

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2009

D. Schwartz
XConnect Global Networks
R. Mahy
Plantronics
A. Duric
Telio
E. Lewis
NeuStar
July 7, 2008

Consolidated Provisioning Problem Statement
draft-ietf-drinks-cons-rqts-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 7, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document describes the type of data provisioned among Voice Service Providers and underscores the need for clearly defined structures and interfaces to facilitate this provisioning. This work

is in support of the service provider peering as defined by the SPEERMINT WG.

Table of Contents

- [1. Introduction](#) [3](#)
- [2. Terminology](#) [4](#)
- [3. Registry Data](#) [4](#)
 - [3.1. Index/Key Data](#) [4](#)
 - [3.2. Resolution Data](#) [5](#)
- [4. Reachability vs. Routing](#) [6](#)
- [5. Operations on the Registry Data](#) [6](#)
- [6. Other Attributes](#) [7](#)
- [7. Data Encoding](#) [7](#)
- [8. Data Management](#) [7](#)
- [9. Data Propegation](#) [8](#)
- [10. Querying The Registry](#) [8](#)
- [11. Control](#) [8](#)
- [12. Existing Technologies](#) [8](#)
- [13. Security Considerations](#) [9](#)
- [14. IANA Considerations](#) [9](#)
- [15. Acknowledgements](#) [9](#)
- [16. References](#) [10](#)
 - [16.1. Normative References](#) [10](#)
 - [16.2. Informational References](#) [10](#)
- [Authors' Addresses](#) [11](#)
- [Intellectual Property and Copyright Statements](#) [12](#)

1. Introduction

VoIP Service Providers (VSPs) engage in peering relationships with other VSPs to create direct IP-to-IP interconnections. These relationships provide network reach, greater technical capabilities and enhanced economic benefit beyond that available with the Public Switched Telephone Network (PSTN), while providing the security benefit perceived to exist in the PSTN.

Because the business and operational management of hundreds or thousands of direct peering relationships is difficult, VSPs often enlist the help of peering exchanges to centralize (or assist [[I-D.ietf-speermint-terminology](#)] with) the management of the relationships. One of the central functions of these peering exchanges is a registry of identifiers (often implemented as a private ENUM tree) based on telephone numbers. This function is called the peering or numbering registry. VSPs participating in the peering exchange must provision their identifiers into the peering registry to allow other VSPs with access to this registry to query and obtain the correct resolution for a given number. Lack of clear standards for this provisioning step has given rise to many proprietary approaches that are rendering open provisioning cumbersome and error prone.

VSP peering is not the only driver for this work, however. It is quite common today for one VSP to receive number resolution data from both authoritative or regulatory sources (e.g. a national telephone number plan administrator) and commercial or private sources. Since eventually all of the information resides locally, the VSP is interested in merging resolution data across potentially different local platforms so as to avoid multiple queries for each call. Today this is virtually impossible to do in an efficient and standard manner due to the proprietary nature of the individual registry components.

In addition, many registry network architectures dictate sub registries for overall resilience and performance. The transfer of resolution data to the sub registries is not yet standardized and results in unnecessary hardware/software component lock in.

This document attempts to describe the most common data that needs to be exchanged in the cases highlighted above with the ultimate goal being that of identifying the data structures and interfaces required to support the data exchange scenarios specified above.

As a final note it is important to stress that while ENUM is not mentioned explicitly in this document so as not to bias the outcome, it is clear that in the minimum all the information present in a

NAPTR RR must be captured as the information therein has been well thought out and deviating from this may curtail future growth. Additionally, while E.164 numbering is not mentioned either for the same reason, it is clear that in almost all cases the prefixes that are being exchanged are in the e.164 format.

2. 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](#)].

Terms used or referred to elsewhere in the document (including in the introduction):

DNS - Domain Name System [[RFC1034](#)]

RR - Resource Record [[RFC1034](#)]

TXT - DNS Text record [[RFC1035](#)]

NAPTR - Naming Authority Pointer record [[RFC3403](#)]

EPP - Extensible Provisioning Protocol [[RFC3730](#)]

ENUM - E.164 to URI DDDS [[RFC3761](#)]

URI - Uniform Resource Identifiers [[RFC3986](#)]

TLS - Transaction Layer Security [[RFC4366](#)]

3. Registry Data

Registry Data consists of both Index or Key data and Resolution data.

3.1. Index/Key Data

The organization of registry data is based on specific phone numbers or phone number prefixes (which could represent blocks of phone numbers, regions, or theoretically whole countries). For generality, we will use the term prefix to include complete phone numbers as well. Prefixes are the index or key used for all registry manipulation and lookups. Even though some of the numbers represented within these prefixes may not be globally reachable, the prefix itself needs to be globally normalized before it can be entered into a registry. These globally normalized prefixes always

begin with a plus (+) and a telephone country code. (Note that prefixes in some countries can contain hexadecimal digits).

