DRINKS
Internet-Draft
Intended status: Informational
Expires: November 28, 2009

S. Channabasappa, Ed.
CableLabs

May 27, 2009

DRINKS Use cases and Protocol Requirements
draft-ietf-drinks-usecases-requirements-00

Status of this Memo

Copyright Notice

Abstract

This document captures the use cases and associated requirements for interfaces to provision session establishment data into SIP Service Provider components that aid with session routing. Specifically, the current version of this document focuses on the provisioning of one such element, termed the registry.

Table of Contents

1. Terminology 3
2. Overview 5
3. Use Cases and Requirements 10
  3.1. Registry Provisioning 10
    3.1.1. Use Cases 10
    3.1.2. Requirements 14
  3.2. Distribution of data into local data repositories 17
  3.3. Miscellaneous Use Cases 17
    3.3.1. Indirect Peering to Selected Destinations 17
    3.3.2. TBD: RN Destinations 17
4. Security Considerations 18
5. IANA Considerations 19
6. Acknowledgments 20
7. References 21
  7.1. Normative References 21
  7.2. Informative References 21
Author's Address 22

1. Terminology

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].

This document reuses terms from [RFC3261] (e.g., SIP) and [RFC5486] (e.g., LUF, LRF). In addition, this document specifies the following additional terms.

Registry: The authoritative source for provisioned session establishment data (SED) and related information.

Local Data Repository: The data store component of an addressing server that provides resolution responses.

Destination Group: A set of public identities that are grouped together to facilitate session setup and routing.

Public Identity: A generic term that refers to a telephone number (TN), an email address, or other identity as deemed appropriate, such as a globally routable URI of a user address (e.g., mailto:john.doe@example.net).

Routing Group: a grouping of SED records.

SED Record: A SED Record contains much of the session establishment data or a 'redirect' to another registry where the session establishment data can be discovered. SED Records types supported are NAPTRs, CNAME, DNAME, and NS Records.

Data Recipient: SP or SSP that receives or consumes SED and related information.

Data Recipient Group: A group of Data Recipients that receive the same set (or subset) of SED and related information.

2. Overview

The SPEERMINT WG specifies Session Establishment Data, or SED, as the data used to route a call to the next hop associated with the called domain's ingress point. More specifically, the SED is the set of parameters that the outgoing signalling path border elements (SBEs) need to complete the call. See [RFC5486] for more details.

The specification of the format and protocols to configure SED is a task taken up by the DRINKS WG. The use cases and requirements that have been proposed in this regard are compiled in this document.

SED is typically created by the terminating SSP and consumed by the originating SSP. For scalability reasons SED is rarely exchanged directly between the intended parties. Instead, it is exchanged via intermediate systems - termed Registries within this document. Such registries receive SED via provisioning transactions from other SSPs, and then distribute the received data into Local Data Repositories. These local data repositories are used for call routing by outgoing SBEs. This is depicted in Figure 1.

