Locator/Id Separation Protocol

CCS Module 2

Factors

- Among some driving factors of today’s Internet are:
  - the widespread availability of wireless (including Wi-Fi and cellular networks) connectivity allowing more non-PC devices perform ad hoc connections;
  - deployment of virtualization increasing the number of logical computing systems;
  - more cloud computing and peer-to-peer applications changing traffic characteristics towards less deterministic and stochastic models of CDNs;
  - reaching the Zettabyte era more quickly due to the overall increase in broadband speeds.
- Issues below are only consequences of Internet usage, which are completely different comparing to Internet conventions and user base 30 years ago.

Problems: Decoupling ID and Loc

- A single IPv4 address space
- IP address serves multiple roles nowadays:
  1) Identification – Identifier is a bit string that is used during the communication's lifetime. It identifies communicating parties in a way that IP address verifies the source of packets;
  2) Localization – Locator is a bit string that specifies packet destination where it should be delivered. It locates the place on the Internet, where a device is attached. Routing protocols interpret IP address as a locator and build up routing tables based on the situation that routers route traffic towards a destination. The locator is also known as **Point of Attachment (PoA)**.

Problems: Multihoming

- Internet's stands multihoming for the situation when the customer is using two or more ISPs for transit services as it is defined in RFC 4116
- Wider definition of **multihoming** covers following use-cases:
  - multihoming of single host attached redundantly to one or more networks;
  - multihoming of single (LAN) network (containing a set of hosts) interconnected redundantly with one or more networks;
  - multihoming of autonomous systems (containing a set of networks) interconnected redundantly with one or more ISPs;

Problems: Multihoming

- The trouble with multihoming is closely connected with IP address semantics described in the previous section – IP addresses is a PoA which is route dependent (i.e., reachability of multihomed networks depends on the chosen/available route).
- However, IP routing should be route independent, but this cannot be satisfied when it takes into account destination and next-hop IP addresses which are route-dependent PoAs.

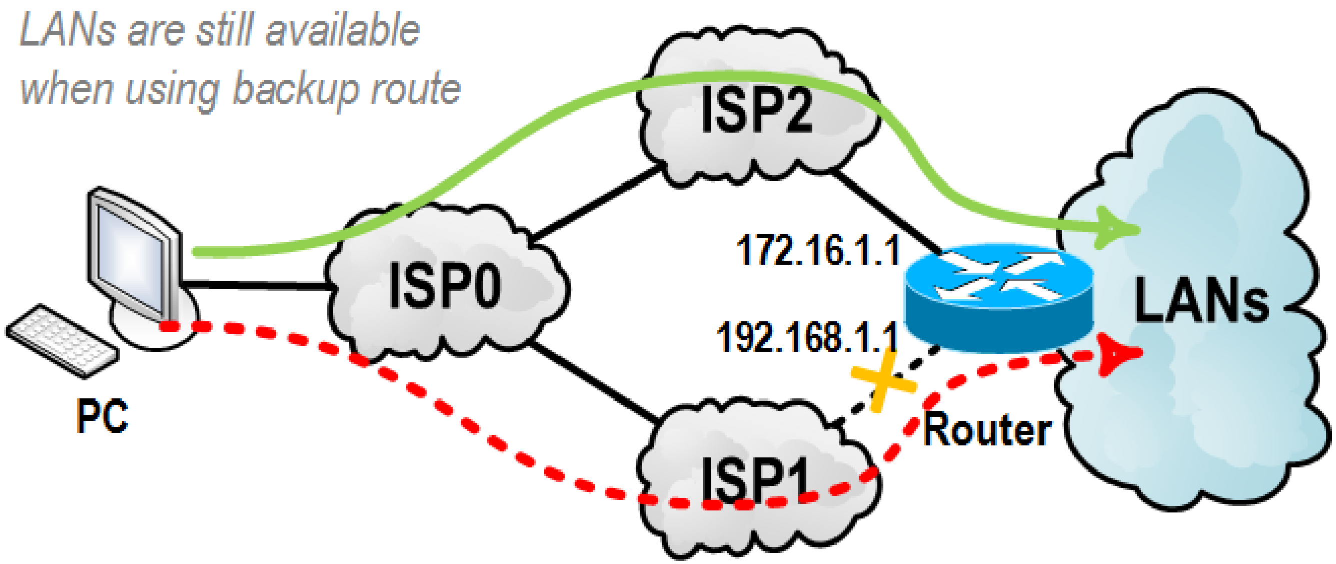

Problems: Mobility

- **Mobility** is the ability of a node or whole network to change its topological connectivity without disruption of ongoing communication
- Solutions like MobileIP, HMIPv6, MP-TCP include:
  - Dynamic renumbering of mobile entity
  - Renumbering and creating a tunnel between old and new location
  - The ability of a mobile entity to actively announce its new location

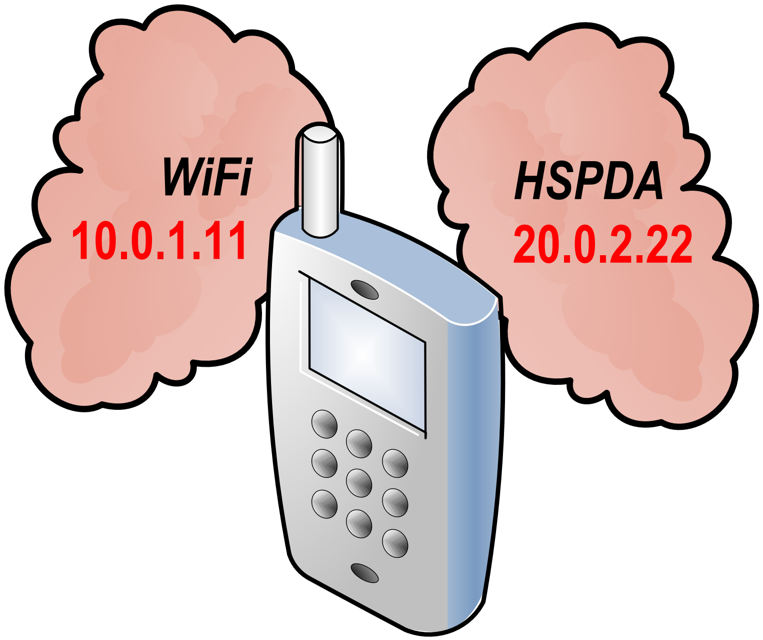

What is unique address???

- A looming current problem (for not just IoT) is how to accommodate possibly billions of devices with the IPv4/IPv6 capability to access the Internet and to provide session survivability when those devices roam.
- NAT is often being used to overcome this limitation by rewriting persistent address to dynamic mobile address.

Problems: Routing Scalability

- The most affected nodes struggling with the situation are Default Free Zone routers.
  - **Default Free Zone (DFZ)**: Backbone of the Internet where routers must keep complete routing tables with all reachable destination networks. In opposite of this are Tier 3 ISP or networks or end customers that are using usually only partial routing information – they have complete knowledge about local connectivity and any other network beyond is available via default route.
- Every year the size of Routing Information Base (RIB) and Forwarding Information Base (FIB) of those routers increases.
  - **Routing Information Base (RIB)**: Basically abstract data structure holding information from a given routing source that holds information about all reachable destination networks and paths to those destinations.
  - **Forwarding Information Base (FIB)**: The FIB is optimized version of RIB. It is consulted most of the time when forwarding packets because it is supported by specialized HW.
- The rate, at which prefix count is growing in the RIB, is the object of discussions but it seems to be slightly faster than linear (sometimes called superlinear) for a couple of last years

Problems: Routing Scalability – FIB

- IPv4

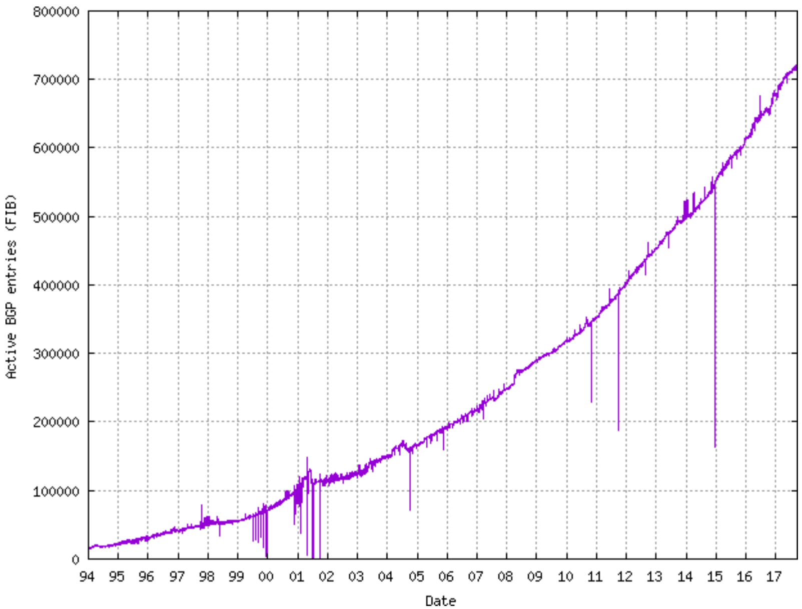

FIB / RIB Table Reports (plots)	Data Sets(txt)
Active BGP entries (FIB)	721791
All BGP entries (RIB)	31967059
RIB/FIB ratio (31967059/721791)	44.2885

- IPv6

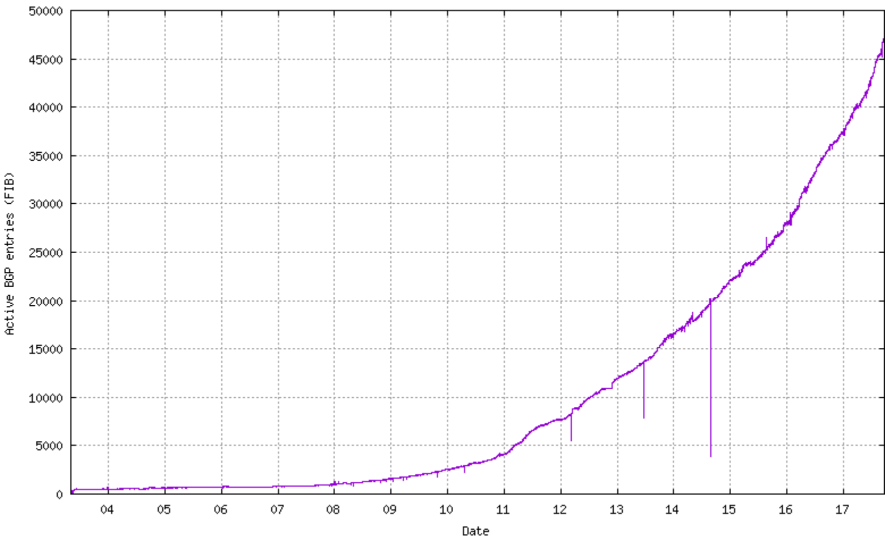

FIB / RIB Table Reports (plots)	Data Sets(txt)
Active BGP entries (FIB)	46693
All BGP entries (RIB)	46694
RIB/FIB ratio (46694/46693)	1.0000

Problems: Routing Scalability – ASN

- IPv4

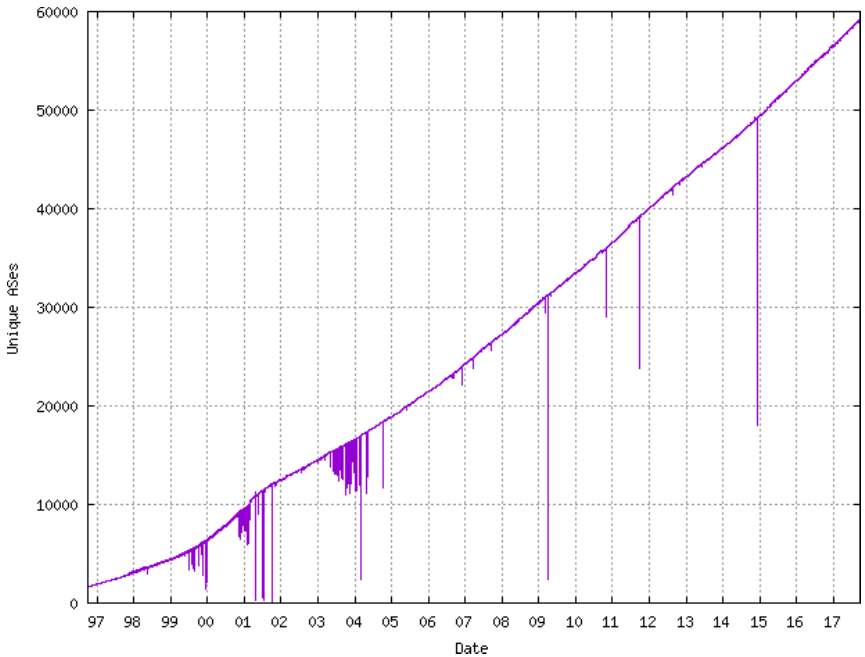

AS Reports (plots)	Data Sets(txt)
Unique ASes	59161
Origin only ASes	49750
Transit only ASes	304
Mixed ASes	9107
Multi-Origin Prefixes	8122
ASes originating a single prefix	22021
Average entries per origin AS	12.2635

- IPv6

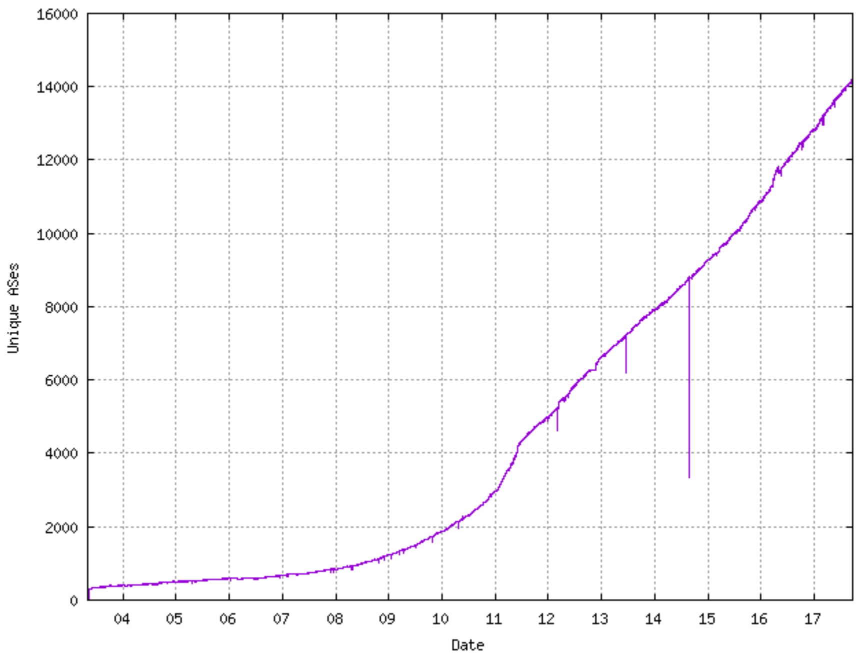

AS Reports (plots)	Data Sets(txt)
Unique ASes	14171
Origin only ASes	11227
Transit only ASes	201
Mixed ASes	2743
Multi-Origin Prefixes	726
ASes originating a single prefix	9618
Average entries per origin AS	3.3424

Approaches

- RFC 6115 clearly states that:
  a) RRG has rough consensus on separating identity and location of devices but does not have consensus how to do it properly;
  b) RRG has consensus that multihoming and traffic engineering issues need to be solved in a scalable manner.
- There are three ways how to decouple identity and locality:
  - **Map-and-encap network-based architecture** – It evolves from Robert Hinden's ENCAPS protocol. When a source sends the packet towards destination outside of source network, the packet must traverse through border router between two address spaces (locator space and identifier space). Here at first border router performs mapping of an identifier to appropriate locator ("map" phase). Then the packet is encapsulated using returned locator address ("encap" phase).
  - **Rewriting hybrid network-based architecture** – Originally this principle comes from papers written by Robert Smart and David Clark 8+8 and later by Mike O'Dell GSE. It utilizes IPv6 field so that upper part of IPv6 address is locator and the lower part stores an identifier. If a source sends packet outside its domain, border router takes addresses containing only identifiers and fills upper bits with appropriate locators.
  - **Host-based architecture** – Decisions in this architecture are purely in the hands of hosts. Thus, hosts prepare and fill all relevant PCI fields (including locators and identifiers) as the packet is being dispatched by the operating system.

