

A Decentralized Mapping System

draft-farinacci-lisp-decent-00

IETF London

March 2018

Dino Farinacci & Colin Cantrell

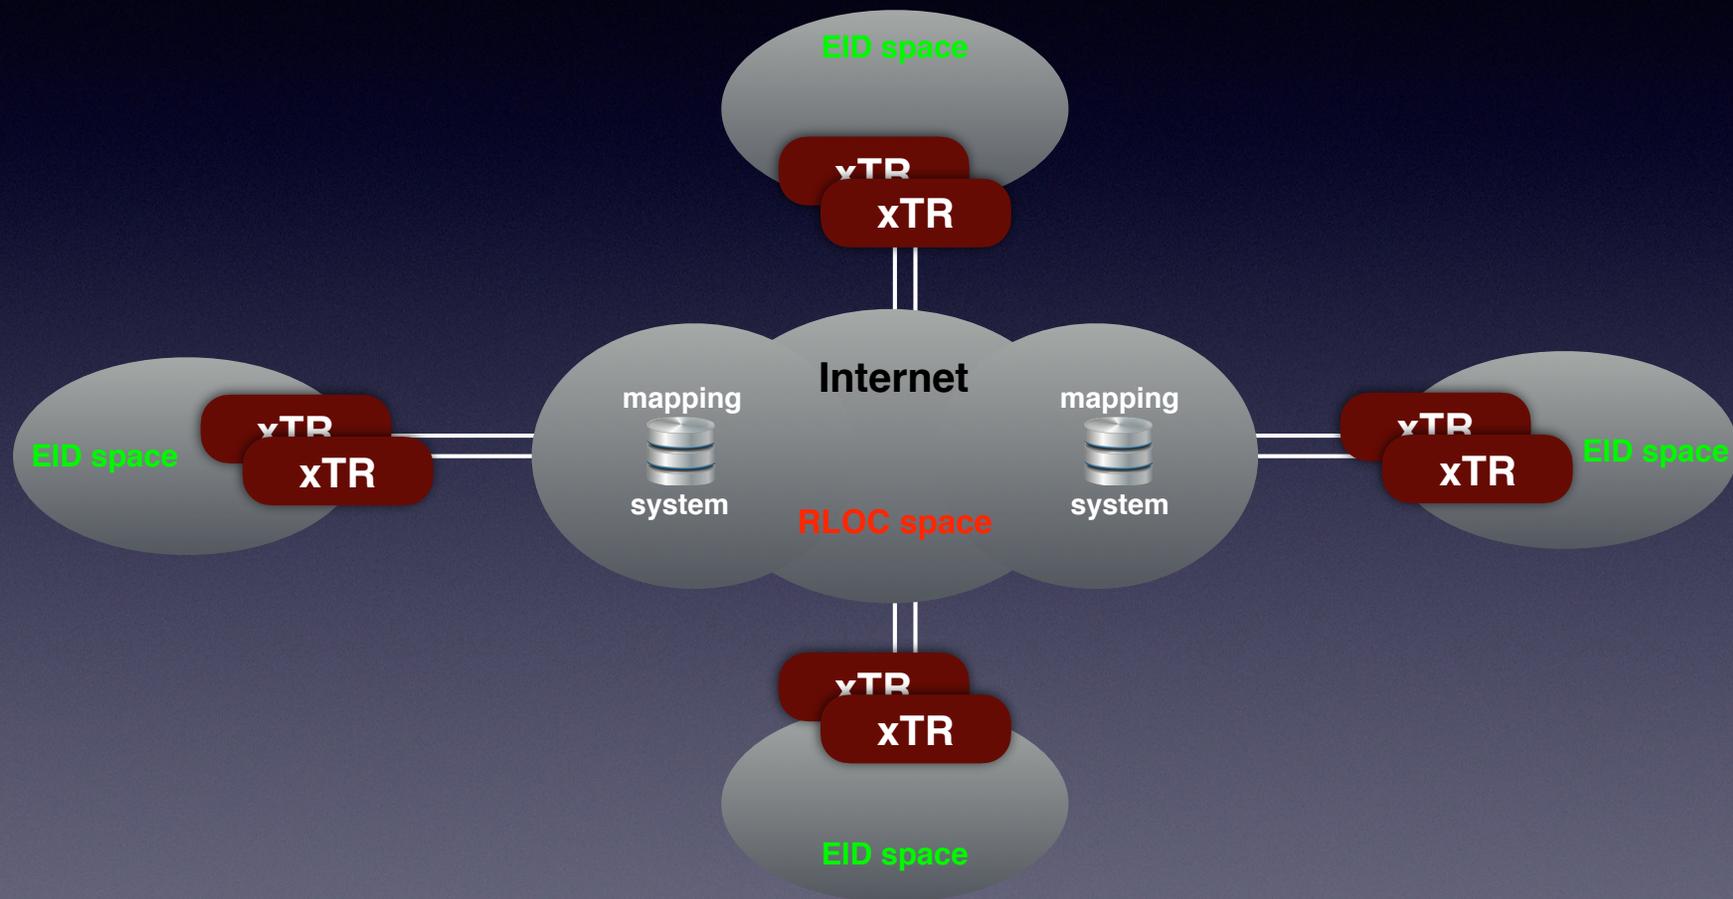
Problem Statement

- 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

