LISP-DDT

Vince Fuller, Cisco

Glen Wiley. Verisign

NANOG-55 Vancouver, BC

Agenda

- Review of LISP: concepts and network elements
- The LISP mapping database today
- Description of LISP-DDT
- Development and deployment status
- Future direction

LISP Review – why do this?

- Initially, to try to find a way to scale the routing system
  - IAB Routing and Address Workshop, October 2006, RFC 4984
  - Separation of ID and location in IP addresses
  - LISP started during workshop, others have followed
- LISP-like indirection turns out to be useful for other things:
  - multi-homing with ingress traffic engineering
  - mobility
  - network virtualization
  - large-scale VPNs
  - ipv6 transition

LISP Review – why do this?

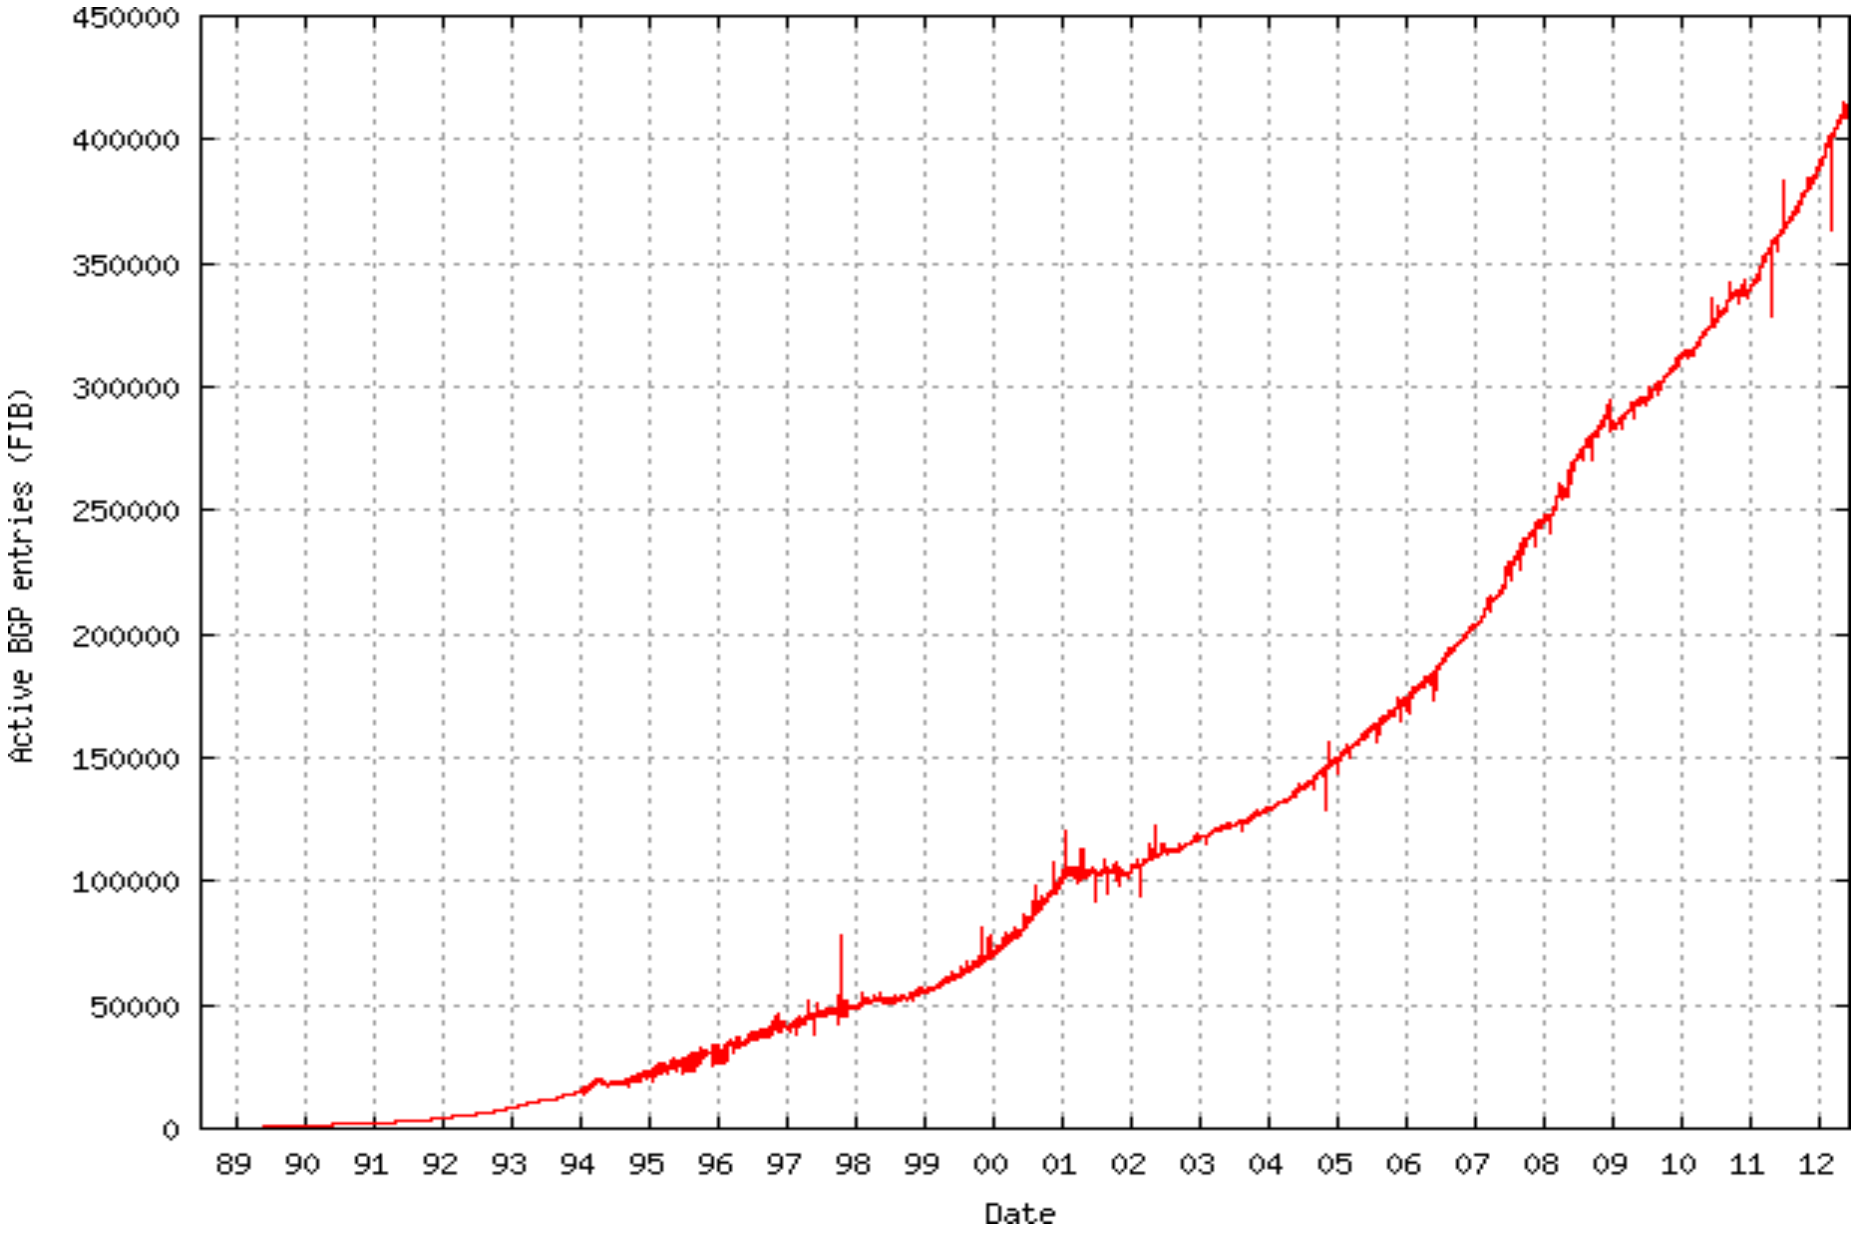

LISP Review – EIDs and RLOCs

- Endpoint IDs (EIDs) are used by hosts
  - assigned to “sites”, portable like “PI” space
  - not propagated by the global routing system
- Routing Locators (RLOCs) are used by the routing system
  - assigned to sites according to topology, like “PA” space
  - aggregated for use by the global routing system
- User data is LISP-encapsulated
  - RLOC source and destination in outer header
  - Source EID is “mapped” to RLOC by edge device
  - Mapping system consulted to find RLOC(s) for destination EID

