Speermint Working Group O. Lendl
Internet-Draft enum.at
Intended status: Informational October 28, 2008
Expires: May 1, 2009

VoIP Peering: Background and Assumptions
draft-lendl-speermint-background-02.txt

Status of this Memo

Copyright Notice

Abstract

This documents provides background for the work on VoIP peering and tries to provide guidance on what kind of work is needed to facilitate widespread SIP-based peering of telephony networks. It is intended to spur discussion on the work about peering (in the SPEERMINT and DRINKS WGs) and should also serve as input to the ongoing discussions on reducing Spam for Internet Telephony (SPIT).

Table of Contents

1. Introduction

The Speermint WG is chartered to help with the interconnection of SIP based layer 7 networks. It should not deal with basic IP connectivity and SIP protocol issues; those are covered by other working groups.

Speermint focuses on what guidelines (and perhaps protocol elements) are needed by service providers and enterprises to move from ad-hoc, manual peerings to a fully standardized, secure, easy to implement, and thus widespread SIP based peering setup.

This document aims solely at the telephony network aspects of SIP and ignores applications like Instant Messaging or Presence which might also be implemented using SIP. The focus here is on the use of SIP in PSTN replacement services.

Version -01 of this draft expands on the implications of the interconnection structure on the SPIT problem. The concepts listed in here should thus also be worthwhile for the RUCUS EG.

Version -02 of this draft provides a rough overview on how the voice interconnection landscape might look like in the future and lists specific work items for the SPEERMINT and DRINKS WGs.

This document was written as discussion input. It is not intended for publication as an RFC.

2. Interconnection Models

In order to understand the VoIP peering world it is necessary to go beyond pure protocol issues and instead talk about the ecosystems in which the protocols operate. This section tries to be purely descriptive and makes no recommendations.

2.1. The PSTN model

The public switched telephone network (PSTN) is built upon the following fundamental assumptions:

- When the system was designed, the numbers of operators was limited. Often, there was just a single incumbent per country. Building a full mesh for international interconnection was possible. This is no longer true.

- Global reachability is achieved by interconnecting individual smaller networks. There is no global lower-layer connectivity: if two networks are not directly interconnected, then calls are passed through transit networks on the application layer.

- There are no ad-hoc connections between networks: all links are manually configured lines (physical, or other transparent bit-pipes).

- There is a clear separation between network operators and network users. This applies both to protocols (e.g., SS7 ISUP versus ISDN), as well as, to regulatory rules.

- Routing information is not directly passed from the destination network to the source network via some global database. Instead, transit networks communicate to their customers which destinations they can handle.

- Accounting and settlement are core features.

2.2. The email model

SIP according to RFC 3261 [2] and RFC 3263 [3] follows an email alike model. It can be summarized as follows:

- Email and SIP addresses are structured as username@domain. For routing purposes, only the domain part is relevant. The username is only interpreted by the machines serving this specific domain.

- The global, public DNS is used to map the domain from the address to a prioritized set of ingress points that handle incoming communication requests for this domain. As the DNS is agnostic to the entity querying data stored there, all senders receive the same set of ingress points.

- In order to achieve global reachability, these ingress points need to accept incoming requests from the open Internet. If they reject, for example, incoming packets from a VoIP provider X from country Y then there is no backup path for this communication, and the destination just will not be reachable from that VoIP provider X.

- As anybody on the Internet can contact any destination domain, no business relationship between sender and destination domain is required. This implies that there is no settlement: No money is changing hands because of such a communication.

o There is no inherent distinction between end users and service providers hard-coded into the protocol. Any client can do the DNS lookups himself and directly contact the destination servers.

o Usually clients do not talk directly to each other: On the source side a SIP INVITE is forwarded to the outbound proxy that applies the routing algorithm, which then contacts its peer on the destination side that also performs additional processing before handing off the communication to the destination side.

The email model has proved to be extremely successful -- for email.

3. Why is SPEERMINT needed?

In theory, the Speermint WG is not necessary: The SIP RFCs envision global reachability of all SIP devices over the public Internet. Source networks just need to resolve the domain from the URI according to RFC 3263 and send the INVITE to the SIP proxy in charge of that destination domain.

Telephone number (TN) based calling is also supported: RFC 3761 [1] provides TN to URI mapping and thus reduces the call routing problem to the already solved case of SIP URI resolution.