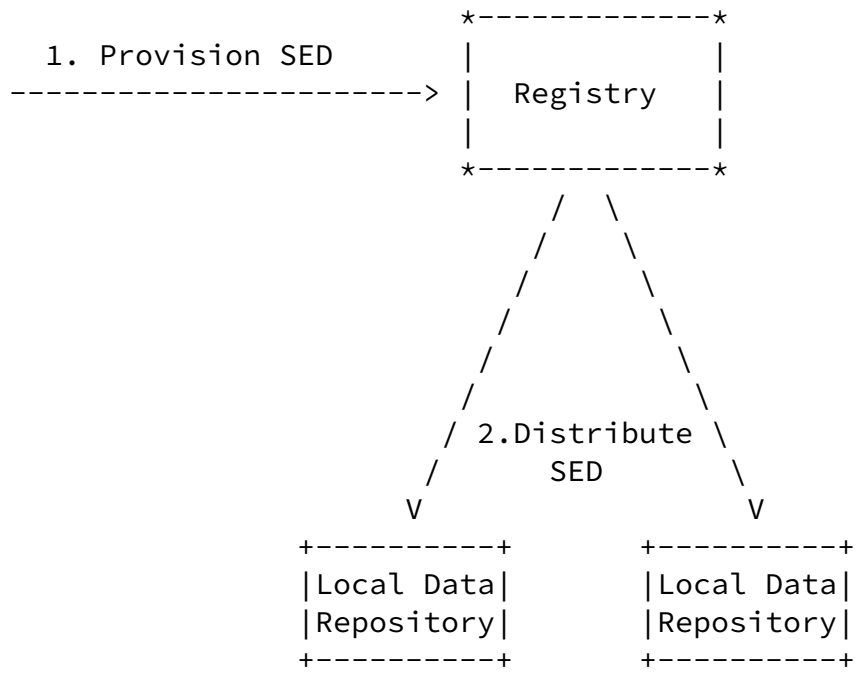

Figure 1: General Diagram

In this version of the document, we primarily address the use cases and requirements for provisioning registries. Future revisions may include data distribution. The resulting provisioning protocol can be used to provision data into a registry, or between registries. This is depicted in Figure 2.

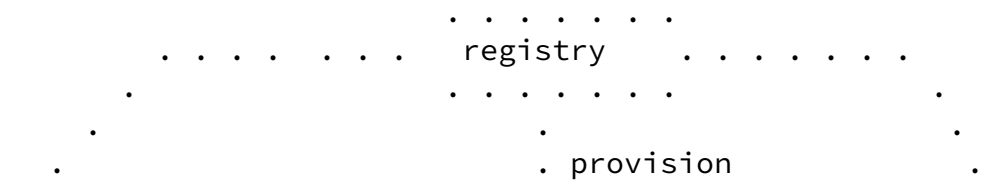

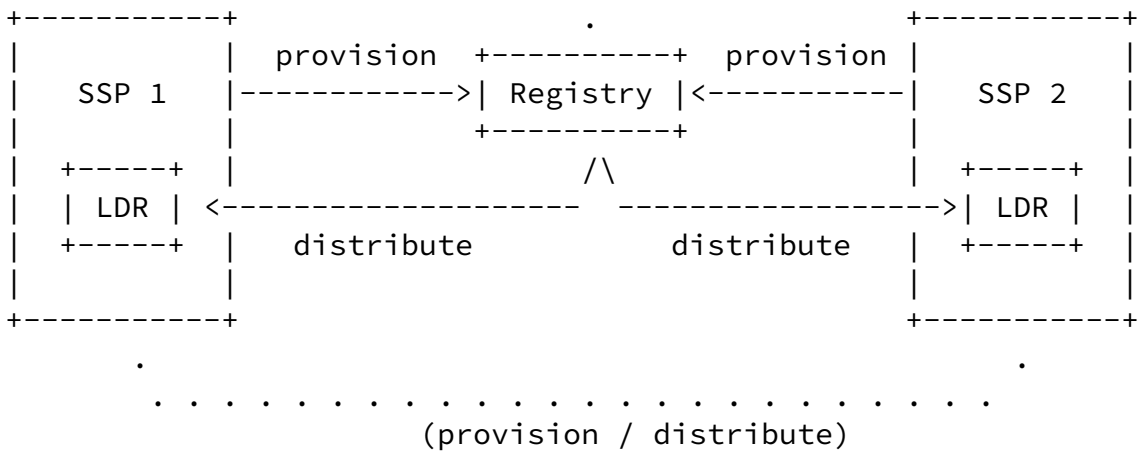

Where, LDR = Local Data Repository

Figure 2: Functional Overview

The following is a summary of the proposed responsibilities for Registries and Local Data Repositories:

- Registries are the authoritative source for provisioned session establishment data (SED) and related information.

- Local Data Repositories are the data store component of an addressing server that provides resolution responses.

- Registries are responsible for distributing SED and related information to the local data repositories.

In addition, this document proposes the following aggregation groups with regards to SED (certain use cases also illustrate these groups):

- Aggregation of public Identifiers: The initial input "key" to a SED lookup is a public identifier. Since many public identifiers will share similar (or identical) destinations, and hence return similar (or identical) SED, provisioning the same set of SED for millions of public identifiers is inefficient, especially in cases where the SED needs to be modified. Therefore, an aggregation mechanism to "group" public identifiers is proposed. This aggregation is called a "destination group".

- Aggregation of SSPs: It is expected that SSPs may want to expose different sets of SED, depending on the receiving SSP. This exposure may not always be unique, in which case an aggregation makes it efficient. Such an aggregation is proposed, and termed "Data Receipient Group".