LISP Review – network elements

- Ingress Tunnel Router (ITR) at LISP site
  - encapsulates user data in LISP datagram
- Egress Tunnel Router (ETR) at LISP site
  - decapsulates LISP datagram, delivers user data
- ETR/ITR functions typically co-located in one device (“xTR”)

- Proxy Ingress Tunnel Router (PITR), infrastructure
  - acts as ITR for non-LISP traffic (“legacy” Internet sources)
- Proxy Egress Tunnel Router (PETR), infrastructure
  - acts as ETR for non-LISP traffic (“hop over” non-LISP routers)

LISP Data Plane
Unicast Packet Forwarding

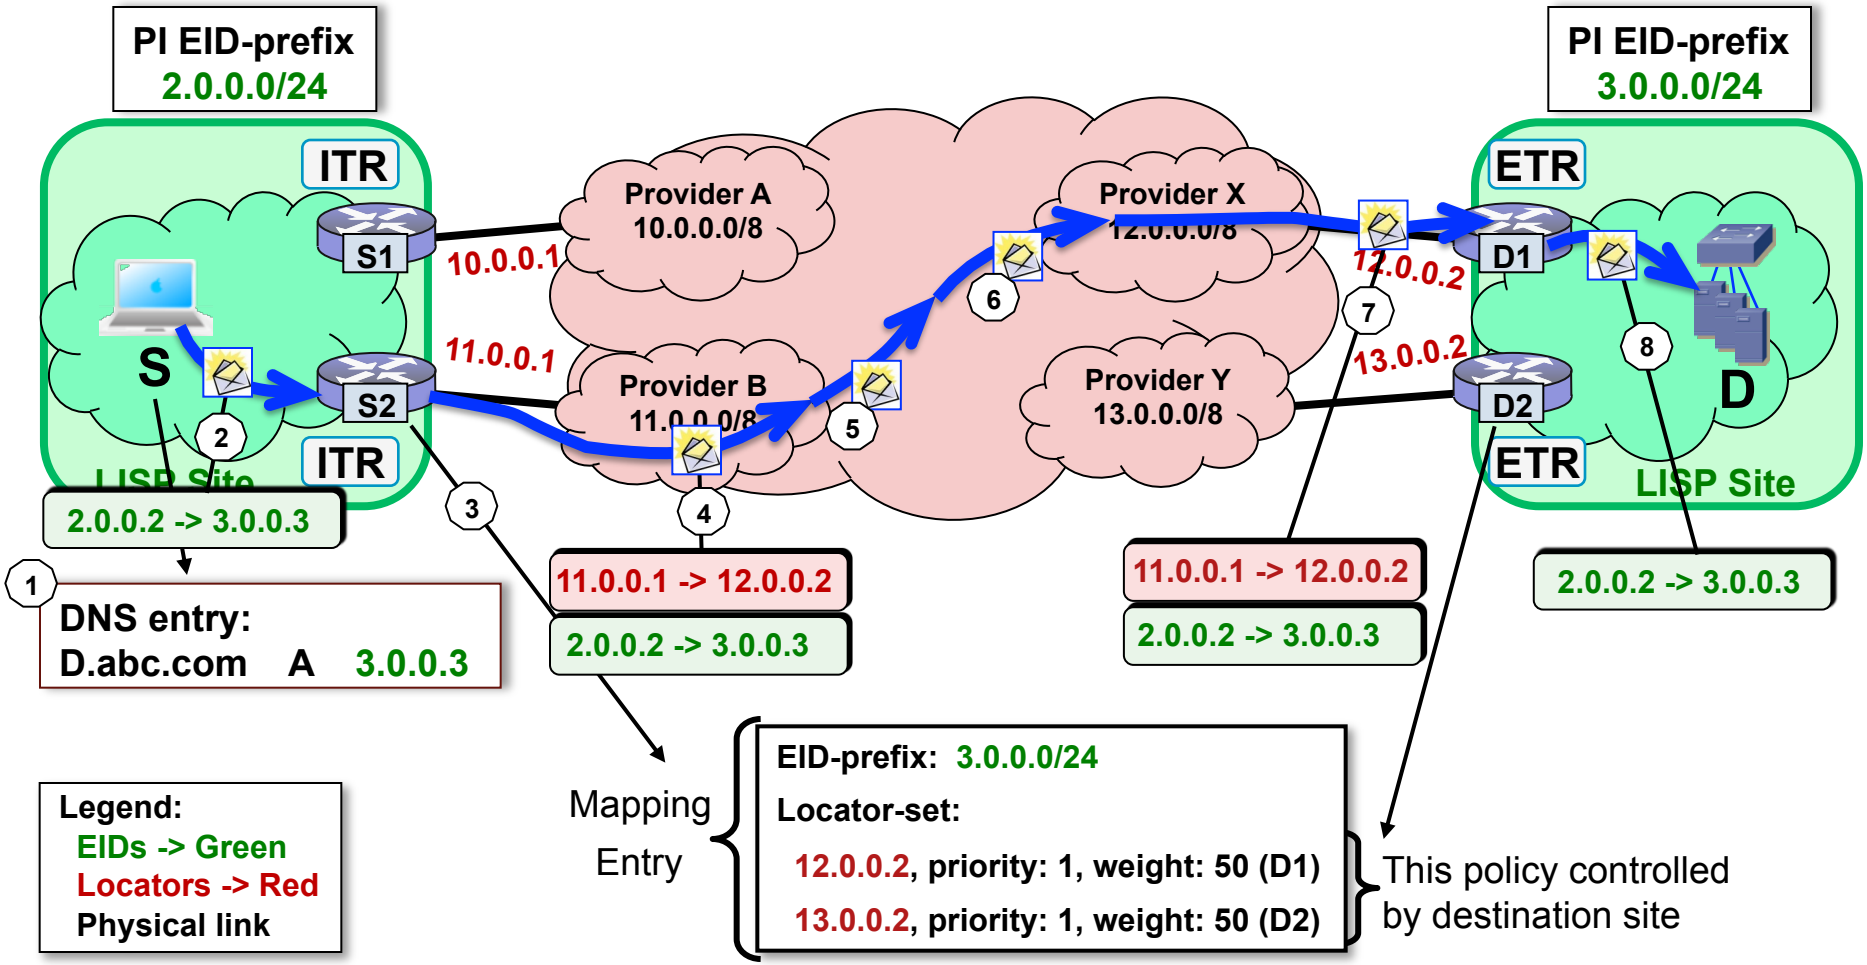

LISP Review - Mapping Database

- EID-to-RLOC Mappings
  - Originated by LISP site ETRs
  - ITR sends Map-Request to ETR (via MR/MS) to obtain mapping
  - Mapping Database enables an ITR to find the ETR(s) for an EID
- Must Support Millions of Sites, Changes
  - Rapid response to connectivity changes
  - Less so for subscription-time database changes
- Prototyped LISP+ALT Using BGP+GRE+VRF
  - Was conceptually simple, operationally complex in practice
  - Lacked flexibility for adding EID types, instance IDs, etc.