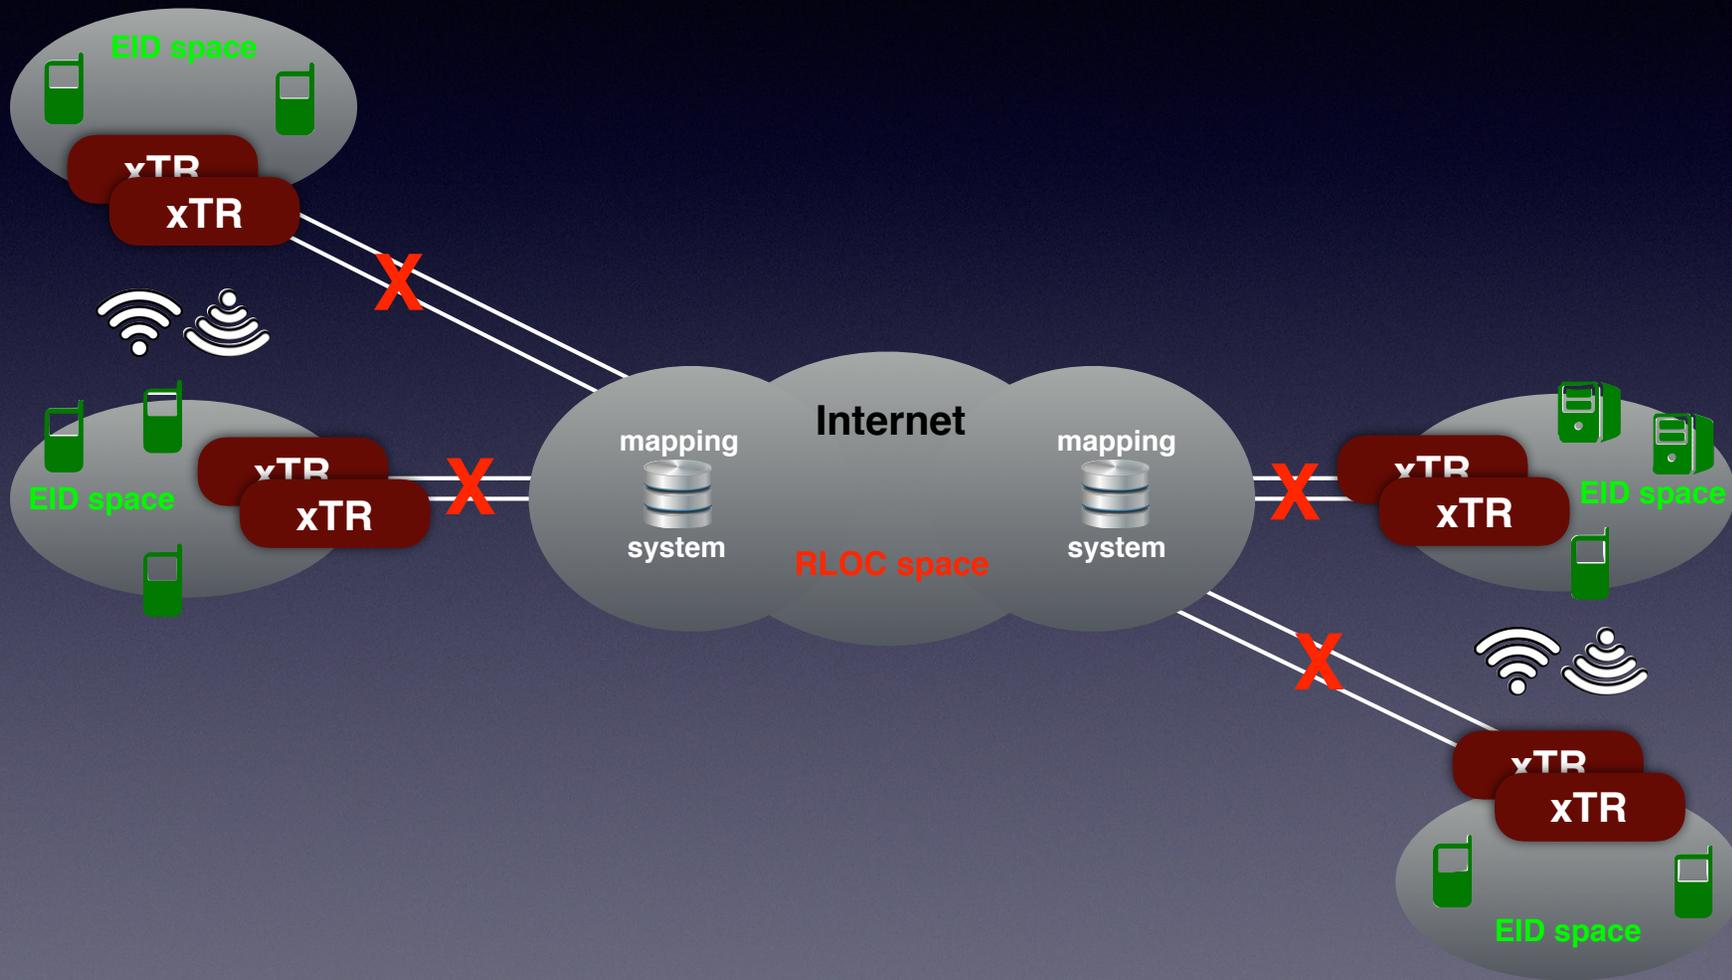
Endpoint IDs (EIDs)

Routing Locators (RLOCs)