- Aggregation of SED records: Finally, it is anticipated that a complete set of routing data will consist of more than just one SED record. To be able to create and use the same set of SED records multiple times (without creating duplicates) an aggregation mechanism at this level is proposed, and called "routing group".

The above aggregations are illustrated in Figure 3.

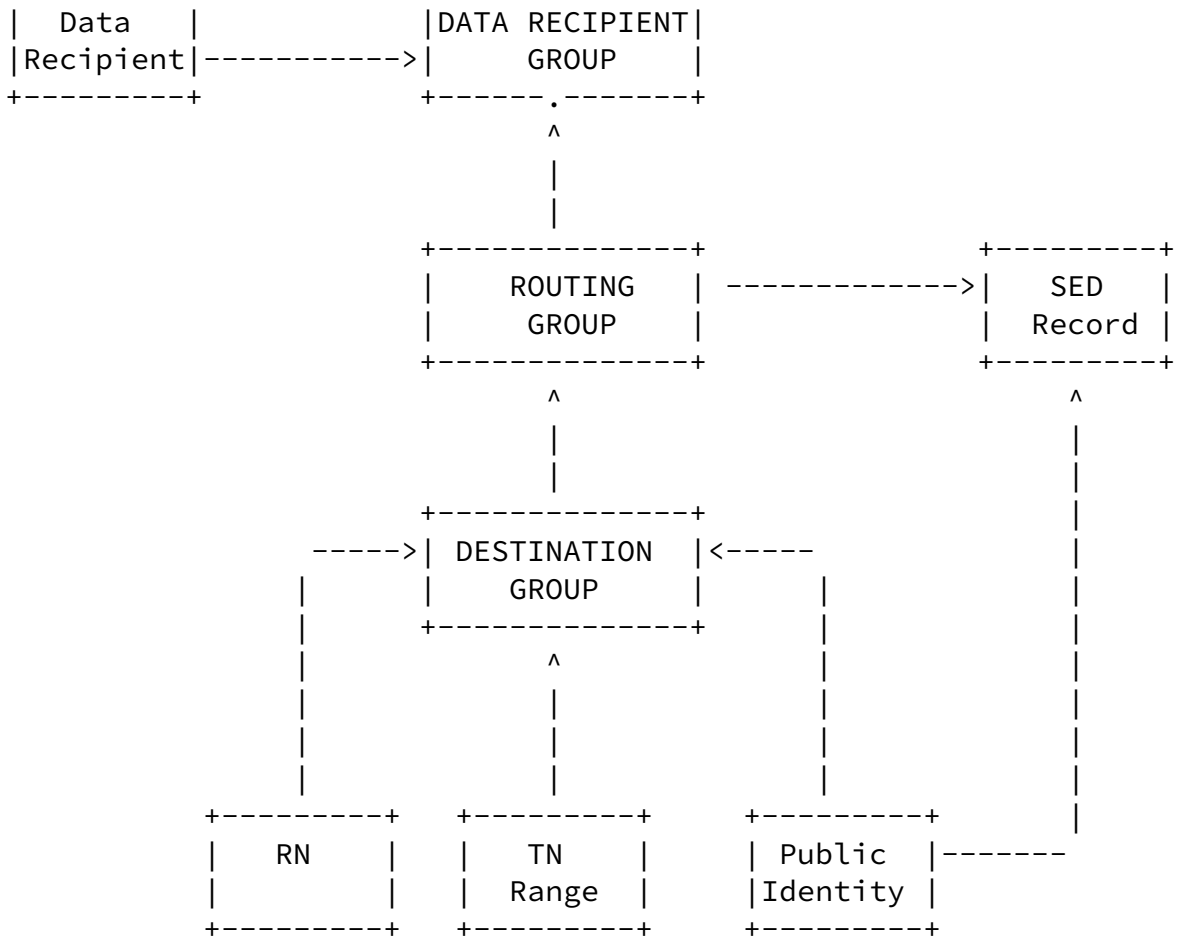

Figure 3: Data Model Diagram

A description of the relationships follows:

- An RN is associated with one or more Destination Groups
- A TN Range object is associated with one or more Destination Group
- A Public Identity is associated with zero or more Destination Group
- A Public Identity is associated with zero or more SED Records

- A Destination Group is associated with zero or more Routing Groups.

