

# A History of Multicast Routing that brings us to today's State of Affairs in Multicast Overlays

*LISP Working Group - Yokohama IETF  
October 2015*

*Dino Farinacci  
farinacci@gmail.com*

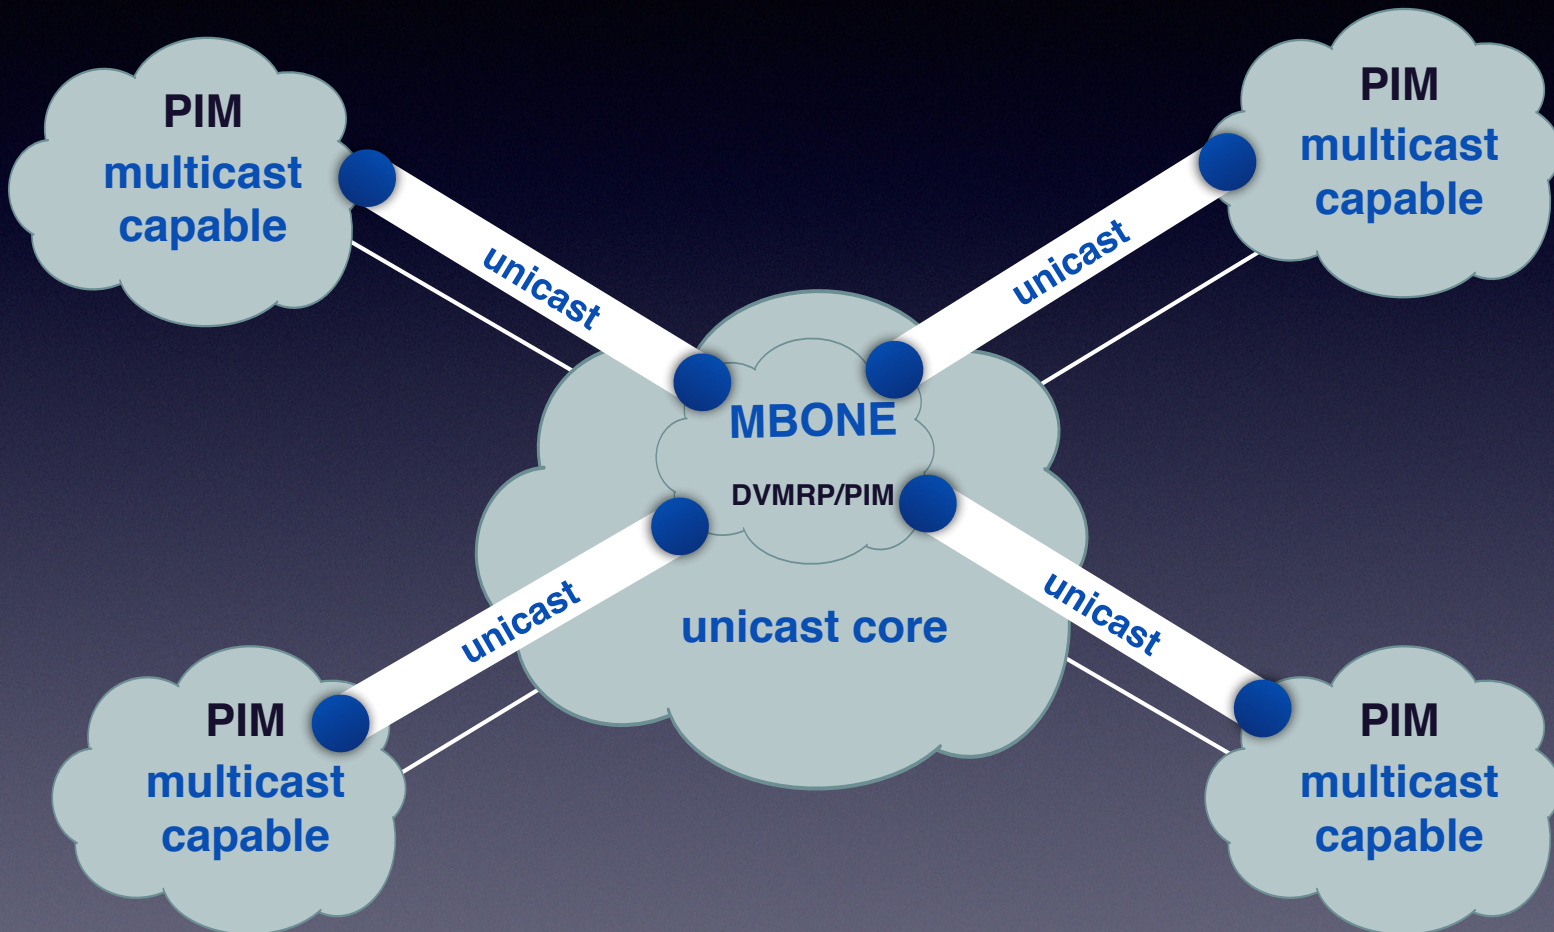


# Agenda - Multicast Delivery Models

- Multicast-over-Unicast (MBONE)
- Multicast Native
- Virtualizing Multicast - MVPNs
- Multicast-over-Unicast (AMT)
- Multicast Map-and-Encap Overlays