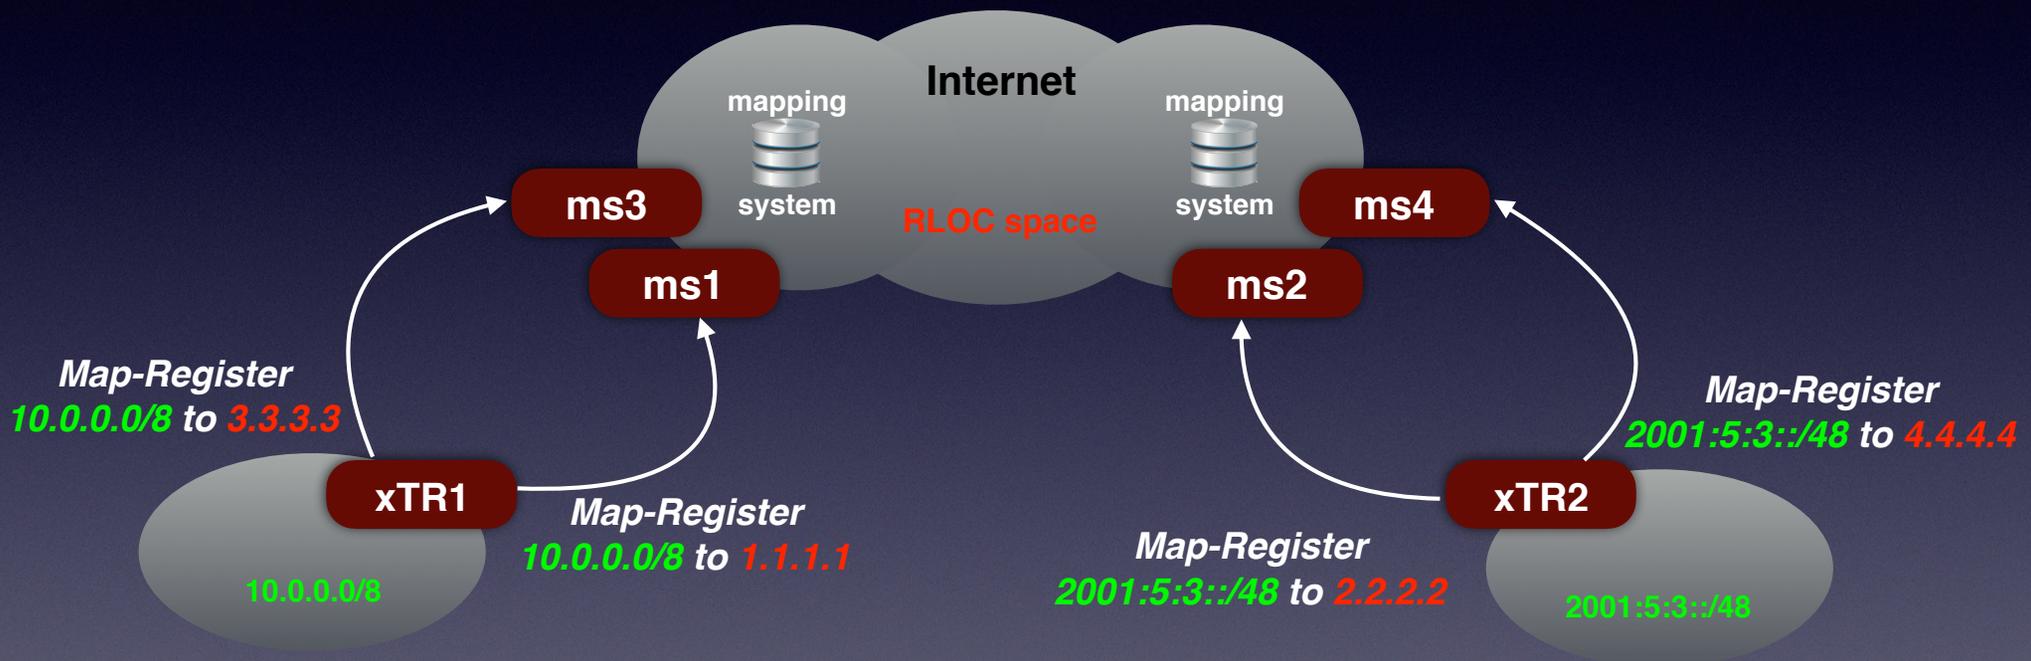
Today's Model Mapping System



Network Connectivity



