


# A Decentralized Mapping System

*draft-farinacci-lisp-decent-03*

***IETF Prague***  
March 2019

*Dino Farinacci & Colin Cantrell*

# Draft Status



*After implementation  
of Pull-Based*

- B.1. Changes to draft-farinacci-lisp-decent-03
  - o Posted March 2019.
  - o Introduce the Hash Mask which is used to grab common bits from a registered prefix and a lookup prefix.
  - o Spec how multicast lookups are done in the pull-based mapping system.
  - o Indicate the hash string includes the unicast EID mask-length and multicast group and source mask-lengths.

*Added Pull-Based  
to Draft*

- B.2. Changes to draft-farinacci-lisp-decent-02
  - o Posted November 2018.
  - o Changed references from peer-group to seed-group to make the algorithms in this document more like how blockchain networks initialize the peer-to-peer network.
  - o Added pull mechanism to compliment the push mechanism. The pull mechanism could be used as a seed-group to bootstrap the push mechanism.

*Colon presented -01  
in Montreal*

- B.3. Changes to draft-farinacci-lisp-decent-01
  - o Posted July 2018.

*Dino presented -00  
in London*

- B.4. Changes to draft-farinacci-lisp-decent-00
  - o Initial draft posted January 2018.

# Problem Statement

- What if LISP xTRs didn't need to depend on a third-party
- What if LISP xTRs could multi-home and roam to inform each other about new RLOCs
- What if LISP xTRs could be their own mapping system
- Let's build a purely democratized and decentralized control-plane

