

LISP update

NANOG-46 Philadelphia, June 2009

Vince Fuller

(for the LISP crew: Noel Chiappa, Dino Farinacci, Darrel Lewis, Dave Meyer, Andrew Partan, and John Zwiebel)

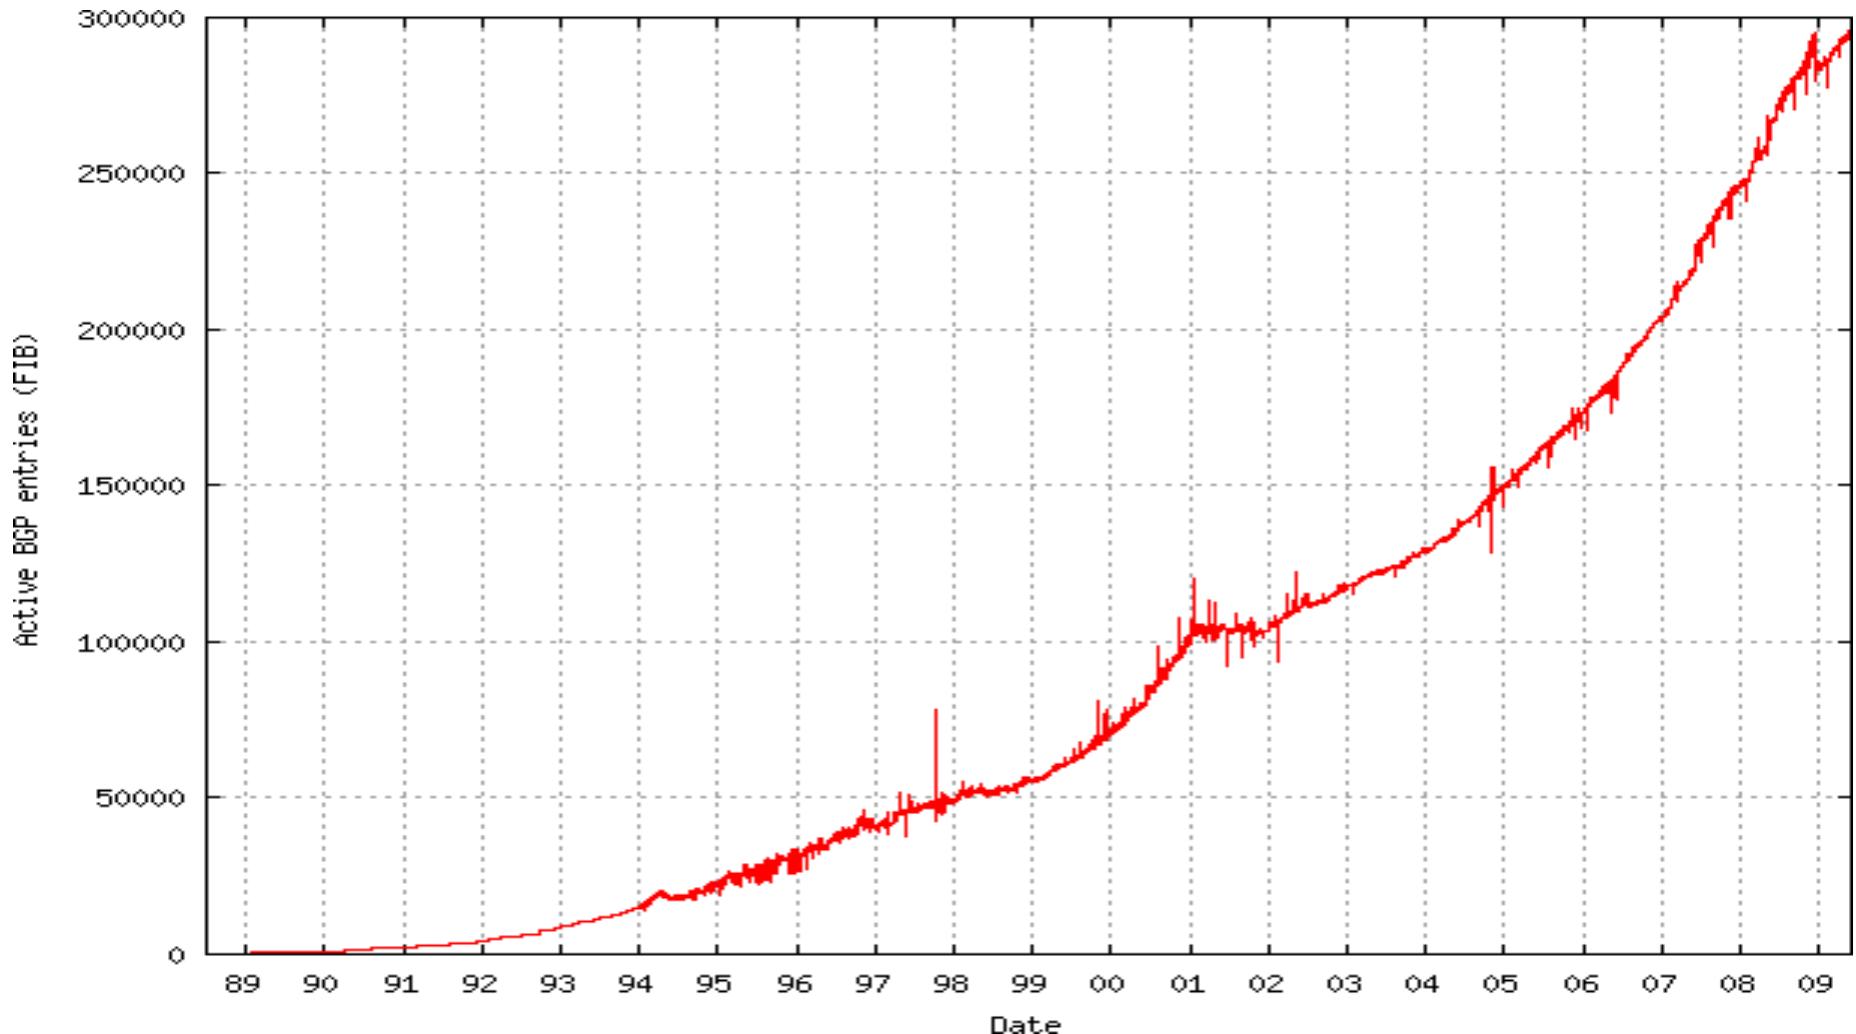
Agenda

- Problem Statement - good news & bad news
- Brief review of LISP
- LISP+ALT mapping database
- LISP-MS map server/map resolver - an easier way
- What's happening in the IETF?
- Implementation & Deployment Status
- Spec References
- Q & A