LISP Review – Database Infrastructure

- Map Server (MS) – publishes EID-prefixes in database
  - ETR registers to one or more Map Servers
  - A Map Server publishes EID-prefixes so an ITR can find them
- Map Resolver (MR) – index for EID to ETR
  - accepts Map-Request from ITR, finds correct ETR(s)
- LISP+ALT – overlay network between MR and MS
  - tunnels and BGP sessions between MR, MS, and ALT routers
  - ALT routers are intermediate nodes that provide EID aggregation
  - Map-Request, not user data, is forwarded through the ALT to MS
  - ETR receives Map-Request from MS and sends Map-Reply to ITR
  - ALT is being replaced by DDT

LISP Control Plane – MR/MS with ALT

Map Request uses BGP/GRE overlay from MR to MS

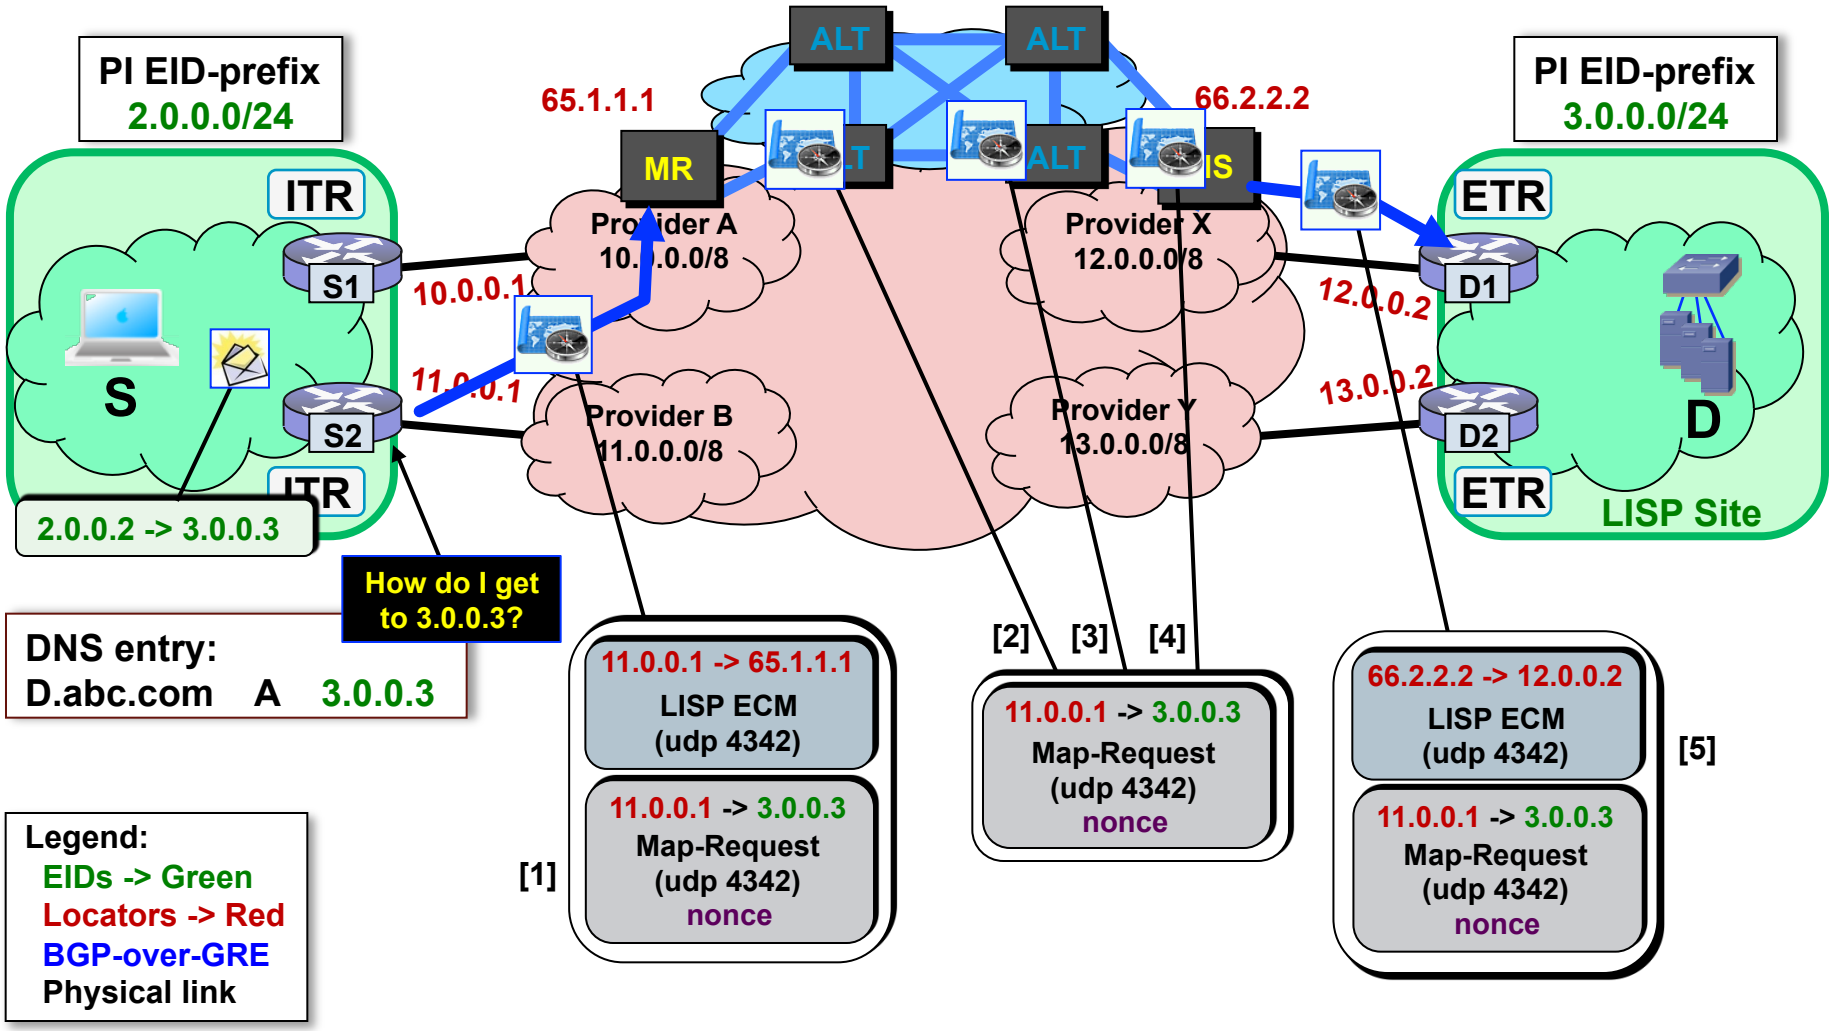

What is LISP-DDT?