Candidates

Name	type	CE	IPv	RS	DIL	MH	Mob	TE	Ren	Dep
LISP	M	CES	v4v6	yes	yes	yes	yes	yes	yes	yes
HIP	H	CEE	v6	yes	yes	yes	yes	no	yes	no
SHIM6	H	CEE	v6	no	yes	yes	no	no	no	yes
RANGI	H	CEE	v6	yes	yes	yes	yes	yes	yes	yes
Ivip	M	CES	v4v6	yes	yes	yes	yes	yes	yes	yes
hIPv4	diff	diff	v4	yes	yes	cond	cond	cond	yes	no
NOL	R	diff	v4v6	yes	yes	yes	yes	yes	no	no
GLI-Split	R	CEE	v6	yes	yes	yes	yes	yes	yes	yes
TIDR	M	CES	v4v6	no	yes	yes	no	yes	yes	yes
ILNP	R	CEE	v6	yes	yes	yes	yes	yes	yes	yes
Evolution	diff	diff	v4v6	yes	no	no	no	no	no	n/a
NBS	diff	CEE	v4v6	yes	yes	cond	cond	cond	no	no
APT	M	CES	v4v6	yes	yes	yes	yes	yes	yes	yes
IRON-RANGER	M	CES	v4v6	yes	yes	yes	yes	yes	yes	yes
TRRP	M	CES	v4v6	yes	no	yes	no	yes	no	yes
Six/One	R	CES	v6	yes	yes	yes	no	no	yes	yes
RINA	diff	diff	v4v6	yes	yes	yes	yes	yes	yes	yes

Core-Edge Separation

- Locator/Identifier split is commonly performed
- Edge networks are separated from DFZ routing tables or are at least highly aggregated. Routing scalability is visible in direct proportion to how widely is CES solution adopted;
- CES benefits are available immediately to adopters – multihoming, inbound TE and if possible also mobility;
- Deployment of CES does not affect DFZ routers, but new devices on the border between core and edge are needed to interconnect this two address spaces together with mapping system;
- CES solutions do not require host stack, API or application changes;
- Tunneling and overlaying impose additional size overhead on fragments, thus introducing MTU concerns when employing CES.

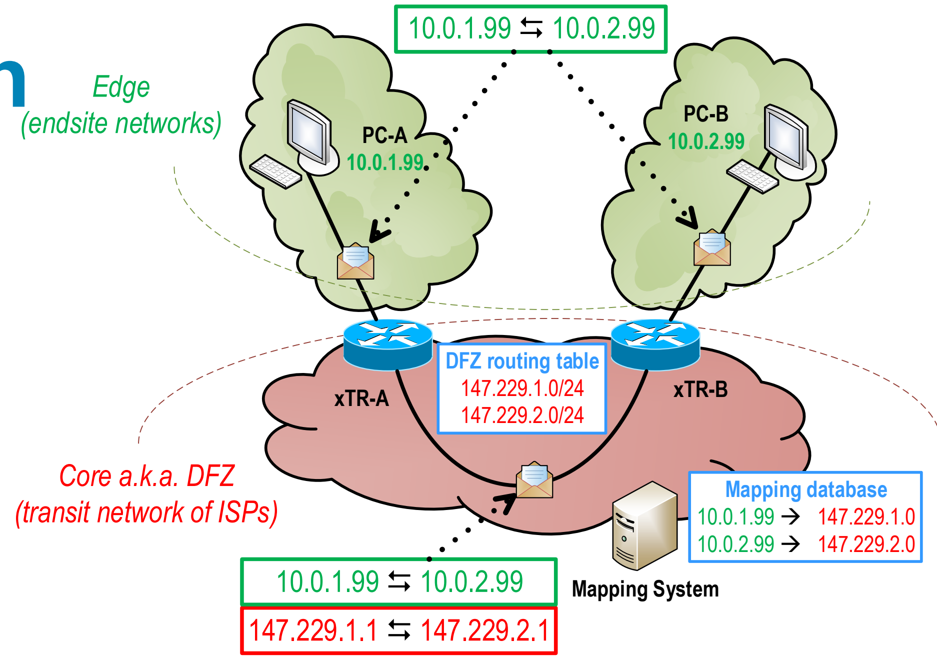

LISP, APT, Ivip, IRON-RANGER, TRRP:

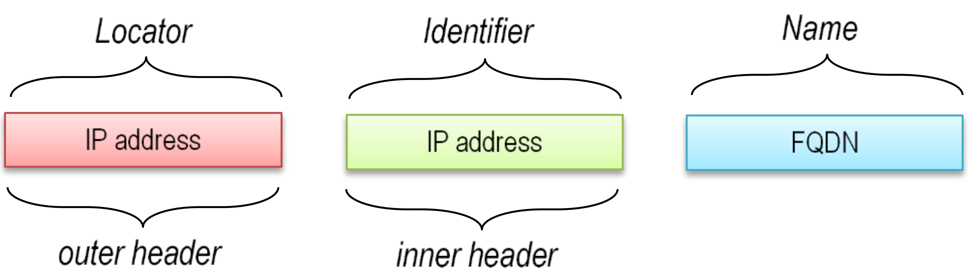

Six/One:

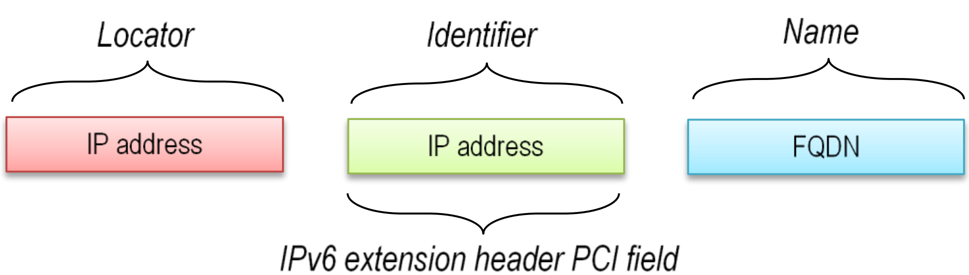

Core-Edge Elimination

- The most of CEE solutions separates locators and identifiers into two completely different namespaces.
- CEE benefits are visible and widely available to adopters only after majority of network migrate.
- Routing scalability is attained in a way that applications are no longer dependent on stable PI (or de-aggregated PA) addresses. Hence, PA addresses could be easily preferred and administratively more available than PI addresses.
- CEE host stack must determine which locator should use. Besides that, potential set of locat could be retrieved, thus implying resolving multihoming, inbound TE issues, and ideally mobility issues.
- DFZ routers are not affected, and no additional tunneling devices are needed, however, a new infrastructure (or at least upgrade of current one, i.e. DNS) must be present to provide mapping between identifiers and locators.
- CEE solutions need host stack changes and applications augmentations.
- The most of CEE solutions do not support IPv4 and have some troubles with NAT so additionally clutches are needed.

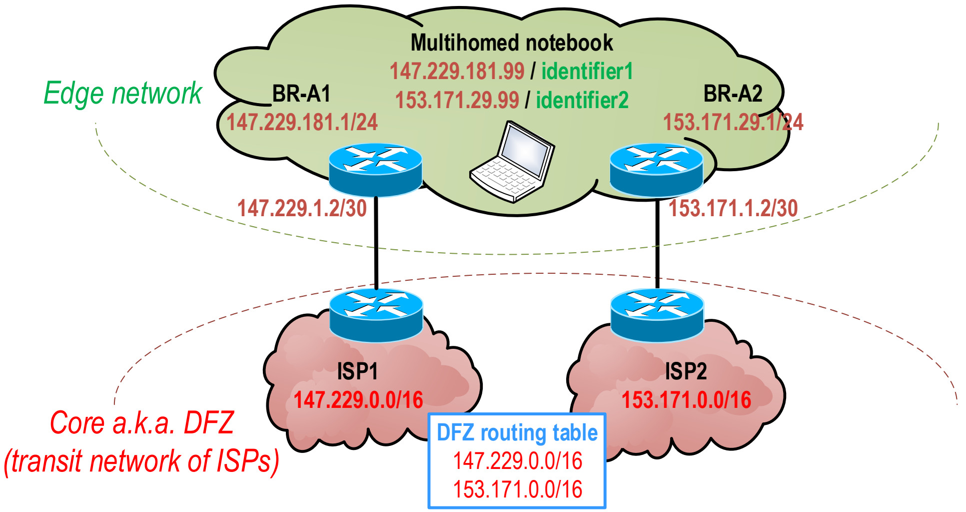

HIP, Shim6, RANGI:

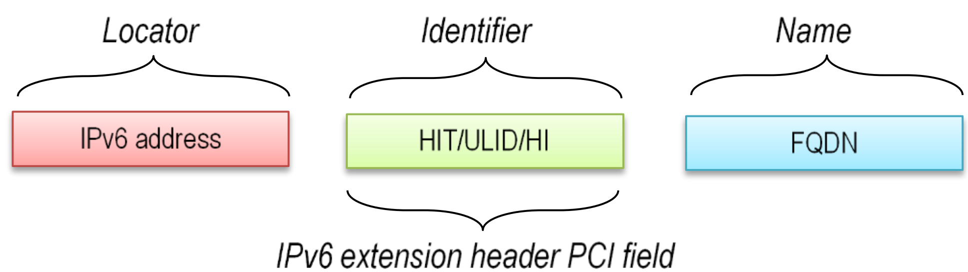

ILNP, GLI-Split:

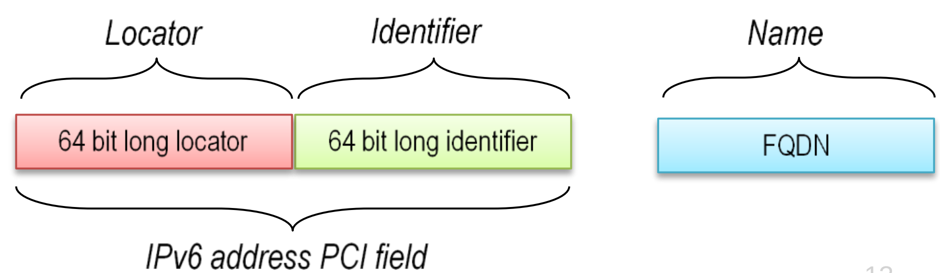

Theory

Basics

- The main idea behind LISP is to separate localization and identification
- LISP accomplishes this by splitting the IP address into two namespaces:
  - Routing Locator (RLOC) namespace where addresses fulfill their localization purposes by telling where is device connected to the network
  - Endpoint Identifier (EID) namespace where each device has a unique name that identifies it from each other
- Also a non-LISP namespace exists (and probably always will exist), where direct LISP communication is (even intentionally) not supported

Basics: Illustration

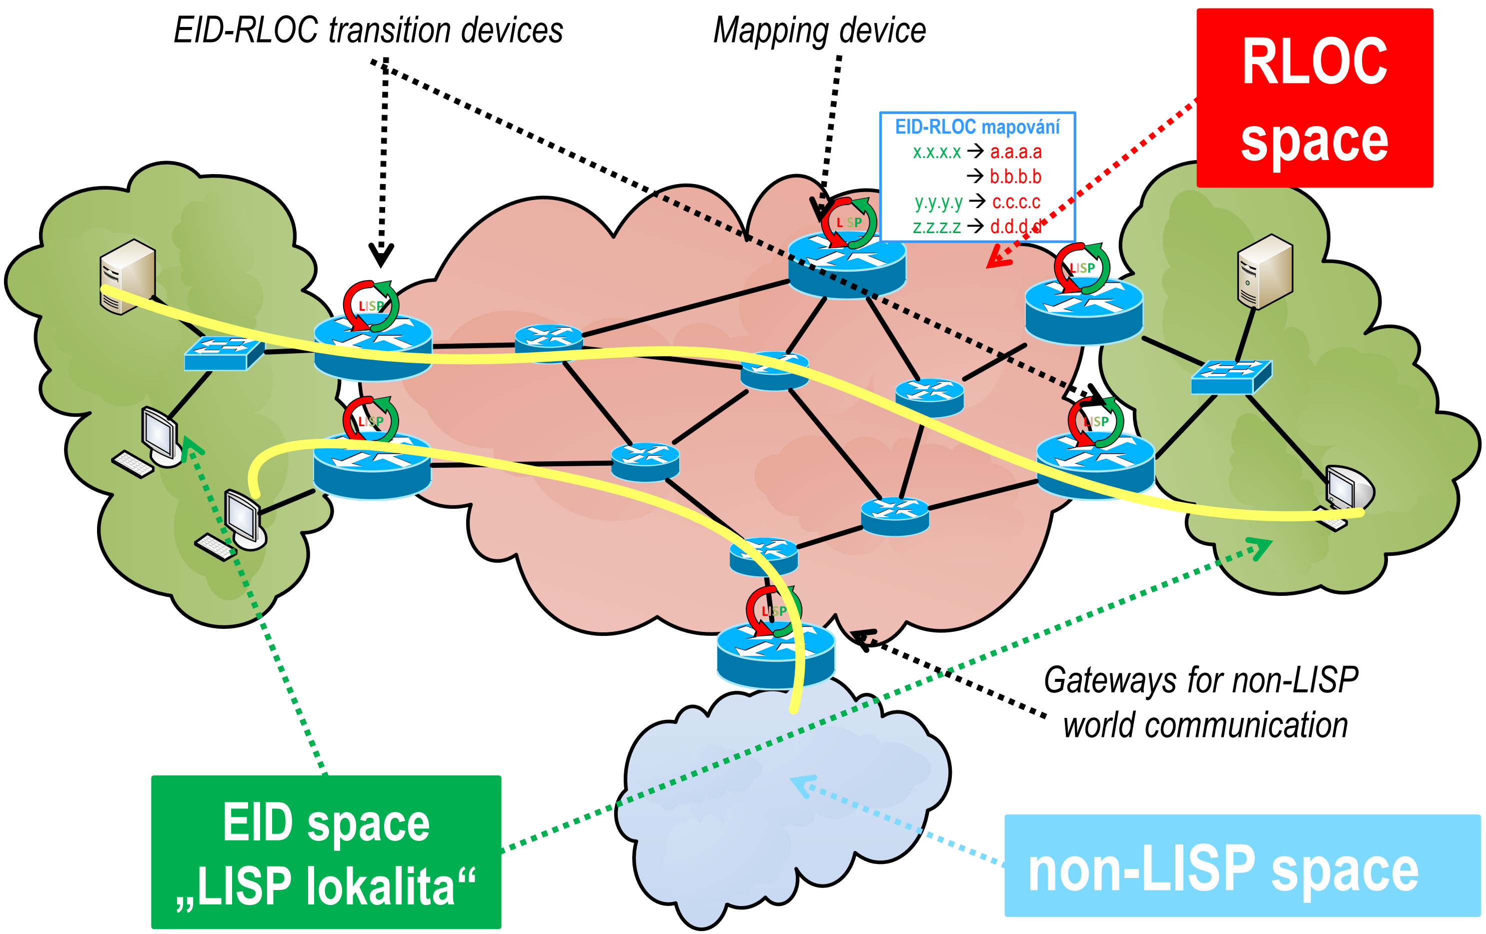

Advantages

- Transparent to end-host without any additional configuration
- Does not offload nor changes anything in DNS
- Independent on address-family
  - concept works with IPv4, IPv6 or any network protocol
- Mobility and inbound traffic engineering is supported by design
- Built-in transition mechanisms
  - protocol spec introduce gateway devices

Encapsulation

- LISP header(36
- UDP header
  - Destination **p**
  - Destination **p**
- Outer IP heade