Problem Statement

- What provoked this?
 - Stimulated by problem statement effort at the Amsterdam IAB Routing Workshop on October 2006
 - RFC 4984
 - More info on problem statement:
 - <http://www.vaf.net/~vaf/apricot-plenary.pdf>
- First and foremost - scale the Internet

Internet Routing State



Good news & bad news

- + Recent analysis work by Geoff Huston suggests that stability (update & withdrawal rate) may be improving

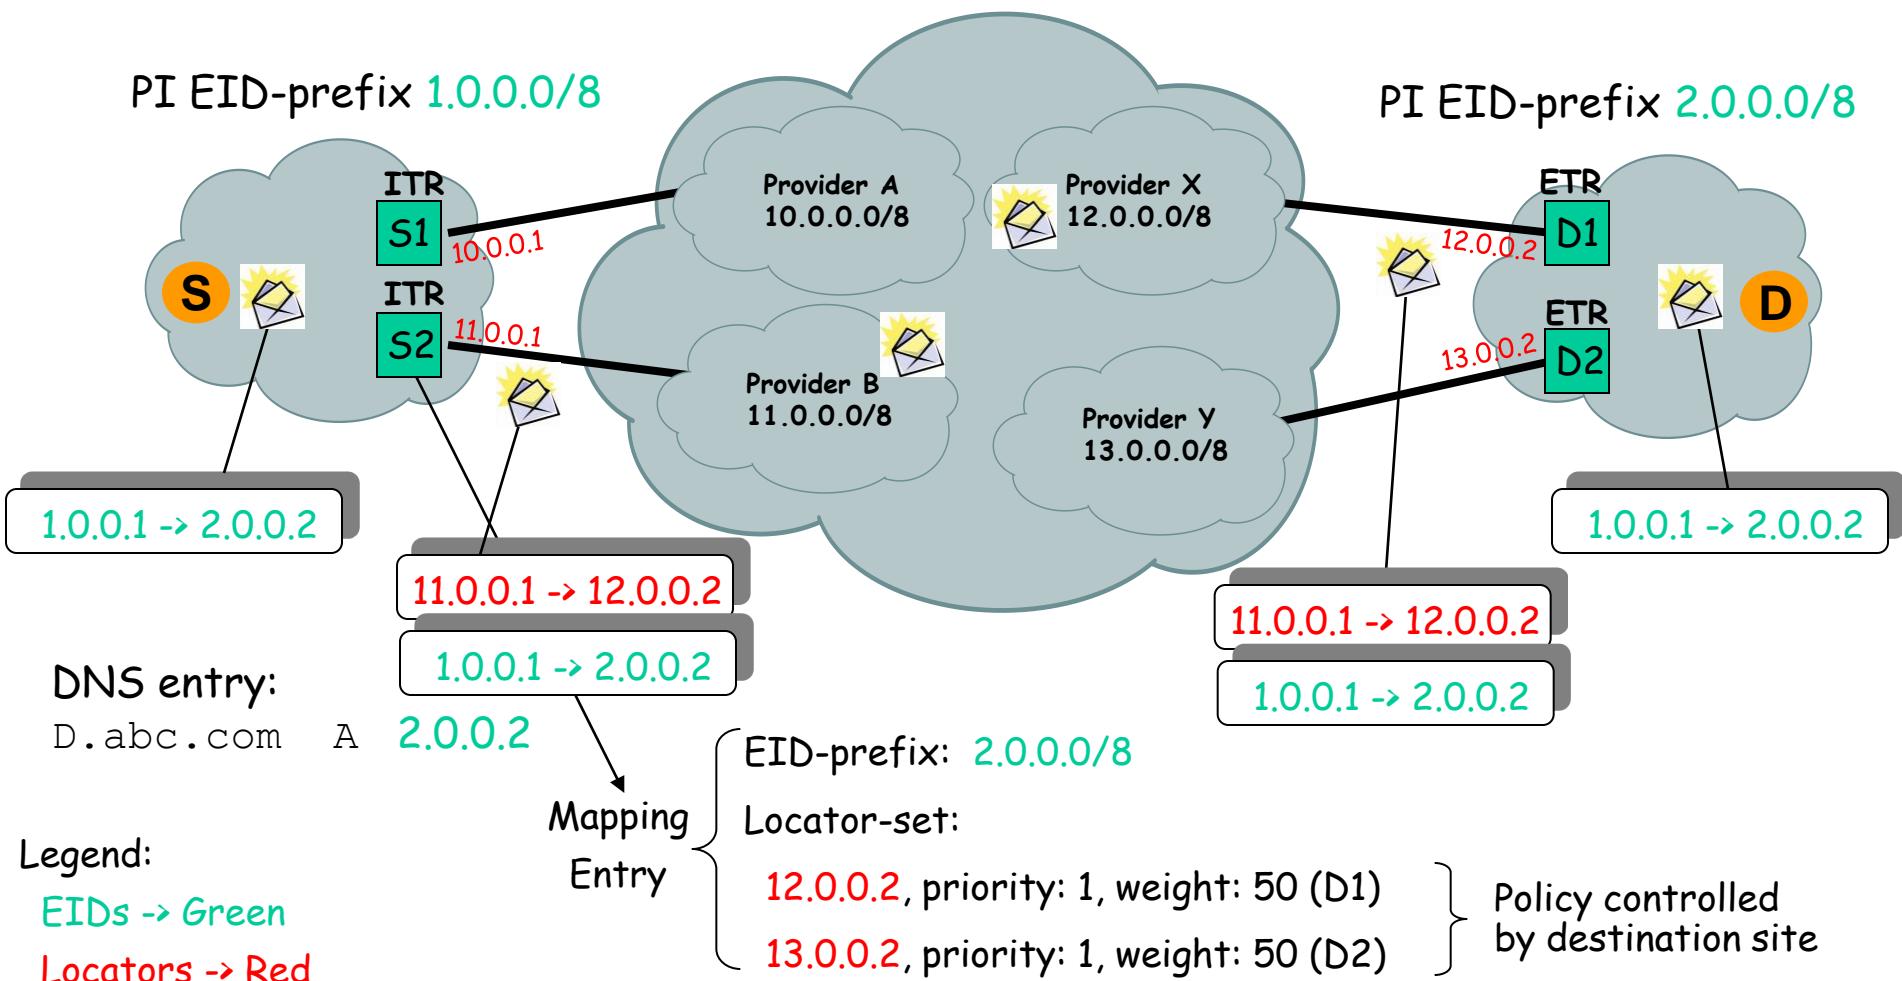
<http://www.potaroo.net/presentations/2009-05-06-bgp2008.pdf>

- Trend is still “up & to the right”
- Multi-homing is still hard
- Long-term trends are unclear

What is LISP?