LISP Control-Plane Messages



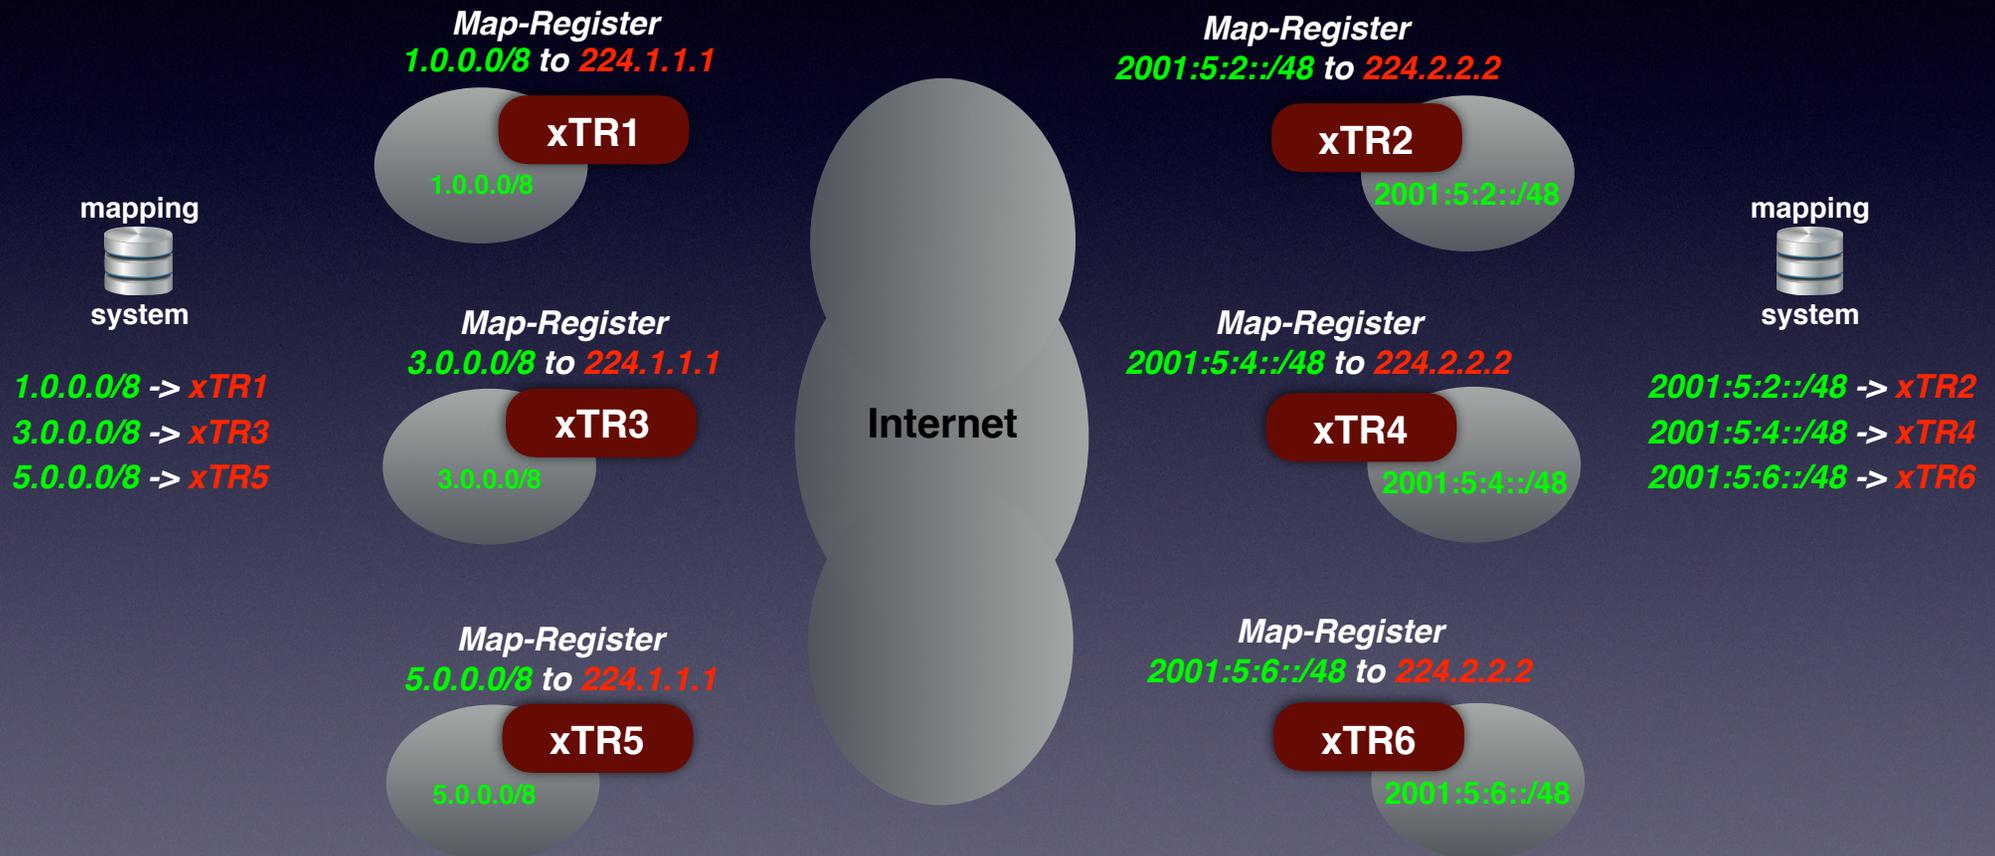
Decentralized Map-Server?

