```

Figure 7: An example: dispatching a BFCP message

6. Security Considerations

TBD...

7. References

[RFC2234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.

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

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, February 2006.

[RFC4582] Camarillo, G., Ott, J., and K. Drage, "The Binary Floor Control Protocol (BFCP)", RFC 4582, November 2006.

[I-D.romano-dcon-requirements]
Romano, S., Amirante, A., Castaldi, T., Miniero, L., and A. Buono, "Requirements for Distributed Conferencing", draft-romano-dcon-requirements-08 (work in progress), December 2010.

[I-D.romano-dcon-xdsp-reqs]
Romano, S., Amirante, A., Castaldi, T., Miniero, L., and A. Buono, "Requirements for the XCON-DCON Synchronization Protocol", draft-romano-dcon-xdsp-reqs-08 (work in progress), December 2010.

[RFC5239] Barnes, M., Boulton, C., and O. Levin, "A Framework for Centralized Conferencing", RFC 5239, June 2008.

[I-D.ietf-xcon-common-data-model]
Novo, O., Camarillo, G., Morgan, D., and J. Urpalainen, "Conference Information Data Model for Centralized Conferencing (XCON)", draft-ietf-xcon-common-data-model-31 (work in progress), June 2011.

Authors' Addresses

Simon Pietro Romano
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: spromano@unina.it

Alessandro Amirante
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: alessandro.amirante@unina.it

Tobia Castaldi
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: tcastaldi@meetecho.com

Lorenzo Miniero
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: lorenzo@meetecho.com

Alfonso Buono
Ansaldo Trasporti e Sistemi Ferroviari
Via Argine, 425
Napoli 80147
Italy

Email: alfonso.buono@atsf.it

Network Working Group
Internet-Draft
Expires: December 22, 2011

S P. Romano
A. Amirante
University of Napoli
T. Castaldi
L. Miniero
Meetecho
A. Buono
Ansaldo Trasporti e Sistemi
Ferroviari

June 20, 2011

Requirements for Distributed Conferencing
draft-romano-dcon-requirements-09

Abstract

This document examines the requirements for Distributed Conferencing (DCON). Separate documents will map the requirements to existing protocol primitives, define new protocol extensions, and introduce new protocols as needed. Together, these documents will provide a guideline for building interoperable conferencing applications. The current works in SIPPING and XCON working groups marginally address the matter, which is nonetheless considered as out-of-scope. The requirements listed in this document are in part based on thoughts derived from the cited working groups activities.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

This document examines the requirements for an architecture capable to provide a distributed conferencing service. It draws inspiration from a number of existing research efforts inside the IETF, mainly in the context of both the SIPPING and the XCON WGs. We will herein present high-level requirements, starting from considerations upon the well-known concept of cascaded conferencing [RFC5239][RFC4575].

2. Conventions

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119] and indicate requirement levels for compliant implementations.

3. Terminology

Distributed conferencing uses, when appropriate, and expands on the terminology introduced in both the SIPPING [RFC2119] and XCON [RFC5239] conferencing frameworks. The following additional terms are defined for specific use within the distributed conferencing work.

Focus Discovery -- this term refers to the capability to detect the presence of new focus entities in a distributed conferencing framework.

Information Spreading -- this term refers to the spreading of conference related information among the focus entities in a distributed environment.

Protocol Dispatching -- this term refers to the capabilty of appropriately forwarding/distributing messages of a natively centralized protocol in order to let them spread across a distributed environment.

DCON Focus -- this term refers to a specific entity enabling communication of a centralized conferencing system with the outside world. A DCON focus allows for the construction of a distributed conferencing system as a federation of centralized conferencing components.

Conferencing Cloud -- this term refers to a specific pair composed of a centralized focus entity (XCON) and its associated distributed focus (DCON). We will herein indifferently use both "cloud" and "island" to refer to a conferencing cloud.