- Locator/ID Separation Protocol
 - EIDs for hosts, topological RLOCs for “core”
 - Separate numbering allows routing to scale
- Ground rules for LISP
 - Network-based solution
 - No changes to hosts whatsoever
 - No new addressing changes to site devices
 - Minimal configuration file changes
 - Imperative to be incrementally deployable
 - Address family agnostic

Unicast Packet Forwarding

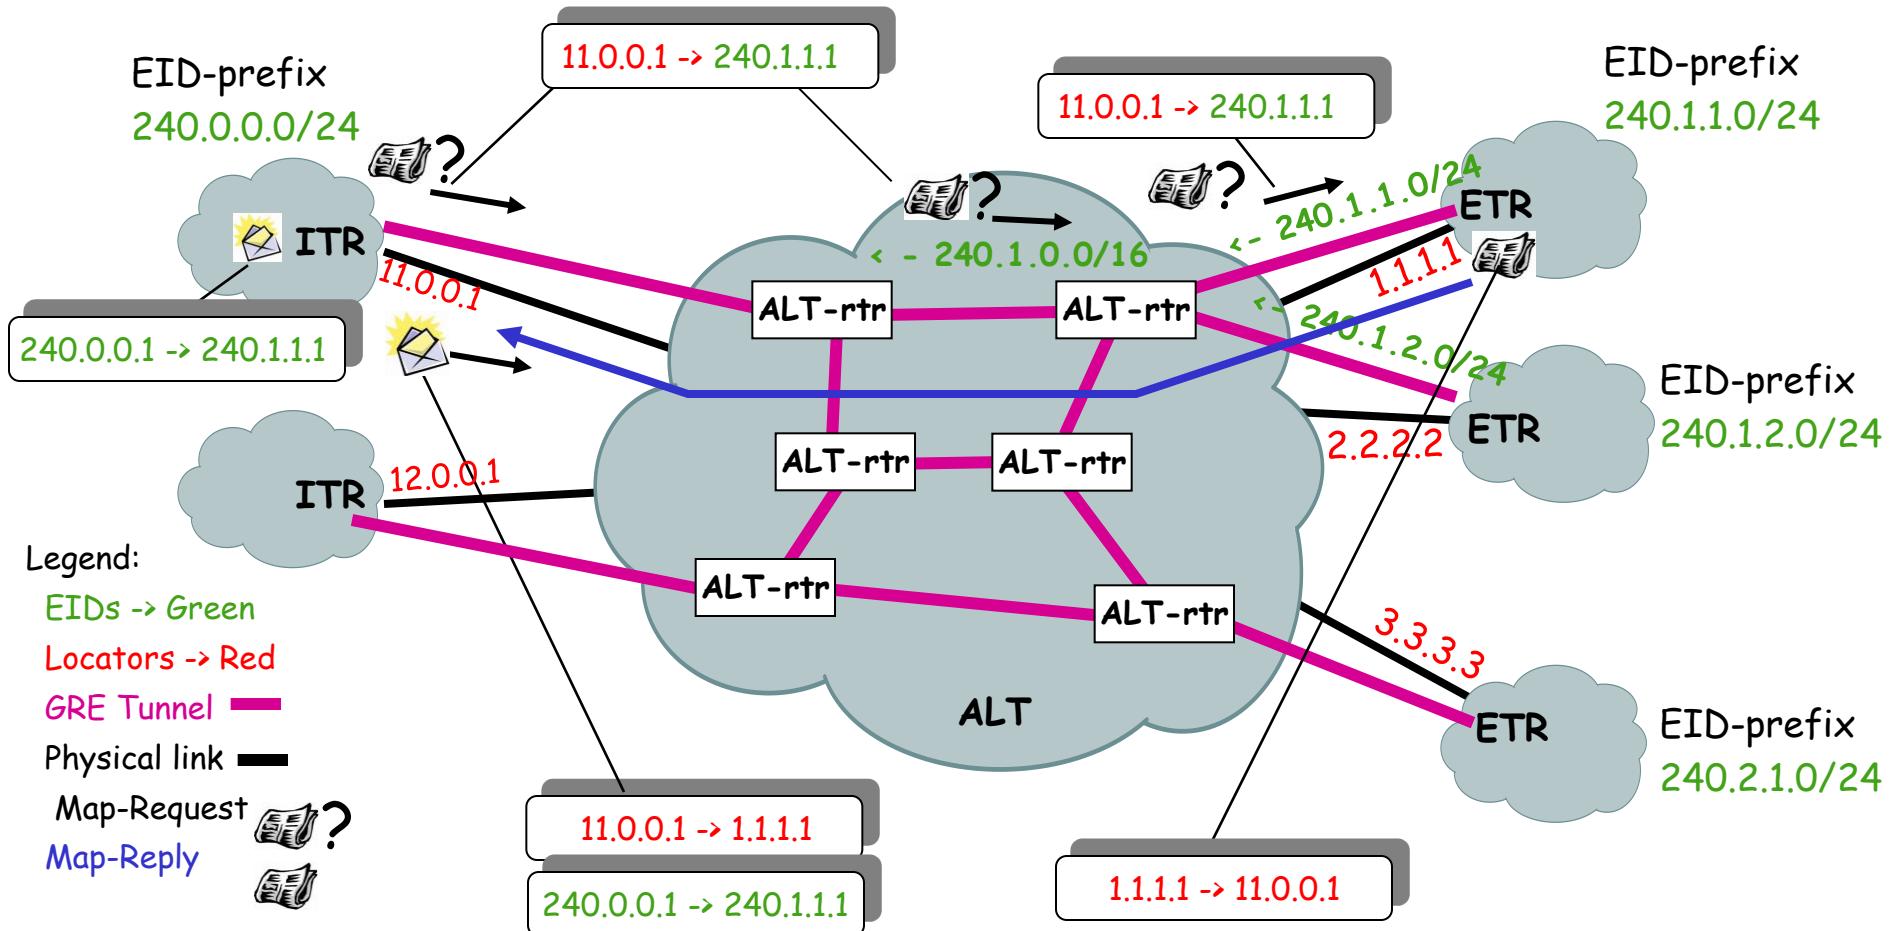


What is LISP+ALT?

