

DRINKS
Internet-Draft
Intended status: Standards Track
Expires: January 6, 2012

D. Schwartz
XConnect
J. Barkan
DigitalShtick
July 5, 2011

Session Peering Distribution Protocol
draft-schwartz-drinks-spdist-00

Abstract

This document defines a protocol for distributing session establishment data from often centralized Session Data Registries or SIP Service Provider data stores into local data repositories for local querying of the data. The distributed data is typically used by various network elements for session peering.

This document describes the Session Peering Distribution Protocol (SPDP) that should be seen as an extension to the session peering provisioning protocol defining the provisioning process itself. The document provides a set of guiding principles for the design of this protocol including an extension of the SPPP basic data model for distribution purposes and an associated extension XML Schema Document.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 6, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.


Table of Contents

1. Introduction	3
2. Rationale	5
2.1. Data Model Changes	5
2.2. Multiple Data Sources	5
2.3. Multiple Local Repositories	5
2.4. Data Versioning	5
2.5. Availability of Local Repository	5
2.6. Partitioning Across Stores	5
2.7. Asynchronous distribution	5
2.8. Bulk Data Transport Protocol	6
3. Protocol Commands	7
4. SPDP Examples	8
4.1. TBD	9
5. Security Considerations	10
6. References	11
6.1. Normative References	11
6.2. Informative References	11
Authors' Addresses	13

1. Introduction

Service providers and enterprises use registries to make call or session routing decisions for Voice over IP, SMS and MMS traffic exchanges. This document is narrowly focused on the distribution protocol, used to take session establishment data (SED) already provisioned into a registry and distribute to a local data repository. This protocol prescribes a way for a registry to distribute to a local cache belonging to a peering entity, data previously provisioned to the registry by a different peering entity willing to share its peering data. The requirements and use cases driving this protocol have been documented in [I-D.ietf-drinks-usecases-requirements] and the aforementioned provisioning protocol has been documented in [I-D.ietf-drinks-spprov]. The reader is expected to be familiar with the the previously mentioned documents.

Three types of provisioning flows have been described in the use case document: client to registry provisioning, registry to local data repository and registry-to-registry. This document addresses a subset (distribution of SED from registry-to-local repository) by defining a Session Peering distribution Protocol (SPDP) for this purpose (arrow "2" in the figure below).



Three Registry Provisioning Flows

Figure 1

The session establishment data that is distributed to local repositories is typically used by various downstream SIP signaling systems to route a call to the next hop associated with the called domain. These systems typically use a local data store ("Local Data Repository") as their source of session routing information. More specifically, the SED data is the set of parameters that the outgoing signaling path border elements (SBEs) need to initiate the session. See [RFC5486] for more details.

A "terminating" SIP Service Provider (SSP) provisions SED into the registry to be selectively shared with other peer SSPs. Subsequently, a Registry may distribute the provisioned data into local Data Repositories used for look-up queries (identifier -> URI) or for lookup and location resolution (identifier -> URI -> ingress SBE of terminating SSP).

While the provisioning protocol [I-D.ietf-drinks-spprov] itself made no determination as to whether or not one common baseline protocol could be used for all three provisioning flows, this document, posits the need for an extension to the provisioning protocol to deal with session data distribution specific issues. To reiterate, as this document is an extension of the provisioning document, it is imperative that the reader be familiar with the provisioning document before starting on this document. Furthermore, and due to this dependence, SPDP will reuse the terminology, high level design and the data model sections found in SPPP and these sections will be omitted from this document. What this document will provide is the rationale for a separate distribution protocol, the potentially different transport requirements, the potentially different security model and the additional protocol commands and resultant additional xml specification.

2. Rationale

While a provisioning protocol must define a data model representing the data being provisioned, it is not required to address issues that arise only once the data is actually provisioned. Data consistency across stores, Data transfer size and data versioning are only some of the complexities that arise when attempting to distribute or replicate data from a master to a slave and the underlying protocol (data model, protocol commands, transport requirements, etc.) must provide support for the added functional complexity. This section takes a look at aspects of data distribution that are different from simple data provisioning.

2.1. Data Model Changes

What is the SPPP model missing? Do we need data structures that capture data in manner more efficient for bulk distribution? Time-stamping (for versioning)?

2.2. Multiple Data Sources

How is this different from multiple provisioning entities? Conflict mechanism?

2.3. Multiple Local Repositories

Consistency across destination stores. push vs pull. scaling considerations.

2.4. Data Versioning

full vs incremental synchronization

2.5. Availability of Local Repository

dealing with downtime when trying to distribute.

2.6. Partitioning Across Stores

How do we deal with data partitioning across local stores under a single administrative domain?