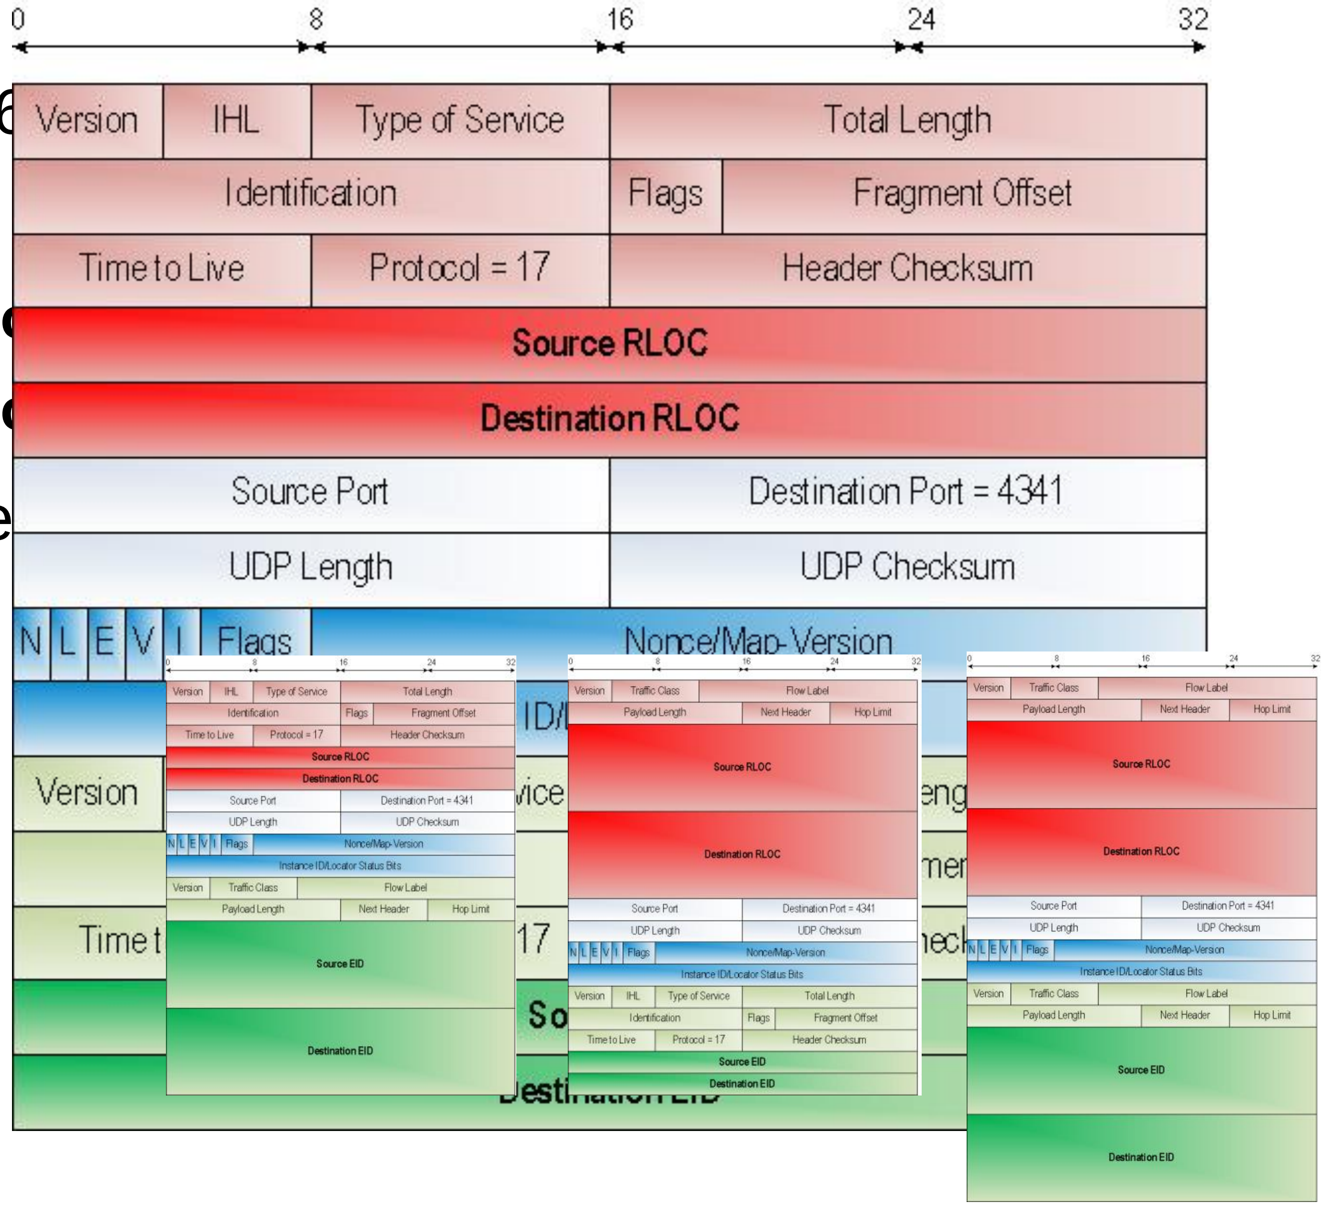

Components: ITR a ETR

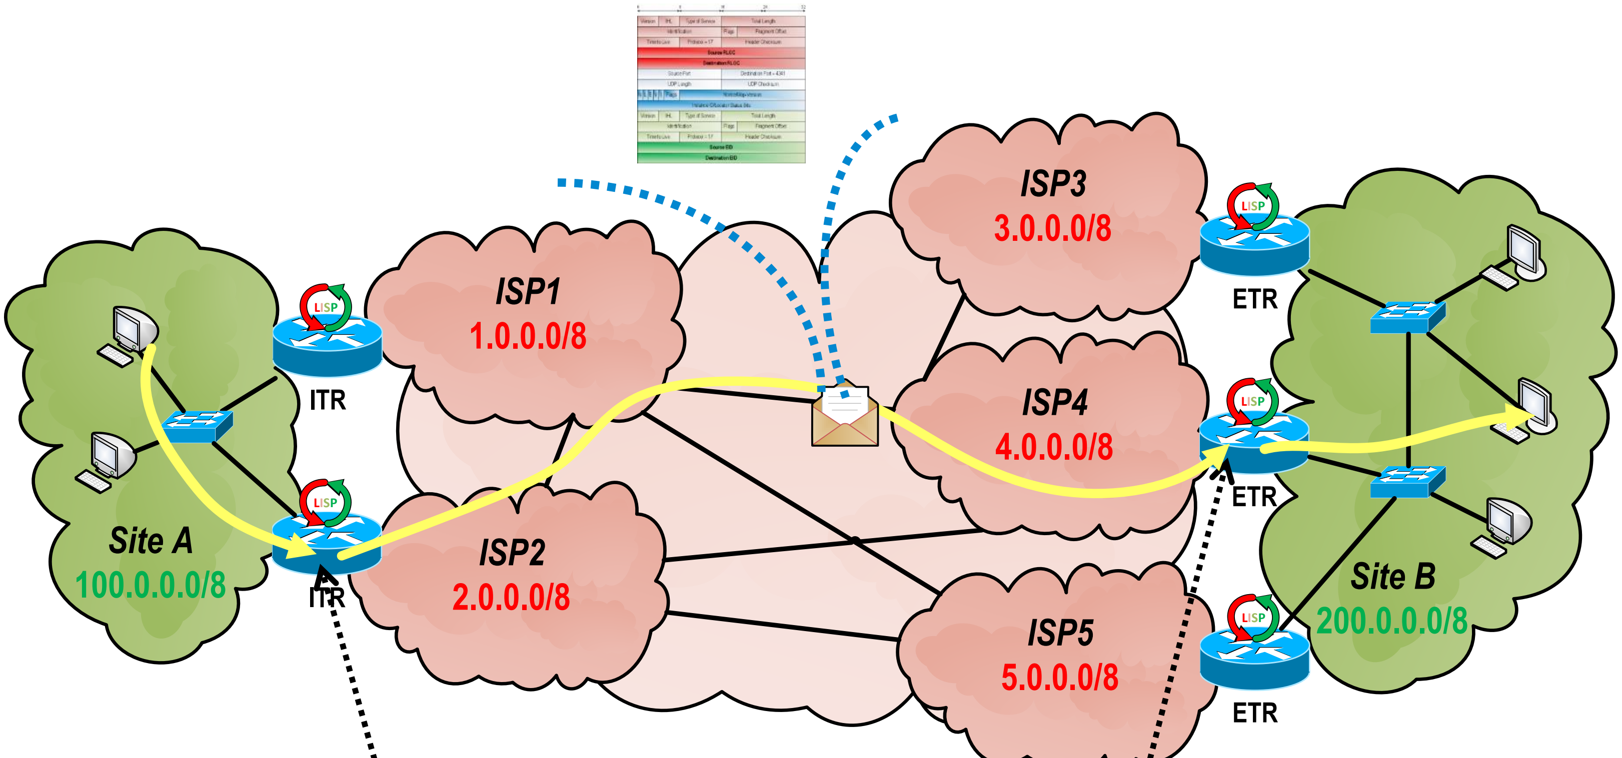

Ingress Tunnel Router (ITR)

- Entry point to RLOC namespace
- Encapsulates outgoing traffic
- Maintains mapping cache

Egress Tunnel Router (ETR)

- Exit point from RLOC namespace
- Decapsulates incoming traffic and relays it to EID hosts

Example: Unicast Routing

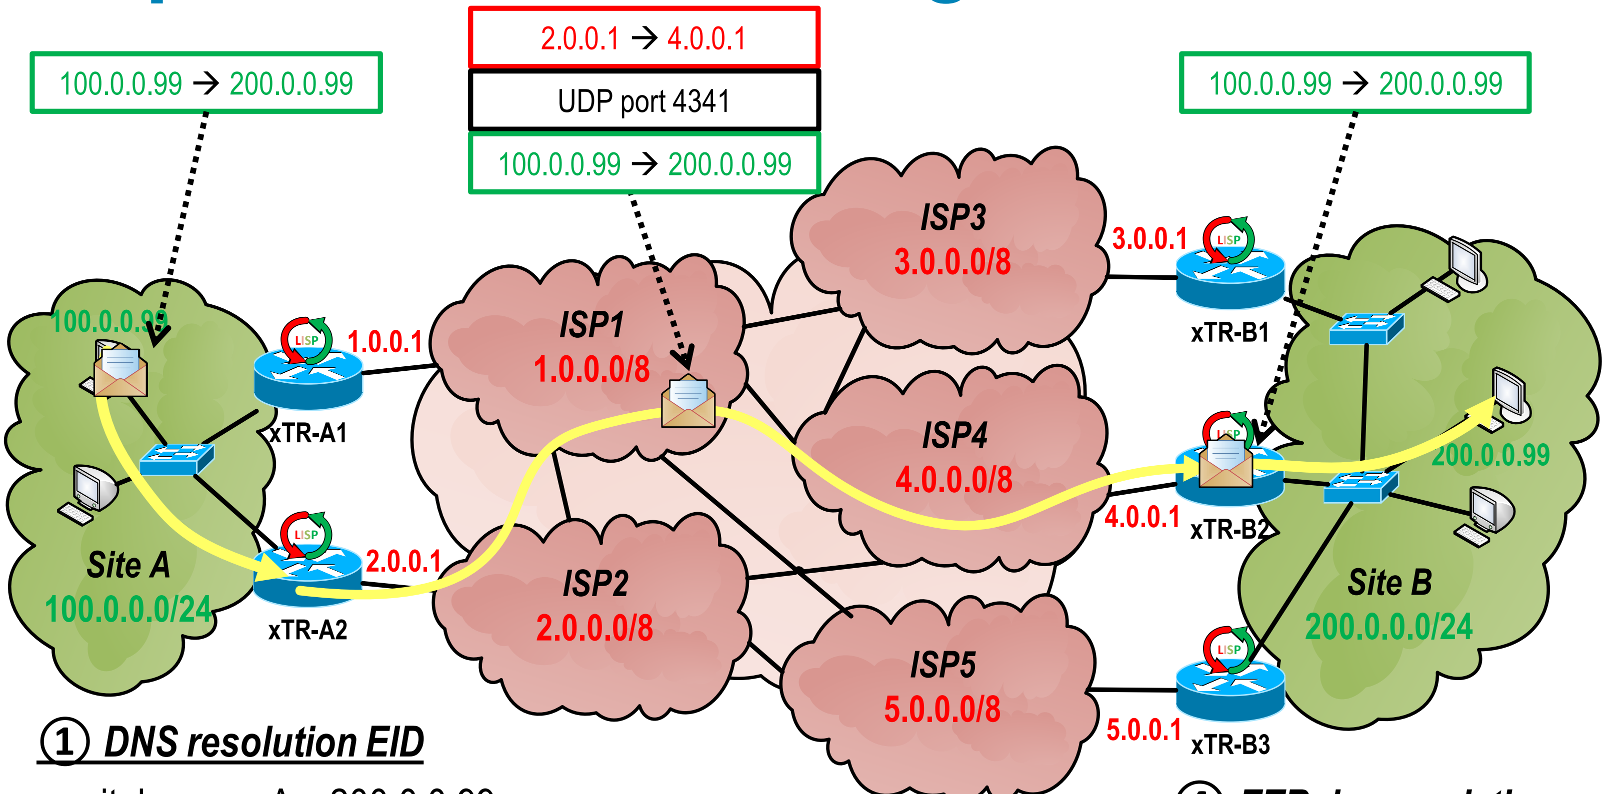

① ***DNS resolution EID***

pc.siteb.com A 200.0.0.99

④ ***ETR decapsulation***

② ***Map-cache lookup***

EID-prefix: 200.0.0.0/24

RLOC:

3.0.0.1, priority 254, weight 50 = xTR-B1

4.0.0.1, priority 1, weight 100 = xTR-B2

5.0.0.1, priority 254, weight 50 = xTR-B3

③ ***ITR encapsulation***

RLOC xTR-B2 chosen as destination due to the lowest priority

Mapping

- *How to obtain RLOC for a given EID?*
- Analogy to DNS

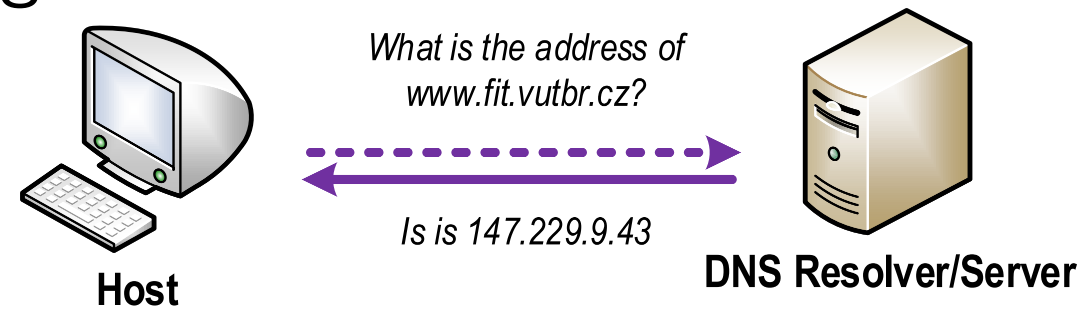

- In case of LISP

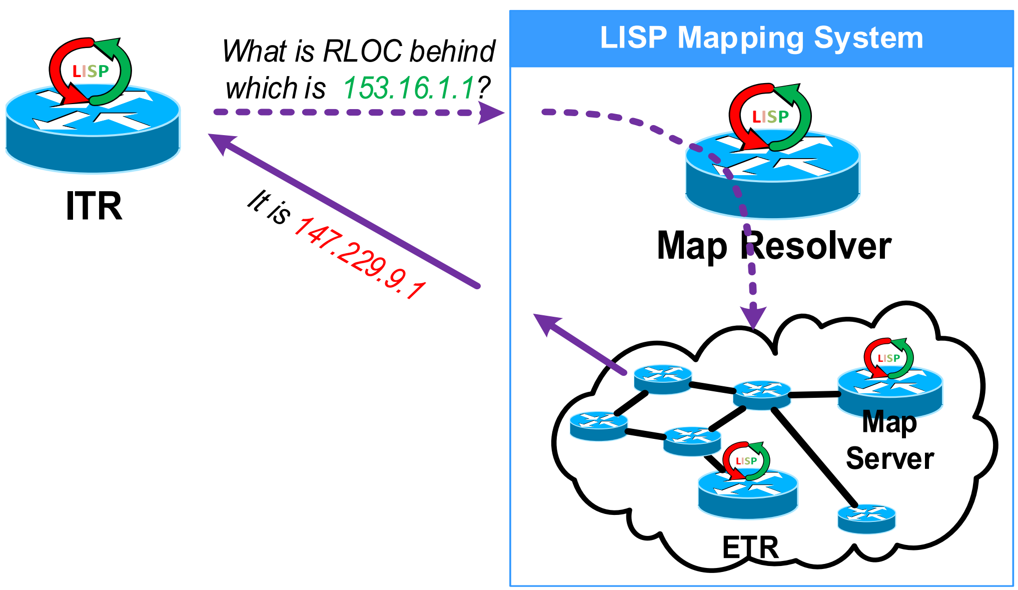

Messages

- **LISP Map-Register**
  - Each ETR announces as authority one or more LISP site(s) to the MS with this message
  - Each registration contains authentication data and the list of mappings and their properties
  - *Notifies others about EID space state*
- **LISP Map-Notify**
  - UDP cannot guarantee message delivery
  - MS may optionally (when the particular bit is set) confirm reception of *LISP Map-Register* with this message
- **LISP Map-Request**
  - ITR generates this request whenever it needs to discover current EID-to-RLOC mapping and sends it preconfigured MR
  - *Asking for a mapping*
- **LISP Map-Reply**
  - This is solicited a response from the mapping system to a previous request and contains all RLOCs to a certain EID together with their attributes
  - Each ITR has its map-cache where reply information is stored for a limited time and used locally to reduce signalization overhead of mapping system
  - Moreover, mapping system generates **LISP Negative Map-Reply** as a response whenever given identifier is not the EID, and thus proxy routing for external LISP communication must occur
  - *Answering the question regarding mapping*