- Mechanism for an ITR to find the ETR for an EID
- Advertise EID-prefixes in BGP on an alternate topology of GRE tunnels
- An ALT Device is:
 - An xTR configured with GRE tunnels
 - A Map-Server
 - A Map-Resolver
 - A pure ALT-only router for aggregating other ALT peering connections
- An ALT-only device can be off-the-shelf gear:
 - Router hardware, commodity Linux host, etc.
 - Just needs to run BGP and GRE

Using ALT to find an ETR

Hard way: extend ALT to xTRs for policy control

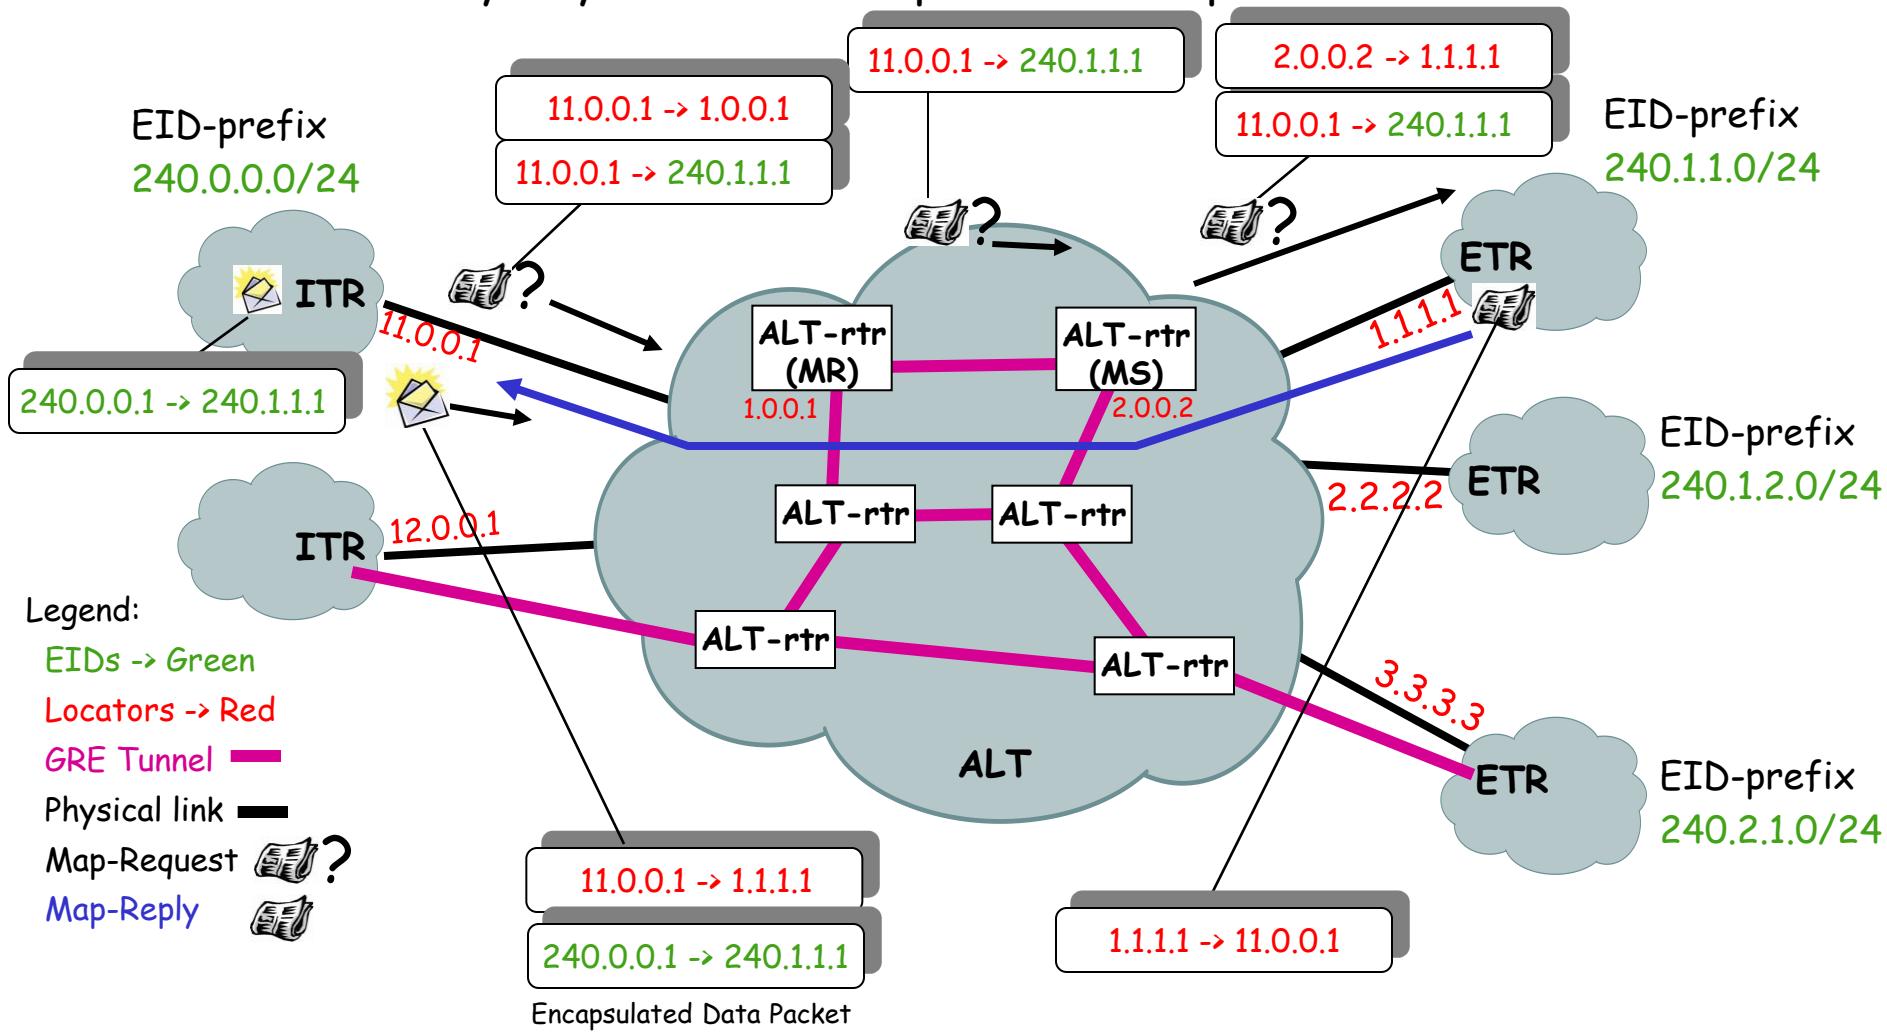


LISP-MS: ALT an easier way

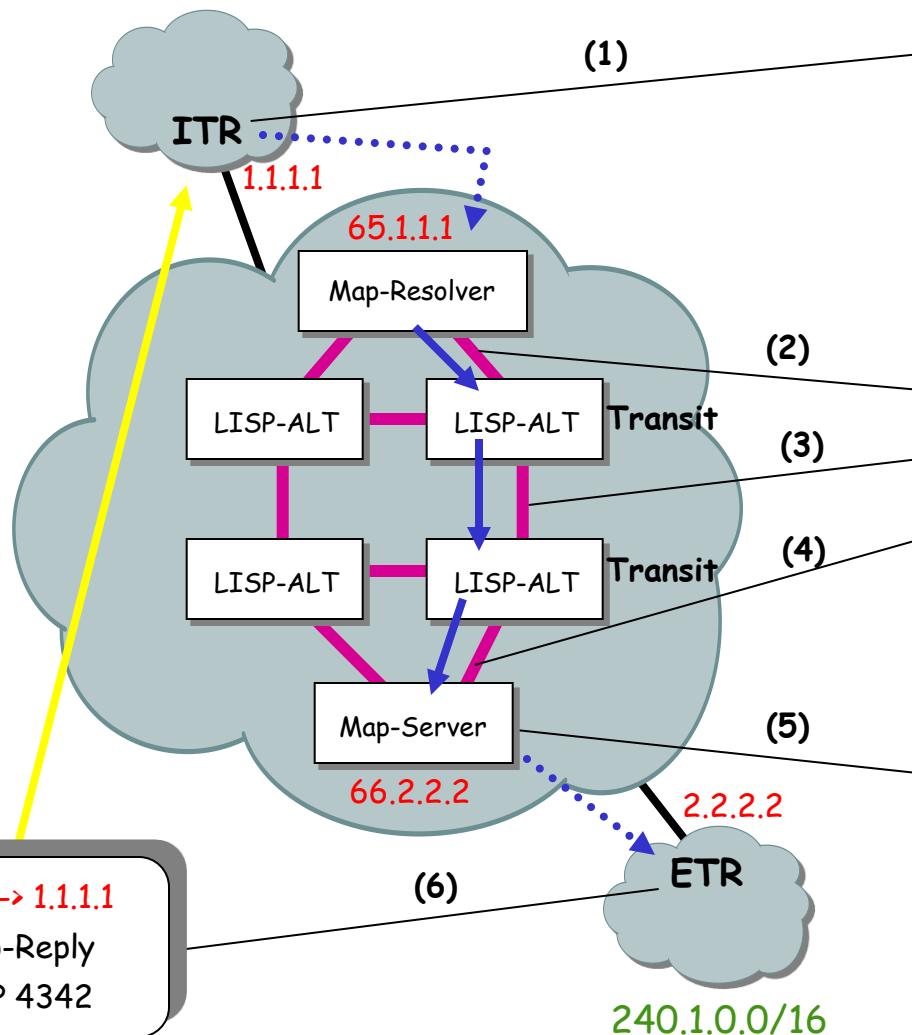
- ETRs register site EID-prefixes with Map-Servers
 - Securely with pair-wise trust model (no PKI needed)
 - Policy can be applied on Map-Servers before EID-prefix accepted into mapping service
 - ETR is still authoritative for its database mappings
- Map-Servers advertise EID-prefixes in to the ALT on behalf of their client ETRs
- ITRs send encapsulated Map-Requests to Map-Resolvers instead of connecting to ALT
- Map-Server/Map-Resolver functionality may be in existing ALT router (most likely) or separate box

Using LISP-MS to attach xTRs

Easy way: xTR uses Map-Server/Map-Resolver



Example in detail



Legend:

- EIDs → Green
- Locators → Red
- BGP-over-GRE —
- Physical link —
- Map-Request path ⚡...