- LISP Delegated Database Tree
  - Hierarchy for Instance IDs and for EID Prefixes
  - Statically Configured
  - Delegations are signed (public-key) and verified when used
- Conceptually, similar to DNS (IN-ADDR hierarchy)
  - but different prefix encoding, messages, etc.
  - we did try using DNS protocol directly, but…
    - prefix encoding was painful/ugly
    - negative map replies were problematic
    - couldn't do it right without protocol modifications

DDT Node

- Statically configured
- Authoritative prefix
  - IID and EID/length for which DDT node is responsible
  - root nodes have IID=any, EID=0/0
- Delegations
  - sub-prefixes delegated to “child” DDT nodes (or Map Servers)
- Accepts DDT Map-Request, returns Map-Referral message
  - contains pointer to node with more-specific information
    - NODE-REFERRAL for another DDT Node
    - MS-REFERRAL for a DDT Map Server
  - “negative” action codes also possible (more on those later)

DDT Map Server

- DDT Node may also be a DDT Map Server
  - accepts ETR registrations (just like any other Map Server)
  - forwards Map-Request to ETR
  - can return proxy Map-Reply (if configured to do so)
  - LISP-SEC authentication sent only following MS-REFERRAL
  - MS-ACK returned when EID fully resolved
  - “negative” action codes described later

DDT Map Resolver

- DDT Map Resolver finds RLOC for authoritative Map Server
  - Cache Map Request from ITR
  - Query the DDT hierarchy iteratively with DDT Map-Requests
  - Detect Loops/Delegation Errors
  - Follow referrals to find the right DDT Map-Server for an EID
- DDT Map Resolvers thus have state:
  - Referral Cache
  - Map-Request Queue
  - Key differences from “ordinary” Map Resolvers

Referrals & Their Actions

- ‘Positive’ referral is a pointer to a DDT node(s) with information about a (usually) more-specific EID-Prefix
  - Type 0, NODE-REFERRAL
  - Type 1, MS-REFERRAL
  - Type 2, MS-ACK
  - DDT MR uses this information to contact the next DDT node/MS
- ‘Negative’ referrals are used to indicate other actions:
  - Type 3, MS-NOT-REGISTERED
  - Type 4, DELEGATION-HOLE
  - Type 5, NOT-AUTHORITATIVE
- Referral includes public-key signature by delegator

Configuration and Setup

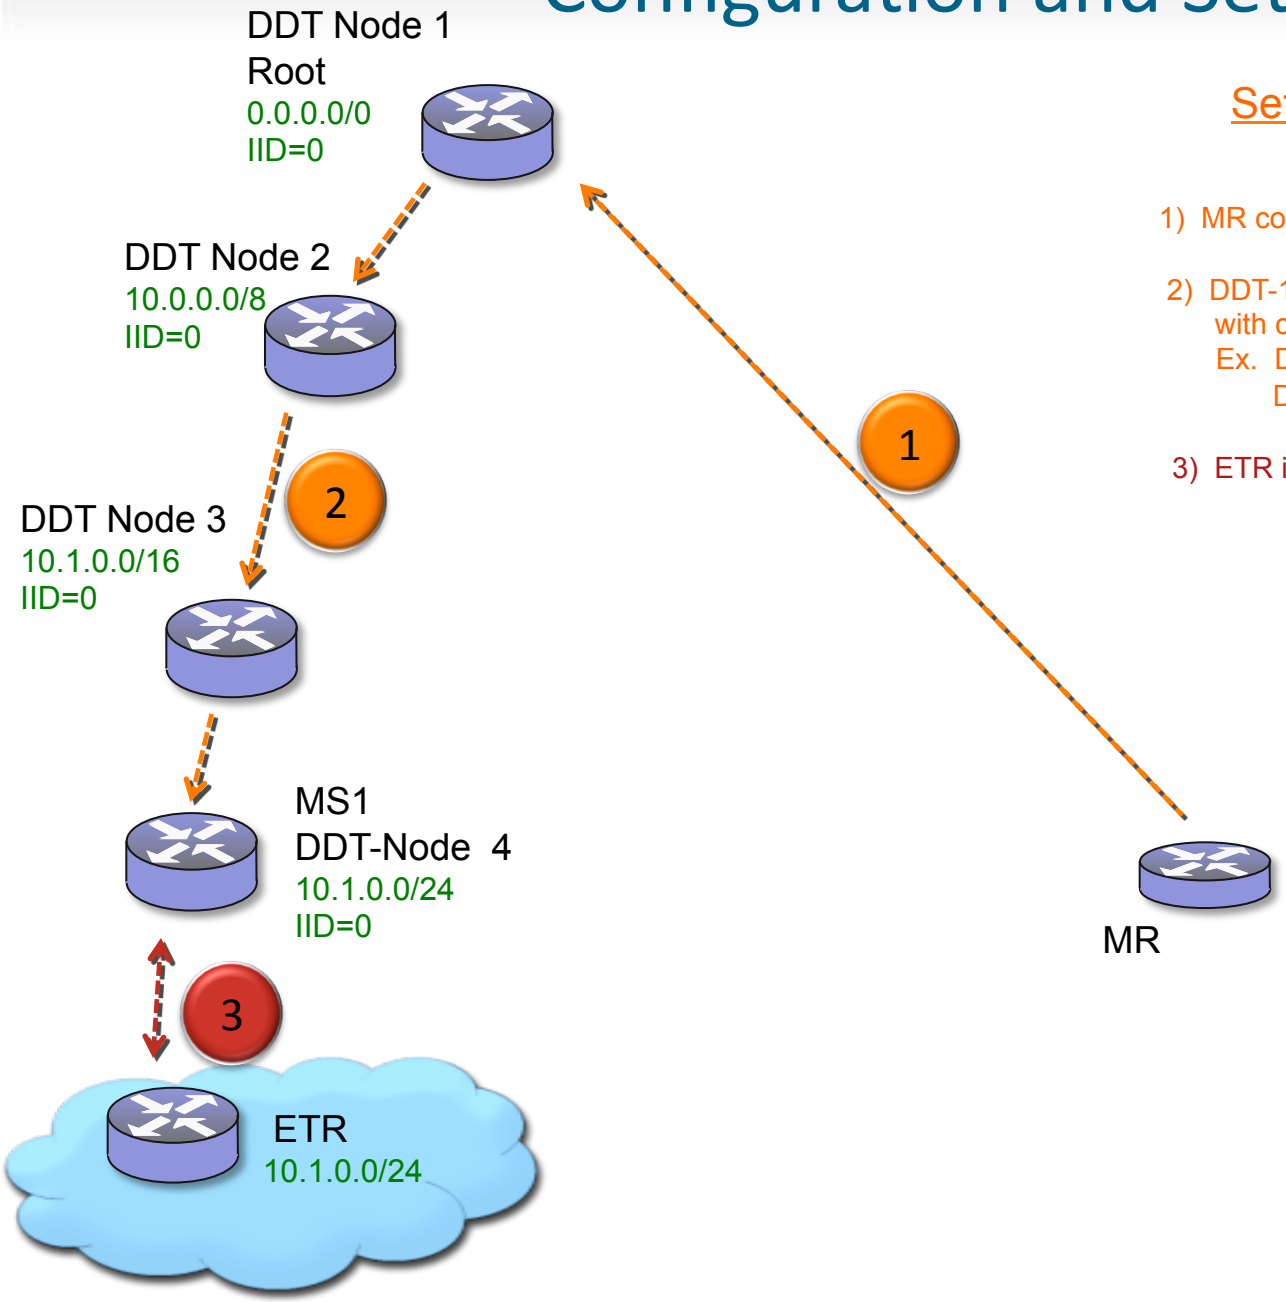

Setup & Configuration

1) MR configured with Root (DDT Node 1) RLOC

2) DDT-1, DDT2, DDT-3, DDT/MS-4 configured children with child prefixes, and authoritative prefixes
   Ex. DDT-2 Delegates child 10.1.0.0/16 to MS3
   DDT-2 configured authoritative for 10/8 in IID0