2.7. Asynchronous distribution

Bulk vs incremental distribution

2.8. Bulk Data Transport Protocol

transport issues associated with bulk transfer

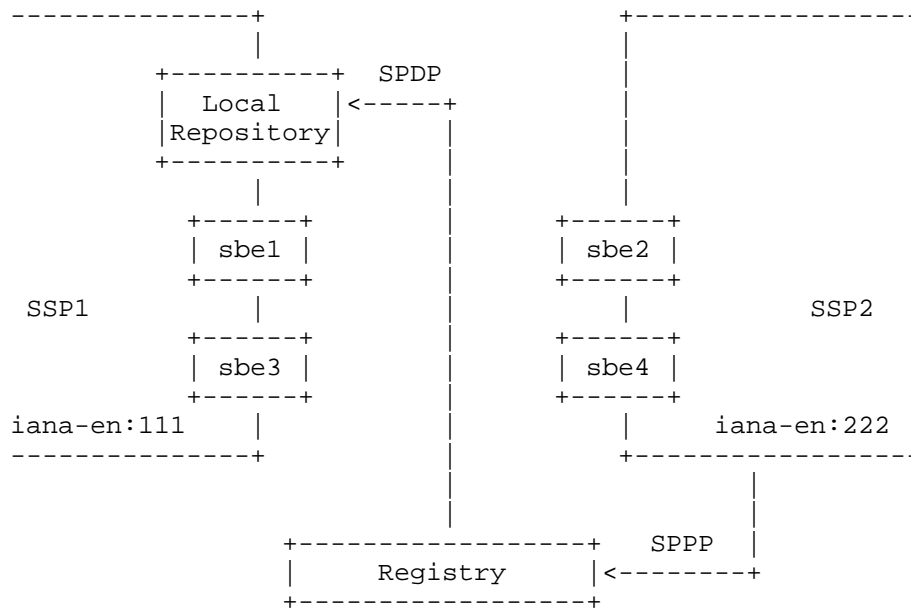
3. Protocol Commands

TBD

4. SPDP Examples

This section shows XML message exchange between a centralized Registry previously provisioned with Session Establishment Data (SED) (using the SPPP protocol) by a peering entity SSP2 willing to share this information, and a local data repository residing at SSP1 (approved by SSP2 to receive this data) needing a copy of SSP2 data. For the sake of simplicity, the transport wrapper for the SPDP protocol is left out as well as most of the provisioning messages already covered in the SPPP document. The SPDP protocol messages in this section are valid XML instances that conform to the SPDP schema version within this document.

In this sample use case scenario, SSP2 provisions resource data to the registry and uses SPPP constructs to selectively define sharing attributes of the route groups. SPDP is subsequently used to have this information distributed to relevant local repository. In the figure below, SSP2 has two ingress SBE instances that are associated with the public identities that SSP2 has the retail relationship with.



Internet-Draft

draft-schwartz-drinks-spdist

July 2011

4.1. TBD

TBD

5. Security Considerations

SPDP implementations manage data that is considered confidential and critical. Furthermore, SPDP implementations can support provisioning activities for multiple registrars and registrants. As a result any SPDP implementation must address the requirements for confidentiality, authentication, and authorization.

With respect to confidentiality and authentication, the transport protocol section contains some security properties that the transport protocol must provide so that authenticated endpoints can exchange data confidentially and with integrity protection.

With respect to authorization, the SPDP server implementation must define and implement a set of authorization rules that precisely address (1) which local repositories will be authorized to get each SPDP object type for given registrant(s). These authorization rules are left as a matter of policy and are not specified within the context of SPDP. However, any SPDP implementation must specify these authorization rules in order to function in a reliable and safe manner.

6. References

6.1. Normative References

- [I-D.ietf-drinks-sppp-over-soap]
Cartwright, K., "SPPP Over SOAP and HTTP",
draft-ietf-drinks-sppp-over-soap-03 (work in progress),
May 2011.
- [I-D.ietf-drinks-spprov]
Mule, J., Cartwright, K., Ali, S., and A. Mayrhofer,
"Session Peering Provisioning Protocol",
draft-ietf-drinks-spprov-07 (work in progress),
April 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and
Languages", BCP 18, RFC 2277, January 1998.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO
10646", STD 63, RFC 3629, November 2003.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688,
January 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform
Resource Identifier (URI): Generic Syntax", STD 66,
RFC 3986, January 2005.

6.2. Informative References

- [I-D.ietf-drinks-usecases-requirements]
Channabasappa, S., "DRINKS Use cases and Protocol
Requirements", draft-ietf-drinks-usecases-requirements-05
(work in progress), March 2011.
- [RFC2781] Hoffman, P. and F. Yergeau, "UTF-16, an encoding of ISO
10646", RFC 2781, February 2000.
- [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.
- [RFC3761] Faltstrom, P. and M. Mealling, "The E.164 to Uniform
Resource Identifiers (URI) Dynamic Delegation Discovery

System (DDDS) Application (ENUM)", RFC 3761, April 2004.

- [RFC4725] Mayrhofer, A. and B. Hoeneisen, "ENUM Validation Architecture", RFC 4725, November 2006.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [RFC5486] Malas, D. and D. Meyer, "Session Peering for Multimedia Interconnect (SPEERMINT) Terminology", RFC 5486, March 2009.

Authors' Addresses

David Schwartz
XConnect
Malcha Technology Park
Jerusalem, 96951
Israel

Email: dschwartz@xconnect.net

Jeremy Barkan
DigitalShtick
Jerusalem,
Israel

Email: jbarkan@digitalshtick.com