Components: MS, MR and ALT

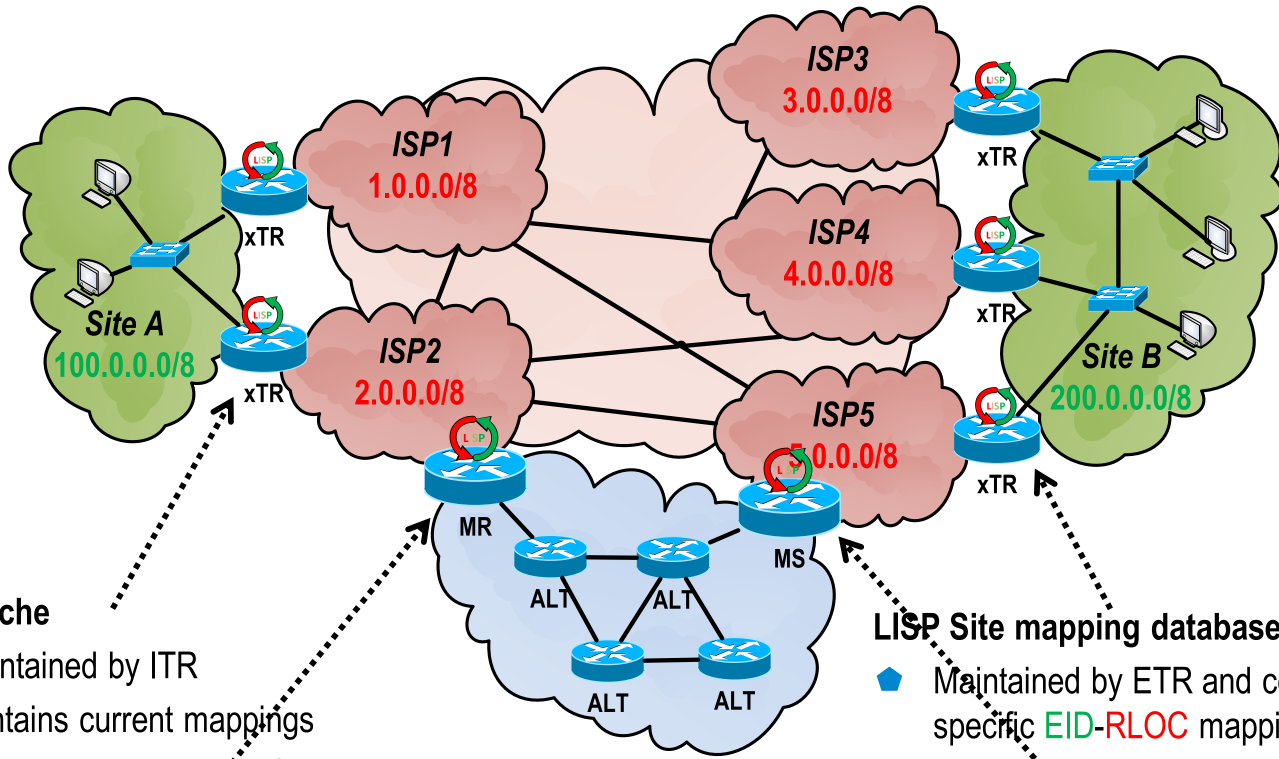

Map Cache

- Maintained by ITR
- Contains current mappings
- Updated by Map-Reply from MS or from ETR
- Routes according to Map-Reply content (namely priority, weight, TTL and RLOC state)

Map-Resolver (MR)

- Accepts ITR's Map-Requests
- Relays to ALT
- Responds with Negative Map-Reply in case of non-LISP namespace

Alternate Topology (ALT)

- Aggregates EID prefix and maintains allocation policy
- EID are announces through BGP inside GRE
- Relays Map-Requests

LISP Site mapping database

- Maintained by ETR and contains site-specific EID-RLOC mappings
- ETR responsible for EID namespace
- Sends Map-Reply to ITR
- Based on Map-Request, ETR changes access policies

Map-Server (MS)

- Consolidates EID-RLOC mappings from known ETRs
- May respond on behalf of ETR

Example: EID Registration

3.0.0.1 → 5.0.0.255

UDP port 4342

Map-Register
SHA-1 hash
local, 3.0.0.1, pri 254, wei 50
4.0.0.1, pri 1, wei 100
5.0.0.1, pri 254, wei 50

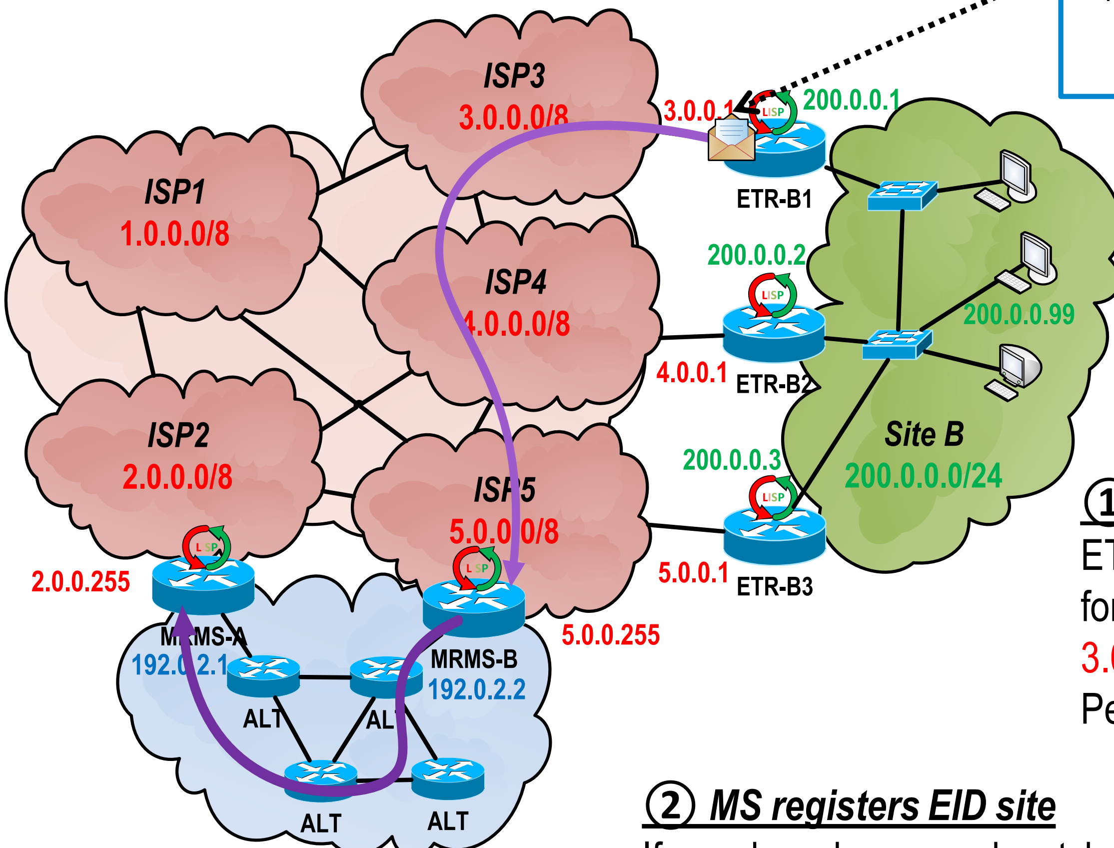

① *EID announcement*
ETR generates Map-Register for 200.0.0.0/24 with 3.0.0.1, 4.0.0.1, 5.0.0.1
Periodically sent every minute

② *MS registers EID site*
If pre-shared password match, then RLOCs to a given EID are updated
Timeout after 3 minutes

③ *MS propagates to ALT*
BGP in GRE
Processed by other MR

Example: Mapping Query

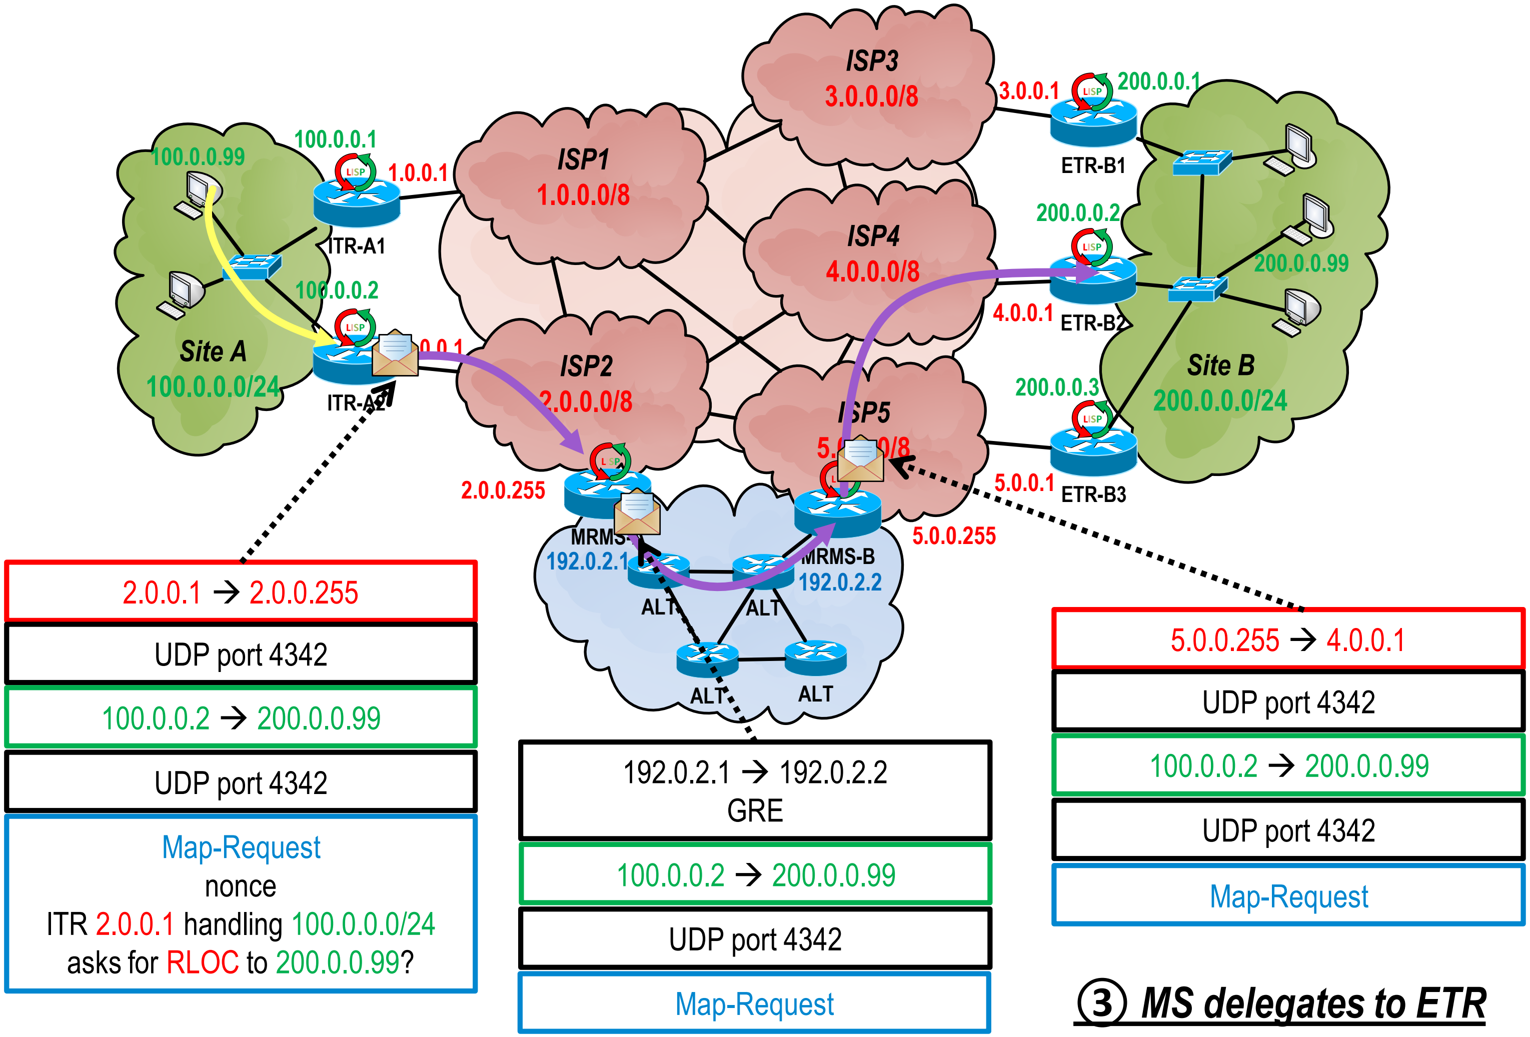

③ *MS delegates to ETR*

Sends to the last ETR, which announced EID-RLOC to mapping database for a given site

① *Map-cache miss*

Generate Map-Request

② *MR relays to ALT*

Encapsulates to GRE and sends through ALT

Rate-Limiting

- Map-Cache records are created on demand
- ITR routing performance depends on map-cache content
  - `IF` mapping exists `THEN` use the best RLOC

    `ELSE` initiate mapping query (map-cache miss)
- However, `WHILE` undergoin mapping query