- A Routing Group is associated with zero or more SED Records; NAPTRs and other SED Record Types associated with Routes are not User or TN specific. For example the user portion of a NAPTR regex will be "\1".

- An Routing Group is associated with zero or more peering organizations to control visibility/access privs to that Routing Group and the Destination Groups they expose.

- A Data Recipient Group is associated with (contains) zero or more Data Recipients to facilitate the allocation of access privileges to Routing Groups.

3. Use Cases and Requirements

This section presents the use cases and associated requirements.

3.1. Registry Provisioning

Registry is the authoritative source for session establishment data (SED). The registry needs to be provisioned with this data to perform its function. This data includes: destination groups, routing groups and data recipient groups. It can also include RNs and TN Ranges. The following sub-sections illustrate the use cases and the requirements, respectively.

3.1.1. Use Cases

USE CASE #1 Near-real-time provisioning: The registry is provisioned with data that is not accompanied by an effective date or time. In such cases, the registry will validate the data and make it effective in near real-time.

USE CASE #2 Non-real-time provisioning: The registry is provisioned via an asynchronous provisioning process. For instance, an SSP has commissioned a new registry and wishes to download a very large number of telephone numbers. Rather than stream individual entities, one at a time, the SSP's back-office system prefers to perform the operation as a set of one or more batches (e.g., via an external data file), instead of the near-real-time provisioning interface.

USE CASE #3 Deferred provisioning: The registry is provisioned with data that is accompanied by an effective date and time. In scenarios such as this, the registry will validate the data and wait until the effective date and time to

make it effective. TBD: What happens if near-real time data overrides data parked for later incorporation?

USE CASE #4 Intra-network SED: SSP wishes to provision their intra-network Session Establishment Data (SED) such that it enables current and future network services to identify and route real-time sessions. For example, in the case

of VoIP calls it allows one SoftSwitch (calling) to discover another (called). The SSP provisions IP addressing information pertaining to each SoftSwitch, which is provisioned to the registry but only distributed to a specific local data repository. This SED may differ from the SED that is distributed to other local data repositories.

USE CASE #5 Destination Groups: An SSP may wish to control the flow of traffic into their network (ingress) based on groupings of Public Identities. Associating each group of Public Identities with a specific set of ingress SBEs or points-of-interconnect. The SSP, for example, sub-divides the country into four regions: North-East, South-East, Mid-West, and West-Coast. For each region they establish points-of-interconnect with peers and provision the associated SED hostnames or IP addresses of the SBEs used for ingress traffic. Against each region they provision the served Public Identities into groups- termed Destination Groups - and associate those destination groups with the appropriate points of ingres.

USE CASE #6 Public Identity is taken out of service: A public identity (or a TN range) is taken out of service because it is no longer valid. The Registry receives a delete operation and removes the public identity from its database. This can also trigger delete operations to keep the local data repositories up-to-date.

USE CASE #7 Assigning a set of public identities to a different Destination Group: A set of public identities are assigned a different Destination Group which effectively changes their routing information. This may be due to a network re-arrangement, a Signaling path Border Element being split into two, or a need to do maintenance, two carriers merging, or other considerations. This scenario can also include an effective date and time.

USE CASE #8 Moving an SSP from one Data Recipient Group to another: An SSP would like to re-assign the Destination Groups it shares with a peer and move the peer SSP from one Data Recipient Group to another. This results in the moved peer seeing a new and different set of routing data.

USE CASE #9 Inter-network SED (Direct and Selective Peering): In this case the SSP is the actual carrier-of-record; the entity serving the end-user. The SSP wishes to communicate different SED data to some of its peers that wish to route to its destinations. So the SSP has implemented direct points-of-interconnect with each peer and therefore would like address-resolution to result in different answers depending on which peer is asking.

USE CASE #10 Separation of Responsibility: An SSP's operational practices can seperate the responsibility of provisioning the routing information, and the associated identities, to different entities. For

example, a network engineer can establish a physical interconnect with a peering SSP's network and provision the associated domain name, host, and IP addressing information. Separately, for each new service subscription, the SSP's back office system provisions the associated public identities.