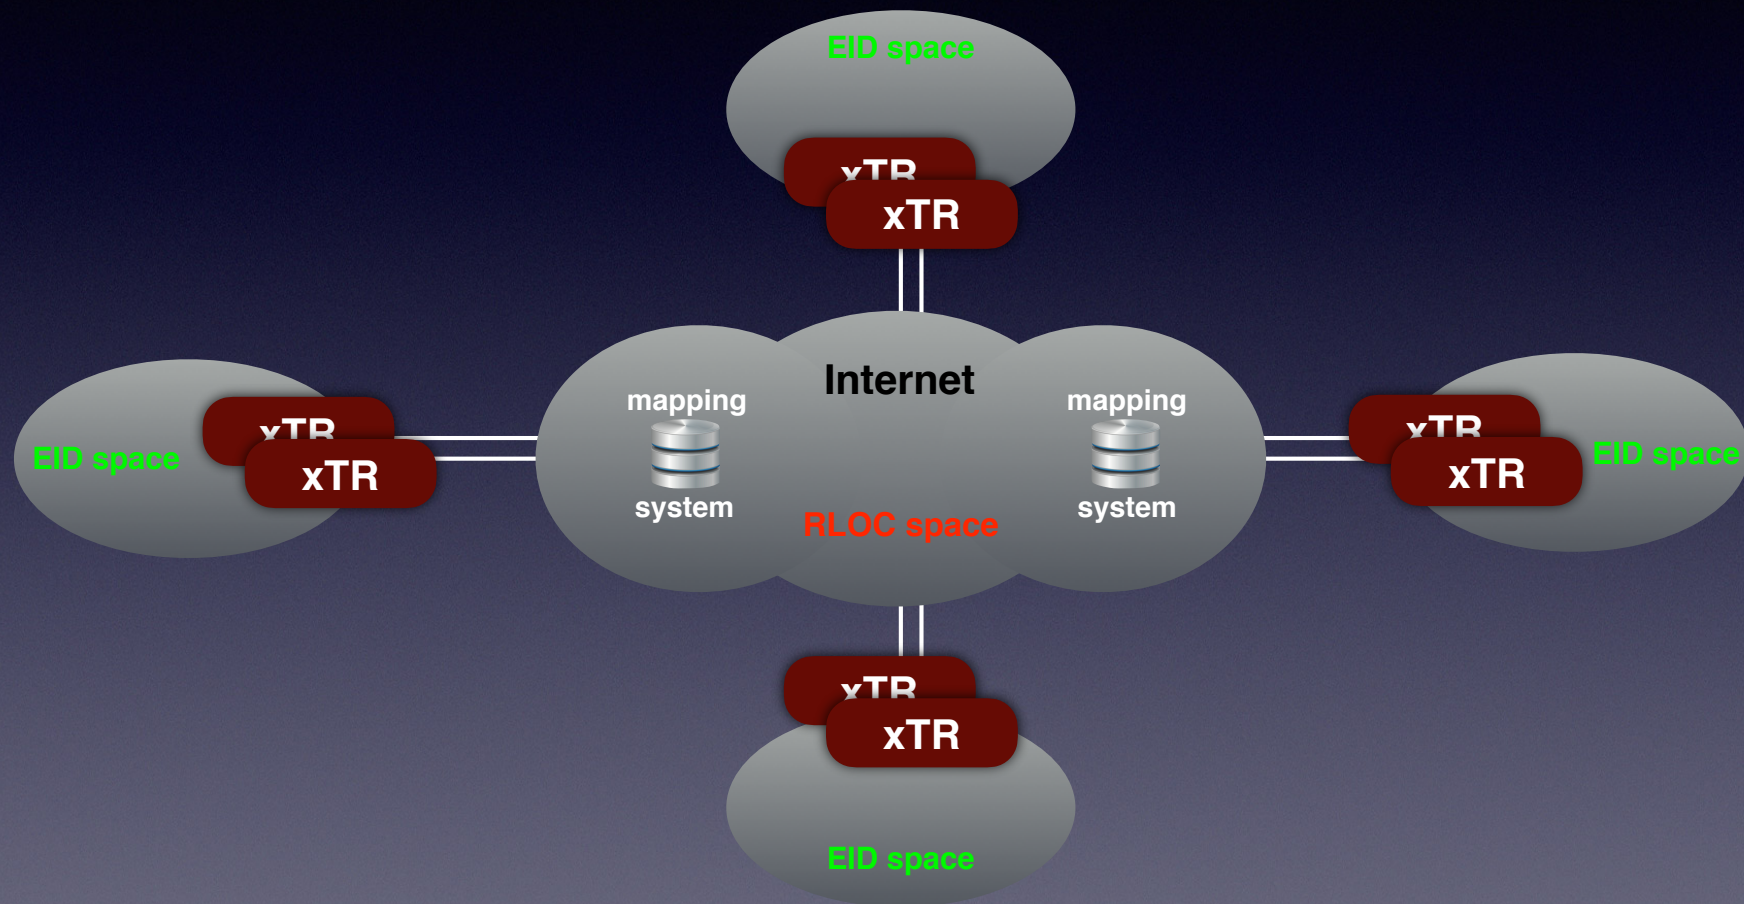
*Endpoint IDs (EIDs)*

*Routing Locators (RLOCs)*

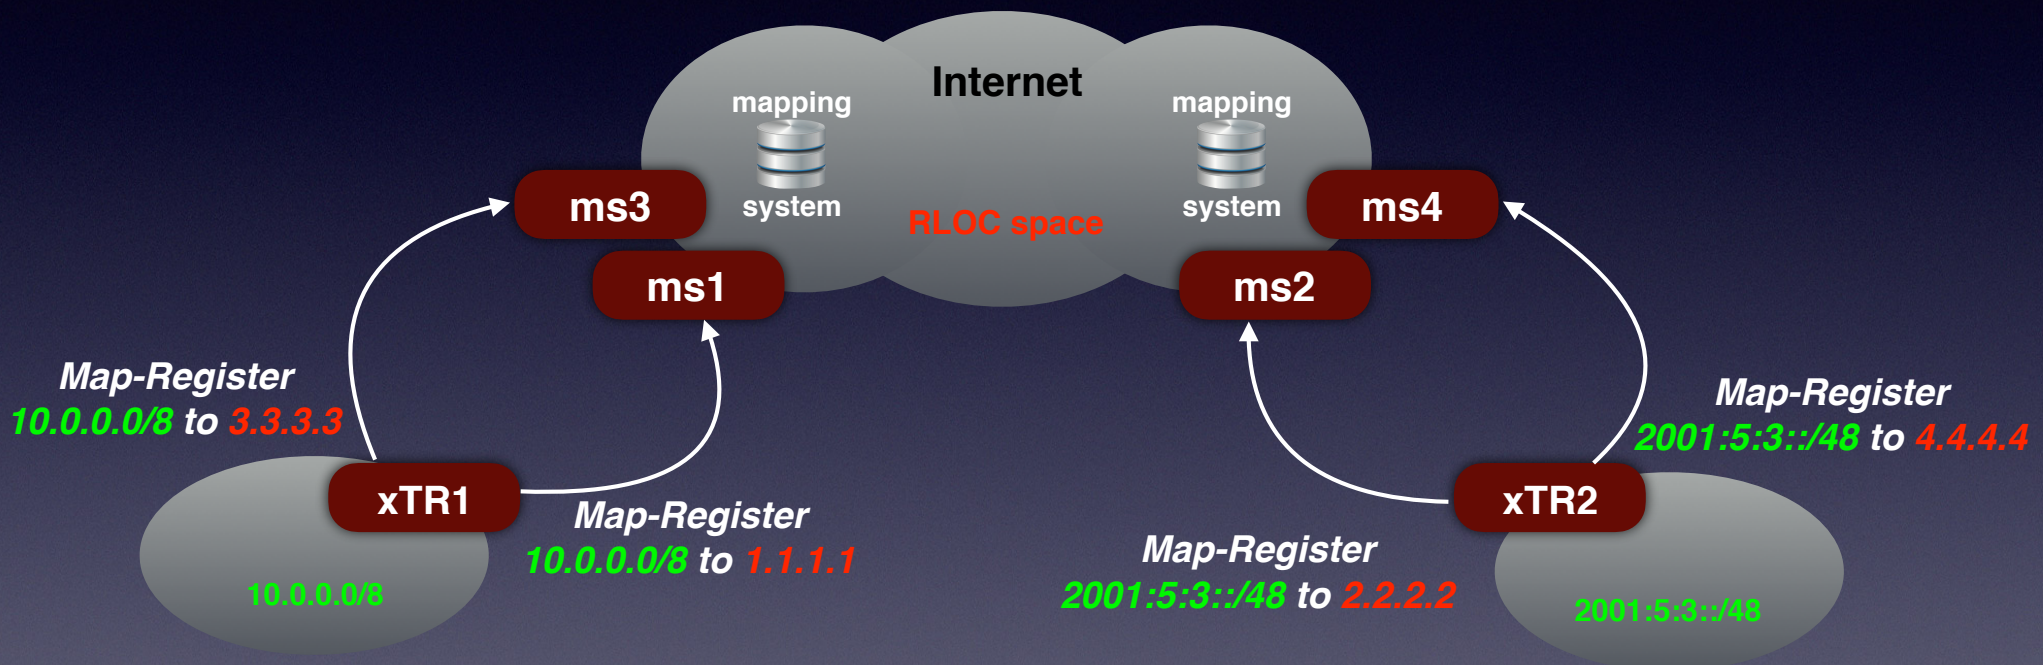
# Use-Cases

- Crypto-Currency & Blockchain Applications
- Emergency Networking (Mesh Networks)
- Plug-and-Play VPN Networking
- Space Networking (Software-Defined Satellites)
- Sharable Economy Apps