  `DO` disc

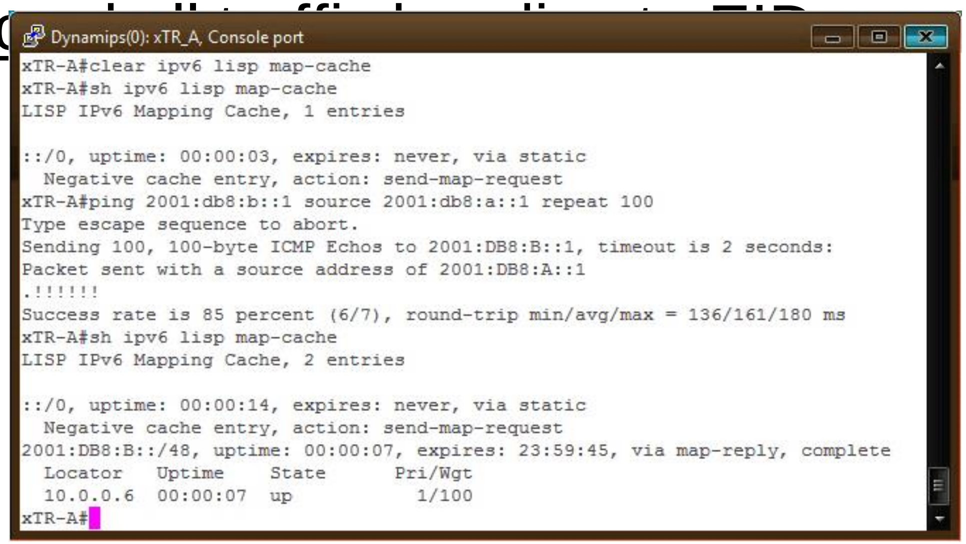

Example: Mapping Response

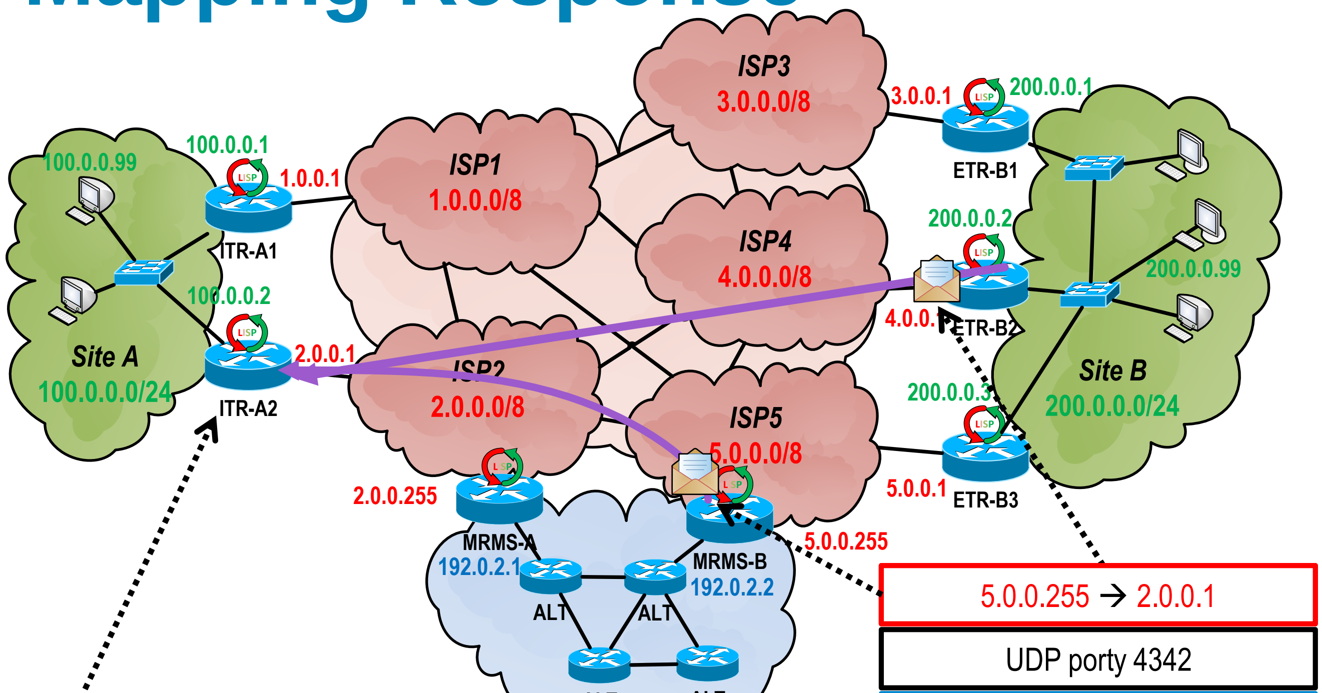

5.0.0.255 → 2.0.0.1

UDP porty 4342

Map-Reply
nonce
3.0.0.1, pri 254, váha 50
lokální, 4.0.0.1, pri 1, váha 100
5.0.0.1, pri 254, váha 50

② *Map-cache update*

EID-prefix: 200.0.0.0/24

RLOCs:

- 3.0.0.1, priority 254, weight 50
- local, 4.0.0.1, priority 1, weight 100
- 5.0.0.1, priority 254, weight 50

❶ *Alternative answer*

MS sends Map-Reply to ITR on behalf of ETR, if EID registration allows this (proxy-reply option)

① *Answering query*

ETR sends Map-Reply to ITR

Distribution of Mapping

- ETR registers itself only to a limited number of MSs.
- It is technically impossible for all ETRs to be registered to the same MS.
- Hence, there must be a way how to distribute mapping database and interconnect different MS between each other in order to guarantee the availability of mapping information to all MRs.
- Three approaches exist
  - LISP-ALT
  - LISP-DDT
  - LISP-DHT

LISP-ALT

- **Alternative Topology (LISP-ALT)**
  - MS are connected via dedicated GRE tunnels across the non-LISP world. LISP routing information are carried as external routes redistributed into BGP. LISP-ALT aggregates EID prefixes and enforces allocation policy. LISP-ALT is not a scalable solution when the number of MSs starts to increase.

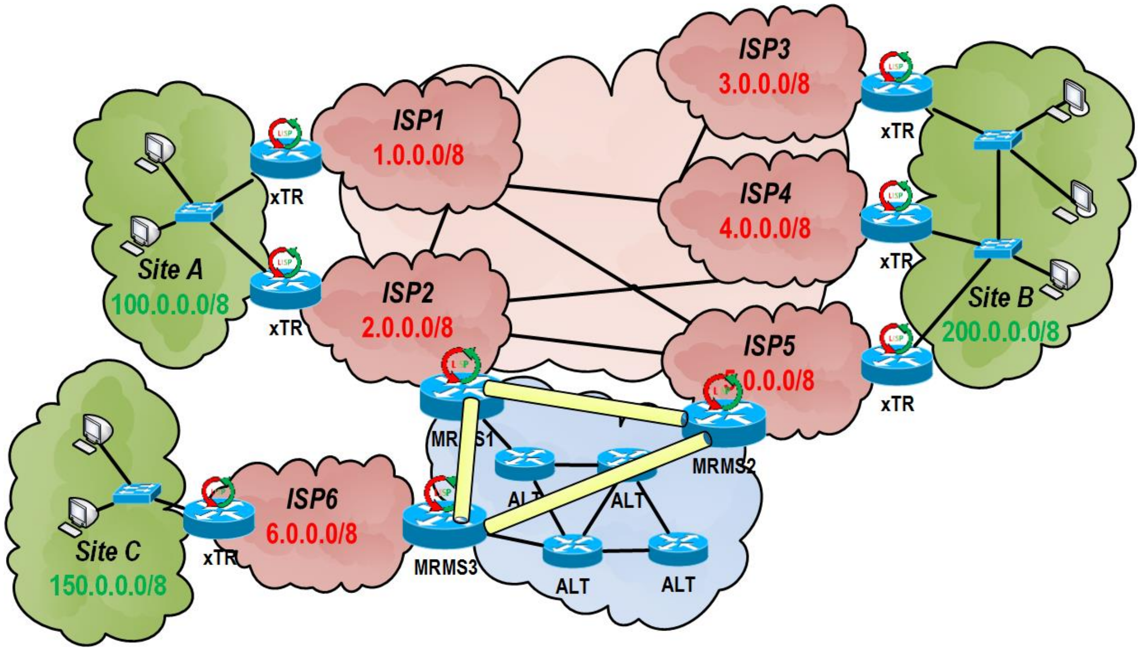

LISP-DDT

- **Delegated Distributed Tree (LISP-DDT)**
  - LISP-DDT is hierarchical distributed database, where each EID block is delegated to some authoritative organization.
  - The concept is similar to DNS with its hierarchy
  - Analogously, mapping request traverses from MR via tree towards the leaf, which is either designated MR, or ETR
  - Iterative query delegation between LISP-DDT nodes is accomplished by special LISP **Map-Referral** message

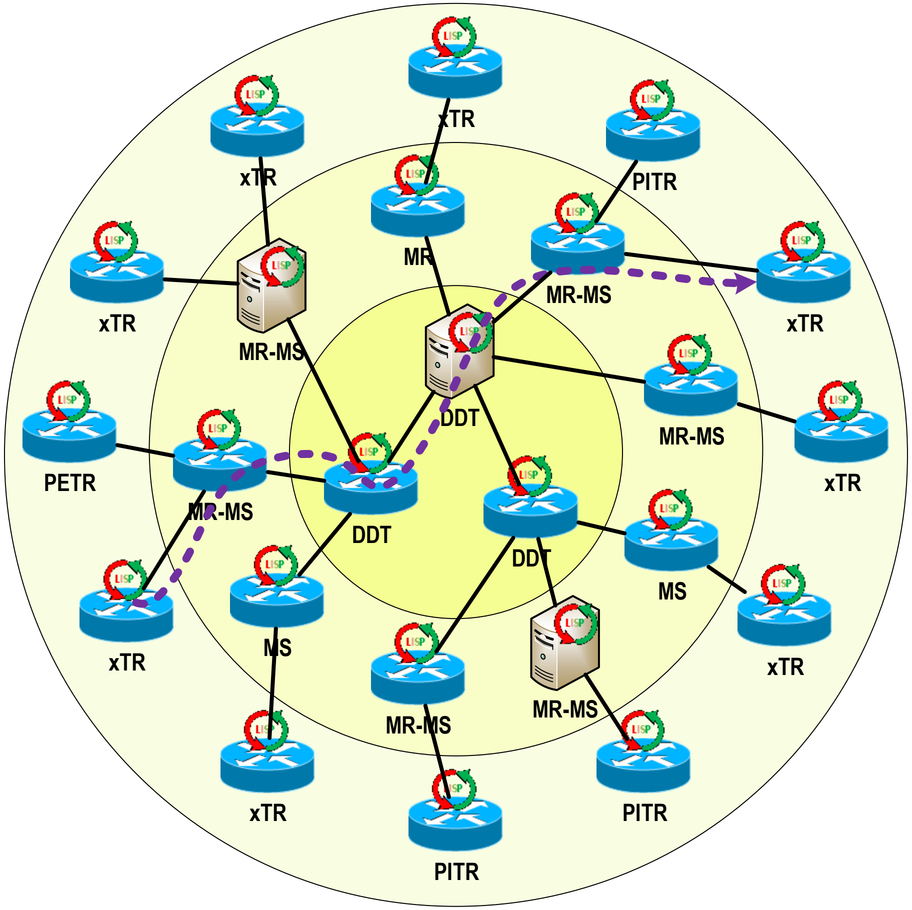

LISP-DHT

- **Distributed Hash Tables (LISP-DHT)**
  - LISP-DHT leverages DHT technology, namely Chord protocol and algorithm.
  - LISP mapping system forms ring-shaped overlay network, where ChordIDs are highest numerical EIDs instead of being randomly chosen.
  - Nodes are divided into two groups:
    - a) MSs as service nodes that are full-fledged DHT nodes;
    - b) xTRs as stealth nodes that can inject messages into DHT but neither do the route nor provide key management.
  - LISP-DHT allows a mapping request to be automatically forwarded to the owner without any previous specific advertisements.

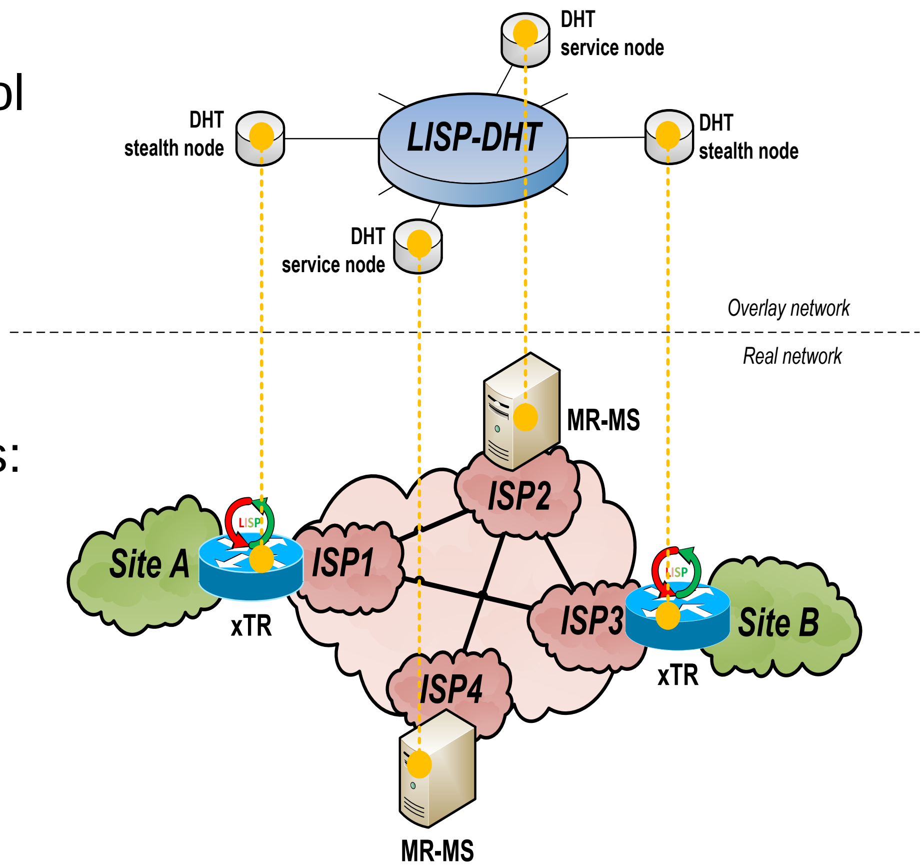

Transiting to LISP

- *Immediate switch over to LISP is impossible!*

- We need to guarantee communication:
  - non-LISP → LISP – Hosts and routers do not know anything about loc/id split. Hence, EIDs are considered as ordinary addresses and natively routed to “EID network entry point”
  - LISP → non-LISP – ITR must recognize that the destination address is not EID. Hence, there are no RLOCs associated with it. The packet is then delivered to “LISP world exit point”

- Transition mechanisms
  - **Proxy ITR** and **Proxy ETR**
  - **LISP-NAT**
    - LISP → non-LISP translates EID onto RLOC

Components: PITR and PETR

Proxy Egress Tunnel Router (PETR)

- Enables LISP sites to reach non-LISP wordl (when uRPF is satisfied)
- Allows cross address-families communication

Proxy Ingress Tunnel Router (PITR)

- Acts as ITR pro non-LISP site by encapsulating non-LISP traffic as a LISP traffic
- Announces aggregated EID prefixes to non-LISP world