4. Related work: Cascaded Conferencing

The requirements for a distributed conferencing framework have already been partially addressed in previous works within the IETF. Specifically, RFC 4245 (High-Level Requirements for Tightly Coupled SIP Conferencing) [RFC4245] introduces the concept of cascading of conferences and illustrates three different scenarios to which it might be applied: (i) peer-to-peer chaining of signaling; (ii) conferences having hierarchal signaling relations; (iii) cascading as a means to distribute the media "mixing". For the three scenarios above, a number of possible requirements are identified, among which the availability of a SIMPLE-based Presence and Instant Messaging architecture plays a major role.

The concept of cascaded conferences is further expanded in RFC 4353 [RFC4353] (A Framework for Conferencing with the Session Initiation Protocol (SIP)), where the term "Cascaded Conferencing" is used to indicate "a mechanism for group communications in which a set of conferences are linked by having their focuses interact in some fashion". In the same document, a specific scenario called "Simplex Cascaded Conferences" is presented as a typical interaction paradigm envisaging that the user agent representing the focus of one conference is a conference-unaware participant in another conference. In other terms, a conference "calls" another conference and gets connected to it as if it were a simple participant. For both such conferences, the peering party is just like any other user participating in the conferencing session. For the sake of completeness, we remark that the previous observation is somehow confuted by RFC 4575 (A Session Initiation Protocol Event Package for Conference State) [RFC4353], which explicitely states:

> "It is possible that a participant in the conference may in fact be another focus. In order to provide a more complete participant list, the focus MAY subscribe to the conference package of the other focus to discover the participant list in the cascaded conference. This information can then be included in notifications by use of the <cascaded-focus> element as specified by this package".

Even though the simplex cascaded conferencing is an established way to concatenate conferences, we claim that it is not flexible enough to effectively cope with a number of potential distributed conferencing scenarios. More precisely, we envisage a situation where an overlay network infrastructure is in charge of managing distributed conferences, whereas the sinlge focus entities keep on managing their own centralized "realm". As it will come out in the next section, this entails that a specific requirement is formulated about the need for explicit management of distributed conference information.

5. Requirements

In the following we are going to list the requirements we have identified for distributed conferencing. Each requirement is presented in general terms and some examples about its applicability are provided.

REQ-1: Overlay Creation and Management

To enable effective operation of the distributed conferencing framework, an overlay network made of all interconnected conferencing clouds MUST be created. As an example, the mentioned overlay MAY be built by interconnecting all focus entities (with each such entity being the root of a local centralized conferencing cloud) through a full-meshed topology. Once the overlay is created, appropriate management of its structure SHOULD be envisaged; this includes, for example, dynamic updating of the topology information at the occurrence of relevant events (grafting/ pruning of new centralized conferencing islands, etc.).

REQ-2: Focus Discovery

An appropriate mechanism for the discovery of peering focus entities SHOULD be provided. Given the sensitive nature of the shared information, an appropriate authentication mechanism SHOULD be adopted. The trigger of the discovery process MAY be related to the concept of "presence"; in such case, an Instant Messaging (IM) based paradigm is RECOMMENDED. Alternatively, a logically centralized, physically distributed repository (e.g. UDDI) MAY be employed as a single reference point for the discovery of peering entities. A pure peer-to-peer approach can also be exploited for the same purpose.

REQ-3: Self-configuration

At the occurrence of events like the grafting of a new cloud onto the overlay distributed conferencing network, the needed configuration steps SHOULD be performed in an automated fashion. This entails that all the news are appropriately exchanged across the overlay and, if needed, notified to the underlying centralized clouds as well.

REQ-4: Information Sharing

The core of the operation of a distributed conferencing framework resides in the possibility to exchange information among all involved entities. The information sharing process SHOULD be made as effective as possible, e.g. by limiting the information that is forwarded outside a single centralized conferencing cloud to the data that are strictly necessary in order to guarantee that the overall state of the overaly is consistent, yet not redundant. Information sharing MAY be achieved either by exploiting a request/response paradigm, or through the adoption of asynchronous notification messages. A combined use of both the aforementioned paradigms is RECOMMENDED.

REQ-5: Dynamic Update

All the clouds participating in the distributed overlay MUST keep the peers updated with respect to worth-noting events happening in their local realm. This MAY be achieved either by exploiting a request/response paradigm, or through the adoption of asynchronous notification messages. A combined use of both the aforementioned paradigms is RECOMMENDED. A pure peer-to-peer approach can also be exploited for the same purpose.