* Apparently, the real world did not choose to implement and deploy SIP and public ENUM as initially envisioned by their inventors. In other words: the motivation for Speermint is the failure of the world to conform to the original IETF vision of SIP based real-time communication. *

3.1. Why did the Email Model fail for SIP?

Although SIP has won the protocol war against H.323 (just as SMTP won against X.400), it failed to establish the same sort of ecosystem in the Internet as SMTP was able to do. The number of SIP users who are reachable via the open Internet using RFC 3263 is minuscule compared to the number of SIP based telephones in operation today.

SIP as a protocol has succeeded; SIP as an ecosystem similar to SMTP has failed.

The need for Speermint arises from that failure: SIP has seen widespread deployment within enterprises and service providers, but the inter-connection part of SIP has not: current deployments usually do not follow RFC 3263, but use either hard-coded IP addresses or private DNS to route calls between SSPs, if they use SIP at all and not the PSTN.

The same applies to ENUM according to RFC 3761: The technology has been successful (as the large number of private ENUM trees demonstrates), but the original vision of ENUM proved to be elusive: the "golden tree" under e164.arpa contains a fraction of entries compared to the numbers found in private trees.

Speermint is chartered to provide solutions for the interconnection problem. It is thus essential to examine why the current standards have failed. Without this gap-analysis there is little chance that Speermint will come up with the missing pieces.

As mentioned before, this analysis cannot be restricted to pure technological aspects, and will thus touch on the business models implied by the technical standards. It is the firm believe of the author that the IETF credo "we don't do business models" has been implicitly violated by existing standards. One approach is thus to identify these implications and augment the protocols to allow them to support a varity of business models and ecosystems.

Business Model

The email model hard-codes a "sender-keeps-all" settlement regime. As anybody is able to connect to anybody, no business relationship is needed between communication partners. Thus, no termination fees can be collected.

The economic model of the current carrier landscape in most countries depends on these charges, and it just does not make sense for any single carrier to allow anonymous incoming SIP-based interconnection as that means lost income. If call patterns are about symmetrical, switching to sender-keeps-all is revenue-neutral for all carriers. There is no clear path on how such a fundamental shift in the bedrock of telco settlement could happen. (other than by regulatory fiat)

The end-state might be a viable business model, but there is no incentive for any individual SSP to start the transition.

This argument applies only to SSPs that are substitutes for PSTN carriers, and not to enterprises operating a SIP infrastructure.

Unwanted Calls

Spam over Internet Telephony (SPIT) is another concern. The free for all nature of the email ecosystem has led to a barrage of unsolicited email (SPAM) which poses a serious threat to the usefulness of email.

Email is non-interactive: filters can be deployed to detect spam by the content of the mail before the recipient is alerted. That is not possible for SPIT; content is only available after the recipient has picked up the phone. A number of SPIT mitigation strategies have been proposed over the past few years, their effectiveness is yet untested. See also [6].

As of 2008, SPIT is not a problem, mainly because the number of open reachable SIP devices is so low. Just as SPAM only started to become a problem after open SMTP servers became common, many SSPs fear that SPIT will appear if they open up their networks.

Identity

Traditional telecom services provide reasonably reliable caller identification. Telcos trust each other's signaling and end users have learned to trust caller-id (even if this trust is somewhat unjustified). Such a trust model is not compatible with the email model of open SIP servers: the INVITE message can come from any host on the Internet and is thus not trusted.

Providing a reliable caller identification is also important for policing: Harassing and abusive calls are more or less under control, as legal and contractual rules can be enforced by tracing calls back to the culprit.

SIP Identity (RFC 4474 [5]) uses a different approach that is based on an authentication service cryptographically asserting the identity of the caller. As such, it is different to the current practice in the telco space, which is based on transitive trust.

QoS and Denial of Service

The email model is not suitable for stringent Quality of Service (QoS) deployments. As there are no pre-arranged relationships with between all communicating SIP servers, there are no mechanisms to guarantee neither network performance on the IP layer for the actual voice transmission, nor can there be comprehensive tests on SIP layer compatibility.

As the ingress points need to be open to anybody on the Internet, they are exposed to Denial of Service attacks.

This combination is at odds with the telco mind-set that thrives on predictable quality and stringent service level guarantees.