1.1.1.1 -> 65.1.1.1

LISP Packet

UDP 4341

1.1.1.1 → 240.1.1.1

Map-Request

UDP 4342

1.1.1.1 → 240.1.1.1

Map-Request

UDP 4342

66.2.2.2 -> 2.2.2.2

LISP Packet UDP 4341

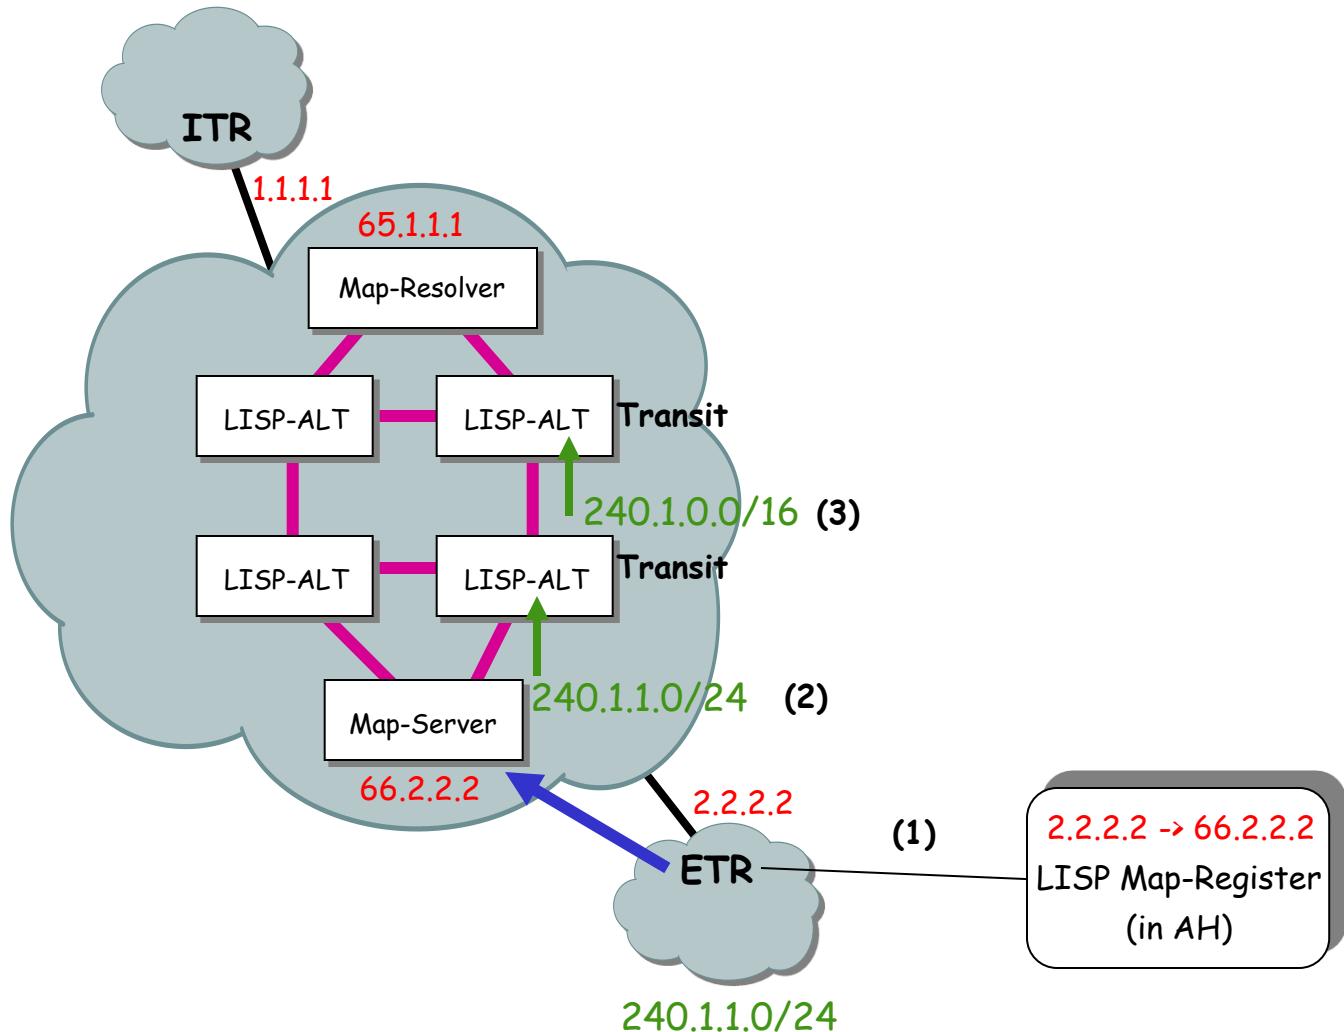
1.1.1.1 → 240.1.1.1

Map-Request

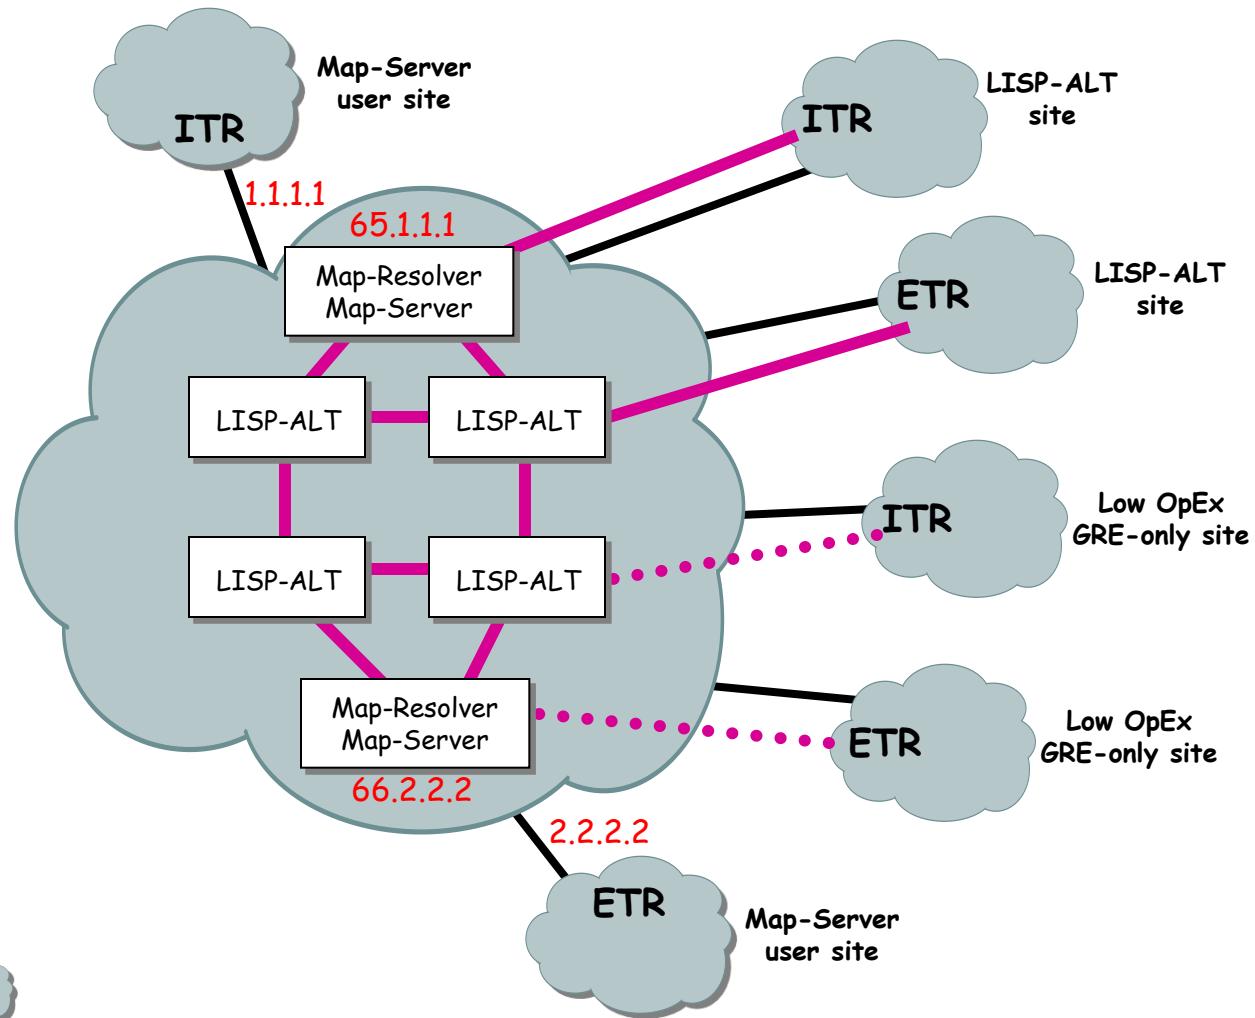
How Map-Server Registration Works

Legend:

- EIDs → Green
- Locators → Red
- BGP-over-GRE —
- Physical link —
- Map-Register —
- BGP update —