3) ETR is registering its EID to the Leaf MS

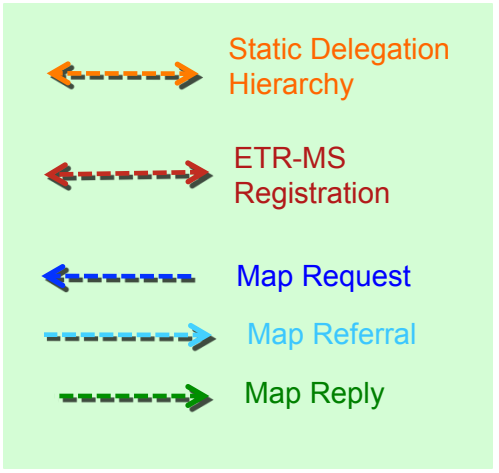

First Request Packet Flow

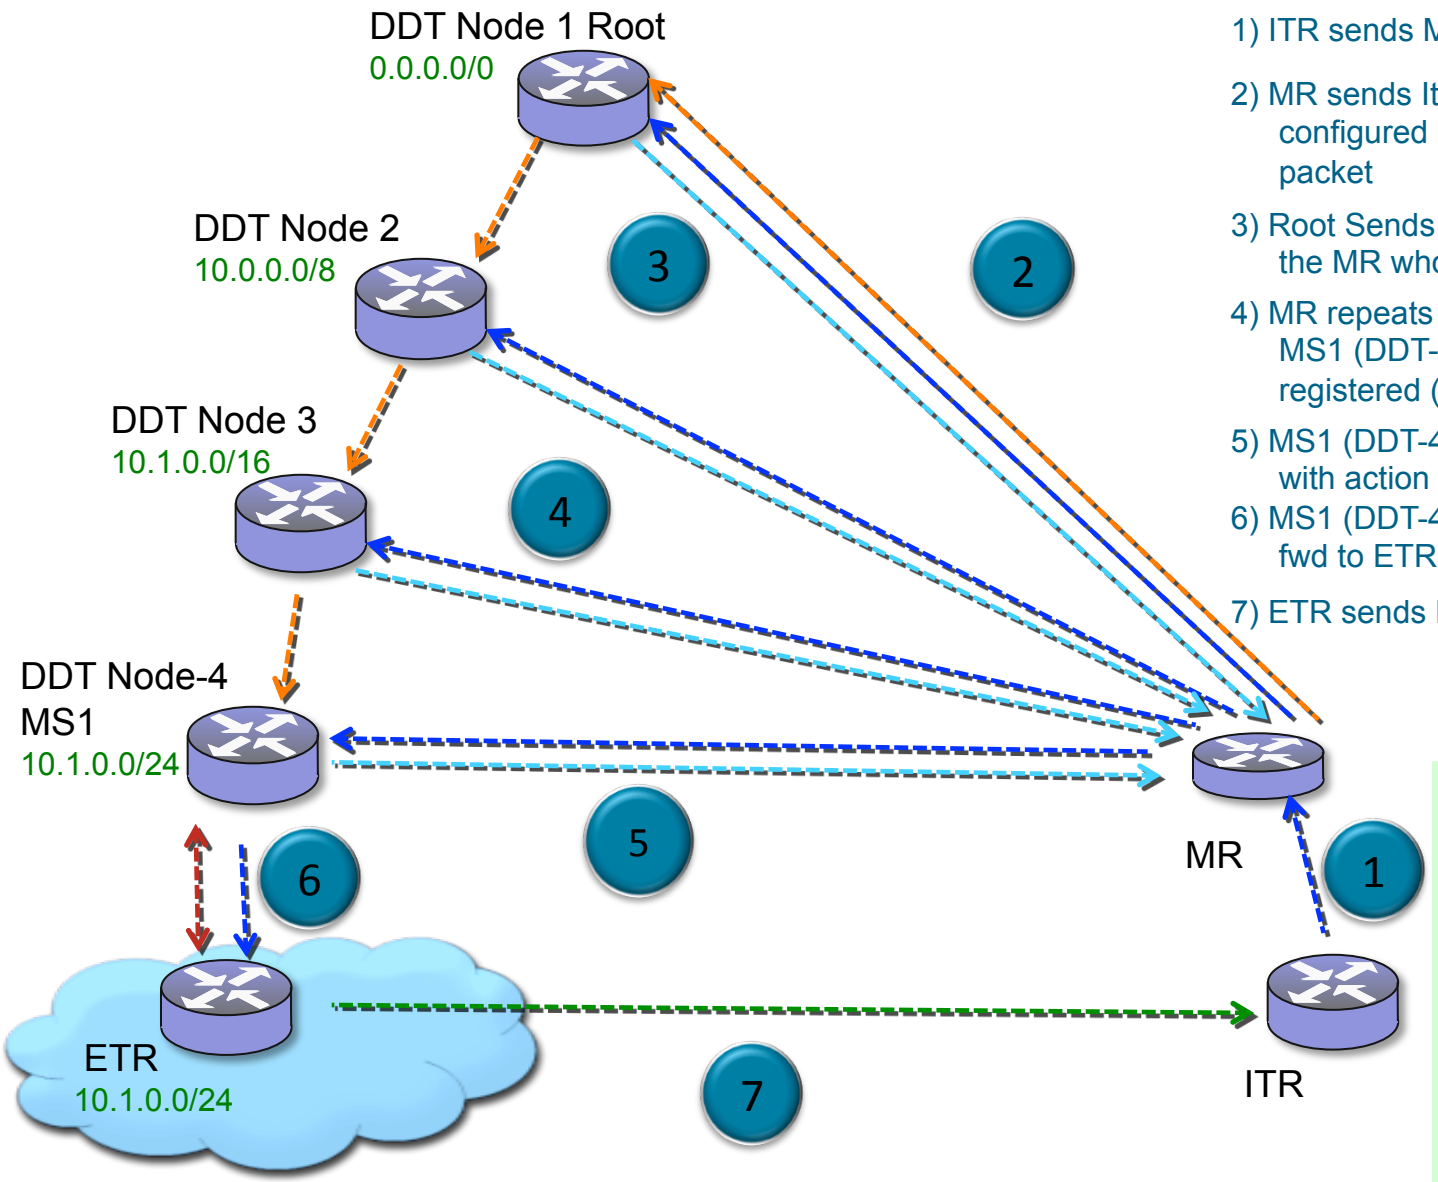

Map Request, Referral, & Reply

1) ITR sends MRQ to MR via ECM
2) MR sends Iterative-MRQ to its statically configured Root DDT-Node via ECM-Like-packet
3) Root Sends a Map Referral to MR informing the MR who is the next DDT-Node (2) to try
4) MR repeats steps 2 & 3 until it gets to leaf MS1 (DDT-Node-4) which has the EID registered (or an error indication)
5) MS1 (DDT-4) sends Map-Referral to MR with action code MS-ACK
6) MS1 (DDT-4) receives, processes MR and fwd to ETR
7) ETR sends Map-Reply to the ITR