Legal Requirements

Operators of public telephony services need to observe a range of regulatory requirements. These rules were written for the PSTN scenario with clearly defined boundaries between operators and users of the telephone network. Changing the interconnection model make these regulations a bad fit for the email model.

For example, if the user requests CLIR (Calling Line Identification Restriction) then its SSPs needs to differentiate the call handling, whether the peering partner is another commercial SSPs (transmit caller-ID, signal CLIR) or an enterprise (suppress caller-ID). Interconnecting with other SSPs that operate under the same rules simplifies compliance.

3.2. The PSTN Model does not fit, either

It is of course possible to rebuild the PSTN based on SIP instead of SS7. Some might argue that this is what the IMS and NGN efforts are all about. This is selling SIP and the Internet short: The basic infrastructure that the Internet offers allows for far more flexible interconnection arrangements than a simple copy of the PSTN structure.

Shared Layer 3 Infrastructure

The PSTN is based on trunk lines connecting voice switches. These trunks are manually established between carriers. Each such link needs physical ports, as well as, dedicated bandwidth. Establishing direct links between carriers is thus only sensible if call volumes can justify the effort.

In contrast to the point-to-point link world of the PSTN, SIP assumes an any-to-any IP based communication model. This has a profound impact on the economics of interconnection: A new peering is not a matter of provisioning a new bit-pipe, but just one of configuring border elements on both sides.

Economic theory states that there must be a optimal number of peerings per SSP given the costs to establish an interconnection versus the costs of transit. As the cost structure is fundamentally different, the mesh density of the optimal SIP based network will deviate significantly from the current PSTN.

Enterprise Peering

As a corollary: Peering between TDM-based enterprise telephony systems is usually limited to very high traffic cross-links. As enterprise-to-enterprise calls do not require settlement, there is a huge potential for additional peering in this space.

Dynamic Routing

Worldwide routing in the PSTN is still based mainly on manually established routes; these routes reflect business relationships. As a consequence, it takes years to a get a new number range routed in the PSTN. The switch from SS7 to SIP must be taken as a chance to upgrade the worldwide call routing to a better routing algorithm.

4. Core Assumptions

The author believes that some working assumptions have to be re-thought based on the feedback from current deployment.

4.1. The Real Problem with SPIT

SPIT and QoS/DoS issues have often been cited as reason why so few people (enterprises and commercial SSPs) run an open SIP service (i.e. accept SIP calls from the public Internet without pre-association). The IETF has taken on the challenge and tried to develop protocol extensions that should help with the SIP adoption. These are often based on the following reasoning:

```
+------------+          +-------------+          +-------------+
|We want to  |          |They need to |          |They have a  |
|interconnect|===(1)===>|run open     |===(2)===>|problem with |
|SSPs        |          |SIP servers  |          |SPIT and DoS.|
+------------+          +-------------+          +-------------+
```

A lot of time has been spent on step (2). Protocols and procedures have been proposed to mitigate the exposure of open SIP proxies. These include the consent framework for SIP, SPIT identification, anti-SPIT policy rules, the Identity: header, etc.

These are all worthwhile proposals that solve certain symptoms. Regrettably, they do not remove the roadblocks to widespread SIP-based peering. For that, they tackle the wrong set of problems. They assume that the email model can be successful and we just need

to make sure that all the associated problems are addressed.

This assumption is wrong. The author believes that it is necessary to tackle step (1) first.

The question therefore should not be "How do we deal with the unpleasant side effects of universal peering?", but "How can we get SSPs to peer at all?". Instead of "How to keep out the unwanted calls?" we should focus on "How can we entice willing partners to a peering?".

4.2. What is a SIP URI?

SIP URIs are used in various contexts: They can specify contact points (sip:user@10.0.0.1), they can specify next hop information in a private interconnection setting (sip:012345678@sbc1.chicago.us.example.net), and they can be public SIP URIs (sip:alice@example.com) for which the responsible SIP proxy can be determined using RFC 3263.

There is yet another interpretation of the SIP URI that may be relevant for Speermint: The URI as a simple identifier of a telephony customer without the commonly implied semantic on how that user can be contacted.

While that is close to the public URI, the difference is important: RFC 3263 does not apply. There is no simple, globally valid set of ingress points for calls towards that URI. The default SIP call routing logic just is not applicable to such URIs.