REQ-6: Distributed Conference Management

In order to allow users' access to remotely created conferences, appropriate mechanisms MUST be provided by the framework. Such mechanisms SHOULD enable transparent management of both locally- and remotely-created conference instances. A pure peer-to-peer approach can be exploited for the same purpose.

REQ-7: Centralized Protocols Routing and Dispatching

Focus entities MUST forward any centralized protocol message to their related peer in the distributed overlay whenever the message is directed to a receiver who does not belong to the local centralized system. Natively centralized protocol messages include, but are not limited to, any protocol defined and specified in the XCON framework (e.g. conference control management and floor control) as well as DTMF messages propagation. An example could be BFCP [RFC4582] messages the local floor control server might need to send to a user who is remotely participating in the conference (because she/he does not belong to the current XCON cloud). Another example concerns BFCP messages a local user might send to the remote floor control server handling the remote distributed conference she/he is involved in. Any message sent by local entities to local entities has to be treated in the usual centralized way according to the relative protocol specifications (i.e. dispatching shall not be involved).

REQ-8: Distributed Mixing

As soon as two or more centralized conferencing islands get connected in order to provide for a distributed conferencing scenario, the need arises to cope with the issue of mixing media flows generated by the conference participants. This challenging issue is currently considered out-of-scope in this document, which mainly focuses on the distribution of conference signalling/control information rather than addressing media management.

6. Security Considerations

The communication between each distributed focus entity contains sensitive information, since it envisages the possibility to spread important data that only authorized parties should know (e.g. the full internal state of the centralized conference objects and relevant privacy information about users authenticated by the system).

Hence it is very important that protocol messages be protected because otherwise an attacker might spoof the legitimate identity of the distributed focus entity or inject messages on its behalf.

To mitigate the above threats, all focus entities SHOULD mutually authenticate upon initial contact. All protocol messages SHOULD be authenticated and integrity-protected to prevent third-party intervention and MITM (Man-In-The-Middle) attacks. All messages SHOULD be encrypted to prevent eavesdropping.

7. References

[RFC2234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.

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

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, "A Session Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, August 2006.

[RFC4245] Levin, O. and R. Even, "High-Level Requirements for Tightly Coupled SIP Conferencing", RFC 4245, November 2005.

[RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, February 2006.

[RFC4582] Camarillo, G., Ott, J., and K. Drage, "The Binary Floor Control Protocol (BFCP)", RFC 4582, November 2006.

[RFC5239] Barnes, M., Boulton, C., and O. Levin, "A Framework for Centralized Conferencing", RFC 5239, June 2008.

Authors' Addresses

Simon Pietro Romano
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: spromano@unina.it

Alessandro Amirante
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: alessandro.amirante@unina.it

Tobia Castaldi
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: tcastaldi@meetecho.com

Lorenzo Miniero
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: lorenzo@meetecho.com

Alfonso Buono
Ansaldo Trasporti e Sistemi Ferroviari
Via Argine, 425
Napoli 80147
Italy

Email: alfonso.buono@atsf.it

Network Working Group S P. Romano
Internet-Draft A. Amirante
Expires: December 22, 2011 University of Napoli
T. Castaldi
L. Miniero
Meetecho
A. Buono
Ansaldo Trasporti e Sistemi
Ferroviari
June 20, 2011

Requirements for the XCON-DCON Synchronization Protocol
draft-romano-dcon-xdsp-reqs-09

Abstract

The Distributed Conferencing (DCON) framework provides the means to distribute Centralized Conference (XCON) information by appropriately orchestrating a number of centralized focus entities (clouds). The mechanism we propose to make each XCON cloud communicate with its related DCON peer is based on the use of some kind of XCON-DCON Synchronization Protocol (XDSP). This document gives the requirements for XDSP.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

The Distributed Conferencing framework [I-D.romano-dcon-requirements] describes the requirements for the overall architecture, terminology, and protocol components needed for distribuited conferencing. DCON is based on the idea that a distributed conference can be setup and accessed by appropriately orchestrating the operation of a number of XCON "focus" elements, each in charge of managing a certain number of participants. Each pair composed of a centralized focus entity (XCON) and its related distributed counterpart (namely, a DCON focus) is called "island", or "cloud". These islands are then made part of an overlay network composed of several inter-communicating clouds.