USE CASE #11 Global TN Destinations: The SSP wishes to add or remove one or multiple fully qualified TN destinations in a single provisioning request.

USE CASE #12 TN Range Destinations: The SSP wishes to add or remove one or multiple TN Range destinations in a single provisioning request. TN Ranges support number ranges that need not be 'blocks'. In other words the TN range start can be any number and the TN range end can be any number that is greater than the TN range start.

USE CASE #13 Non-TN Destinations: An SSP chooses to peer their messaging service with another SSPs picture/video mail service. Allowing a user to send and receive pictures and/or video messages to a mobile user's handset, for example. The important aspect of this use case is that it goes beyond simply mapping TNs to IP addresses/ hostnames that describe points-of-interconnect between peering network SSP's. Rather, for each user the SSP provisions the subscriber's email address (i.e. jane.doe@example.com). This mapping allows the mobile multimedia messaging service center (MMSC) to use the subscriber email address as the lookup key and route messages accordingly.

USE CASE #14 Tier 2 Name Server: An SSP maintains a Tier 2 name server that contains the NAPTR records that constitute

the terminal step in the LUF. The SSP needs to provision an registry to direct queries for the SSPs numbers to the Tier 2. Usually queries to the registry should return NS records, but, in cases where the Tier 2 uses a different domain suffix from that used in the registry, CNAME and NS records may be employed instead.

USE CASE #15 Peering Offer/Acceptance: An SSP offers to allow terminations from another SSP by adding that SSP to a Data Recipient Group it controls. This causes notification of the offered SSP. An SSP receiving a peering offer should be able to accept or decline the offer. If the offer is rejected the Registry should not provision corresponding SED to the rejecting SSP. It is expected that this capability will apply mainly in the transit case where non-authoritative parties (in the sense of not being the terminating SSP for an identity) wish to offer the ability to reach the identity and originating SSPs may wish to restrict the routes that are provisioned to their local data repositories.

USE CASE #16 Points of Egress: An SSP has a peering relationship with a peer, and when sending messages to that peer's point of interconnection, the originating SSP wishes to use one or more points of egress. These points of

egress may vary for an given peer. This capability is supported by allowing an originating SSP to provision SED for identities terminating to other SSPs where the originating SSP is itself the data recipient. The provisioning SSP may make use of multiple data recipient identities if it requires different sets of egress points be used for calls originating from different parts of its network. Routing from egress points to ingress points of the terminating SSP may be accomplished by static routing from the egress points or by the egress points using data provisioned to the

Registry by the terminating SSP.

3.1.2. Requirements

The following data requirements apply:

DREQ1: The registry provisioning data model MUST support the following entities: public identities, destination groups, routing groups and data recipient groups.

DREQ2: The registry provisioning data model MUST support the grouping and aggregation of public identities within destination groups.

DREQ3: The registry provisioning data model SHOULD support the grouping and aggregation of TN Ranges within destination groups.

DREQ4: The registry provisioning data model SHOULD support the grouping and aggregation of RNs within destination groups.

The following functional requirements apply:

FREQ1: The registry provisioning interface MUST support the creation and deletion of: public identities, destination groups, routing groups and data recipient groups.

FREQ2: The registry provisioning interface MUST support near-real-time, non-real-time and deferred provisioning operations.

FREQ3: The registry provisioning interface MUST support the following types of modifications:

- reassignment of one or more public identities from one destination group to another;

- reassignment of one data recipient from one destination group to another;

- association and disassociation of a "Default Routing Group" with a Data Recipient; and,

- identification of a destination group as a "Primary Provider" destination group or a "Transit" destination group.

FREQ4: When an entity with a different client identifier than that of the carrier of record for a public identity in a destination group adds a new SSP to a destination recipient group associated with that destination group, the registry provisioning interface MUST: a) notify the new SSP of the updated routing information (which constitutes a peering offer) b) not provision the SED to the new SSP's LDR unless the new SSP signals acceptance.

FREQ5: The registry provisioning interface MUST separate the provisioning of the routing information from the associated identities.

FREQ6: The registry provisioning protocol MUST define a discrete set of response codes for each supported protocol operation. Each response code MUST definitively indicate whether the operation succeeded or failed. If the operation failed, the response code MUST indicate the reason for the failure.