In other words: It is a very useful concept to use the SIP protocol and the URIs from RFC 3261 without also adopting RFC 3263, because the latter more or less implies the email model that has not seen a lot of deployment yet. It is thus expected that any SIP URI published in a public infrastructure ENUM will not imply the applicability of RFC 3263.

In order to place calls, some alternative to RFC 3263 needs to be developed that accommodates the needs of carriers.

4.3. Peering vs. Reachability

Whatever the interconnection setup, subscribers of a telephony service expect to be able to call all subscribes of all other SSPs. When the email model cannot be assumed, this requires the use of transit networks and thus some sort of routing mechanism to find a path to the destination SSP.

4.4. The Key to Routing Data

Currently, the PSTN side is using telephone numbers (TN) as the key to the routing information, whereas the RFC 3263 SIP uses the domain name.

The TN used to be the perfect identifier for routing as the hierarchical structure of the number corresponded to the network topology in the PSTN. The emergence of alternative carriers, number porting, and service numbers (free-phone, premium rate, ...) changed that: This is a form of "locator" / "identifier" split.

Prefix-based routing used to be the way to aggregate routes to telephone numbers in order to keep the routing tables and their updates manageable. While that is still useful to encode policies like "Route all of +43 to Carrier XYZ", within a country number portability made prefix-based routing increasingly inefficient.

Looking at the routing information from a database design point of view, it does not make sense to store the set of possible routes (incl. all meta-data like prices, capacity, quality,...) for every individual number, as these will be identical for at least all numbers operated by a single carrier in some area.

Any routing protocol will thus scale by several orders of magnitude better if it is based on some sort of "Destination-Group" that comprises a carrier identification plus optionally a service or region-specific tag.

While the telephone number is the starting point of the routing information lookup, it is not a good identifier to use as the key for storing routes.

4.5. Lookups vs. Announcements

Generally speaking, there are two ways how to distribute routing information:

On Demand

Whenever routing information is needed some external database is queried.

Example: DNS (including ENUM)

Pro-active Distribution

The information for all possible destinations is distributed before the first routing decision is made.

Example: BGP, OSPF

There are of mixed models as well, e.g., when an organization gathers routing information pro-actively to load an internal database that is then queried on an "on-demand" basis by network elements. Alternatively, some systems might pro-actively fetch the fraction of the global routing information covering the most likely destinations, and only fall back to on-demand queries for the rest.

The on-demand model requires a lower level transport infrastructure to contact the external database. It's thus clear that Layer 3 routing cannot use that model as this leads to a chicken-and-egg problem. However, for VoIP peering basic Internet connectivity can be assumed and the same constraints do not apply.

The pro-active model on the other hand operates under a different constraint: Distributing all information needed for the routing decisions to all carriers requires that the aggregate dataset size of these routing information tables does not exceed sensible size limits. For example, it is not feasible to replace the MX record lookup of mail-servers by a routing protocol which replicates the domain-name to mail-server mapping information to all ISPs and enterprise mail-servers. There are just too many domains in use and thus the "mail-routing tables" would exceed all practical limits.

With regards to TN based calls, both options are possible. ENUM according to RFC 3761 is a clear on-demand approach. On the PSTN side, downloads of database dumps are a common method to distribute routing information.

A scalable routing system is needed, up to the set of all active TN. They number in the billions. Installing a "full routing table" into a core telephony (soft-)switch is thus not feasible. Current PSTN implementations cope by crude aggregation of routes to foreign countries.

The number of reachable IP addresses is roughly of the same order of magnitude, but the aggregation properties of IP addresses reduces to global routing table to under 500000 entries without any impact on the quality of the routing decisions. Telephone numbers do not aggregate as well and therefore makes a TN-based protocol in the style of BGP infeasible (that is one reason why TRIP [4] failed).

The obvious solution is to add an on-demand mapping step ahead of the routing protocol. That on-demand mapping should include the option to seed a cache with the most likely destination TNs.

4.6. No National Solutions

Telecom regulation, especially concerning number assignments and interconnection rules, is a national matter. Calling patterns favor local destinations: local and national calls make up the majority of all calls. It is thus not surprising that number based PSTN (and some of the emerging VoIP-based) routing exchanges only deal with numbers from a single country code.

On the other hand, international voice termination markets deal usually not with individual numbers, but with routes to number prefixes.