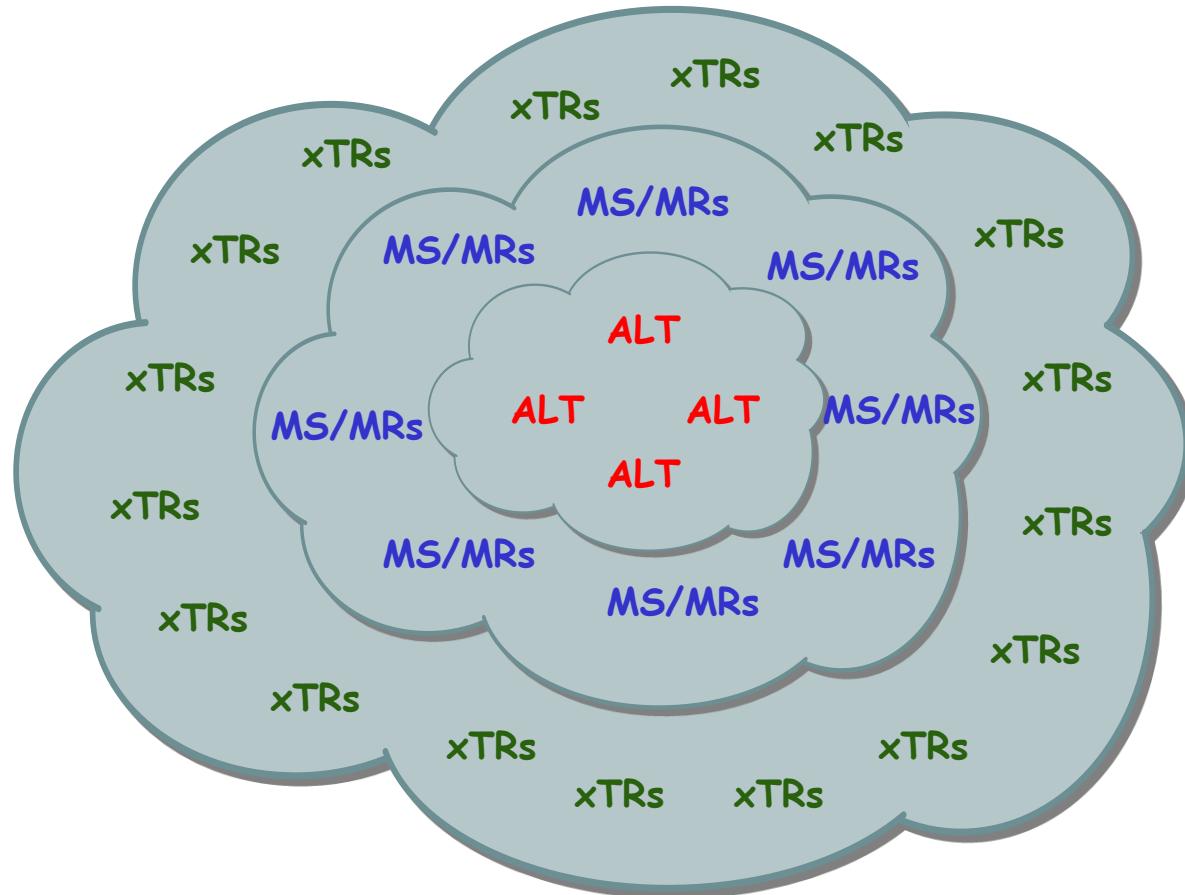
Interoperates with LISP-ALT Sites



Mapping Database Infrastructure

Legend:

- LISP Sites -> green
- 1st layer access infrastructure -> blue
- 2nd layer core infrastructure -> red



LISP in the IETF

- Started presenting in RRG and IETF ops/routing WGs in 2007-2008
- "EXPLISP" BOF in Dublin (July '08)
 - Basically a disaster due to process issues
- Much deliberation on RRG list, etc.
- LISP BOF in San Francisco (March '09)
- New WG at Stockholm (July '09)
 - Darrel Lewis & Sam Hartman co-chairs
 - Core LISP documents are now WG I-D's

Prototype Implementation

- Cisco NXOS, on NX7000 and Titanium
 - Underlying Linux code base
- Includes LISP, ALT, Interworking, and Map-Server/Map-Resolver functionality
 - “lig” diagnostic tool
- Software switching only
- Supports LISP for both IPv4 and IPv6
 - ITR, ETR, and PTR
 - LISP-NAT for IPv4 only