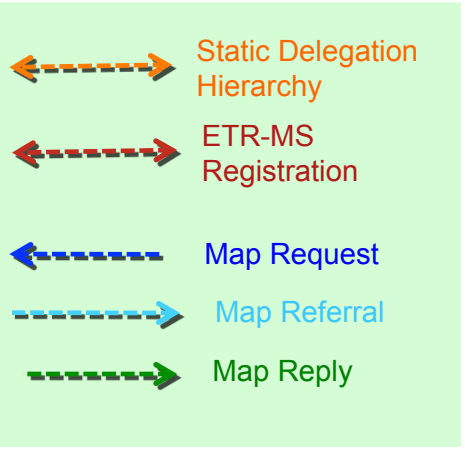

Steady State

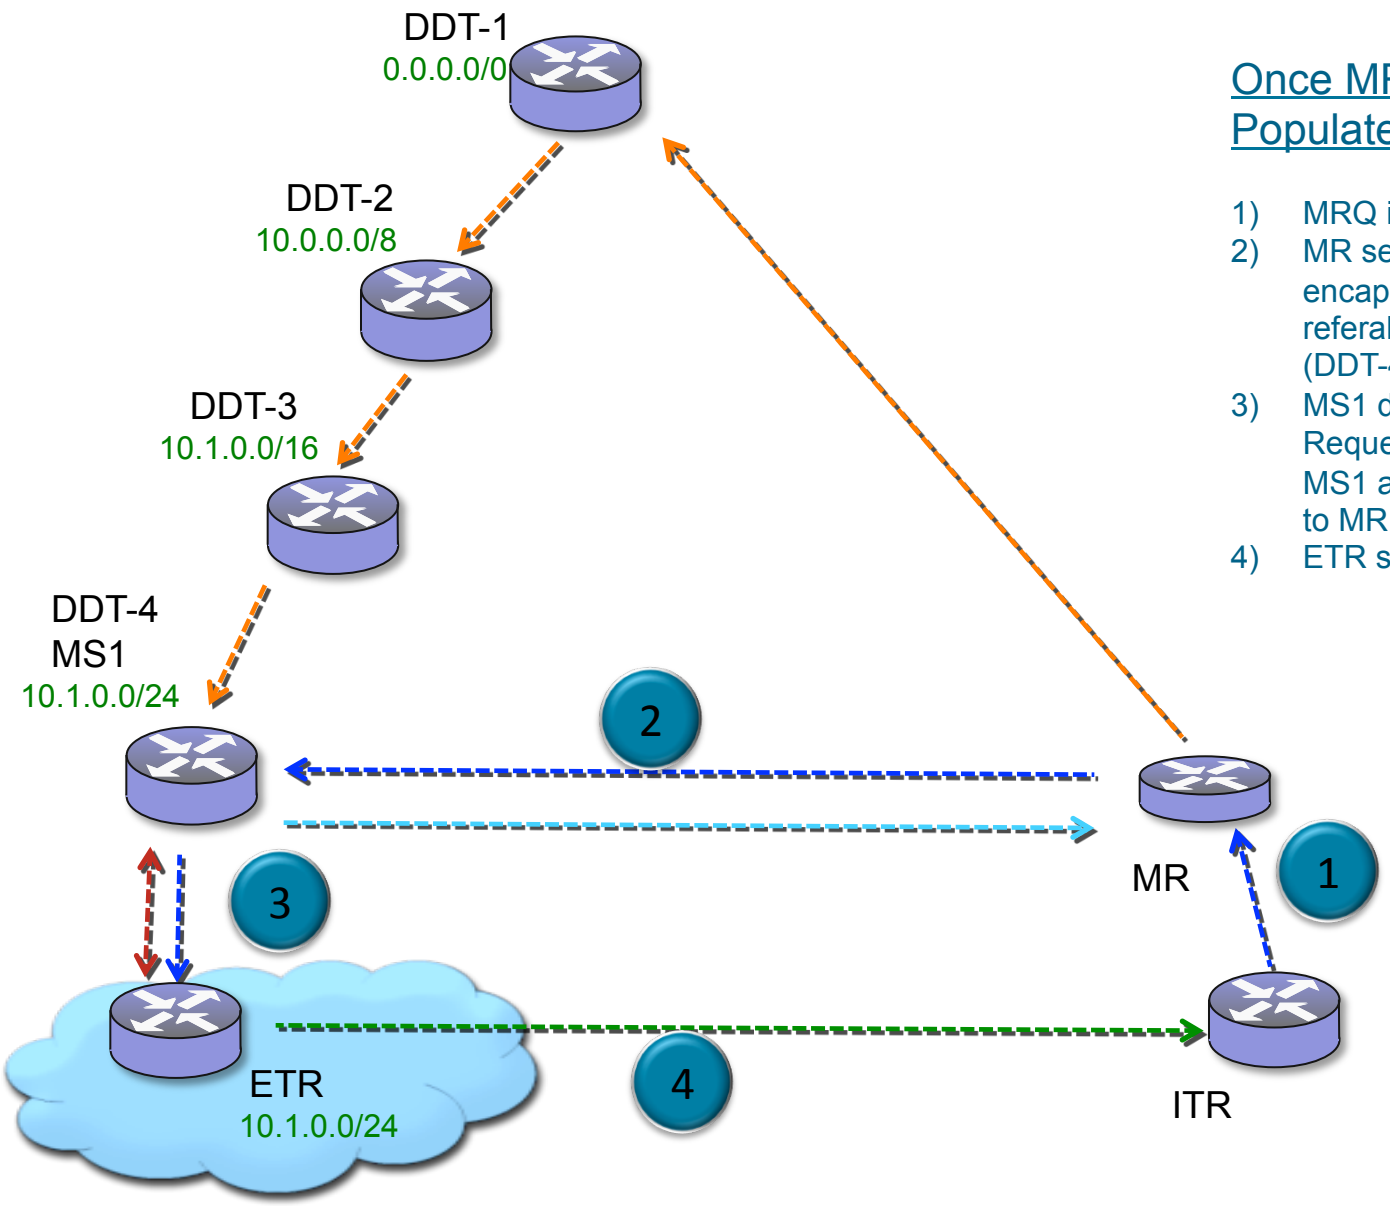

Once MR's Referal-Cache is Populated

1) MRQ in ECM arrives on MR
2) MR sends MRQ in ECM (possibly double encaped if lisp-sec is used to secure referal path) to Cached Map-Server MS1 (DDT-4)
3) MS1 decaps ECM and then sends Map-Request in new ECM to ETR
   MS1 also sends a Map-Referal ACK back to MR
4) ETR sends Map-Reply to ITR

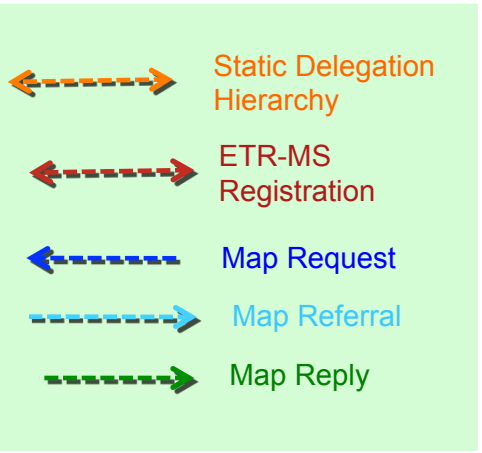

Implementation Status

- IOS and NXOS implementations complete
- OpenLISP implementation nearly complete
- Verisign implementation in progress
- Development and interoperability testing going on now
- Configuration is pretty simple
  - much easier to configure and operate than LISP+ALT!
- Does not include proposed DDT-SEC extensions

Organization and Operational Status

- Collaboration among Cisco, Verisign, Intouch NV
  - discussions with others, more welcome
- Common root: ddt-root.org
- Running on LISP pilot network
  - transition from LISP+ALT in March, 2012
  - ALT configurations removed in April, 2012
- Looking at various options for organizational structure
  - emphasis on transparency, scalability, efficiency, simplicity