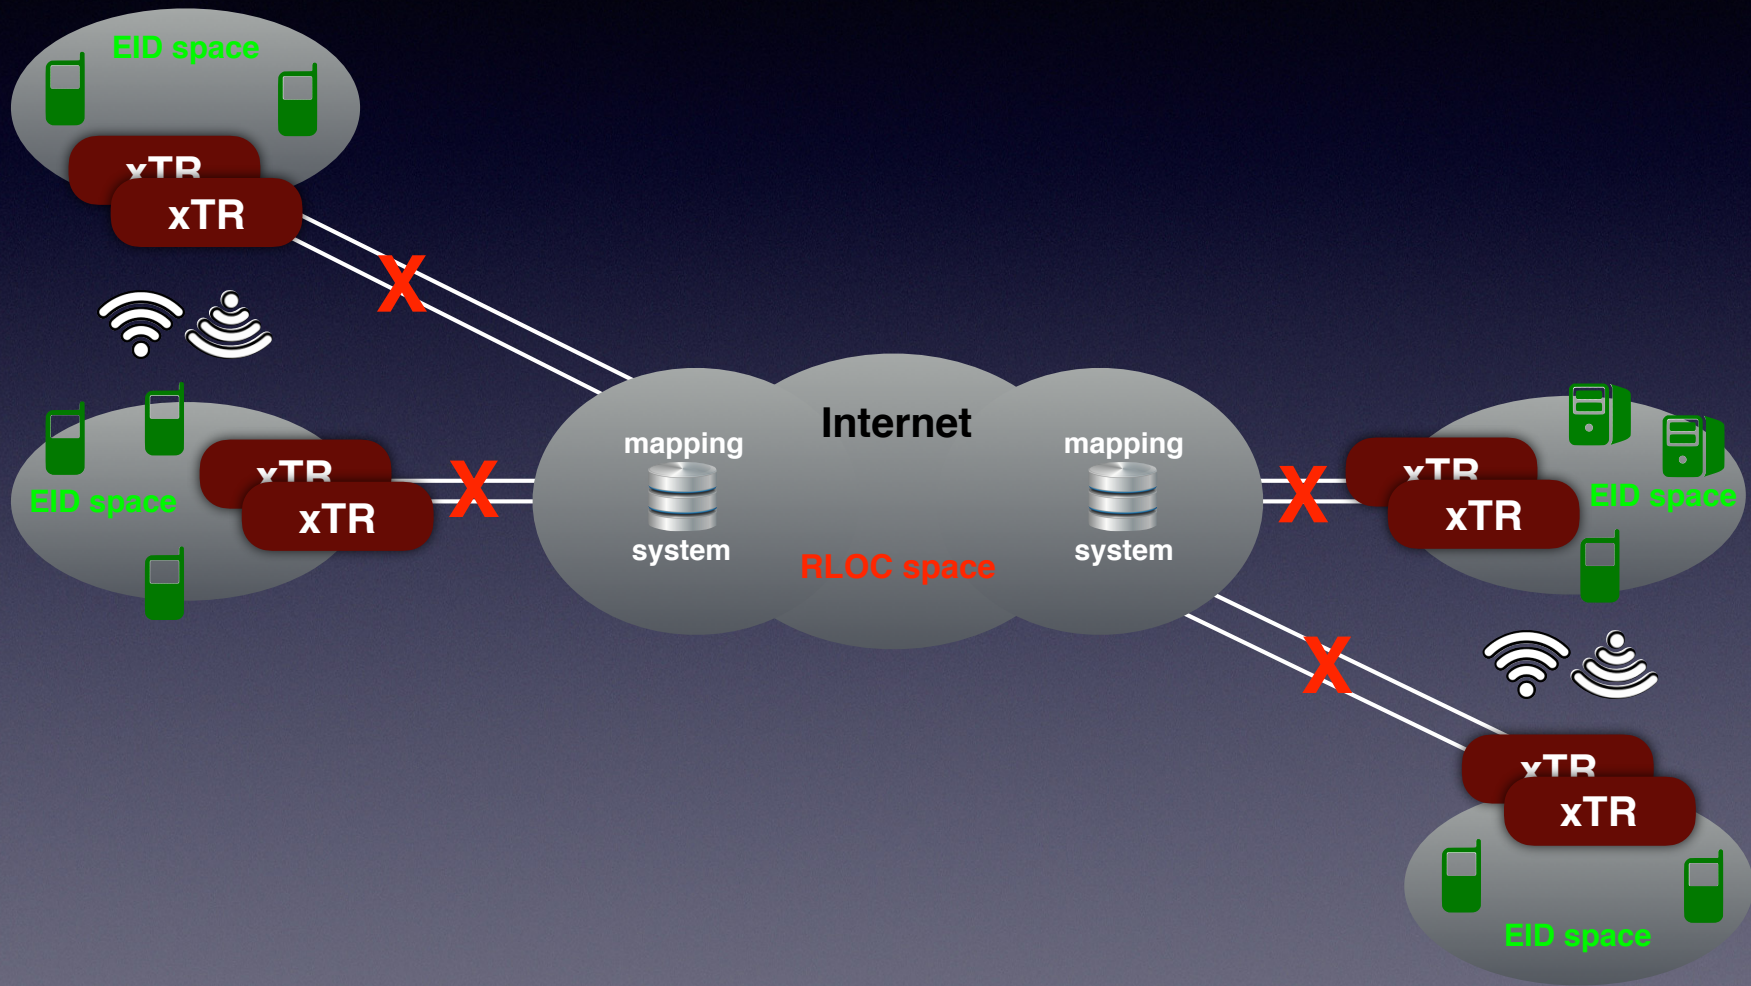
# Today's Model Mapping System



# LISP Control-Plane Messages



# Network Connectivity



# Decentralized Map-Server?

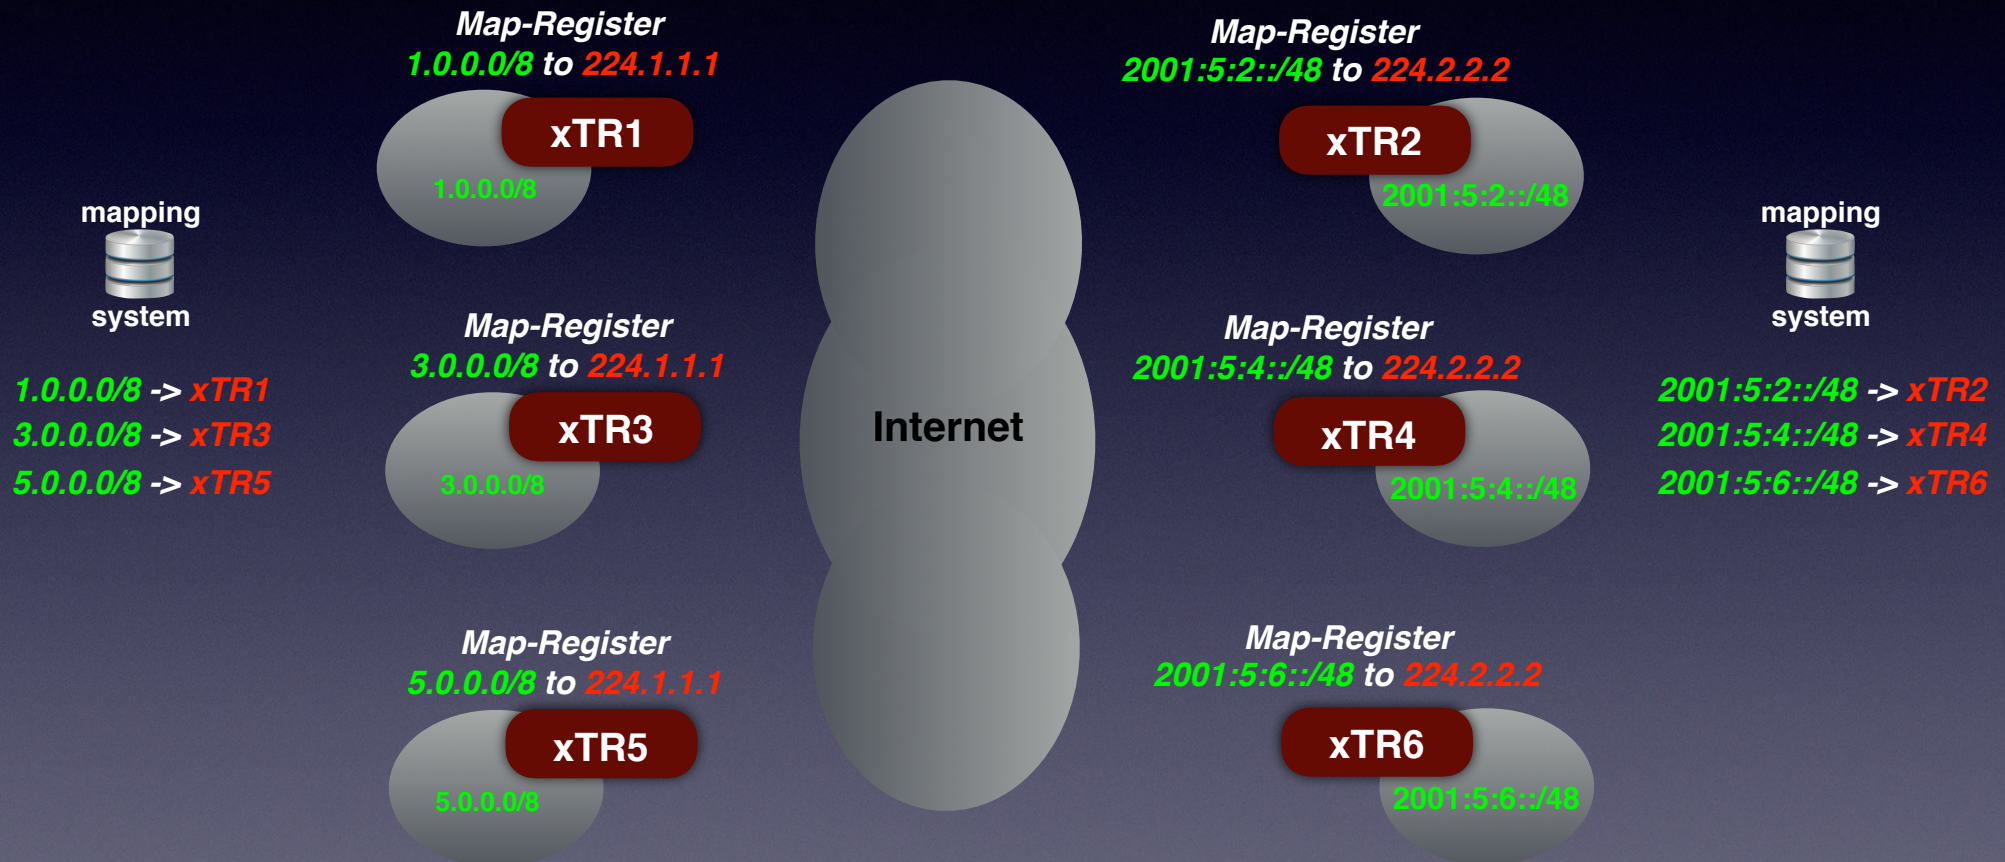
- **Push-Based** Model
  - What if each xTR was a Map-Server
  - What if each xTR could Map-Register to each xTR
  - The mapping system would be synchronized
  - An xTR could be a Map-Resolver for itself
- **Pull-Based** Model
  - Subset of xTRs are co-located Map-Resolvers & Map-Servers
  - Registrations are sharded (just like today)
  - The EID as input to a hash function define where Map-Requests and Map-Registers go
  - Redundancy achieved through DNS level of indirection