Interaction between each participant and the corresponding conference focus is based on the standard XCON framework [RFC5239], whereas inter-focus interaction is based on a peer-to-peer paradigm. The interaction between the centralized conference focus and the distributed conference focus, instead, has requirements that are defined in this document.

2. Conventions

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119] and indicate requirement levels for compliant implementations.

3. Terminology

Distributed conferencing uses, when appropriate, and expands on the terminology introduced in the both the SIPPING [RFC4353] and XCON [RFC5239] conferencing frameworks. The following additional terms are defined for specific use within the distributed conferencing work.

Conferencing Cloud:

> A specific pair composed of a centralized focus entity (XCON) and its associated distributed focus (DCON). We will herein indifferently use both "cloud" and "island" to refer to a conferencing cloud.

DCON Focus:

A specific entity enabling communication of a centralized conferencing system with the outside world. A DCON focus allows for the construction of a distributed conferencing system as a federation of centralized conferencing components.

Focus Discovery:

The capability to detect the presence of new focus entities in a distributed conferencing framework.

Information Spreading:

The spreading of conference related information among the focus entities in a distributed environment.

Protocol Dispatching:

The capabilty of appropriately forwarding/distributing messages of a natively centralized protocol in order to let them spread across a distributed environment.

Label Swapping:

The opportune swap of labels assigned to a specific resource, needed to avoid conflicts in the assignment of labels across several point-to-point communications regarding the same resource.

4. XDSP Requirements

This section describes requirements for the XCON-DCON synchronization protocol (XDSP).

4.1. General Protocol Requirements

REQ-A1:

XDSP protocol MUST be a reliable client-server protocol. Hence, it MUST have a positive response indicating that the request has been received, or an error response in case an error has occurred.

REQ-A2:

It MUST be possible for the XCON focus entity, the server, to authenticate the related DCON focus entity, the client.

REQ-A3:

It MUST be possible for the DCON focus entity to be authenticated by the server, the related XCON focus entity.

REQ-A4:

It MUST be possible to ensure message integrity between each pair of XCON and DCON focus entities.

4.2. Requests and responses

REQ-B1:

It MUST be possible for the involved XCON and DCON entities to communicate on a stateless synchronous request-response based mechanism.

REQ-B2:

An error message MUST be sent back to the entity placing the request, in case the message couldn't be processed for any reason.

REQ-B3:

An authentication mechanism SHOULD be made possible on the basis of such stateless synchronous request-response based interaction between the two involved entities.

REQ-B4:

It SHOULD be possible for the XCON focus entity to request access to remote (e.g. avaliable on different islands) resources by means of an answer sent to the related DCON focus entity. This includes requesting a join to a remote conference for a local user, setting up distributed conferences, actively requesting the list of all the remote conferences and/or the list of all users (remote and local) in a currently running conference, etc..

REQ-B5:

The DCON focus entity SHALL forward any request directed to resources available in the related XCON cloud to the related XCON focus entity which will manage it and properly answer the request.

4.3. Updates and asynchronous notifications

REQ-C1:

It SHOULD be possible for the DCON focus entity to subscribe to the related XCON focus entity for events related to the conference system state, in order to receive asynchronous notifications.

REQ-C2:

The XCON focus entity SHALL generate new asynchronous notifications every time there is any change in the state of any of the conferences it is currently handling.

REQ-C3:

It SHOULD be possible for the DCON focus entity to receive full state updates from the related XCON focus entity, in case some of the events were missed, to make the known state consistent with the actual conference system internal state.

REQ-C4:

Both partial notifications and full updates SHOULD be sent through the same authenticated channel used for XDSP communication. In case a separate channel and/or a separate protocol are used (e.g. by means of the XCON event package, when it is available, or of the already available SIPPING conference event package [RFC4575]), the same issues about security and integrity SHOULD be addressed to avoid attacks and exploits by unauthenticated users.