Since prefixes have variable lengths, a provisioning protocol must be able to enter data for a sub-prefix or super-prefix of an existing record. For example, it must be possible to enter records about "+1202555" and "+12025551234" at the same time. For lookups information about the most specific prefix will be returned. This allows for some measure of aggregation. Also, in order to provide maximal flexibility a provisioning protocol must provide a mechanism for specifying both minimum and maximum number of digits in a prefix.

3.2. Resolution Data

For each prefix, there is a variety of data that can be exchanged. The most important set of data identifies that a specific VSP is responsible for the prefix and in most cases the VSP provides a SIP URI through which this prefix can be reached.

In complex cases, several VSPs may claim some form of responsibility for the same prefix. We can use the term "last hop" VSP to describe the VSP closest to the end-user of a phone number. The provider who was assigned a prefix by the national numbering authority is the "first hop" VSP. The first hop VSP may have no way of knowing if the last hop VSP will include itself in the registry. Therefore it is important that the querier can obtain both records and use the most specific one which contains reachability information.

In many cases, commercial registries also contain information used for Local Number Portability. Even if a prefix is not reachable for IP peering, it is useful to provide the "dipped" number and carrier code. This information could be provided as a tel URI with the number portability attributes defined in [RFC 4769](#) [[RFC4769](#)]. Likewise it is very useful to know that a prefix is known not to exist anywhere.

Like stated in the introduction it is imperative that the minimal set of resolution data contain all the information required for a DNS NAPTR RR.

Additionally, as a future proofing mechanism it is recommended that resolution data include a catchall (similar to a DNS TXT RR) for data that is not currently present in any ENUM service definitions.

One final note. One of the common "rewrite" rules for the URI in NAPTR RRs for example is of the form "\1@somehost.company.example" where the "\1" refers to the telephone number being processed. This kind of shorthand should be available when processing prefixes in

bulk.

4. Reachability vs. Routing

The goal of the registry is to provide sufficient information to identify a responsible VSP for a prefix. The responsible VSP can provide a SIP URI that can be proxied or redirected as desired by the VSP. It is important to note that the registry is expected to return the same responsibility data for all parties that query it.

Routing information is also out of scope of the registry provisioning problem. Once a responsible VSP is contacted, that VSP can use a variety of techniques to route that request. For example, the VSP could use TRIP [3], a private database, or a SIP location server. Since this information is highly dynamic, it is inappropriate to provision in a registry.

5. Operations on the Registry Data

Below is a list of logical operations on the registry data. Please note that as stated above a provisioning protocol MUST apply these operations equally to individual prefixes as well as prefix blocks.

Add: Add (responsible VSP) data about a new prefix to the registry.

Delete: Remove a prefix from the registry. Semantically it means that the prefix no longer exists anywhere.

Port-out: A port-out is similar to a delete, but could be logged differently. The semantics are that the prefix still exists, but that the previous VSP is no longer responsible for it.

Port-In: A port-in is similar to an add, but the semantics are that the prefix was previously assigned to a different provider.

Transfer: A transfer is a port-out and port-in operation performed atomically. This operation insures that lookups do not fail for the transferred prefix during the transfer.

Renumber: A renumber operation occurs when a specific prefix is completely changed, but the data corresponding to the prefix has not changed. This typically happens when a prefix is lengthened. For example, when France moved from an 8-digit to a 10-digit dial plan, the corresponding globally normalized prefix for a Parisian phone number had a 1 inserted between the country code and the old prefix. Renumbering could also accomplish prefix shortening (although this is

quite unlikely to occur) or prefix splitting (in the past United States area codes where split when they were exhausted).

Modify: A modify operation occurs when some other attribute of a prefix is modified (for example the target URI used for reachability).

6. Other Attributes

All the registry records will need to include some kind of validity information. The provisioning protocol can indicate that a record is not valid before one date and time and no longer valid after another date and time.

In addition to responsibility data, we have identified the following extra attributes as important or interesting:

Regardless of which protocol is used, issues that should be addressed include

Number type (unknown, IP, PSTN, both)

PSTN carrier code (for numbers with no IP reachability)

Fee category (free, landline, mobile, pay)

Media types supported (voice, video, message)

There MUST be a mechanism for an audit trail as issues of ownership are bound to surface and a clearly defined mechanism for tracking changes to registry data is essential.

7. Data Encoding

Since data may need to traverse administrative domain boundaries some thought needs to be given to the rendering of the messages in transmission. The encoding mechanism MUST be robust and easy to design and troubleshoot with a strong bias towards an existing and widely recognized standard.