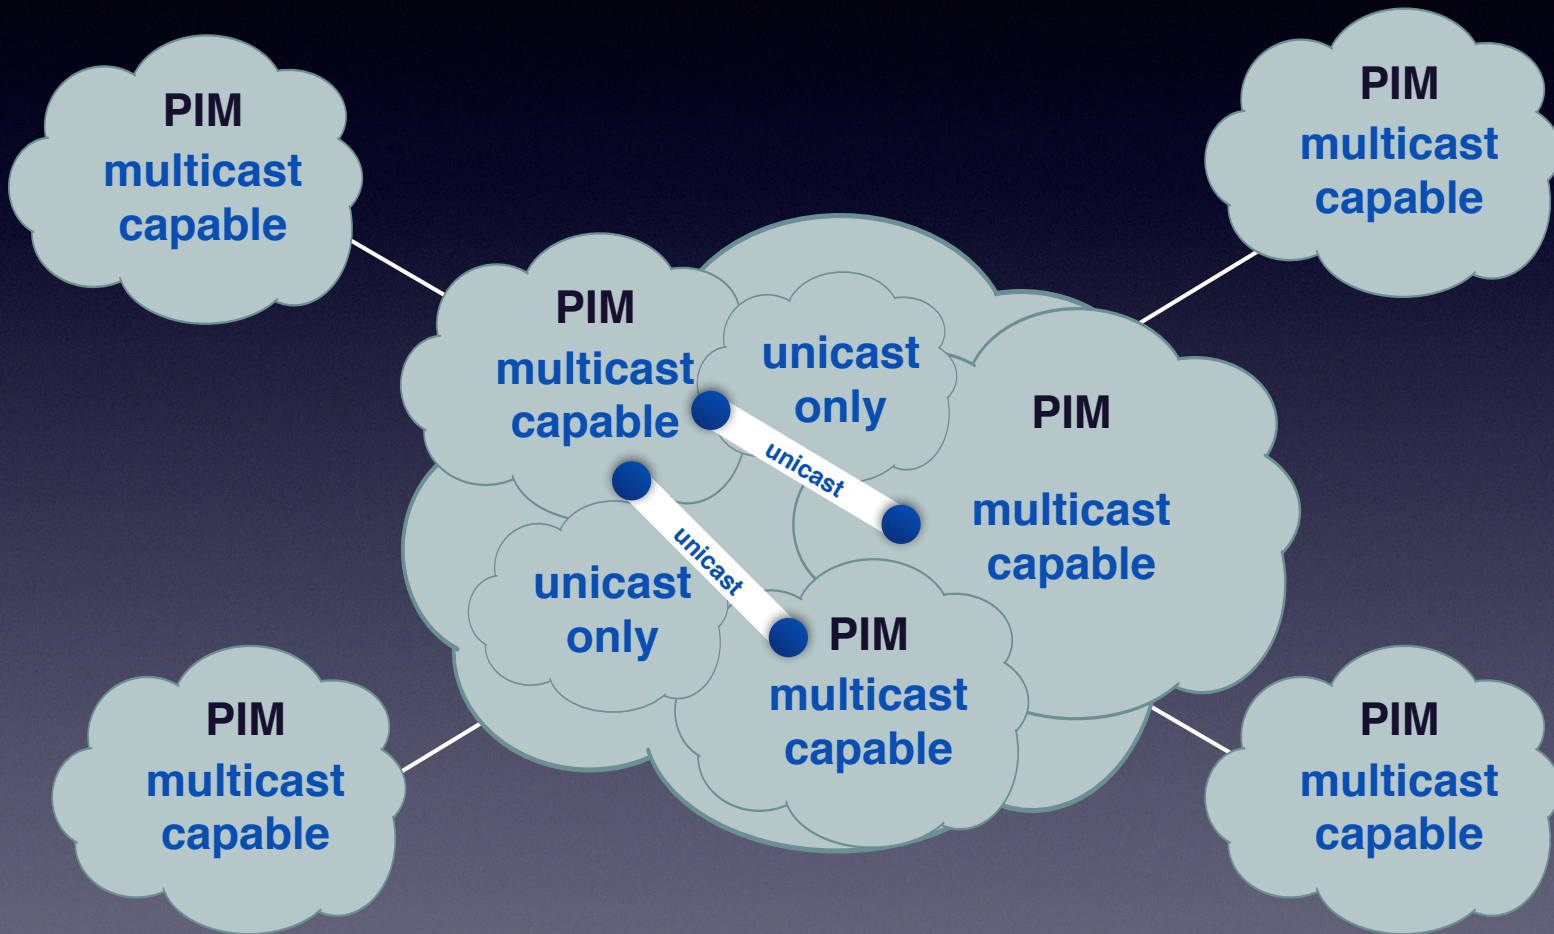


# 1995



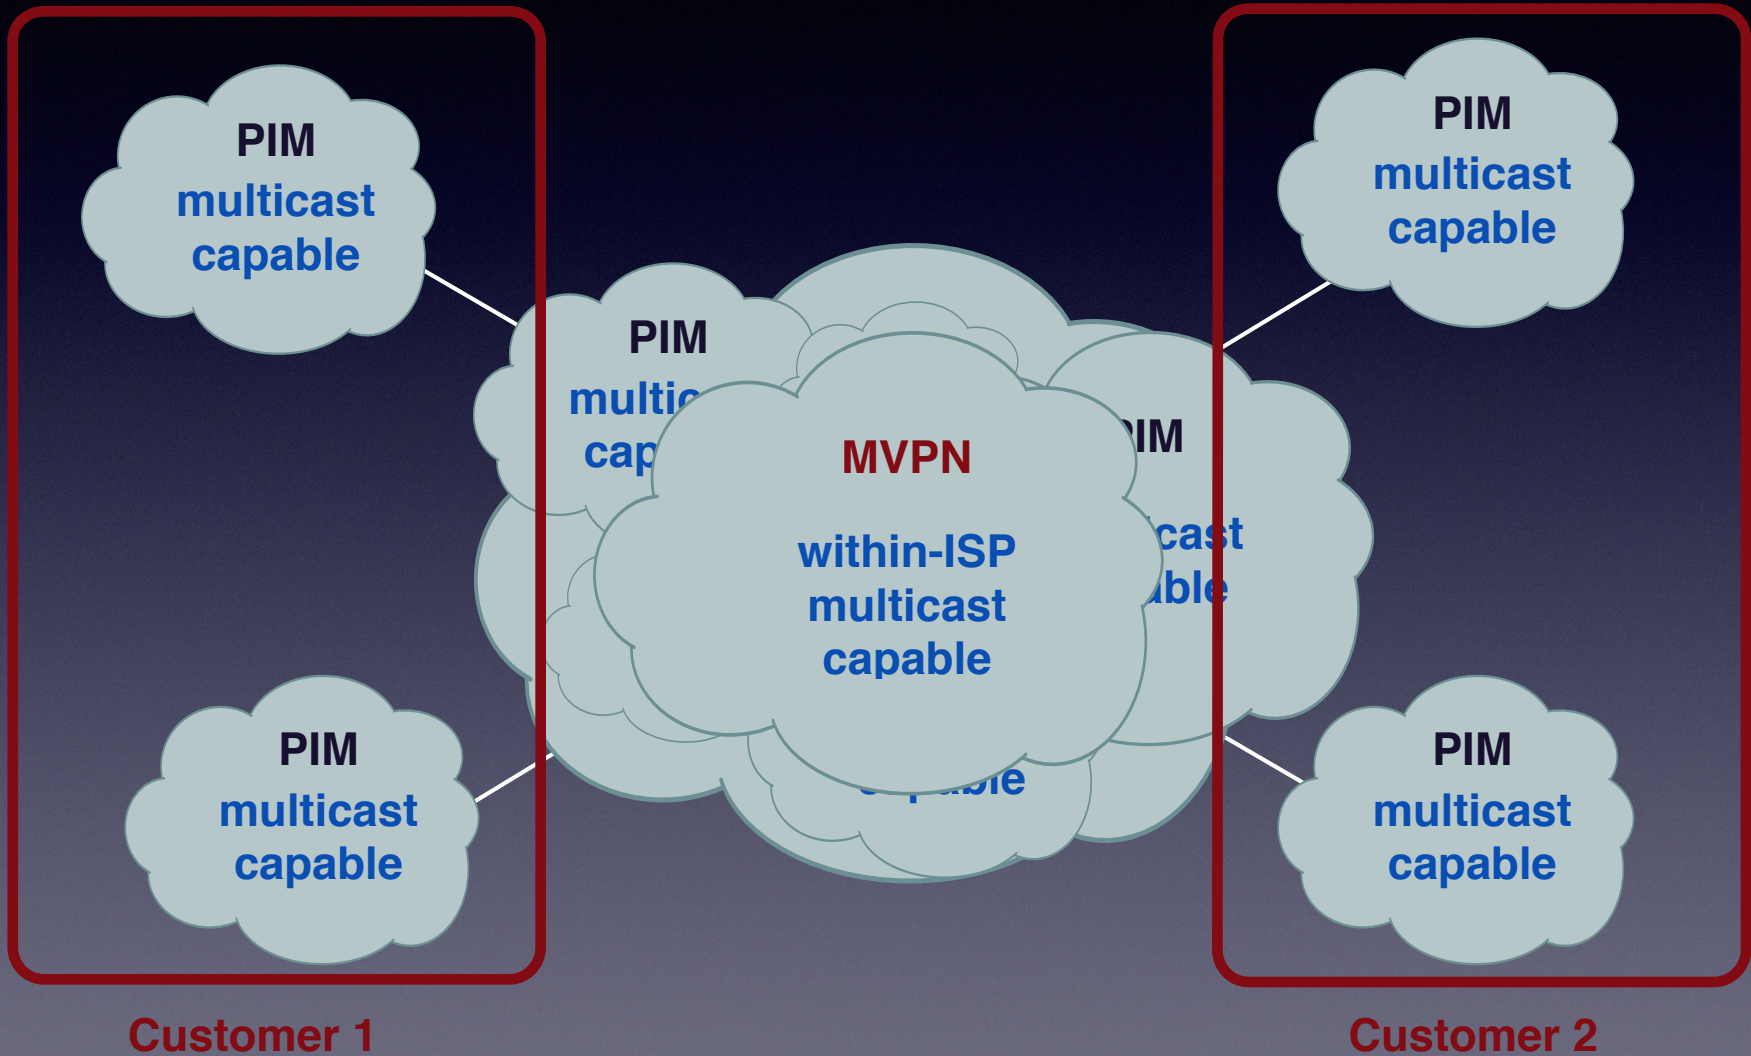


# Native Arrived (kind of)



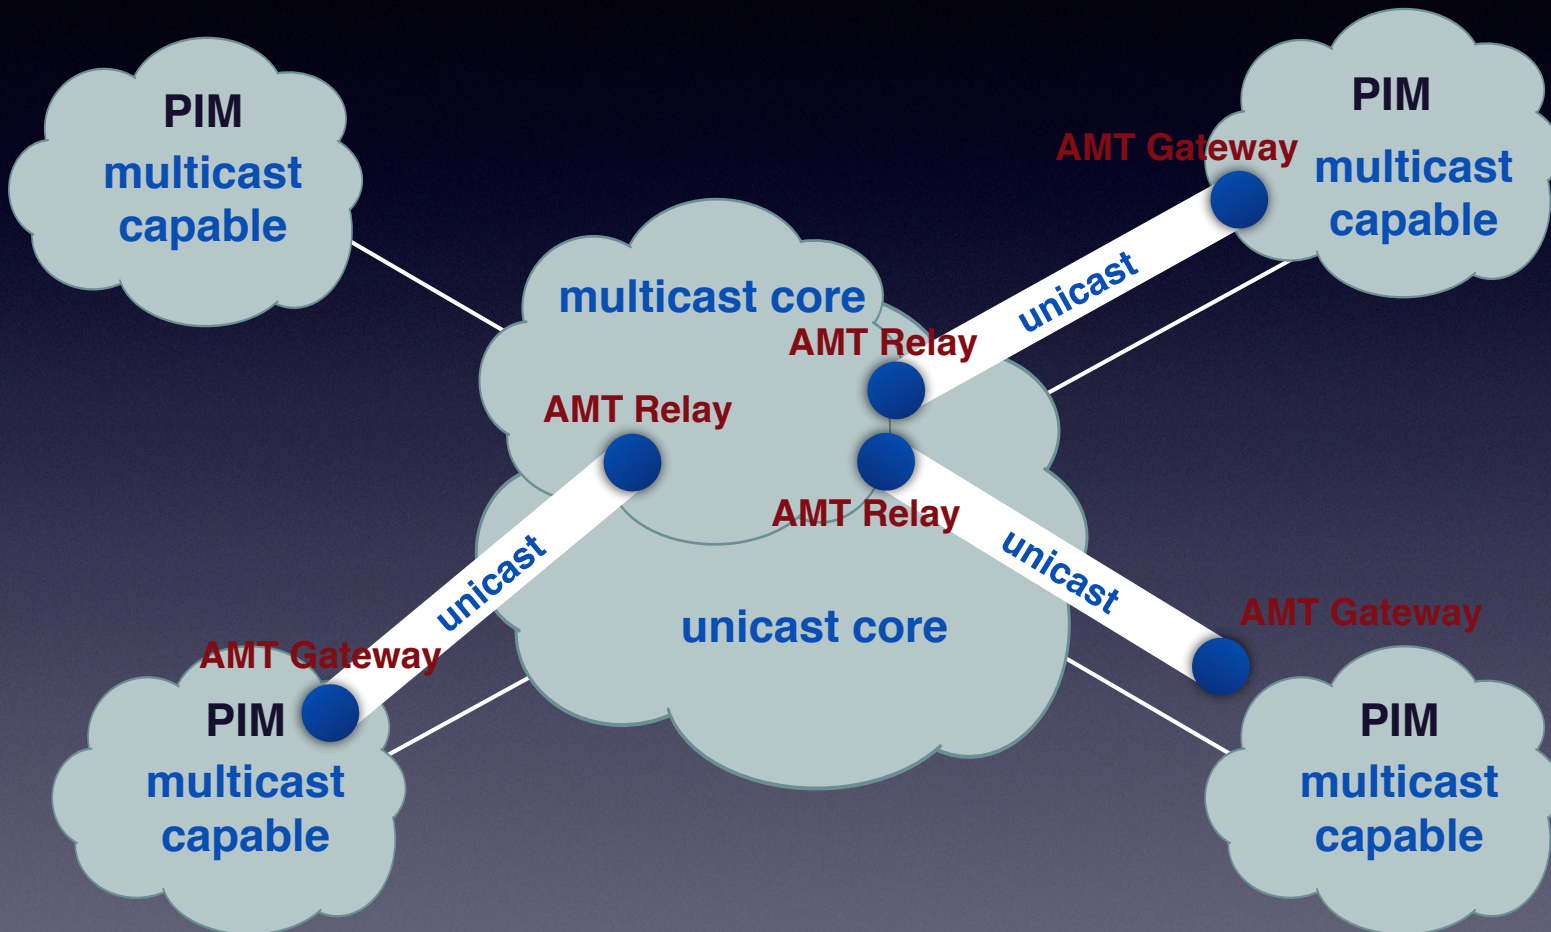


# ISP Multicast Service



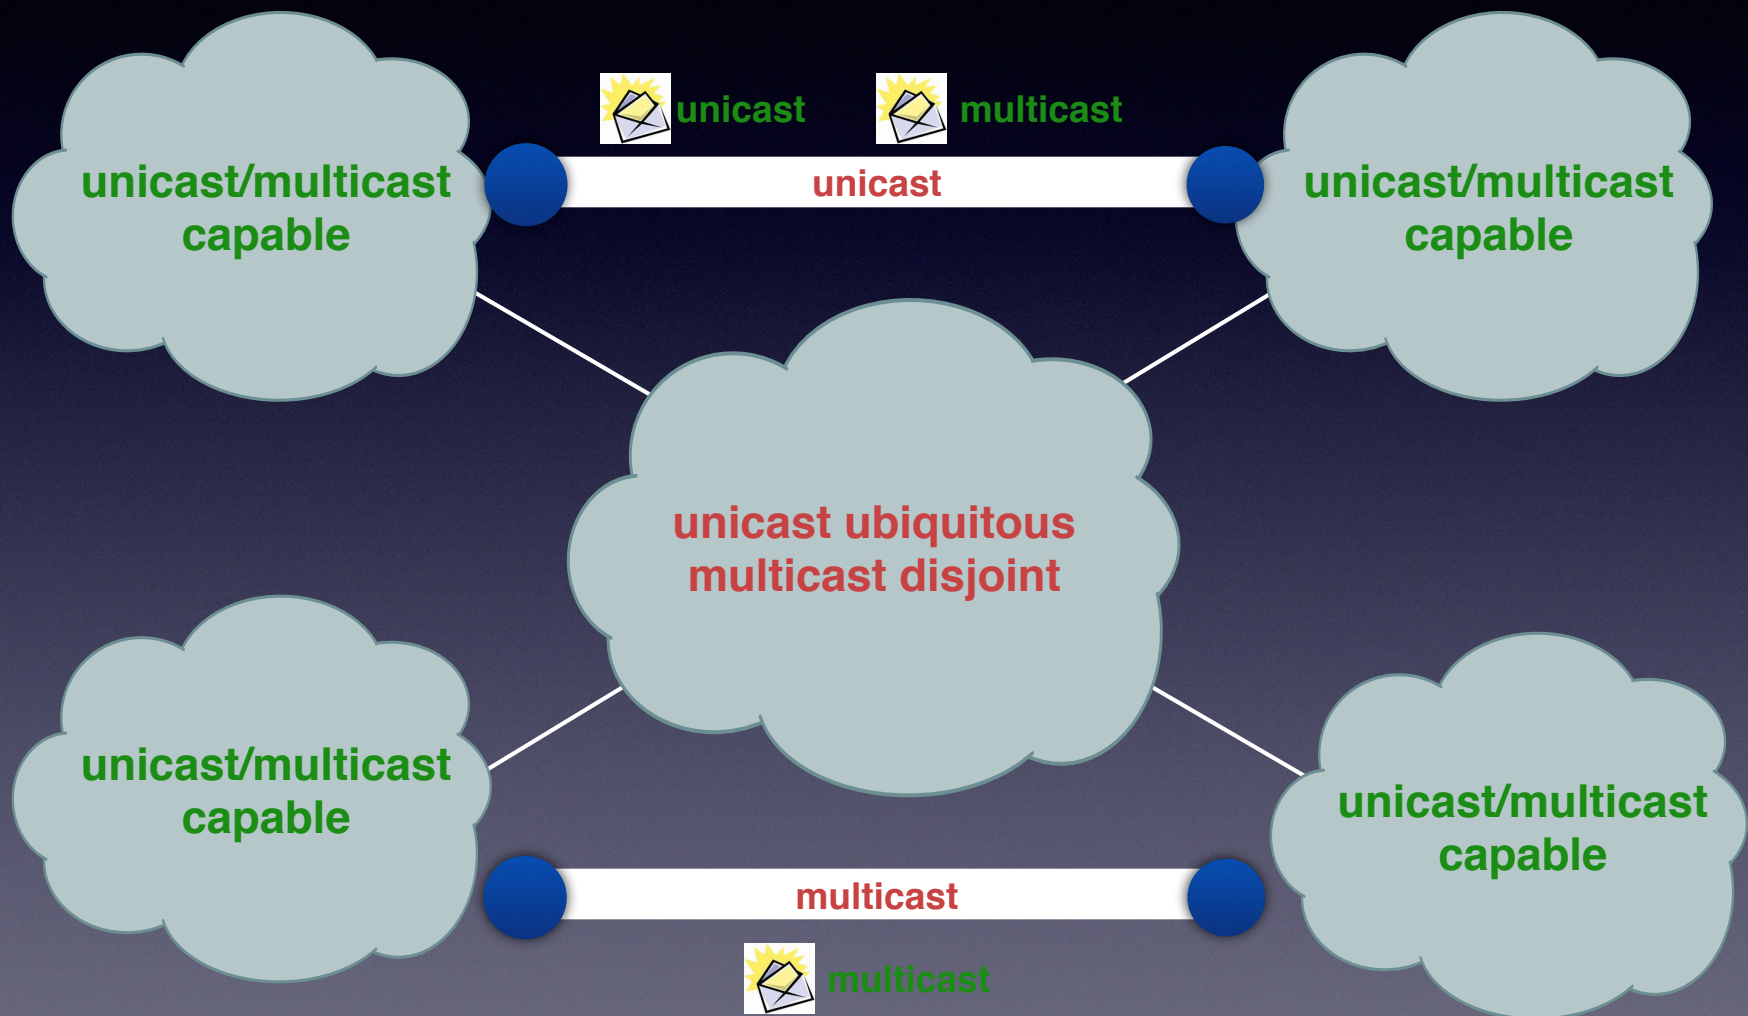


# We Wanted Multicast Anywhere





# Now We Have Overlays



EIDs -> green

RLOCs -> red




# LISP-Multicast Today

- RFC 6831 - *"The Locator/ID Separation Protocol (LISP) for Multicast Environments"*
  - Defines how to encaps multicast into multicast or unicast
  - Defines use of unicast PIM J/P message exchange between ETRs and ITRs
  - Defines how to work with native PIM at source and receiver multicast sites
  - Enumerates various combinations and recommends how to avoid combinatoric nightmares



# Core Supports Native Multicast



LISP xTRs

Hosts in EID space

Routers in EID space

(EID, G) tree state

Link with no tree state

(RLOC, DG) tree state

Replicating router

Routers in RLOC space




# Multicast Overlay Signaling Mechanisms

- **In-the-Network Signaling** - Mechanisms of Today
  - RFC 6831 - LISP-Multicast (**PIM** for signaling)
  - *draft-farinacci-lisp-mr-signaling* (**LISP** for signaling)
- **Out-of-the-Network Signaling** - Next-Gen Mechanisms
  - Mapping Database Based
    - *draft-farinacci-lisp-signal-free-multicast*
    - *draft-coras-lisp-re* & *draft-ietf-lisp-lcaf*
  - Programmable Interfaces
    - i2rs, OpenFlow, RESTful