Given the increasingly international footprint of voice operators the country-specific ways of handling inter-carrier routing is an anachronism. Just as the Internet routing does not care about national borders, there is no inherent reason why a single set of TN mapping and voice routing protocols cannot be seamlessly deployed in an international setting. There should be no need for special handling per country-code in the routing logic.

Consider the case of a pan-European mobile operator Foo. If Foo has signed a peering agreement with a local Austrian VoIP operator Bar, then Bar should pass all calls over this link that terminate in any GSM network that Foo operates. Ideally, Bar should notice when a number was ported to Foo's network in Germany and adapt the routing. If Foo acquires a new network in, say Bulgaria, then Bar should automatically route all calls to that set of numbers over the peering with Foo. All this should happen without Bar having to participate in Germany- or Bulgaria-specific TN database exchanges.

All this has been standard in Internet-based communication: both BGP, as well as, application layer protocols like SMTP or HTTP do not care about national borders. The protocol to resolve a .com name is the same as the one to resolve a domain under .cn. BGP speakers announce routes without any regard for national borders. Speermint should strive to achieve the same level of universality.

This does not preclude local optimizations. For example, if the mapping from TN to some sort of routing identifier is done by Infrastructure ENUM, then it makes sense to pro-actively prime the SSP name-servers with the data for all local numbers.

5. A Vision

This section provides a "view from 20,000 feet" on how the interconnection of voice networks might look like in the future.

5.1. The Players

Historically, voice interconnection used to be the domain of commercial voice network operators. Direct interconnection between corporate PBX systems was the exception as such links required dedicated circuits. With the move to IP based phone networks, this is bound to change.

Right now, there is no protocol support for controlled, secure, and dynamic peering between enterprises. Once this is made possible, the voice interconnection landscape is bound to change dramatically.

The future might look like a mixture between the email and the IP world: The majority of end-users will still contract voice service from large operators (either the same providers where they buy basic connectivity, or independent, voice-only operators). Smaller companies might run their own PBX (which buy upstream "minutes" from operators) or just buy hosted voice services. Larger PBX installations will be (from a technology point of view) indistinguishable from smaller SSP.

5.2. Interconnection fabrics

How will these voice networks interconnect? An "interconnection fabric" will be any physical or logical arrangement where two or more voice operators interconnect. There will be a number of very different such fabrics. No single one will be dominant.

5.2.1. Old-style TDM based interconnection

The SS7 technology will be here to stay for quite some time.

5.2.2. Private SIP interconnection

Whenever two SSPs use a dedicated IP link between their networks to exchange calls, this is a trivial interconnection fabric. Such links will be used for PBX to Upstream-SSP connections, as well as for carrier to carrier peering, just like private peering links in the Layer 3 world.

5.2.3. Commercial Peering Fabrics

Just as Internet Exchange Points have sprung up to simplify interconnection on layer 3, similar dedicated SIP interconnection facilities have appeared. These provide optimal conditions for secure, QoS enabled L3 connections, coupled with support infrastructure like directory services or SIP scrubbing.

5.2.4. Overlay networks

Instead of building a physical peering fabric, the public Internet can be used to build a logical peering fabric. That can been done with VPN technologies or just by implementing suitable access control on SBEs.

5.2.5. RFC3263-style SIP

The set of all SSPs implementing SIP according to RFC 3263 also forms a peering fabric.

5.3. An internet of Voice Networks

Just as the Layer 3 Internet is built by linking individual networks together, the world-wide voice network will similarly emerges as the meta-network that comprises all the member networks.

Such a meta-network (or a lower-case 'I' internet) needs as precondition

- a common addressing scheme
- an inter-domain routing protocol

5.4. The Lookup Function (LUF)

The first step in all call processing is the mapping from telephone number to the Destination-Group (see Section 4.4).

This mapping step is performed not only by commercial voice operators, but also by PBX installations, if they have more than just one upstream interconnection provider. As argued in Section 4.5, this needs to be an on-demand lookup into an external database. That database needs to accept queries from such a large number of legitimate questioners that it is unfeasible to keep that data private.

Regarding who can provision entries into that database, some more restrictions seem plausible, e.g. all registered service providers who directly provide service to the customer using that number.

The logical choice for the LUF is thus public Infrastructure ENUM.

In the case of URI-based dialing, ENUM cannot be used. We thus also need some sort of LUF which maps domain-names to Destination-Group. This assumes that all SIP URIs with the same domain-name will be routed the same path. That is a reasonable design decision.