Other Coding Efforts

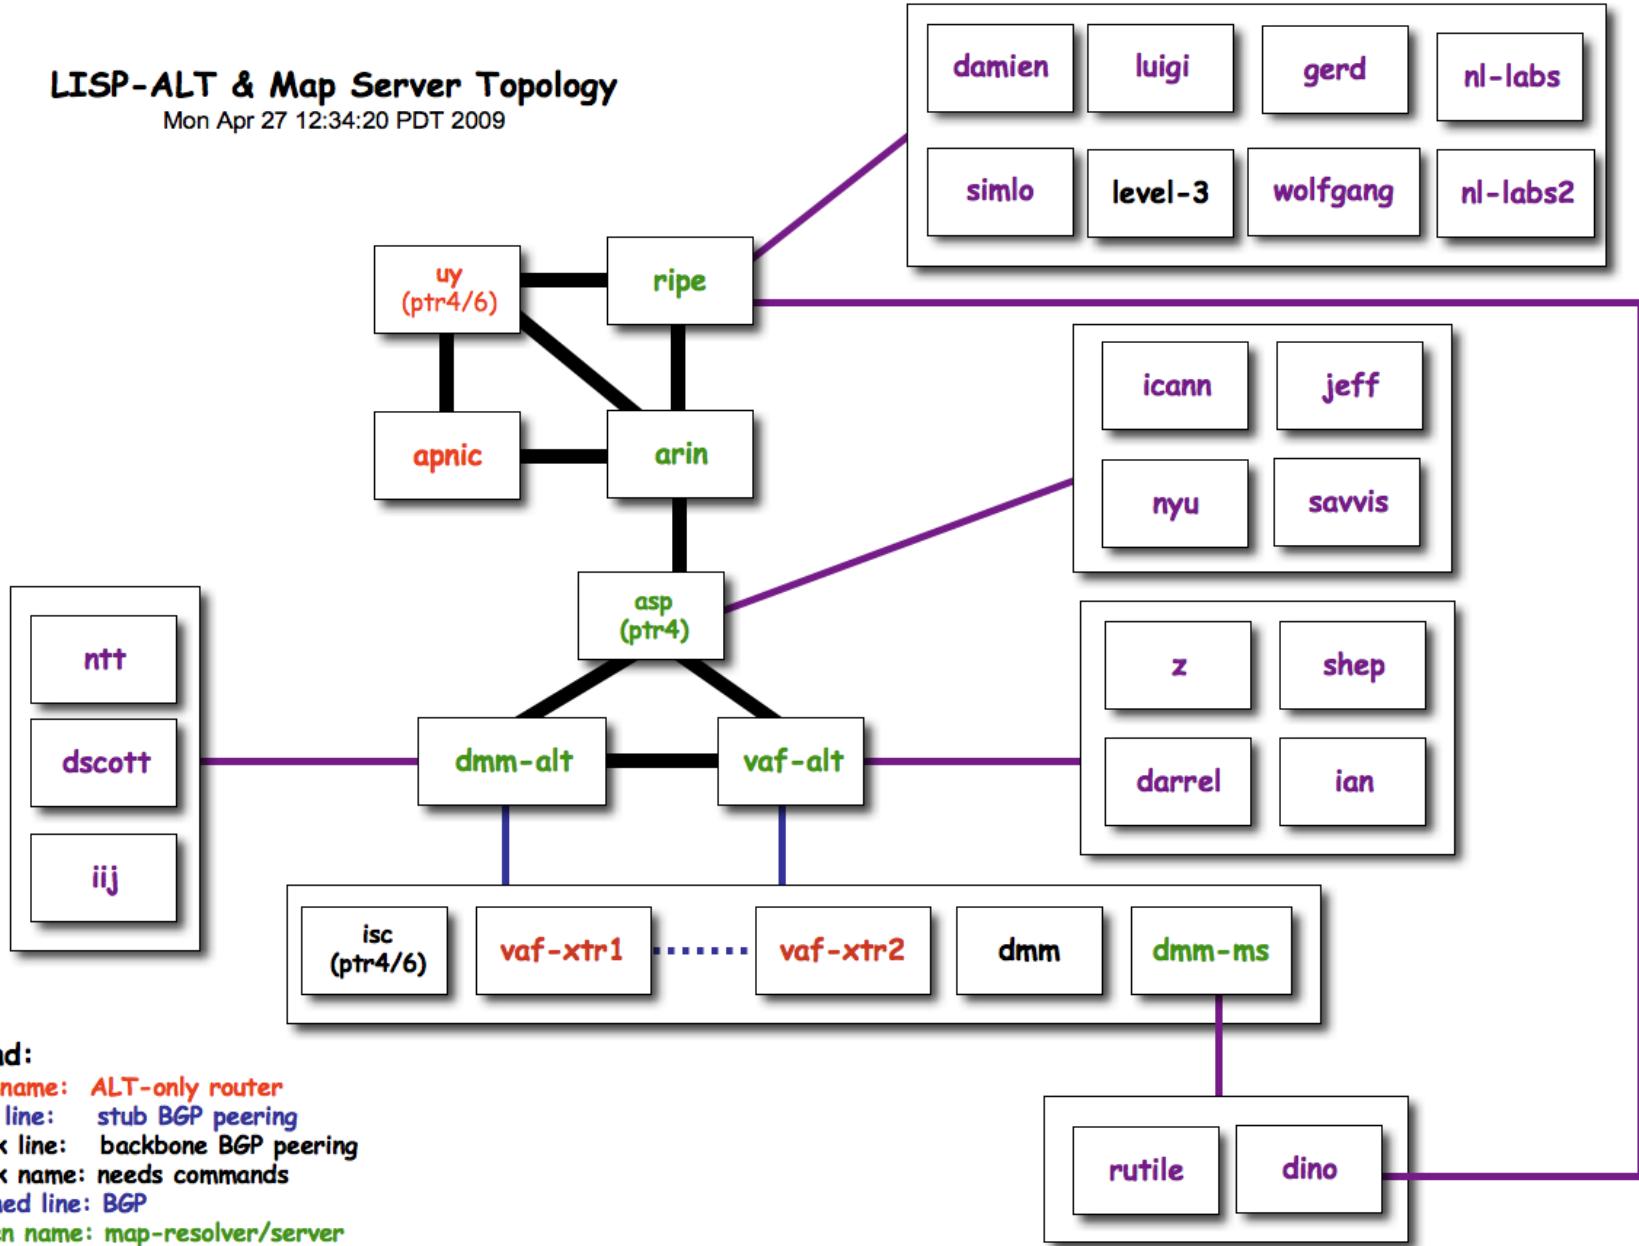
- IOS implementation under-way
 - Loc/ID split functionality
- Considering IOS-XR implementation
 - TE-ITR/TE-ETR functionality
- OpenLISP implementation been available for FreeBSD a while and being updated
 - For testing the specs
- Considering native Linux implementation (and recently learned of preliminary work on one)
- Any other efforts?

LISP Deployment

- LISP Pilot Network Operational
 - Deployed for nearly 2 years
 - ~32 sites across 7 countries
 - US, UK, BE, JP, UY, AU, DE
 - Uses the NX-OS Titanium Platform
 - IOS and OpenLISP platforms to be added
 - EID-Prefixes used
 - 153.16.0.0/16 and 2610:00d0::/32
 - GRE tunnels out of 240.0.0.0/4, 32-bit ASNs
 - RLOCs used
 - Current site attachment points to the Internet

LISP-ALT & Map Server Topology

Mon Apr 27 12:34:20 PDT 2009



Legend:

- Red name: ALT-only router
- Blue line: stub BGP peering
- Black line: backbone BGP peering
- Black name: needs commands
- Dashed line: BGP
- Green name: map-resolver/server
- Violet name: client of mr/ms

LISP Deployment

- LISP Interworking Deployed
 - Have LISP 1-to-1 address translation working
 - <http://www.translate.lisp4.net>
 - Proxy Tunnel Router (PTR)
 - IPv4 PTRs:
 - Andrew, ISC, and UY
 - IPv6 PTRs:
 - Dave (UofO), ISC, and UY
 - <http://www.lisp6.net> reachable through IPv6 PTR
 - <http://www.ptr.lisp4.net> reachable through IPv4 PTR

Open Policy for LISP

- It's been >2 years since the IAB Routing & Addressing Workshop
- This is not a Cisco only effort
 - We have approached and recruited others
 - There are no patents (cisco has no IPR on this)
 - All documents are Internet Drafts
- We need and seek designers, implementors, testers, and researchers
- As always, please let us know if you are interested

Internet Drafts

- `draft-ietf-lisp-01.txt`
- `draft-ietf-lisp-multicast-01.txt`
- `draft-ietf-lisp-alt-01.txt`
- `draft-ietf-lisp-ms-01.txt`
- `draft-ietf-lisp-interworking-00.txt`
- `draft-meyer-lisp-eid-block-01.txt`
- `draft-meyer-loc-id-implications-01.txt`
- `draft-farinacci-lisp-lig-00.txt`

- `draft-mathy-lisp-dht-00.txt`
- `draft-iannone-openlisp-implementation-02.txt`
- `draft-brim-lisp-analysis-00.txt`
- `draft-meyer-lisp-cons-04.txt`
- `draft-lear-lisp-nerd-04.txt`
- `draft-curran-lisp-emacs-00.txt`

References

- Public mailing list:
lisp@ietf.org
- Core LISP team:
lisp-dddvaz@external.cisco.com
- More info at:
<http://www.lisp4.net>
<http://www.lisp6.net>