8. Data Management

Due to the entrenchment of legacy software development processes (e.g. SOAP/XML, WSDL, TLS) it is of utmost importance that whatever emerges from this WG should be easy to implement with existing

methodologies.

9. Data Propegation

The transport layer is strictly point-to-point, with no caching or forwarding.

10. Querying The Registry

The definition of the registry query mechanism is out of scope for the PEPPERMINT activity. However, it is helpful to know what mechanisms are in popular use to understand the necessary properties for a provisioning interface. At present, there are two well-known methods used by VSPs to query a peering registry. These are ENUM [[RFC3761](#)] and SIP [[RFC3261](#)].

ENUM is simply a method of using DNS. It allows a VSP to query a registry with a telephone number, an E.164 number in most cases, and retrieve a list of URIs, each with a specific preference order.

When SIP is used for the query mechanism, the numbering registry functions as a SIP redirect proxy [[RFC3261](#)]. The input to this mechanism is a URI or more importantly the UserInfo part of the URI containing the number to be queried. The response returned by the proxy is either an error code indicating the absence of information about the number queried or a redirect response (3XX) containing 1 (or more) Contact Headers representing available routing options.

11. Control

Since it may be possible to both push and pull data it is imperative that a throttling mechanism exist to control the flow of data exchange at all levels.

12. Existing Technologies

there has been some consideration to EPP extensions for ENUM [[RFC4114](#)], and why it has not been adopted and why a requirements document is now being produced to cover what would seemingly be addressed by that solution.

There are two reasons for EPP not being adopted. One is that it isn't compatible with legacy participants. The other reason is that it requires more implementation work.

Legacy participants have an existing base of software development built around SOAP/XML and WSDL, and are familiar with TLS. Approaches to ENUM registry interfaces that use these tools will blend more easily into the software products already in use to manage telephone numbers.

The use of SOAP permits automatic generation of software to handle the client side of the exchange. Domain name registries had to provide software tool kits to give to registrars to match this functionality. When a change is made to EPP, there will be a lot of software exchanged.

From experience with both EPP and SOAP based approaches to registry software, the SOAP based approach is much easier on the software engineering process. The difference between the approaches is not seen in a protocol analysis, but in an analysis of software engineering.

13. Security Considerations

Security is to be implemented in the applications exchanging data, the requirements here are meant to say that relevant security data will be exchanged in the building of the transport. Specifically, data integrity, authentication and secrecy. (Please note that all three of these can be provided by using TLS, with the certificate handshake being used by the application to complete the security needs).

14. IANA Considerations

NA

15. Acknowledgements

Many of the points of information in this document are summarizations of objectives and statements made by individuals on the PEPPERMINT and SPEERMINT mailing lists. We would also like to acknowledge Jeremy Barkan and Otmar Lendl for providing insightful inputs to the original draft. Information on the PEPPERMINT mailing list can be found at <http://www.ietf.org/mailman/listinfo/peppermint>.

16. References

[16.1.](#) Normative References

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

[16.2.](#) Informational References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC3403] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", [RFC 3403](#), October 2002.
- [RFC3730] Hollenbeck, S., "Extensible Provisioning Protocol (EPP)", [RFC 3730](#), March 2004.
- [RFC3761] Falstrom, P. and M. Mealling, "The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)", [RFC 3761](#), April 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4114] Hollenbeck, S., "E.164 Number Mapping for the Extensible Provisioning Protocol (EPP)", [RFC 4114](#), June 2005.
- [RFC4366] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", [RFC 4366](#), April 2006.
- [RFC4769] Livingood, J. and R. Shockey, "IANA Registration for an Enumservice Containing Public Switched Telephone Network (PSTN) Signaling Information", [RFC 4769](#), November 2006.
- [I-D.ietf-speermint-terminology] Malas, D. and D. Mayer, "SPEERMINT Terminology", [draft-ietf-speermint-terminology-15.txt](#), November 2007.

Authors' Addresses

David Schwartz
XConnect Global Networks
Malcha Technology Park
Building # 1
Jerusalem 90961
Israel

Phone: +972 52 347 4656
Email: dschwartz@xconnect.net
URI: www.kayote.com

Rohan Mahy
Plantronics

Email: rohan@ekabal.com
URI: www.plantronics.com

Alan Duric
Telio

Email: alan.duric@telio.no
URI: www.telio.no

Edward Lewis
NeuStar
46000 Center Oak Plaza
Building # 1
Sterling, VA 20166
USA

Phone: +1-571-434-5468
Email: ed.lewis@neustar.biz
URI: www.neustar.biz

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

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