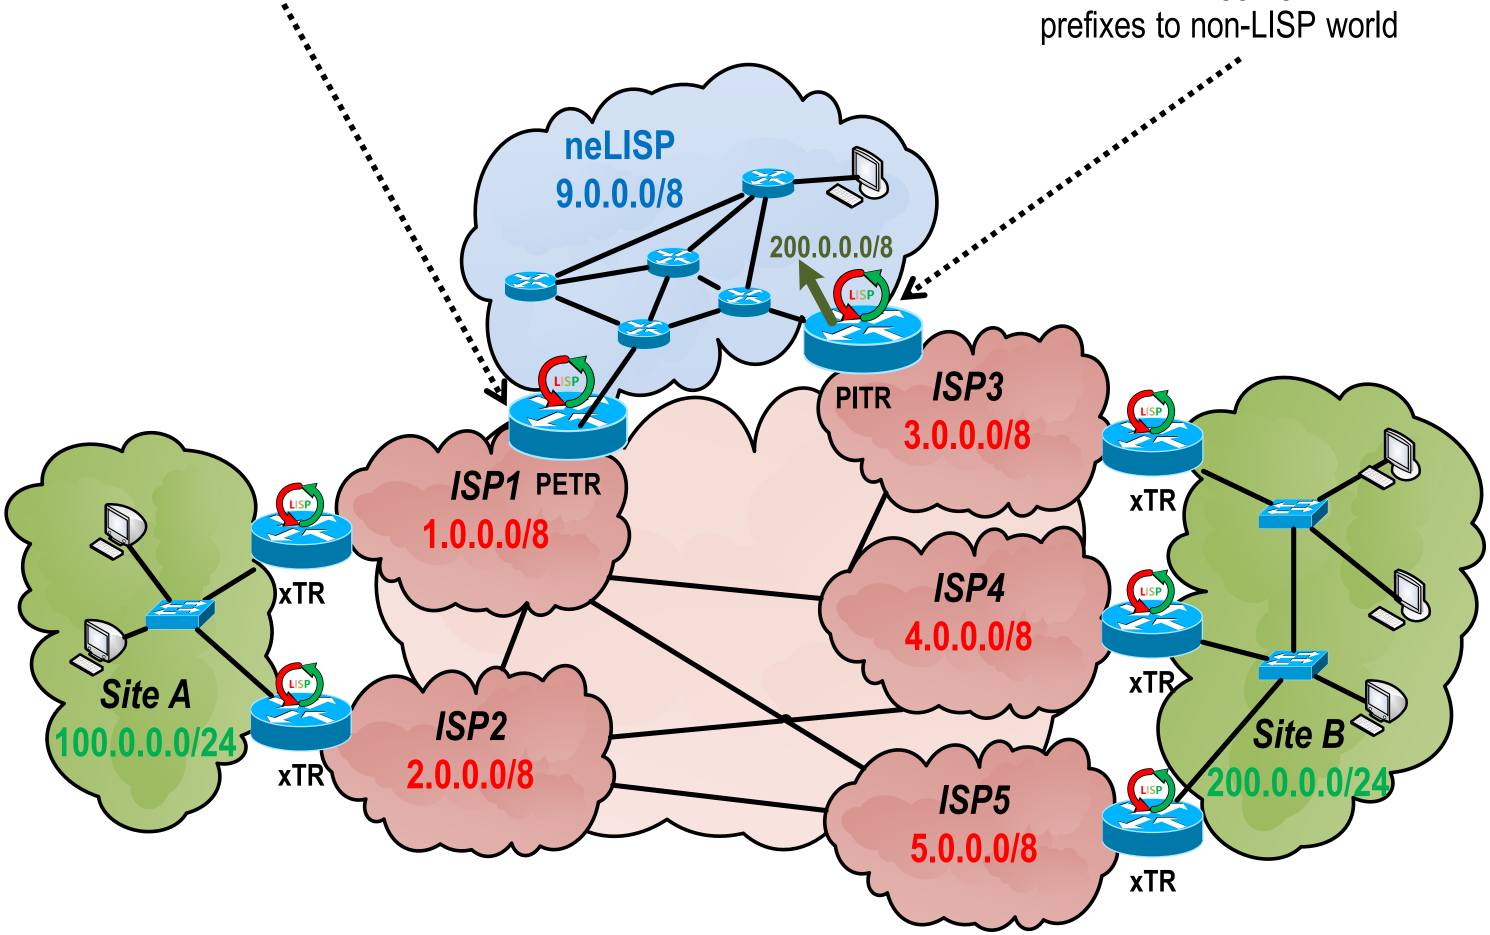

Example: Communication with non-LISP

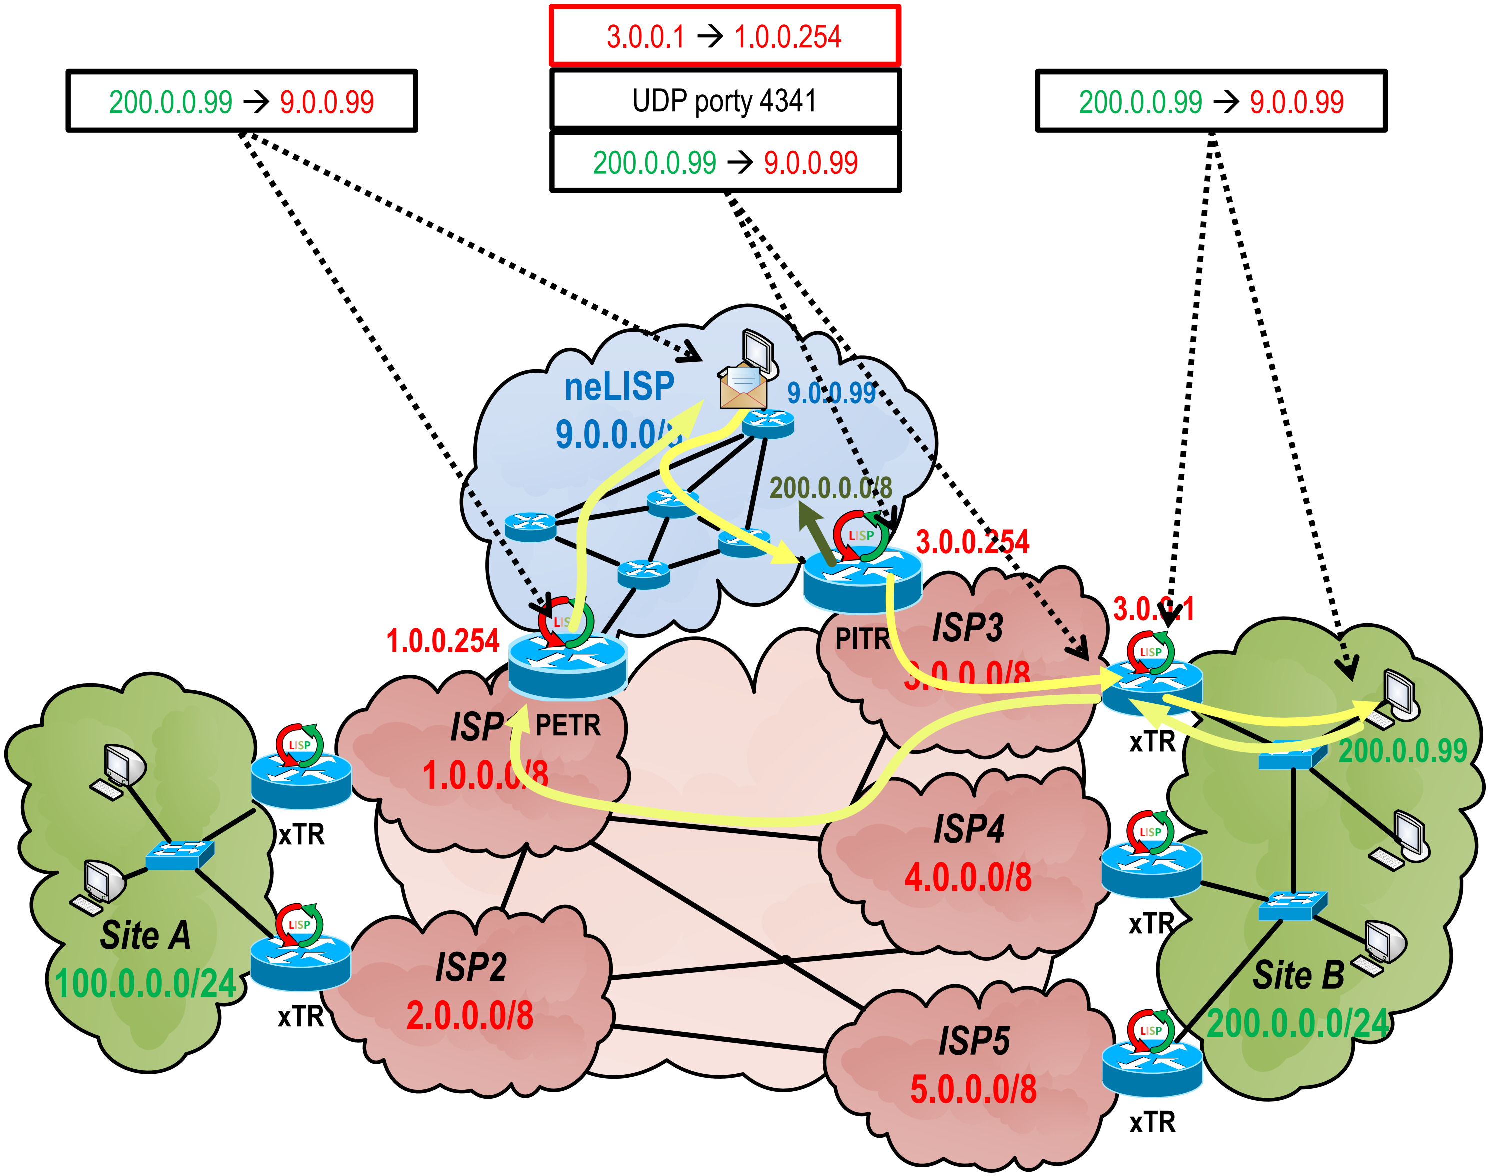

Configuration

Support

- Functionality of all components available since IOS 12.4

Serie	Platforms
ISR 1800 Series	1801, 1802, 1803, 1805, 1811, 1812, 1841, 1861
ISR 1900 Series	1941, 1941W
ISR 2800 Series	2801, 2811, 2821, 2851
ISR 2900 Series	2901, 2911, 2921, 2951
ISR 3800 Series	3825, 3845
ISR 3900 Series	3925, 3945
Cisco 7200 Series	7200, 7200-NPE-G2, 7201, 7301
Cisco ASR 1000 Series	1002, 1002-F, 1004, 1006

Cookbook

1) Basic connectivity
2) ITR and ETR functionality
   a) Enable xTR role
   b) Create mapping
   c) Configuring appropriate MS and MR
3) MS and MR functionality
   a) Enable MS and MR
   b) Assign mappings to sites
4) ALT functionality
   a) Enable ALT
   b) VRF configuration for IPv4 and IPv6
   c) Creating tunnel interface
   d) BGP peering through GRE tunnel with LISP redistribution

Scenario

- C7200 with ADVIPSERVICESK9-M, version 15.2(4)M
- IPv6 ping from one site to another

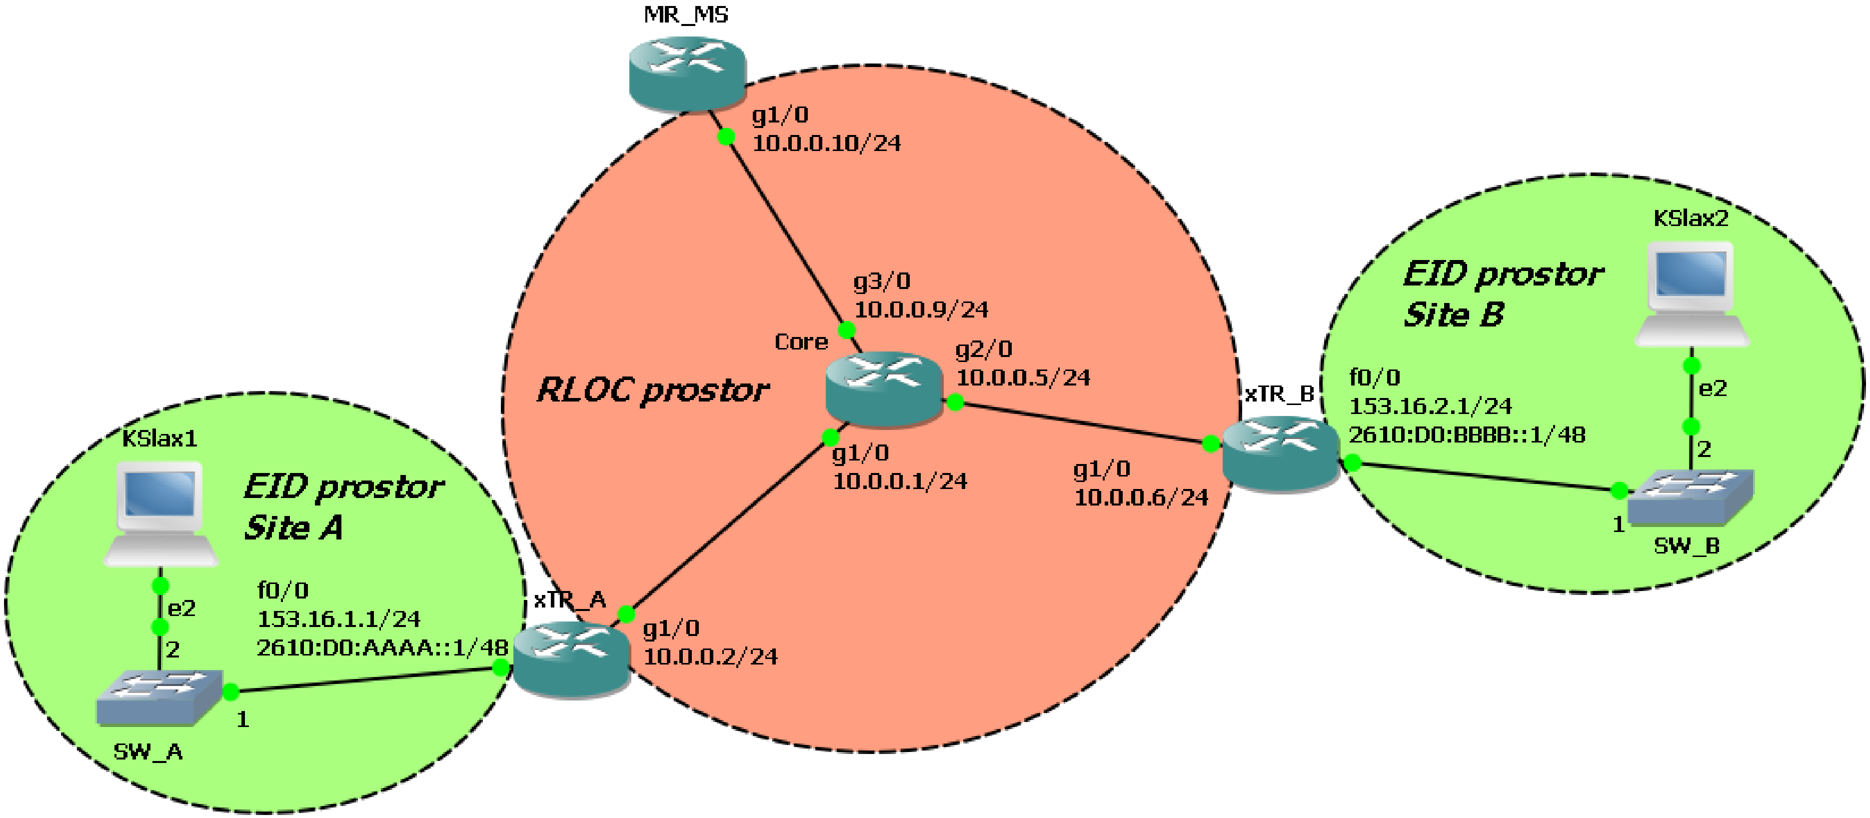

Step 2) ITR and ETR Functionality

a) Enable xTR role

```
(config-router)# router lisp
(config-router-lisp)#
   ipv4 itr
   ipv4 etr
   ipv6 itr
   ipv6 etr
```

b) Create mapping

```
(config-router-lisp)#
   database-mapping EID RLOC1 priority [0-255] weight [0-100]
   database-mapping EID RLOC2 priority [0-255] weight [0-100]
```

c) Configuring appropriate MS and MR

```
(config-router-lisp)#
   ipv4 itr map-resolver address
   ipv4 etr map-server address key heslo {proxy-reply}
   ipv6 itr map-resolver address
   ipv6 etr map-server address key heslo {proxy-reply}
```

Step 3) MS and MR Functionality

a) Enable MS and MR

```
(config-router-lisp)#
   ipv4 map-server
   ipv4 map-resolver
   ipv6 map-server
   ipv6 map-resolver
```

b) Assign mappings to sites

```
(config-router-lisp)# site jmeno
(config-router-lisp-site)#
   authentication-key heslo
   eid-prefix EID
```

Step 4) ALT Functionality

a) Enable ALT role

```
(config-router-lisp)# ipv4 alt-vrf vrf
                      ipv6 alt-vrf vrf
```

b) VRF configuration for IPv4 and IPv6

```
(config)# vrf upgrade-cli multi-af-mode
(config)# vrf definition vrf
(config-vrf)# address-family ipv4
(config-vrf)# address-family ipv6
```

c) Creating tunnel interface

```
(config)# interface tunnel number
(config-interface)#
   ip address addr mask
   tunnel source [iface|addr]
   tunnel destination addr
   vrf forwarding vrf
```

d) BGP peering through GRE tunnel with LISP redistribution

```
(config)# router bgp ASN
(config-router)# address-family [ipv4|ipv6]
(config-af)#
   neighbor addr remote-as ASN
   neighbor addr activate
   redistribute lisp
```

Verification ①

- Successful registration of EID-to-RLOC mapping

Dynamips(3): MS_MR, Console port

```
MS-MR#sh lisp site detail
LISP Site Registration Information

Site name: Site-A
Description: LISP Site A
Allowed configured locators: any
Allowed EID-prefixes:
  EID-prefix: 192.168.1.0/24
    First registered:     00:39:51
    Routing table tag:    0x0
    Origin:               Configuration
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
    ETR 10.0.0.2, last registered 00:00:28, no proxy-reply, no map-notify
                  TTL 1d00h
      Locator    Local   State      Pri/Wgt
      10.0.0.2   yes     up           1/100
  EID-prefix: 2001:DB8:A::/48
    First registered:     00:39:51
    Routing table tag:    0x0
    Origin:               Configuration
    Registration errors:
      Authentication failures:   0
      Allowed locators mismatch: 0
    ETR 10.0.0.2, last registered 00:00:20, no proxy-reply, no map-notify
                  TTL 1d00h
      Locator    Local   State      Pri/Wgt
      10.0.0.2   yes     up           1/100

MS-MR#
```

Verification ②

- LISP-to-LISP IPv4/IPv6 communication over IPv4 core
- Ping from xTR_A to xTR_B with EID as src a dst