DDT Beta (IID0) Network Deployment

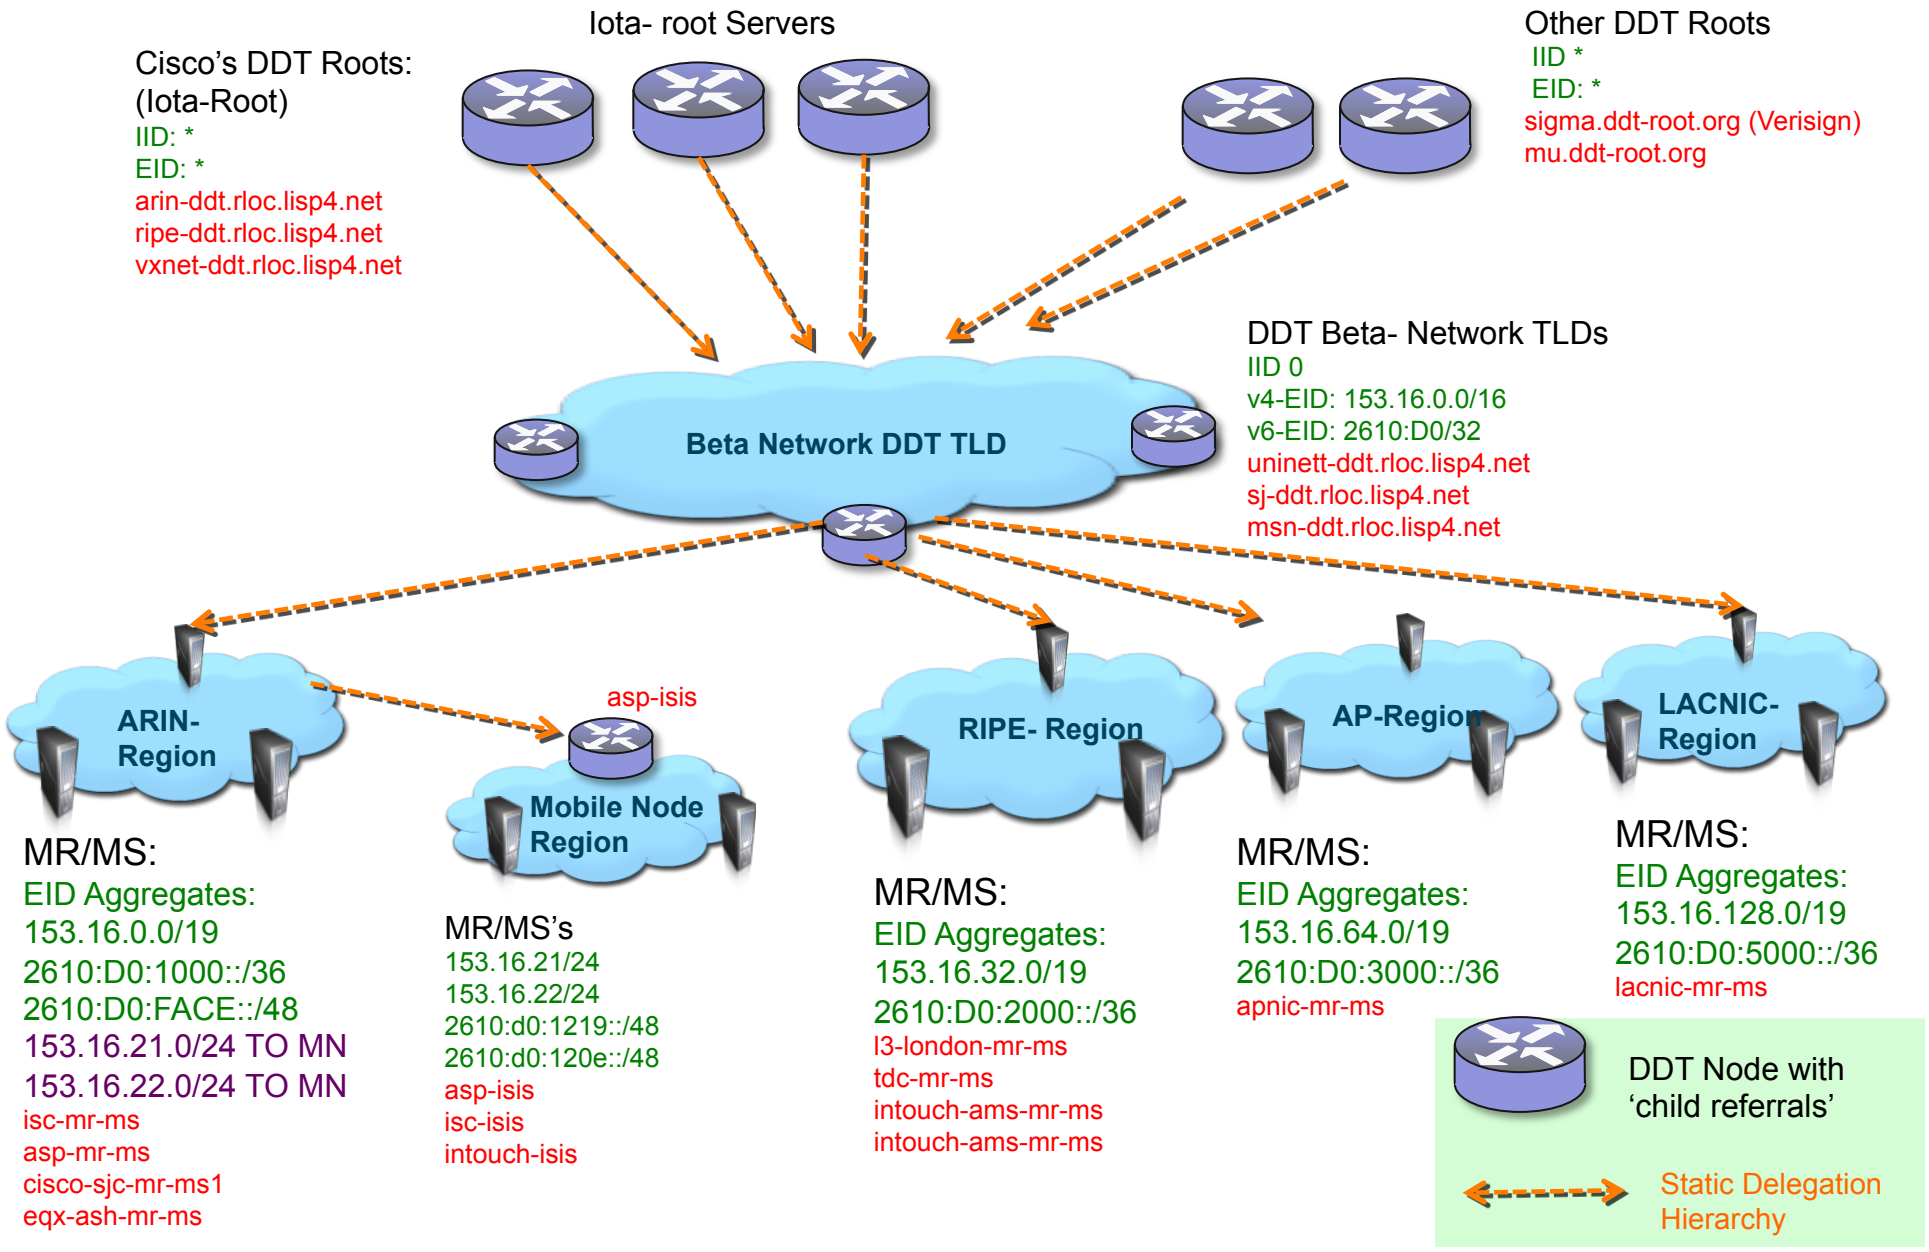

LISP Development at Verisign

- Verisign has operationalized a global public LISP mapping database system
- Support OpenLISP development
- Multiple phases for Verisign LISP rollout
  - Pilot program to provide a secure and reliable EID to RLOC Mapping Service (LISP-DDT currently)
  - Currently testing Customer Portal used for provisioning and managing external customers
  - In-house developed LISP service using the same infrastructure currently in place to support COM/NET DNS (Q4 2012)