- 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

How to Define a 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

LISP-Decent Control-Plane Messages



Benefits

- xTRs only depend on each other - they do so already if they want to talk to each other
- No third-party trust or dependency exists
- 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
- Pretty much same benefits as peer-to-peer networking

Use-Cases

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

Brief LISP-Decent 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-M

lispers.net

Scalable Open Overlay Networking

n2

Enter EID for Site-Cache lookup:

LISP-MS Site Info

Site Name	EID-P
EIDs	[1]
	[1]2.2
	[1]'ne
	[1]1.1
	[1]'ne
	[1]3.3
	[1]'ne
peer-groups	[1](0.0
	[1](0.0

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

LISP-Decent in Action

lispers.net

Scalable Open Overlay Networking

n1

Site name: peer-groups, EID-prefix: [1](0.0.0.0/0, 224.1.1.1/32) registered: **yes**, dynamic
Description:
Last registerer: [0]172.17.0.6, xTR-ID: 0xab64c2b4849e579d, site-ID: 0
First registered: 0:10:39, last registered: 0:00:01, auth-type: sha1, registration flags: P-s-I-t-R-m-n
Default registration timeout TTL: 180 seconds
Forcing proxy Map-Reply: yes
Forcing proxy Map-Reply for xTRs behind NATs: no
Send drop-action proxy Map-Reply to PITR: no
Proxy Map-Reply action: not configured
Allowed RLOC-set: any

Registered RLOC-set (merge-semantics):
[0]no-address, state: up-state, up/uw/mp/mw: 255/0/1/100
rle: 172.17.0.5 (L128), 172.17.0.6 (L128), 172.17.0.7 (L128)

LISP-ITR Map-Cache:

Enter EID for Map-Cache lookup:

EID-Prefix or (S,G)	Uptime TTL	RLOC Record RLOC Keys	Map-Reply Source	RLOC Send Stats	RLOC State RLOC Action	Unicast Priority/Weight Multicast Priority/Weight
[1](0.0.0.0/0, 224.1.1.1/32)	0:05:05 24 hours	rle: 172.17.0.5(L128), 172.17.0.7(L128), 172.17.0.6(L128)	map-notify	packet-count: 35 byte-count: 5315	up-state since 0:05:05 encapsulate	255/0 1/100

Questions/Reactions/Tomatoes?