# PIM Control-Plane Everywhere





# LISP as Control-Plane





# Next-Gen LISP-Multicast

- Eliminate the need for PIM over-the-top
  - Less protocols mean lower OpEx and less complexity
- Use the existing mapping system for ETRs to find the ITRs of source multicast sites
- At the same time allow for encap of multicast into unicast
  - To allow multicast service over partner unicast-only network



# Out-of-the-Network Signaling

- Use the Mapping Database
  - Replication list of ETRs or DGs are stored per (S-prefix, G-prefix) EID entry
- Use a Programmable Interface
  - Have network controller monitor ETRs for joined state
  - Then network controller programs ITRs with replication state
  - Network controllers can program RTRs inside of network to optimize distribution trees



# Using the Mapping Database

(S-EID, G)  
encoding

Multicast Info Canonical Address Format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = 16387                               | Rsvd1 | Flags |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type = 9 | Rsvd2 | R|L|J |                               4 + n |
+-----+-----+-----+-----+-----+-----+-----+
|                               Reserved                               | Source MaskLen | Group MaskLen |
+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = x                               | Source/Subnet Address ... |
+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = x                               | Group Address ... |
+-----+-----+-----+-----+-----+-----+-----+

```

Replication List Entry Address Format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = 16387                               | Rsvd1 | Flags |
+-----+-----+-----+-----+-----+-----+-----+
| Type = 13 | Rsvd2 |                               4 + n |
+-----+-----+-----+-----+-----+-----+-----+
|                               Rsvd3                               | Rsvd4 | Level Value |
+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = x                               | RTR/ETR #1 ... |
+-----+-----+-----+-----+-----+-----+-----+
|                               Rsvd3                               | Rsvd4 | Level Value |
+-----+-----+-----+-----+-----+-----+-----+
|                               AFI = x                               | RTR/ETR #n ... |
+-----+-----+-----+-----+-----+-----+-----+

```

RLOC or DG  
encoding



# Mapping Database Based Example


Map-Cache entry in A:

(SEID, G) ->

multicast DG {which B & C joined (RLOC-A, DG)}

unicast RLOC D

unicast RLOC E



unicast   
multicast 



# Programmable Interface Example

Network controller reads from:

RLOC D wants (SEID, G) via unicast

RLOC E wants (SEID, G) via unicast

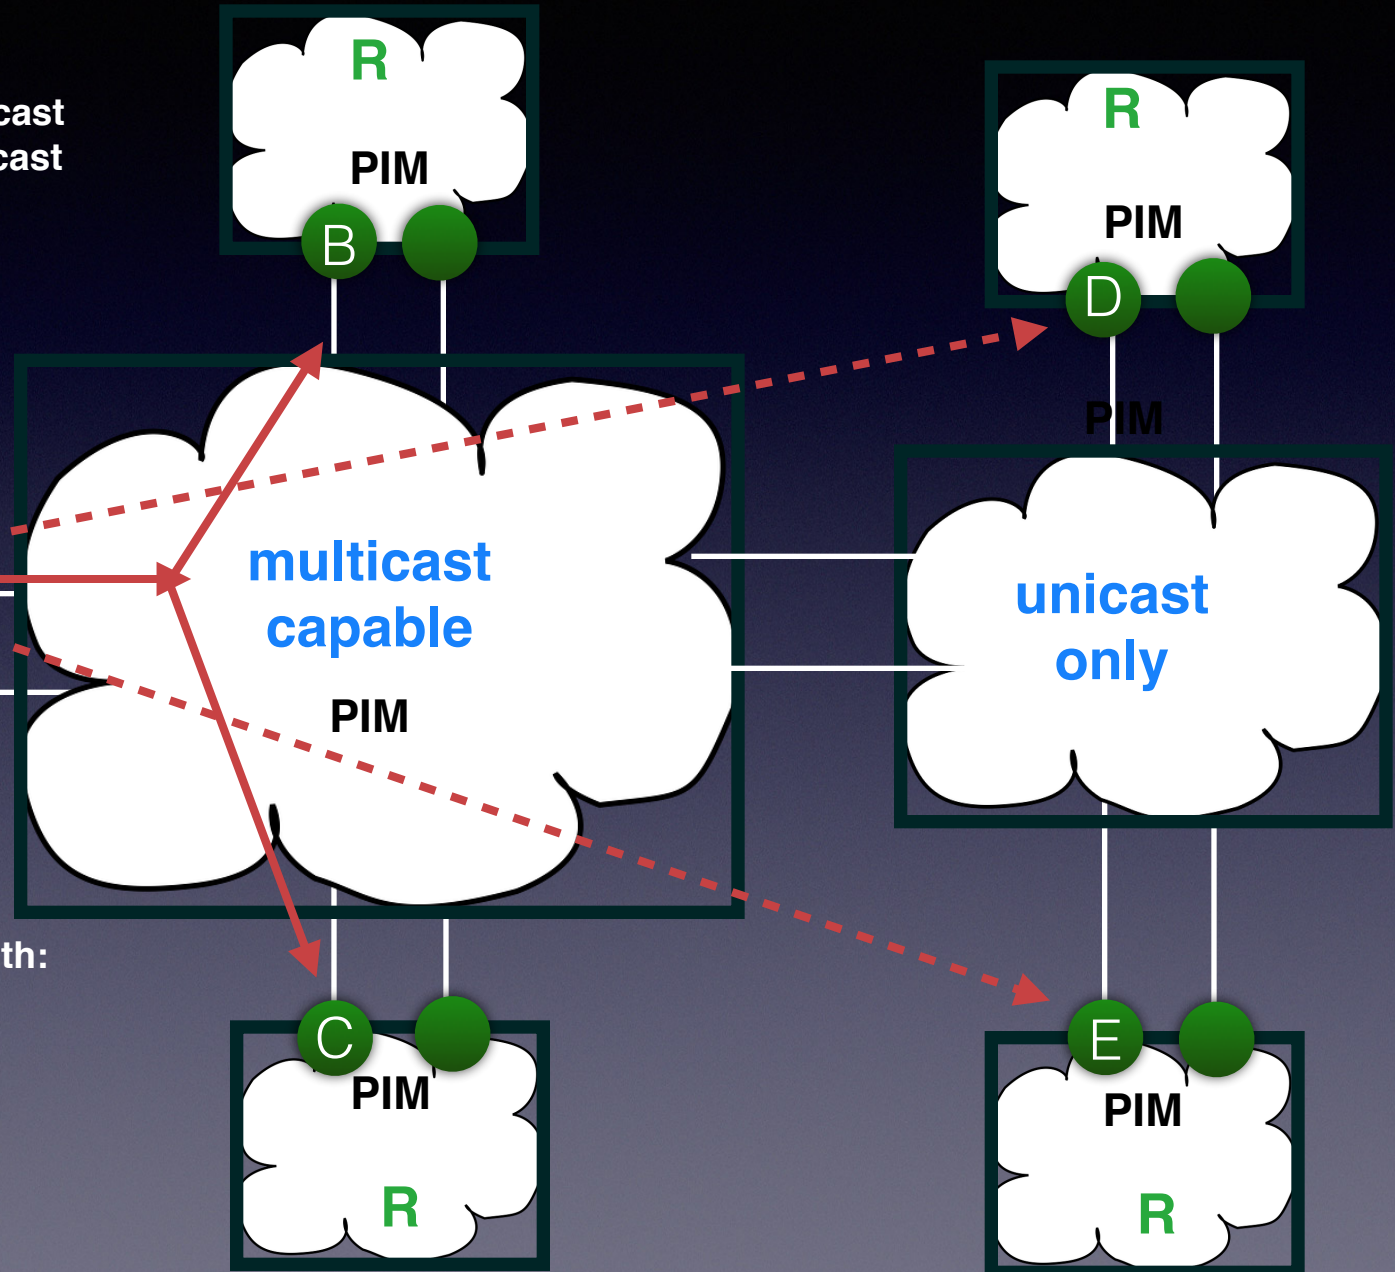
RLOC B wants (SEID, G) via DG

RLOC C wants (SEID, G) via DG

Network controller writes to:

RLOC B to join (RLOC-A, DG)

RLOC C to join (RLOC-A, DG)



Network controller writes to

RLOC A to add (SEID, G) entry with:

multicast DG

unicast RLOC D

unicast RLOC E

unicast   
multicast 



# Recommendations

- RFC 6830 is necessary for underlay-only multicast to transition to overlay-underlay multicast - **keep it active**
- We let *draft-farinacci-lisp-mr-signaling* expire to reduce the number of signaling approaches - **kill it**
- Make *draft-farinacci-lisp-signal-free-multicast* working group draft - **keep it active**
  - Focus future multicast efforts on mapping database based designs for access-control, policy, and replication-engineering (RE)