FREQ7: The registry provisioning interface MUST allow an SSP to define multiple sub-identities that can be used in data recipient groups

FREQ8: The registry provisioning interface MUST define the concurrency rules, locking rules, and race conditions that underlie the implementation of that protocol operation and that result from the coexistence of protocol operations that can operate on multiple objects in a single operation and bulk file operations that may process for an extended period of time.

FREQ9: The registry provisioning interface MUST support the ability for a Data Recipient to optionally define a Routing Group as their Default Routing Group, such that if the Data Recipient performs a resolution request and the lookup key being resolved is not found in the Destination Groups visible to that Data Recipient then the SED Records associated with the Default Routing Group shall be returned in the resolution response.

FREQ10: The registry provisioning interface MUST support the ability for the owner of a Routing Group to optionally define Source Based Routing Criteria to be associated with their Routing Group(s). The Source Based Routing Criteria will include the ability to specify zero or more of the following in association with a given Routing Group: Resolution Client IP Address(es) or Domain Names, Calling Party URI(s). The result will be that the resolution server would evaluate the characteristics of the Source, compare them against Source Based Routing Criteria associated with the Routing Groups visible to that Data Recipient, and return any SED Records associated with the matching Routing Groups.

FREQ11: The registry provisioning interface MUST track, via a client identifier, the entity provisioning each data object (e.g. Destination Group or Routing Group). This client identifier will identify the entity that is responsible for that data object from a protocol interface perspective. This client identifier SHOULD be tied to the session authentication credentials that the client uses to connect into to the registry.

The registry provisioning interface MUST incorporate a data recipient identifier that identifies the organization responsible for each data object from a business perspective. This organization identifier may or may not ultimately refer to the same organization that the client Identifier refers to. The separation of the data recipient identifier from the client identifier will allow for the

separation of the two entities, when the need arises.

Exactly one client identifier MUST be allowed to provision objects under a given data recipient identifier. But a client identifier MUST be allowed to provision objects under multiple data recipient identifiers.

Objects provisioned under one "Protocol Client Identifier" MUST NOT be alterable by a provisioning session established by another "Protocol Client Identifier".

3.2. Distribution of data into local data repositories

This section targets use cases concerned with the distribution of SED to local data repositories. This is considered out-of-scope for this version of the document.

3.3. Miscellaneous Use Cases

This section contains additional use cases for consideration.

3.3.1. Indirect Peering to Selected Destinations

The SSP transit provider wishes to provide transit peering points for a set of destinations. But that set of destinations does not align with the destination groups that already exist. The SSP wishes to establish its own destination groups for the destinations that it provides transit to.

3.3.2. TBD: RN Destinations

The SSP does not wish to provision individual TNs, but instead, for ease of management, wishes to provision Routing Numbers ((e.g., as in some number portability implementations). Each RN effectively represents a set of individual TNs, and that set of TNs is assumed to change 'automatically' as TNs are ported in and ported out. Note that this approach requires a query to resolve a TN to an RN prior to using the provisioned data to route.

4. Security Considerations

Session establishment data allows for the routing of SIP sesions within, and between, SIP Service Providers. Access to this data can compromise the routing of sessions and expose a SIP Service Provider to attacks such as service hijacking and denial of service. The data can be compromised by vulnerable functional components and interfaces identified within the use cases.

5. IANA Considerations

This document does not register any values in IANA registries.

6. Acknowledgments

This document is a result of various discussions held by the DRINKS requirements design team, which is comprised of the following individuals, in alphabetical order: Deborah A Guyton (Telcordia), Gregory Schumacher (Sprint), Jean-Francois Mule (CableLabs), Kenneth Cartwright (Verisign), Manjul Maharishi (Verisign), Penn Pfautz (AT&T Corp), Ray Bellis (Nominet), the co-chairs (Richard Shockey, Nuestar; and Alexander Mayrhofer, enum.at GmbH), and the editors of this document.

7. References

7.1. Normative References

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

7.2. Informative References

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

June 2002.

[RFC5486] Malas, D. and D. Meyer, "Session Peering for Multimedia Interconnect (SPEERMINT) Terminology", RFC 5486, March 2009.

Author's Address

Sumanth Channabasappa
CableLabs
858 Coal Creek Circle

Louisville, CO 80027
USA

Email: sumanth@cablelabs.com