REQ-C5:

Since state changes might happen in both the involved focus entities (even though related to different situations) the same requirements just described for notifications generated by XCON focus entities should be addressed for their related DCON focus entities. It SHOULD be possible for the XCON focus entity to subscribe to the related DCON focus entity for events related to the conference system state, in order to receive asynchronous notifications.

REQ-C6:

The DCON focus entity SHALL generate new asynchronous notifications every time there is any change in its internal state, e.g. whenever new remote conferences have been created or become active, etc.

REQ-C7:

It SHOULD be possible for the XCON focus entity to receive full state updates from the related DCON focus entity, in case some of the events were missed, to make the known state consistent with the actual conference system internal state.

4.4. Centralized protocols routing and dispatching

REQ-D1:

The XCON focus entity MUST forward any centralized protocol message to its related DCON focus entity whenever the message is to be received by a peer who is not a local entity of the centralized system. Natively centralized protocol messages include, but are not limited to, any protocol defined and specified in the XCON framework (e.g. conference control management and floor control) as well as DTMF messages propagation. An example is represented by BFCP messages the local floor control server might need to send to a user who is remotely (i.e. a user who does not belong to the current XCON cloud) participating in the conference. Another example concerns BFCP messages a local user might want to send to the remote floor control server handling the remote, distributed, conference the user is participating in. Any message sent by local entities to local entities has to be treated in the usual centralized way according to the relative protocol specifications (i.e. dispatching shall not be involved).

REQ-D2:

The DCON focus entity MUST forward any natively centralized protocol message it receives from DCON focus peers in the distributed overlay (routing) to the related XCON focus entity (dispatching), whenever the message is addressed to any of the local entities of the centralized cloud.

REQ-D3:

The XCON and DCON focus entities MUST establish and mantain opportune labels to correctly address and identify local entities involved in routed and dispatched messages. These labels MUST be appropriately swapped whenever they leave a DCON focus entity and reach a foreign one, so to avoid conflicts upon assigned labels in different islands.

REQ-D4:

Message dispatching between the two involved focus entities SHOULD occur on an request-response based communication mechanism, and opportune errors should be generated in case any exceptional condition happens while processing the messages.

5. Security Considerations

The communication between each DCON focus entity and its related XCON entity contains sensitive information, since it envisages the possibility to spread important information that only authorized entities should be aware of (e.g. the full internal state of the centralized conference objects and relevant privacy information about users authenticated by the system).

Hence it is very important that protocol messages be protected because otherwise an attacker might spoof the legitimate identity of the DCON focus entity and/or inject messages on his behalf. Many obvious consequences could come out of such an undesirable situation.

To mitigate the above threats, both the DCON focus entity and the XCON focus entity SHOULD be authenticated upon initial contact. All protocol messages SHOULD be authenticated and integrity-protected to prevent third-party intervention and MITM (Man-In-The-Middle) attacks. All messages SHOULD be encrypted to prevent eavesdropping.

6. Acknowledgements

7. References

[RFC2234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax

Specifications: ABNF", RFC 2234, November 1997.

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

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, February 2006.

[RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, "A Session Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, August 2006.

[I-D.romano-dcon-requirements]
Romano, S., Amirante, A., Castaldi, T., Miniero, L., and A. Buono, "Requirements for Distributed Conferencing", draft-romano-dcon-requirements-08 (work in progress), December 2010.

[RFC5239] Barnes, M., Boulton, C., and O. Levin, "A Framework for Centralized Conferencing", RFC 5239, June 2008.

Authors' Addresses

Simon Pietro Romano
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: spromano@unina.it

Alessandro Amirante
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Email: alessandro.amirante@unina.it

Tobia Castaldi
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: tcastaldi@meetecho.com

Lorenzo Miniero
Meetecho
Via Carlo Poerio 89
Napoli 80100
Italy

Email: lorenzo@meetecho.com

Alfonso Buono
Ansaldo Trasporti e Sistemi Ferroviari
Via Argine, 425
Napoli 80147
Italy

Email: alfonso.buono@atsf.it