A trivial approach would be just to re-use the domain-name as the Destination-Group. That works fine for carrier-owned domains, but raises significant scaling issues if customer-owned domains are used in AoRs. It is thus more sensible to define a simple DNS Resource Record which indicates the Destination-Group for SIP URIs with this domain-name.

As a corollary, the I-ENUM lookup could just as well return AoR SIP URIs and only a second lookup into the DNS yields the Destination-Group.

5.5. The Location Routing Function (LRF)

The second step in the call routing algorithm is the next-hop selection based on the Destination-Group the LUF generated.

For each individual call, this is simply a lookup in a routing-table which maps Destination-Group to one (or more) sets of "session establishment data (SED)" records which describe the next SIP hop for this call. This SED might contain such elements like IP address of an SBE, TLS parameters to use, QoS parameters, transcoding instructions, bandwidth limits and whatnot else.

This is straight-forward, and very much like the IP routing lookup into the forwarding routing table. The tricky question is how this routing table is constructed. There are multiple possibilities:

Static / Manual routing

A human operator can configure individual routing rules into the system.

A dynamic routing protocol spoken between two connected networks

Here the human operator just configures basic rules concerning which neighbors his systems should talk to and under which constraints they should accept and prioritize routing information, as well as what routes they should announce. This would be the voice meta-network's equivalent to the BGP of the Layer 3 Internet.

Transit Routing is possible with such a protocol.

Route registries, e.g. run by Peering Fabrics

Support systems built into fabrics might act as a routing clearinghouse which can reduce the n-squared complexity of a full routing mesh on peering exchanges. The Layer 3 analogy would be route-reflectors run by Internet exchange points. Such clearinghouses work fine for pure peering relationships, but have a hard time handling transit routing.

A dynamic routing protocol is vastly preferable to manual routing. The data exchanged in such a protocol could to be structured roughly in the following way:

Key

The key into the routing table is the Destination-Group (the result of the LUF).

SED

This is the information the device handling the call needs to generate SIP messages.

Route metadata

This information is needed to determine whether this route is acceptable for the SSP receiving the route as well as for priorizing alternative routes. It might contain:

* Network path (equivalent to the BGP AS-path)
* Re-distribution restrictions
* What media-types are supported
* Capacity-related data
* Regulatory data (what kind of networks are involved?)
* QoS offered by the route
* Trust (e.g. does the destination require some reputation network score for the caller? SIP-Identity?)
* Commercial terms

Poor-man's multihoming of an enterprise PBX could be implemented by allowing the LUF to return multiple Destination-Groups.

6. Design Pitfalls

Reading documents from various sources, a few common design mistakes

need to mentioned and cautioned against:

6.1. Reliance on fabric internals for multi-hop routing

The cardinal mistake here is to assume that a single peering / interconnection fabric will be the sole means of interconnection between SSPs. A good example is the design of the IPX (see http://www.gsmworld.com/documents/ireg/ir3444.pdf) and the idea of transparent SIP proxies within that network. From the scarce public documentation available it seems like the L3 paths selected via BGP routing should determine which (transparent) SIP proxies need to take care of transit agreements.

This can work in theory as long as *all* possible destination networks partake (i.e. the IP addresses associated with SBEs of the target network) in the BGP routing cloud of the IPX. Even simple non-IPX cross-links between two carriers mess up that scheme, let alone the idea that enterprises will also want to do their own call routing.

Generally speaking, a routing protocol within an interconnection fabric cannot replace a routing protocol which finds the right fabric over which to route the call.

6.2. Mistaking a protocol for softswitch/SBE control for LUF/LRF

The LUF and LRF steps need to happen somewhere in the SSPs network. They don't necessarily need to be done directly on the device which actually generates SIP messages. Neither need the inter-domain routing data exchanges be done between the actual SBEs. There certainly are good arguments for that, but this is not a given. (In the Layer 3 Internet, this is the common practice: the actual routers which do the IP packet routing speak BGP.)

For voice routing, it may well be a good choice to centralize the routing protocol handling and call routing decision processing to dedicated servers which are not in the SIP call path themselves. In such a setup, the soft-switch / SBE needs to query this central instance on what to do with each individual call.

That protocol is neither LUF, nor LRF. It is something completely different.