# Definition of a **Push-Based** Mapping System

- A consolidated mapping system is identified by a multicast group address
- The xTRs that are part of a mapping system join the same multicast group
- Map-Registers are sent to the group - all xTRs receive all mappings
- Efficient distribution when underlay supports multicast or head-end replication at each xTR
- Map-Request lookup has low latency
- xTRs build and send **1** Map-Register for ***n*** xTRs
- Management simplified by accessing one xTR to get all mappings

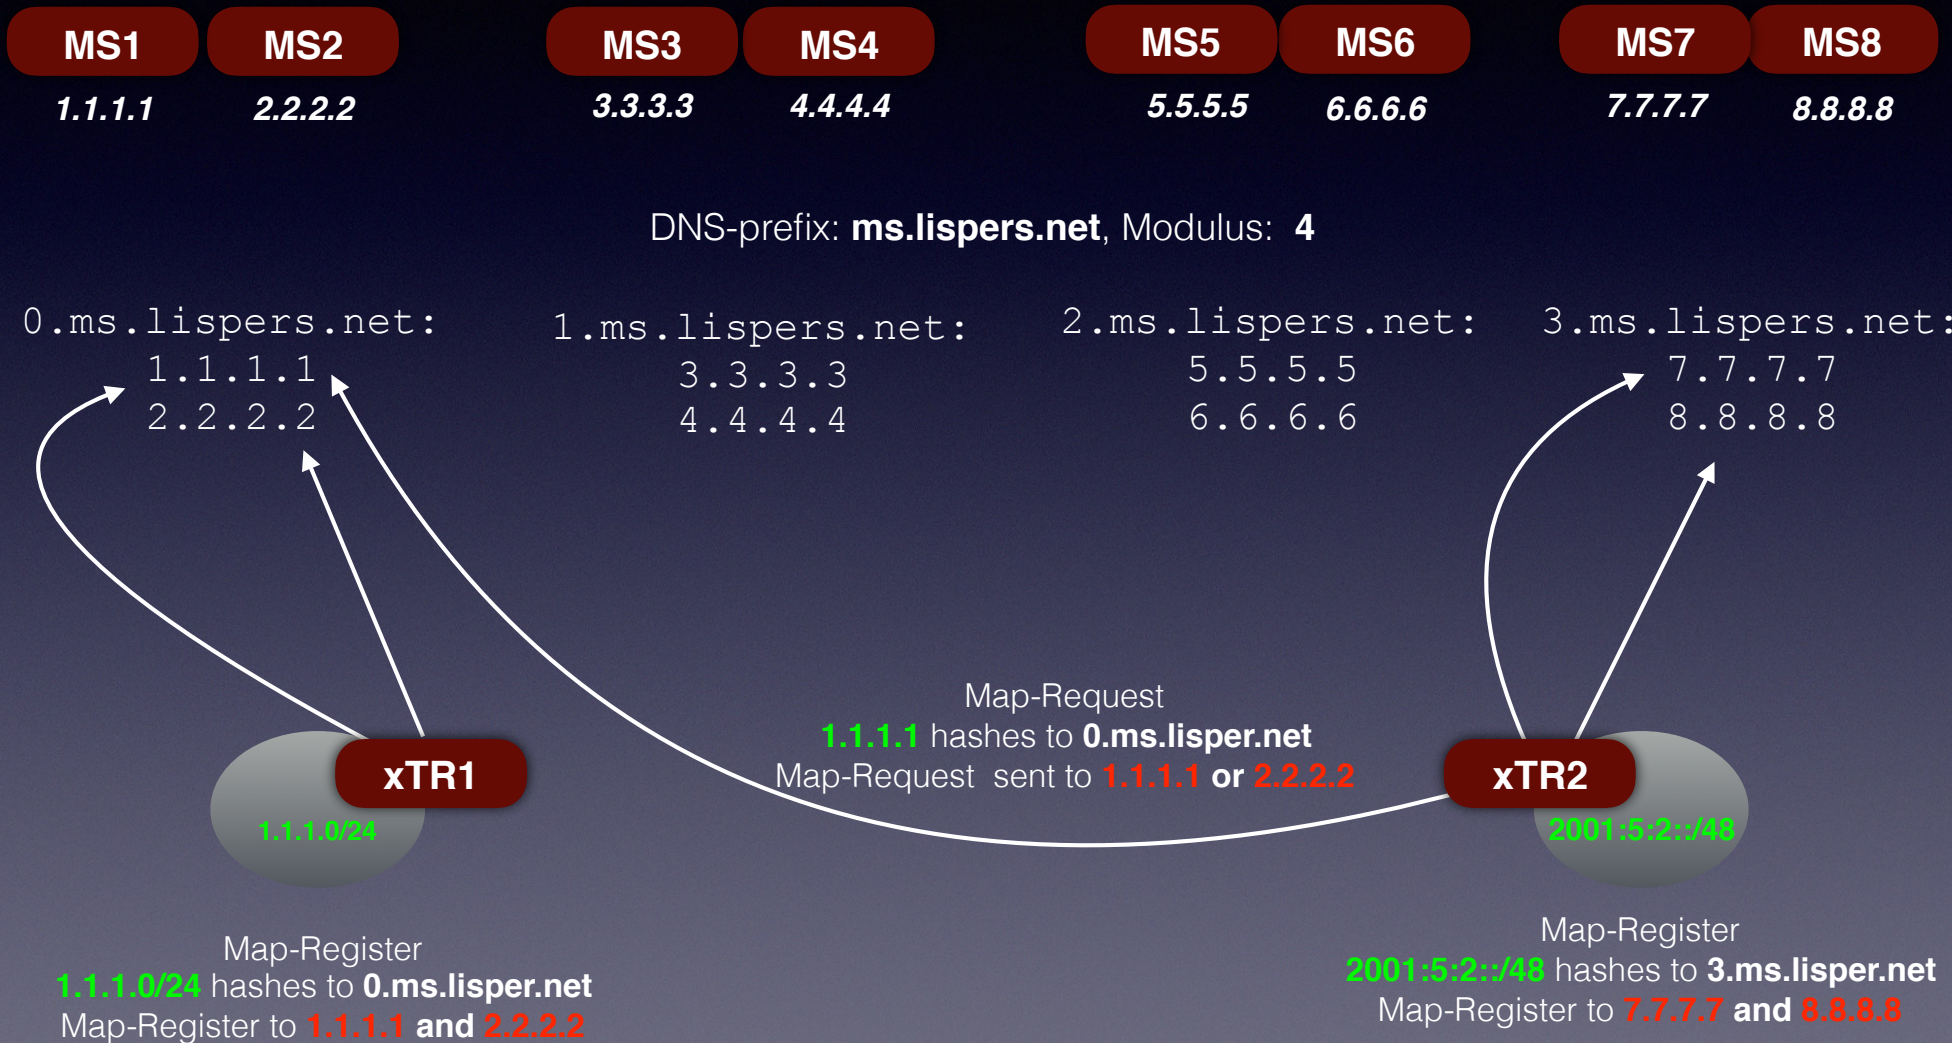
# LISP-Decent (Push-Based) Control-Plane Messages



