draft-lendl-speermint-background-02

Thanks to Alex for presenting!

IETF 73, November 2008

Otmar Lendl

Why thís draft? Because I think that Speermint is in danger of just tinkering on details without first stepping back a bit and looking at the whole problem.

-01

- Content:
  - Analysis of the problem space we face
    - email model
    - PSTN interconnection model
  - Some general consideration on the possible solution space
    - Name / Address / Next-Hop
    - Key to routing
    - Lookups vs. Announcements

This first (and old) part offers some analysis of the problem space.

•A comparison of Interconnection models. Neither the pure email model, nor a simple re-implementation of the PSTN with SIP will do for Speermint.

•Some very high level design question

New in -02 (1)

- “view from 20,000 feet”: Based on this analysis, a vision on how voice interconnection might look like
  - Actors
  - Interconnection fabrics
  - A layer-7 internet of voice networks
  - Public Infrastructure-ENUM as LUF
  - BGP-like protocol as LRF

I’ve often complained (esp. in the drinks formation) that we as a WG don’t share a common vision how the world-wide voice routing setup might look like post-PSTN. That significantly hampered discussions, as everybody interprets drafts and terminology within their own mental framework of the solution.

So here I present what my vision and high-level view on how a pure IP worldwide telephony network could operate.

This includes:

•what are the actors in this game. (enterprises, small SSPs, large SSPs)

•How will these small voice networks interconnect?

•“internet” in the sense of a “large meta-network emerging from the interconnection of smaller networks”. We know what it takes to build such a network on layer3: solve addressing and routing.

•LUF: “public”, as too many players need to use it. Especially if enterprise PBXs participate in the global routing protocol. ENUM as TNs get translated to destination groups

•LRF: we have a full blown routing protocol to solve (remember: we’re not operating in email-model / RFC 3263 world)

New in -02 . Design Pitfalls

- Reliance on fabric internals for multi-hop routing
(That only works if that fabric becomes a monopoly to which everybody connects.)

  "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."

Now, given the thoughts from above, what mistakes do we need to avoid?

Example: the IPX wants to re-use the BGP L3 paths as implicit L7 routing. Yes, this works as long as

•Every single entity involved in voice routing get connected to the IPX (and thus all SBE names resolve to IP addresses in the IPX), and

•No other interconnections between two SSPs outside the IPX (and its L3 address space) exist. (Those are invisible to the BGP inside the IPX and are thus not available for transit traffic)

•No distinction between e.g. IM and voice routing

More pitfalls …

- Mistaking a protocol for softswitch/SBE control for LUF/LRF
  – "please tell me what I should do with this call" vs. "map this number to a URI"
  – draft-ietf-enum-trunkgroup, draft-kaplan-enum-source-uri
- Mixing LUF and LRF
  – That's a layering violation, pure and simple. These shortcuts will bite you.

I've read over and over again that the actual devices handling the SIP requests want to receive all necessary information regarding a call with a single query, thus a need for an integrated LUF/LRF lookup.

This is wrong on a few counts:

•LUF and LRF are abstract functions that may be co-located on the SIP device or not.

•Who ever said that LRF is an external lookup? Most likely, the LRF exchanges data (routing updates) independently of an actual call in order to build an in-memory routing table. A SIP call will this just trigger a lookup in the local memory and not an over-the-network LRF query.

•The fact that the box handling SIP requests might be rather dumb and delegate all high-level call routing to a centralized routing engine is completely orthogonal to the LUF / LRF design. Such a "routing-decision delegation"-protocol need other features than LUF/LRF, e.g. it must be able to transmit all information the dumb device knows (dialed number/URI, call source (trunk, SIP peer IP address, source URI), media requested, codecs offered, …) about the call towards the central intelligence and receive back detailed instructions on what to do. That looks just like RADIUS controlling an dial-in access server or a DSL termination device.

ENUM is the wrong protocol for that, although some people think it can be strong-armed into that service. See e.g. draft-ietf-enum-trunkgroup, draft-kaplan-enum-source-uri and that encrypted ENUM stuff.

•Mixing LUF and LRF is like putting Ethernet MAC addresses in the DNS in order to skip the ARP step.

After all, we know that the whole Internet is just a single switched Ethernet.

See last slide regarding "a single fabric achieves monopoly state".

Even more pitfalls …

- Provisioning vs. routing
  - This is DRINKS stuff: LUF needs provisioning, LRF needs a routing protocol. That doesn’t mix.
- RFC 3263 is part of the problem
  - It’s assumptions just don’t apply.
- Disregarding enterprises

Ceterum Censeo: If we manage to be not so dumb to mix LUF and LRF then we will notice that we have to completely different beasts on our hand. That will also have an effect on DRINKS. That simple requirements document planned there is a big mistake.

RFC 3263 considered harmful: It’s a good solutions for the email interconnection-model. It is a dead-end for Speermint.

Speermint has been dominated by carriers. We really need to make sure that enterprises can operate within the same routing framework as telcos. It works with BGP. I see no *technical* reasons why it can’t also work for SIP.