What does this protocol need? The device needs to pack up all information it has on that call (source URI, source "trunk", media requested, destination TN / URI, softswitch ID) into a query, send that to the central routing instance. The answer needs to contain basically the SED (which set of attribute-value pairs) which tells

the softswitch exactly what do to.

Good role models are access servers for dial-in application and the RADIUS protocol used there. ENUM is not suited for this. A number of recent I-Ds proposing extensions to the basic ENUM algorithm (e.g. source variability, encoding calling-party-ID in the query) and a long list of URI parameters show that people have tried to force ENUM to do something it wasn't design to do. One should use Radius, Diameter, or even MGCP as inspiration on how this can be implemented.

6.3. Mixing LUF and LRF

There is a good reason for the LUF/LRF separation. The two functions solve different problems and require different approaches (simple lookup vs. routing protocol).

Just as with normal network layering, it is sometimes tempting to violate the layering principles. For example, when the focus is not so much on full-scale transit-capable VoIP routing as on very simple SIP peering, then folding the LRF into the LUF can seem like a good idea.

As with all layering-violations, the perceived simplification quickly turns into a mess of workarounds once the scenario expands beyond the trivial case. In this case, just consider the effect of number portability on the routing table size.

6.4. Provisioning vs. routing

With the establishment of the DRINKS working group, data-exchange and "provisioning" in the context of VoIP peering are now subject to IETF work. The data-exchanges needed for LUF and LRF are very different:

To provision and implement the LUF, we need:
- o a provisioning protocol which adds TN->Destination-Group mappings to a registry
- o a query protocol into that registry
- o a database replication protocol to make local copies of LUF registries possible

The data-exchange for the LRF is:
- o Routing data updates between peering SSPs or a fabric's route aggregator (a routing protocol)

These are two completely different types of data exchanges. The LUF side is close to what needed to run the DNS, the second is a routing protocol like BGP. Trying to find a unified solution to these two problems is futile.

6.5. RFC 3263 is part of the problem

RFC 3263 hardcodes the email model of interconnection. With respect to speermint, it is part of the problem, and not part of the solution.

6.6. Disregarding enterprises

As mentioned above (Section 3.2), PBX interconnection used to be rare and utilized completely different technology than carrier to carrier interconnection.

This used to make sense in the TDM world, but in a VoIP landscape it is essential that enterprises also get the tools to simplify peering. As this is pure settlement-free peering, there is a huge potential for ad-hoc, dynamic peering if the protocols coming out of speermint and DRINKS support that in a secure manner.

The IETF should take great care that these two workinggroups do not intentionally exclude the requirements from enterprise peering from their work.

7. What building-blocks are missing?

A mentioned before, the LUF lookup could be implemented via Infrastructure ENUM, so no new protocol work is needed in that case, only a new enumservice-type needs to be defined. More work is needed for other parts of the framework:

7.1. LUF / I-ENUM provisioning

EPP combined with the ENUM extension is a possible solution. Some enhancements (e.g. block provisioning) and better support for the typical number portability operations might be needed, though. There has been good input into the DRINKS WG on that topic.

7.2. LRF Routing announcements

There currently is no IETF protocol suitable for this data-exchange.

7.3. A call-control protocol

As mentioned above (Section 6.2), Radius or similar protocol might be suitable.

8. Security Considerations

Not applicable at this stage of the discussion.

9. IANA Considerations

Not applicable.

10. Acknowledgements

The ideas expressed in this draft evolved during discussions with a large number of people. Version -01 includes significant input from Hannes Tschofenig.

11. References

11.1. Normative References

[1] Faltstrom, P. and M. Mealling, "The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)", RFC 3761, April 2004.

[2] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[3] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, June 2002.

11.2. Informative References

[4] Rosenberg, J., Salama, H., and M. Squire, "Telephony Routing over IP (TRIP)", RFC 3219, January 2002.

[5] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[6] Rosenberg, J. and C. Jennings, "The Session Initiation Protocol (SIP) and Spam", RFC 5039, January 2008.

Author's Address

Otmar Lendl
enum.at GmbH
Karlsplatz 1/9
Wien A-1010
Austria

Phone: +43 1 5056416 33
Email: otmar.lendl@enum.at
URI: http://www.enum.at/

Full Copyright Statement

Intellectual Property

Acknowledgment

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