LISP Footprint at Verisign

- LISP Pilot Network
  - Redundant Servers & Resolvers at each site
  - Redundant Transit from each site
  - Supports both IPv4 and IPv6
- Vendor Diverse Software Deployed
  - Dell R710s running NX-OS
  - Dell R710s running OpenLISP (June ’12)
  - Dell R710s running Verisign Mapping Servers/Resolvers (Q4 ‘12)
- Root Servers Deployed at multiple Datacenters
  - Dell R710s running NX-OS
  - Root IPs:
    - 72.13.36.141
    - 69.58.178.141
  - Efforts within our Labs, Engineering organizations in cooperation with Cisco for some time now

LISP on Verisign Infrastructure

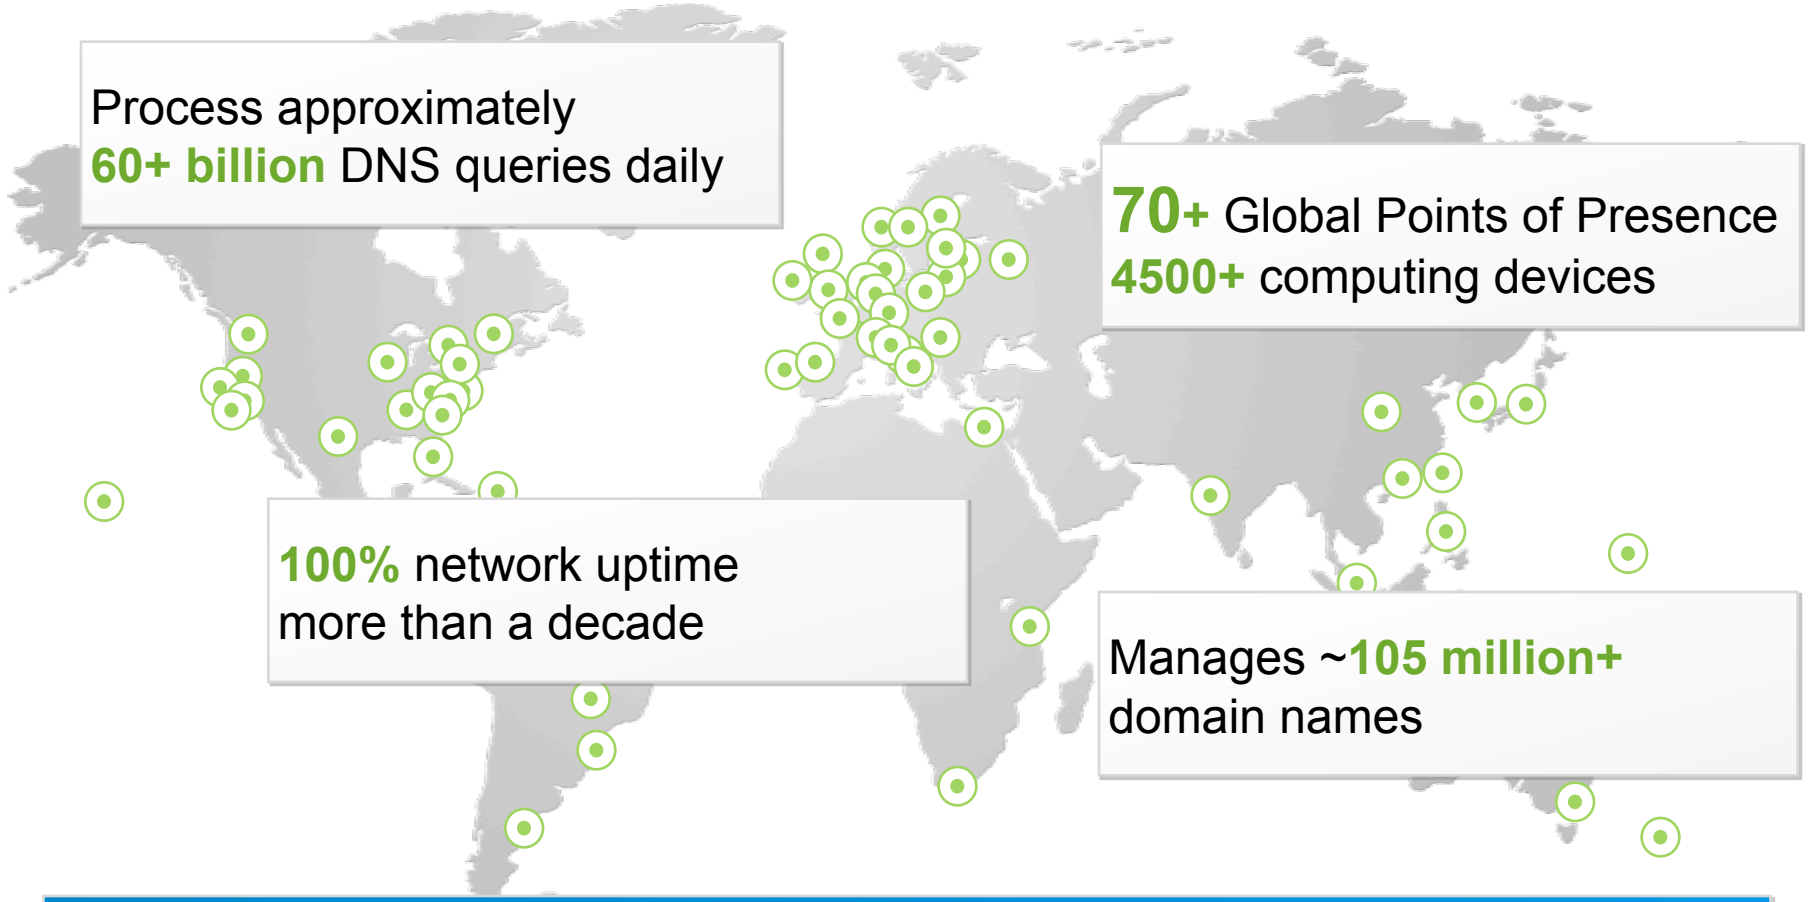

The World Trusts Verisign to Operate the Internet's Critical Infrastructure

Specification & Standardization

- Individual Submission in IETF LISP WG
  - draft-fuller-lisp-ddt-01.txt
- Work In Progress, -02 draft will be out soon
  - better integration of action code descriptions
  - additional work needed on security extensions
  - finite state machine for pending request processing
- Consistent With WG Charter
  - planned adoption as WG draft
  - expected replacement for LISP+ALT

Future Work

- LISP Mapping Provider “eco-system”(?)
- Internet-Scale Deployment(?)

LISP Resources

- www.lisp4.net
  - background information, pointers to other presentations
  - pilot network topology, traffic, etc.
  - LISP Network Operators Group (LNOG)
- lisp.cisco.com
  - Cisco implementation info, image downloads, etc.
- lisp@ietf.org - IETF LISP working group
  - https://datatracker.ietf.org/wg/lisp
  - “core” documents scheduled for RFC publication “RSN”
- LISP-DDT route operation - http://ddt-root.org

Questions and Comments?

Contacts: lisp-support@cisco.com

lisp@verisign.com

lisp-interest@cisco.com

(especially if you're interested in participating ☺)