# Definition of a **Pull-Based** Mapping System

- Map-Server sets are defined by a DNS-name prefix
- EID is hashed to produce a Modulus Index
- Modulus Index selects a DNS-name
- DNS-name has A-records for each Map-Server
- Map-Registers and Map-Requests go to **same** Map-Server Set
- A Pull-Based mapping system can bootstrap a Push-Based mapping system
- A Pull-Based mapping system does not need LISP-DDT or LISP-ALT

# LISP-Decent (Pull-Based) Control-Plane Messages



# Brief Push-Based Demo

- 3 containers each running a *lispers.net* xTR
- Docker bridge NOT doing multicast
- xTRs are doing head-end replication
- xTRs register an IPv4 EID-prefix and a Name EID

# LISP-Decent in Action

**lispers.net**

Scalable Open Overlay Networking

n1

Enter EID for Site-Cache lookup:

**LISP-MS Site Information:**

Site Name	EID-Prefix or (S,G)	Registered	Last Registerer	Last Registered	First Registered	Registration Flags
EIDs	[1]					
	[1]1.1.1.1/32					
	[1]'nexus-n1'					
	[1]2.2.2.2/32					
	[1]'nexus-n2'					
	[1]3.3.3.3/32					
	[1]'nexus-n3'					
peer-groups	[1](0.0.0.0/0, 224.0.0.0/4)					
	[1](0.0.0.0/0, 224.1.1.1/32)					

**lispers.net**

Scalable Open Overlay Networking

n2

Enter EID for Site-Cache lookup:

**LISP-MS Site Information:**

Site Name	EID-Prefix or (S,G)	Registered	Last Registerer	Last Registered	First Registered	Registration Flags
EIDs	[1]					
	[1]2.2.2.2/32					
	[1]'nexus-n2'					
	[1]1.1.1.1/32					
	[1]'nexus-n1'					
	[1]3.3.3.3/32					
	[1]'nexus-n3'					
peer-groups	[1](0.0.0.0/0, 224.0.0.0/4)					
	[1](0.0.0.0/0, 224.1.1.1/32)					

**lispers.net**

Scalable Open Overlay Networking

n3

Enter EID for Site-Cache lookup:

**LISP-MS Site Information:**

Site Name	EID-Prefix or (S,G)	Registered	Last Registerer	Last Registered	First Registered	Registration Flags
EIDs	[1]	(ams)	--	never	never	--
	[1]3.3.3.3/32	yes (dynamic)	[0]172.17.0.7	0:00:29	0:03:34	p-s-l-t-r-m-n
	[1]'nexus-n3'	yes (dynamic)	[0]172.17.0.7	0:00:29	0:03:34	p-s-l-t-r-m-n
	[1]1.1.1.1/32	yes (dynamic)	[0]172.17.0.5	0:00:33	0:01:33	p-s-l-t-r-m-n
	[1]'nexus-n1'	yes (dynamic)	[0]172.17.0.5	0:00:33	0:01:33	p-s-l-t-r-m-n
	[1]2.2.2.2/32	yes (dynamic)	[0]172.17.0.6	0:00:30	0:01:30	p-s-l-t-r-m-n
	[1]'nexus-n2'	yes (dynamic)	[0]172.17.0.6	0:00:30	0:01:30	p-s-l-t-r-m-n
peer-groups	[1](0.0.0.0/0, 224.0.0.0/4)	(ams)	--	never	never	--
	[1](0.0.0.0/0, 224.1.1.1/32)	yes (dynamic)	[0]172.17.0.7	0:00:05	0:03:18	P-s-l-t-R-m-n

# Brief Pull-Based Demo

- 1 **mrms** container simulating all 8 map-servers/map-resolvers
- 1 **etr** container Map-Registering a variety of EID-types
- 1 **itr** container Map-Requesting
- DNS-prefix is **ms.lispers.net**, Modulus is **4**

# DNS Configuration

```
[>>> import socket
[>>> socket.gethostbyname_ex("0.ms.lispers.net")
('0.ms.lispers.net', [], ['1.1.1.1', '2.2.2.2'])
[>>>
[>>> socket.gethostbyname_ex("1.ms.lispers.net")
('1.ms.lispers.net', [], ['3.3.3.3', '4.4.4.4'])
[>>>
[>>> socket.gethostbyname_ex("2.ms.lispers.net")
('2.ms.lispers.net', [], ['5.5.5.5', '6.6.6.6'])
[>>>
[>>> socket.gethostbyname_ex("3.ms.lispers.net")
('3.ms.lispers.net', [], ['7.7.7.7', '8.8.8.8'])
>>> █
```



# MS Map-Registers Received

***lispers.net***

Scalable Open Overlay Networking

mrms

Enter EID for Site-Cache lookup:

## LISP-MS Site Information:

Site Name	EID-Prefix or (S,G)	Registered	Last Registerer	Last Registered	First Registered	Registration Flags
registrations	[0]	(ams)	--	never	never	--
	<a href="#">[0]2.2.2.2/32</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]+000140871820010/64</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]1.1.1.1/32</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]'dino'</a>	yes (dynamic)	[0]172.17.0.3	0:00:07	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0](1.1.1.1/32, 224.1.1.1/32)</a>	yes (dynamic)	[0]172.17.0.3	0:00:07	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]0000-1111-2222/48</a>	yes (dynamic)	[0]172.17.0.3	0:00:07	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]2001::2/128</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]48-51-12-N-2-20-55-E/100</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]+1223344556677880/64</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]+0000000000000000/8</a>	yes (dynamic)	[0]172.17.0.3	0:00:08	0:11:08	p-s-l-t-r-m-n
	<a href="#">[0]2001::1/128</a>	yes (dynamic)	[0]172.17.0.3	0:00:07	0:11:08	p-s-l-t-r-m-n

Mon Mar 25 11:56:57 UTC 2019 - Uptime 0:11:29, Version 0.496

Copyright 2013-2019 - all rights reserved by [lispers.net](#) LLC

Features/Bugs go to [support@lispers.net](mailto:support@lispers.net)

# ETR Sending Map-Registers

```
03/25/19 23:07:32.419: etr: EID-prefix [0]240.1.1.1/32 for ms-name 'all', decent-index 3
03/25/19 23:07:32.419: etr: EID-prefix [0]1.1.1.1/32 for ms-name 'all', decent-index 1
03/25/19 23:07:32.420: etr: EID-prefix [0]2.2.2.2/32 for ms-name 'all', decent-index 0
03/25/19 23:07:32.420: etr: EID-prefix [0]2001::1/128 for ms-name 'all', decent-index 2
03/25/19 23:07:32.420: etr: EID-prefix [0]2001::2/128 for ms-name 'all', decent-index 3
03/25/19 23:07:32.421: etr: EID-prefix [0]'dino' for ms-name 'all', decent-index 1
03/25/19 23:07:32.421: etr: EID-prefix [0](1.1.1.1/32, 224.1.1.1/32) for ms-name 'all', decent-index 1
03/25/19 23:07:32.421: etr: EID-prefix [0]48-51-12-N-2-20-55-E/100 for ms-name 'all', decent-index 3
03/25/19 23:07:32.422: etr: EID-prefix [0]+1223344556677880/64 for ms-name 'all', decent-index 3
03/25/19 23:07:32.422: etr: EID-prefix [0]+000140871820010/64 for ms-name 'all', decent-index 0
03/25/19 23:07:32.422: etr: EID-prefix [0]+0000000000000000/8 for ms-name 'all', decent-index 3
03/25/19 23:07:32.422: etr: Send Map-Register to map-server [0]2.2.2.2, ms-name 'all', decent-index 0.ms.lispers.net
03/25/19 23:07:32.425: etr: Send Map-Register to map-server [0]1.1.1.1, ms-name 'all', decent-index 0.ms.lispers.net
03/25/19 23:07:32.427: etr: Send Map-Register to map-server [0]4.4.4.4, ms-name 'all', decent-index 1.ms.lispers.net
03/25/19 23:07:32.429: etr: Send Map-Register to map-server [0]8.8.8.8, ms-name 'all', decent-index 3.ms.lispers.net
03/25/19 23:07:32.431: etr: Send Map-Register to map-server [0]6.6.6.6, ms-name 'all', decent-index 2.ms.lispers.net
03/25/19 23:07:32.432: etr: Send Map-Register to map-server [0]7.7.7.7, ms-name 'all', decent-index 3.ms.lispers.net
03/25/19 23:07:32.434: etr: Send Map-Register to map-server [0]3.3.3.3, ms-name 'all', decent-index 1.ms.lispers.net
03/25/19 23:07:32.436: etr: Send Map-Register to map-server [0]5.5.5.5, ms-name 'all', decent-index 2.ms.lispers.net
```