Dynamips(0): xTR_A, Console port

```
xTR-A#clear ipv6 lisp map-cache
xTR-A#sh ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 1 entries

::/0, uptime: 00:00:03, expires: never, via static
  Negative cache entry, action: send-map-request
xTR-A#ping 2001:db8:b::1 source 2001:db8:a::1 repeat 100
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 2001:DB8:B::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:A::1
.!!!!!!
Success rate is 85 percent (6/7), round-trip min/avg/max = 136/161/180 ms
xTR-A#sh ipv6 lisp map-cache
LISP IPv6 Mapping Cache, 2 entries

::/0, uptime: 00:00:14, expires: never, via static
  Negative cache entry, action: send-map-request
2001:DB8:B::/48, uptime: 00:00:07, expires: 23:59:45, via map-reply, complete
  Locator    Uptime    State      Pri/Wgt
  10.0.0.6   00:00:07  up           1/100
xTR-A#
```

Results ①

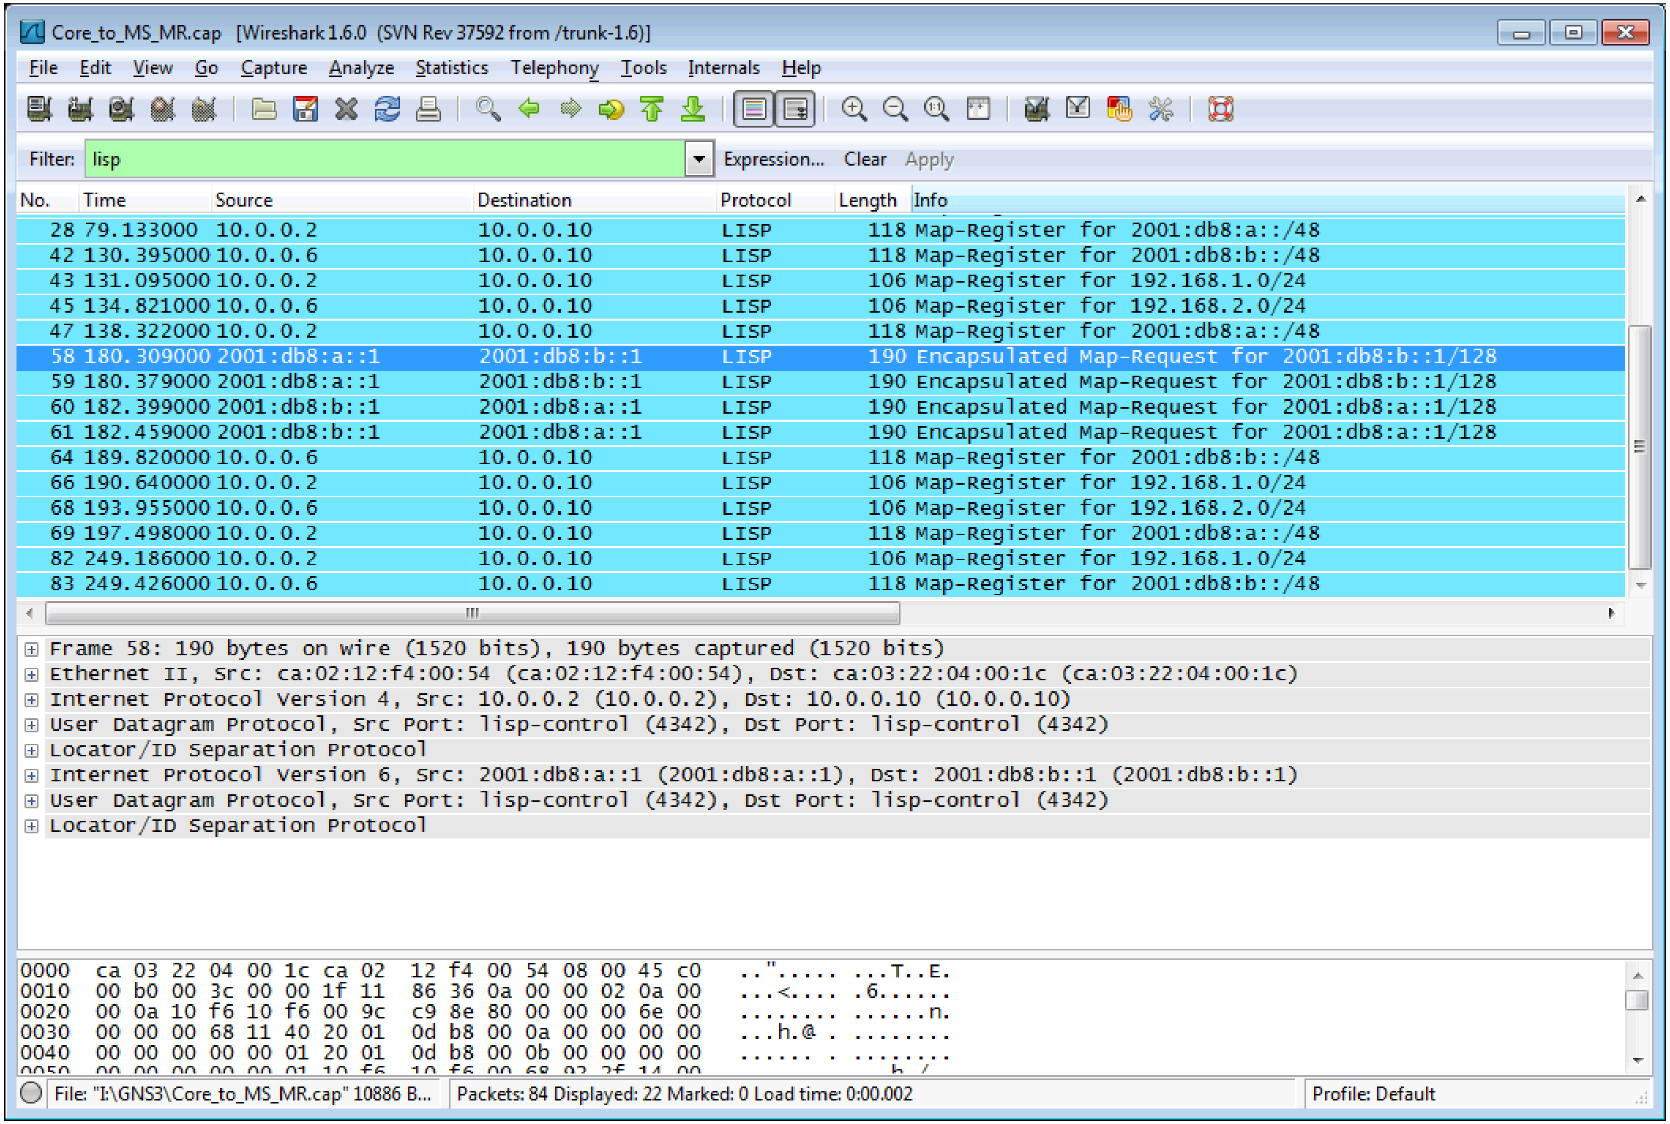

Core_to_MS_MR.cap [Wireshark 1.6.0 (SVN Rev 37592 from /trunk-1.6)]
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter: lisp
Expression... Clear Apply
No. Time Source Destination Protocol Length Info
28 79.133000 10.0.0.2 10.0.0.10 LISP 118 Map-Register for 2001:db8:a::/48
42 130.395000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
43 131.095000 10.0.0.2 10.0.0.10 LISP 106 Map-Register for 192.168.1.0/24
45 134.821000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
47 138.322000 10.0.0.2 10.0.0.10 LISP 118 Map-Register for 2001:db8:a::/48
58 180.309000 2001:db8:a::1 2001:db8:b::1 LISP 190 Encapsulated Map-Request for 2001:db8:b::1/128
59 180.379000 2001:db8:a::1 2001:db8:b::1 LISP 190 Encapsulated Map-Request for 2001:db8:b::1/128
60 182.399000 2001:db8:b::1 2001:db8:a::1 LISP 190 Encapsulated Map-Request for 2001:db8:a::1/128
61 182.459000 2001:db8:b::1 2001:db8:a::1 LISP 190 Encapsulated Map-Request for 2001:db8:a::1/128
64 189.820000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
66 190.640000 10.0.0.2 10.0.0.10 LISP 106 Map-Register for 192.168.1.0/24
68 193.955000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
69 197.498000 10.0.0.2 10.0.0.10 LISP 118 Map-Register for 2001:db8:a::/48
82 249.186000 10.0.0.2 10.0.0.10 LISP 106 Map-Register for 192.168.1.0/24
83 249.426000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
Frame 58: 190 bytes on wire (1520 bits), 190 bytes captured (1520 bits)
Ethernet II, Src: ca:02:12:f4:00:54 (ca:02:12:f4:00:54), Dst: ca:03:22:04:00:1c (ca:03:22:04:00:1c)
Internet Protocol Version 4, Src: 10.0.0.2 (10.0.0.2), Dst: 10.0.0.10 (10.0.0.10)
User Datagram Protocol, Src Port: lisp-control (4342), Dst Port: lisp-control (4342)
Locator/ID Separation Protocol
Internet Protocol Version 6, Src: 2001:db8:a::1 (2001:db8:a::1), Dst: 2001:db8:b::1 (2001:db8:b::1)
User Datagram Protocol, Src Port: lisp-control (4342), Dst Port: lisp-control (4342)
Locator/ID Separation Protocol
0000 ca 03 22 04 00 1c ca 02 12 f4 00 54 08 00 45 c0 .."..... ...T..E.
0010 00 b0 00 3c 00 00 1f 11 86 36 0a 00 00 02 0a 00 ...<.... .6......
0020 00 0a 10 f6 10 f6 00 9c c9 8e 80 00 00 00 6e 00 ....n.
0030 00 00 00 68 11 40 20 01 0d b8 00 0a 00 00 00 00 ...h.@ . ...
0040 00 00 00 00 00 01 20 01 0d b8 00 0b 00 00 00 00 ....
File: "I:\GNS3\Core_to_MS_MR.cap" 10886 B...
Packets: 84 Displayed: 22 Marked: 0 Load time: 0:00.002
Profile: Default

Results ②

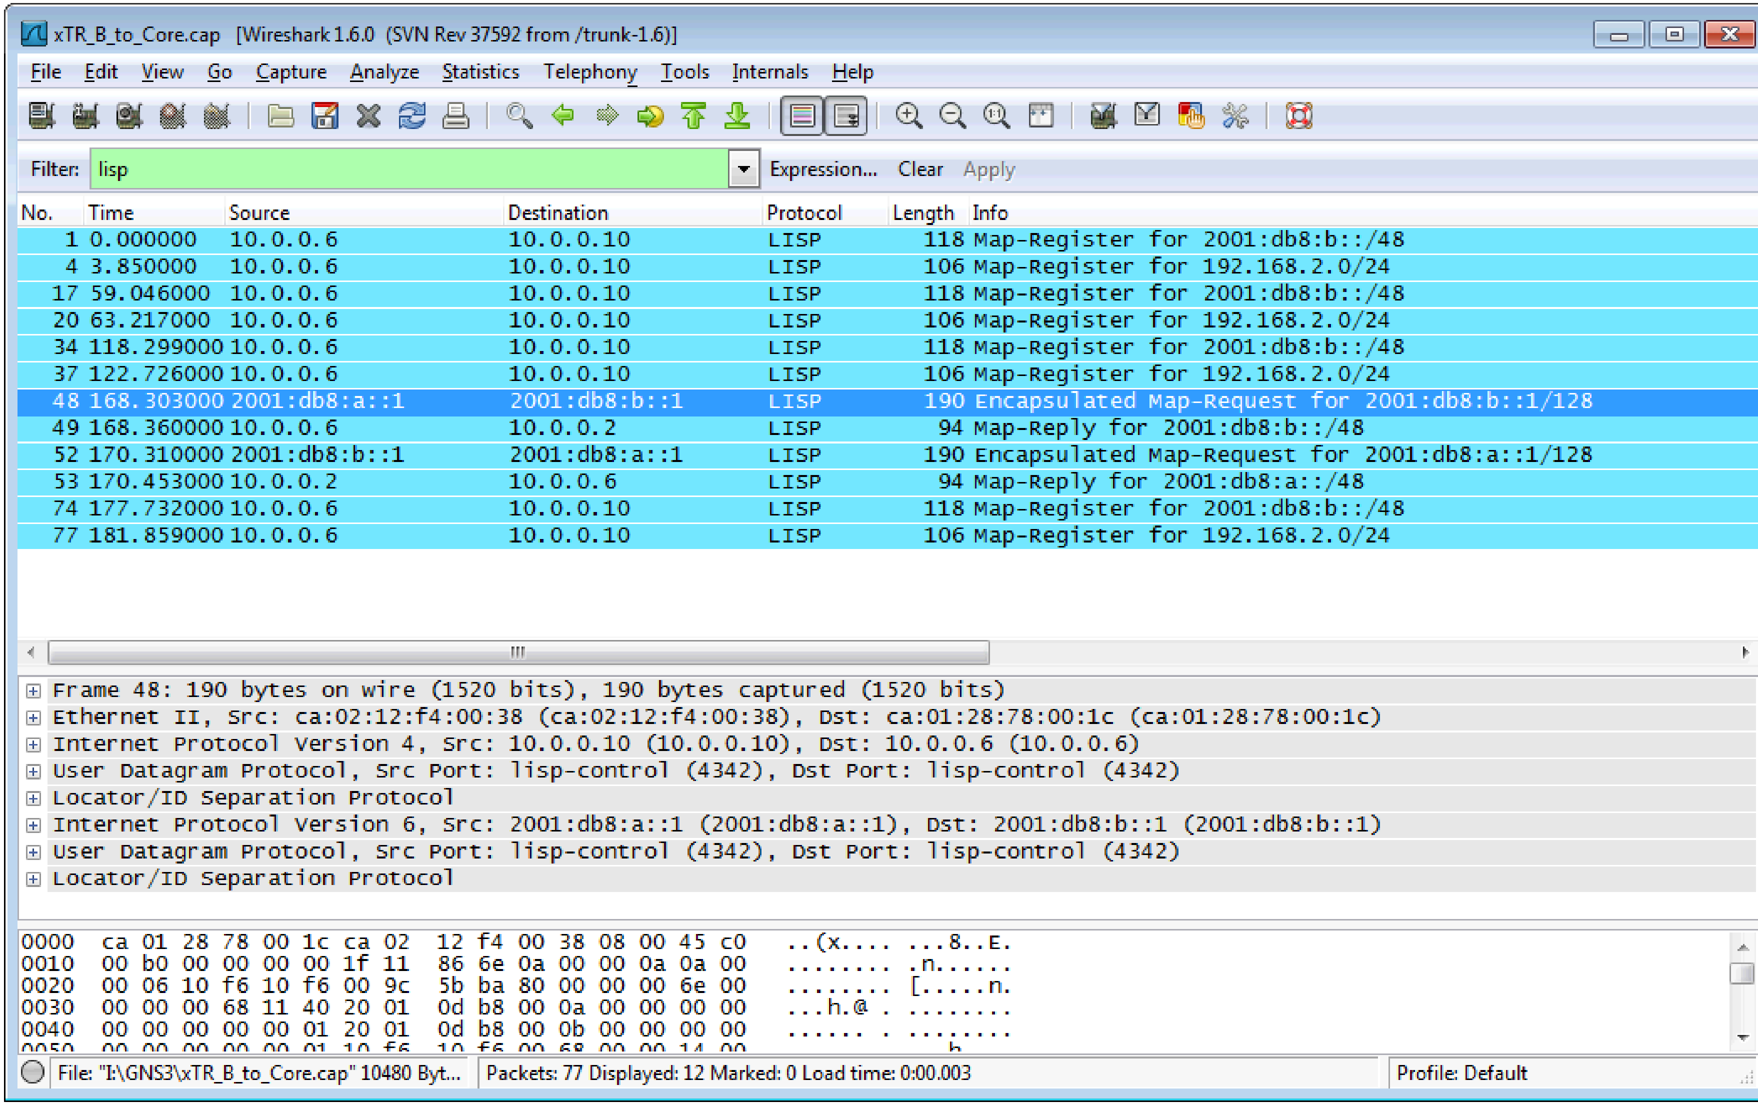

xTR_B_to_Core.cap [Wireshark 1.6.0 (SVN Rev 37592 from /trunk-1.6)]
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter: lisp
Expression... Clear Apply
No. Time Source Destination Protocol Length Info
1 0.000000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
4 3.850000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
17 59.046000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
20 63.217000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
34 118.299000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
37 122.726000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
48 168.303000 2001:db8:a::1 2001:db8:b::1 LISP 190 Encapsulated Map-Request for 2001:db8:b::1/128
49 168.360000 10.0.0.6 10.0.0.2 LISP 94 Map-Reply for 2001:db8:b::/48
52 170.310000 2001:db8:b::1 2001:db8:a::1 LISP 190 Encapsulated Map-Request for 2001:db8:a::1/128
53 170.453000 10.0.0.2 10.0.0.6 LISP 94 Map-Reply for 2001:db8:a::/48
74 177.732000 10.0.0.6 10.0.0.10 LISP 118 Map-Register for 2001:db8:b::/48
77 181.859000 10.0.0.6 10.0.0.10 LISP 106 Map-Register for 192.168.2.0/24
Frame 48: 190 bytes on wire (1520 bits), 190 bytes captured (1520 bits)
Ethernet II, Src: ca:02:12:f4:00:38 (ca:02:12:f4:00:38), Dst: ca:01:28:78:00:1c (ca:01:28:78:00:1c)
Internet Protocol Version 4, Src: 10.0.0.10 (10.0.0.10), Dst: 10.0.0.6 (10.0.0.6)
User Datagram Protocol, Src Port: lisp-control (4342), Dst Port: lisp-control (4342)
Locator/ID Separation Protocol
Internet Protocol Version 6, Src: 2001:db8:a::1 (2001:db8:a::1), Dst: 2001:db8:b::1 (2001:db8:b::1)
User Datagram Protocol, Src Port: lisp-control (4342), Dst Port: lisp-control (4342)
Locator/ID Separation Protocol
0000 ca 01 28 78 00 1c ca 02 12 f4 00 38 08 00 45 c0 ..(x.... ...8..E.
0010 00 b0 00 00 00 00 1f 11 86 6e 0a 00 00 0a 0a 00 .... .n......
0020 00 06 10 f6 10 f6 00 9c 5b ba 80 00 00 00 6e 00 .... [.....n.
0030 00 00 00 68 11 40 20 01 0d b8 00 0a 00 00 00 00 ...h.@ . ...
0040 00 00 00 00 00 01 20 01 0d b8 00 0b 00 00 00 00 ....
File: "I:\GNS3\xTR_B_to_Core.cap" 10480 Byt...
Packets: 77 Displayed: 12 Marked: 0 Load time: 0:00.003
Profile: Default

Use-case: IPv6 over IPv4

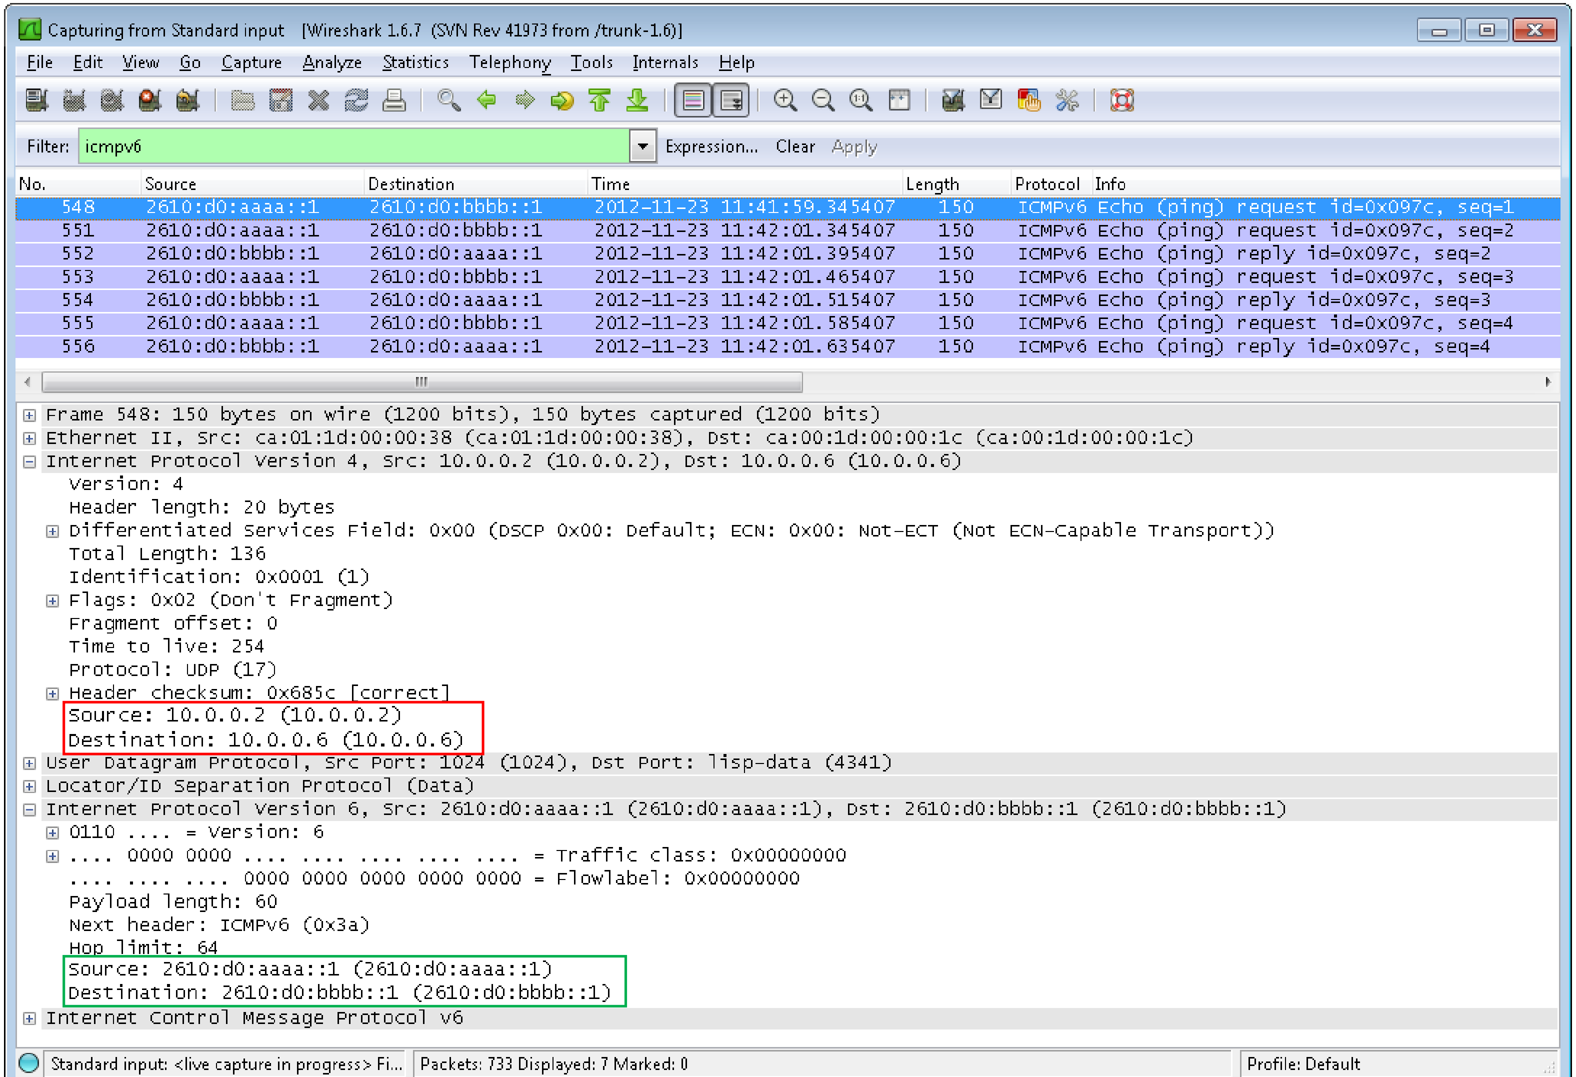

Useful Commands

- xTR

```
show ip lisp map-cache
show ip lisp database
show ip lisp statistics
show ip[v6] lisp
```

- MRMS

```
show lisp site
show lisp
show ip[v6] lisp
```

Conclusion

Where can you use it?

- Multihoming
- IPv6 transition mechanism
- VM mobility

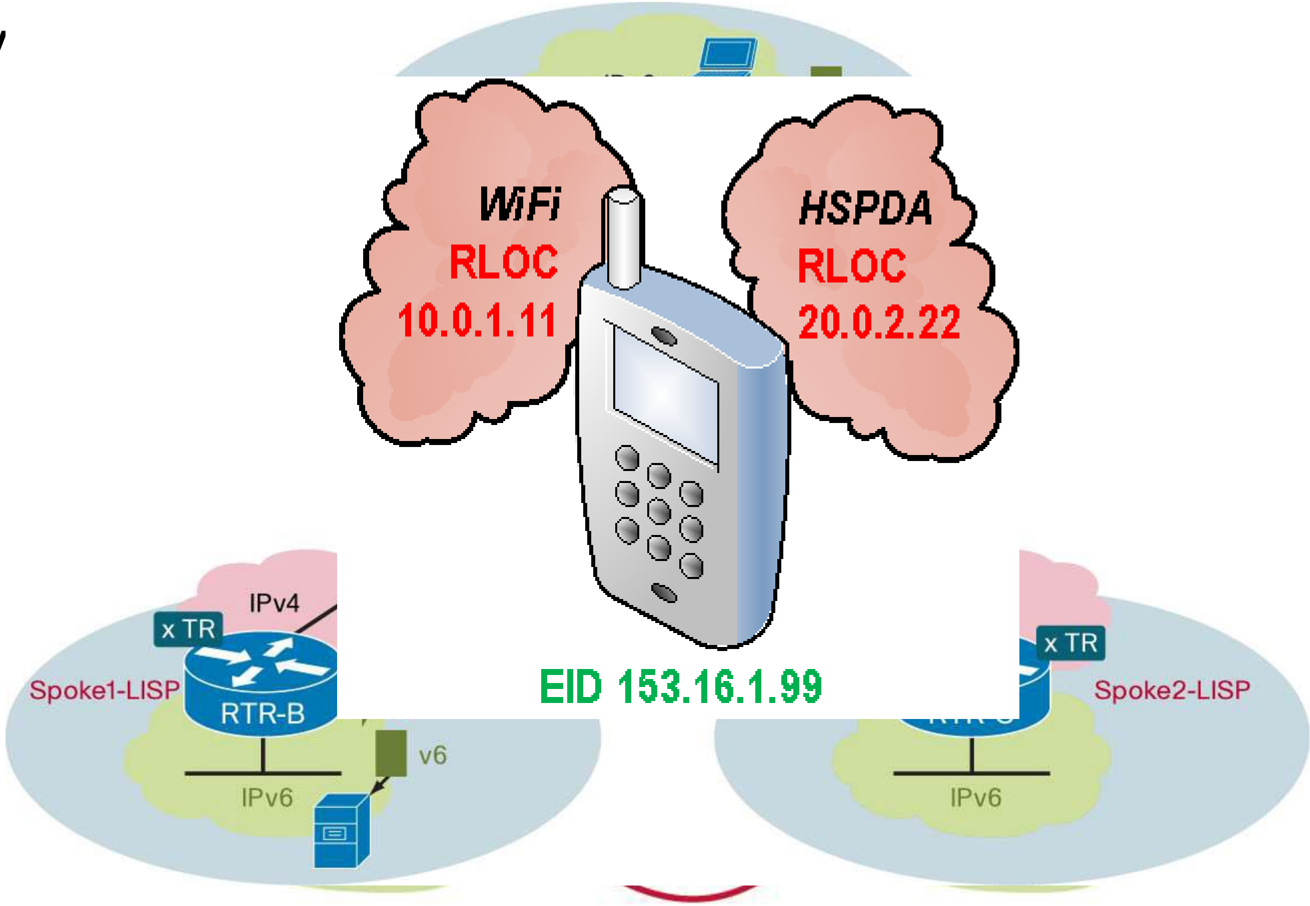

State of The Art

- Reserved prefix for EID (http://lispmon.net/)
  - IPv4 - 153.16.0.0/16
  - IPv6 - 2610:00d0::/32
- ALT infrastructure numbering
  - IPv4 addreses of GRE tunnels 240.0.0.0/4
  - ASN 32656.x
- **LISP Beta Network** (https://www.lisp4.net/beta-network/)
  - Google, Facebook, Cisco, Qualcomm, AT&T, Lufthansa, Microsoft,
- Implementation
  - Cisco v IOS, IOS-XE, IOS-XR, NX-OS
  - OpenLISP onf FreeBSD
  - LISPmob, Aless, OpenWRT for Linux-based system
  - Gingerbread for Android

Standardization Efforts

- https://datatracker.ietf.org/wg/lisp/documents/

Document	Date
Active Internet-Drafts	
draft-ietf-lisp-eid-anonymity-00 **LISP EID Anonymity**	2017-08-17 9 pages
draft-ietf-lisp-eid-mobility-00 **LISP L2/L3 EID Mobility Using a Unified Control Plane**	2017-05-11 23 pages
draft-ietf-lisp-introduction-13 **An Architectural Introduction to the Locator/ID Separation Protocol (LISP)**	2015-04-02 27 pages
draft-ietf-lisp-mn-00 **LISP Mobile Node**	2017-04-28 22 pages
draft-ietf-lisp-predictive-rlocs-00 **LISP Predictive RLOCs**	2017-06-07 13 pages
draft-ietf-lisp-rfc6830bis-05 **The Locator/ID Separation Protocol (LISP)**	2017-08-29 53 pages
draft-ietf-lisp-rfc6833bis-05 **Locator/ID Separation Protocol (LISP) Control-Plane**	2017-05-10 39 pages
draft-ietf-lisp-sec-13 **LISP-Security (LISP-SEC)**	2017-09-20 23 pages Ne
draft-ietf-lisp-signal-free-multicast-06 **Signal-Free LISP Multicast**	2017-08-01 23 pages
draft-ietf-lisp-te-00 **LISP Traffic Engineering Use-Cases**	2017-04-28 16 pages
draft-ietf-lisp-vendor-lcaf-00 **Vendor Specific LCAF**	2017-08-17 5 pages
draft-ietf-lisp-vpn-00 **LISP Virtual Private Networks (VPNs)**	2017-05-11 16 pages
draft-ietf-lisp-yang-05 **LISP YANG Model**	2017-07-03 56 pages
RFCs	
RFC 6830 *(was draft-ietf-lisp)* **The Locator/ID Separation Protocol (LISP)**	2013-01 75 pages
RFC 6831 *(was draft-ietf-lisp-multicast)* **The Locator/ID Separation Protocol (LISP) for Multicast Environments**	2013-01 28 pages
RFC 6832 *(was draft-ietf-lisp-interworking)* **Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites**	2013-01 19 pages
RFC 6833 *(was draft-ietf-lisp-ms)* **Locator/ID Separation Protocol (LISP) Map-Server Interface**	2013-01 13 pages
RFC 6834 *(was draft-ietf-lisp-map-versioning)* **Locator/ID Separation Protocol (LISP) Map-Versioning**	2013-01 21 pages
RFC 6835 *(was draft-ietf-lisp-lig)* **The Locator/ID Separation Protocol Internet Groper (LIG)**	2013-01 12 pages
RFC 6836 *(was draft-ietf-lisp-alt)* **Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)**	2013-01 25 pages
RFC 7052 *(was draft-ietf-lisp-mib)* **Locator/ID Separation Protocol (LISP) MIB** Errata	2013-10 66 pages
RFC 7215 *(was draft-ietf-lisp-deployment)* **Locator/Identifier Separation Protocol (LISP) Network Element Deployment Considerations**	2014-04 30 pages
RFC 7834 *(was draft-ietf-lisp-impact)* **Locator/ID Separation Protocol (LISP) Impact**	2016-04 18 pages

Reference

- GoogleTech Talks
  - http://www.youtube.com/watch?v=WSI1RAlFU3s
  - http://www.youtube.com/watch?v=_bz4cRuAcak
  - http://www.youtube.com/watch?v=fxdm-Xouu-k
  - http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_11-1/111_lisp.html
- IETF WG
  - https://datatracker.ietf.org/wg/lisp/charter/
- Additional info
  - Official webpage - http://lisp.cisco.com/lisp_tech.html
  - LISP Beta Network - http://www.lisp4.net a http://www.lisp6.net
  - LISP DDT Root - http://www.ddt-root.org
  - LinkedIn - http://www.linkedin.com/groups/LISP-3776183