# ITR Sending Map-Requests

```
03/25/19 23:30:54.245: itr: Use LISP-Decent map-resolver 1.ms.lispers.net for EID [0]1.1.1.1/32
03/25/19 23:30:54.245: itr: Send 82 bytes to ::ffff:4.4.4.4 4342, packet: 80000000 4500004e 00000000 80114699 f0020202 01010101 e9f510f6 003a0000 10000
101 8e090d54 88990ab9 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00010101 0101
03/25/19 23:31:04.264: itr: Use LISP-Decent map-resolver 0.ms.lispers.net for EID [0]2.2.2.2/32
03/25/19 23:31:04.264: itr: Send 82 bytes to ::ffff:1.1.1.1 4342, packet: 80000000 4500004e 00000000 80114497 f0020202 02020202 e9f510f6 003a0000 10000
101 c9a7dd71 aac2a3cc 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00010202 0202
03/25/19 23:31:14.286: itr: Use LISP-Decent map-resolver 3.ms.lispers.net for EID [0]3.3.3.3/32
03/25/19 23:31:14.286: itr: Send 82 bytes to ::ffff:7.7.7.7 4342, packet: 80000000 4500004e 00000000 80114295 f0020202 03030303 e9f510f6 003a0000 10000
101 a686b807 8323e29e 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00010303 0303
03/25/19 23:31:24.304: itr: Use LISP-Decent map-resolver 0.ms.lispers.net for EID [0]4.4.4.4/32
03/25/19 23:31:24.304: itr: Send 82 bytes to ::ffff:2.2.2.2 4342, packet: 80000000 4500004e 00000000 80114093 f0020202 04040404 e9f510f6 003a0000 10000
101 aa2de2e2 82e346db 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00010404 0404
03/25/19 23:31:34.314: itr: Use LISP-Decent map-resolver 2.ms.lispers.net for EID [0]44.44.44.44/32
03/25/19 23:31:34.314: itr: Send 82 bytes to ::ffff:6.6.6.6 4342, packet: 80000000 4500004e 00000000 8011f042 f0020202 2c2c2c2c e9f510f6 003a0000 10000
101 f38c5956 793007c9 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00012c2c 2c2c
03/25/19 23:31:44.223: itr: Use LISP-Decent map-resolver 2.ms.lispers.net for EID [0]128.1.1.1/32
03/25/19 23:31:44.223: itr: Send 82 bytes to ::ffff:5.5.5.5 4342, packet: 80000000 4500004e 00000000 8011c798 f0020202 80010101 e9f510f6 003a0000 10000
101 b0e4ff28 d700b86c 0001f002 02020001 ac110007 0002fe00 00000000 00000000 00020002 00020020 00018001 0101
03/25/19 23:31:54.245: itr: Use LISP-Decent map-resolver 2.ms.lispers.net for EID [0]2001::1/128
03/25/19 23:31:54.245: itr: Send 126 bytes to ::ffff:6.6.6.6 4342, packet: 80000000 60000000 00521180 fe000000 00000000 00000002 00020002 20010000 0000
0000 00000000 00000001 e9f510f6 00520000 10000101 a61f6610 88e7992d 0002fe00 00000000 00000000 00020002 00020001 ac110007 0002fe00 00000000 00000000 00
020002 00020080 00022001 00000000 00000000 00000000 0001
03/25/19 23:32:04.295: itr: Use LISP-Decent map-resolver 3.ms.lispers.net for EID [0]ff02::1:ff00:2/128
03/25/19 23:32:04.295: itr: Send 146 bytes to ::ffff:8.8.8.8 4342, packet: 80000000 60000000 00661180 fe000000 00000000 00000002 00020002 ff020000 0000
0000 00000001 ff000002 e9f510f6 00660000 10000101 63f676ed b3be8417 00000001 ac110007 0002fe00 00000000 00000000 00020002 00020080 40030000 0900002c 00
000000 00000000 80800002 fe000000 00000000 00000002 00020002 0002ff02 00000000 00000000 0001ff00 0002
03/25/19 23:32:17.395: itr: Use LISP-Decent map-resolver 0.ms.lispers.net for EID [0]dfdf::2/128
03/25/19 23:32:17.396: itr: Send 126 bytes to ::ffff:1.1.1.1 4342, packet: 80000000 60000000 00521180 fe000000 00000000 00000002 00020002 dfdf0000 0000
0000 00000000 00000002 e9f510f6 00520000 10000101 827db308 c22dcff0 0002fe00 00000000 00000000 00020002 00020001 ac110007 0002fe00 00000000 00000000 00
020002 00020080 0002dfdf 00000000 00000000 00000000 0002
```

# Questions/Reactions/Tomatoes?

