

From Map-and-Encap to BIER

Observations on Network Routing Scalability

Tianyuan Yu · Lan Wang · Beichuan Zhang · Lixia Zhang

Talk Roadmap

1

The Root Problem

IP's dual role as identifier & locator

2

Unicast Story

Map-and-Encap → LISP vs MPLS

3

Multicast Story

Stateful multicast → BIER

4

The Scaling Principle

Scale with topology, not applications

5

Key Observations & Lessons

What this tells us about future design

The Root Problem: IP's Dual Role

IDENTIFIER

- Stable name for an endpoint
- Embedded in configs, contracts, ACLs
- Independent of provider
- → Must not change

CONFLICT

LOCATOR

- Encodes position in topology
- Assigned by provider
- Enables prefix aggregation
- → Must track attachment point

Rekhter's Law : "Addressing can follow topology, or topology can follow addressing. Choose one."

The Consequence: Unbounded Routing Table Growth

1M+

DFZ routing table entries today

Up from ~10,000s in the early 1990s

What's driving this growth?

✓ Multi-homed end sites wanting provider-independent addresses.

✓ Administrative & Comm

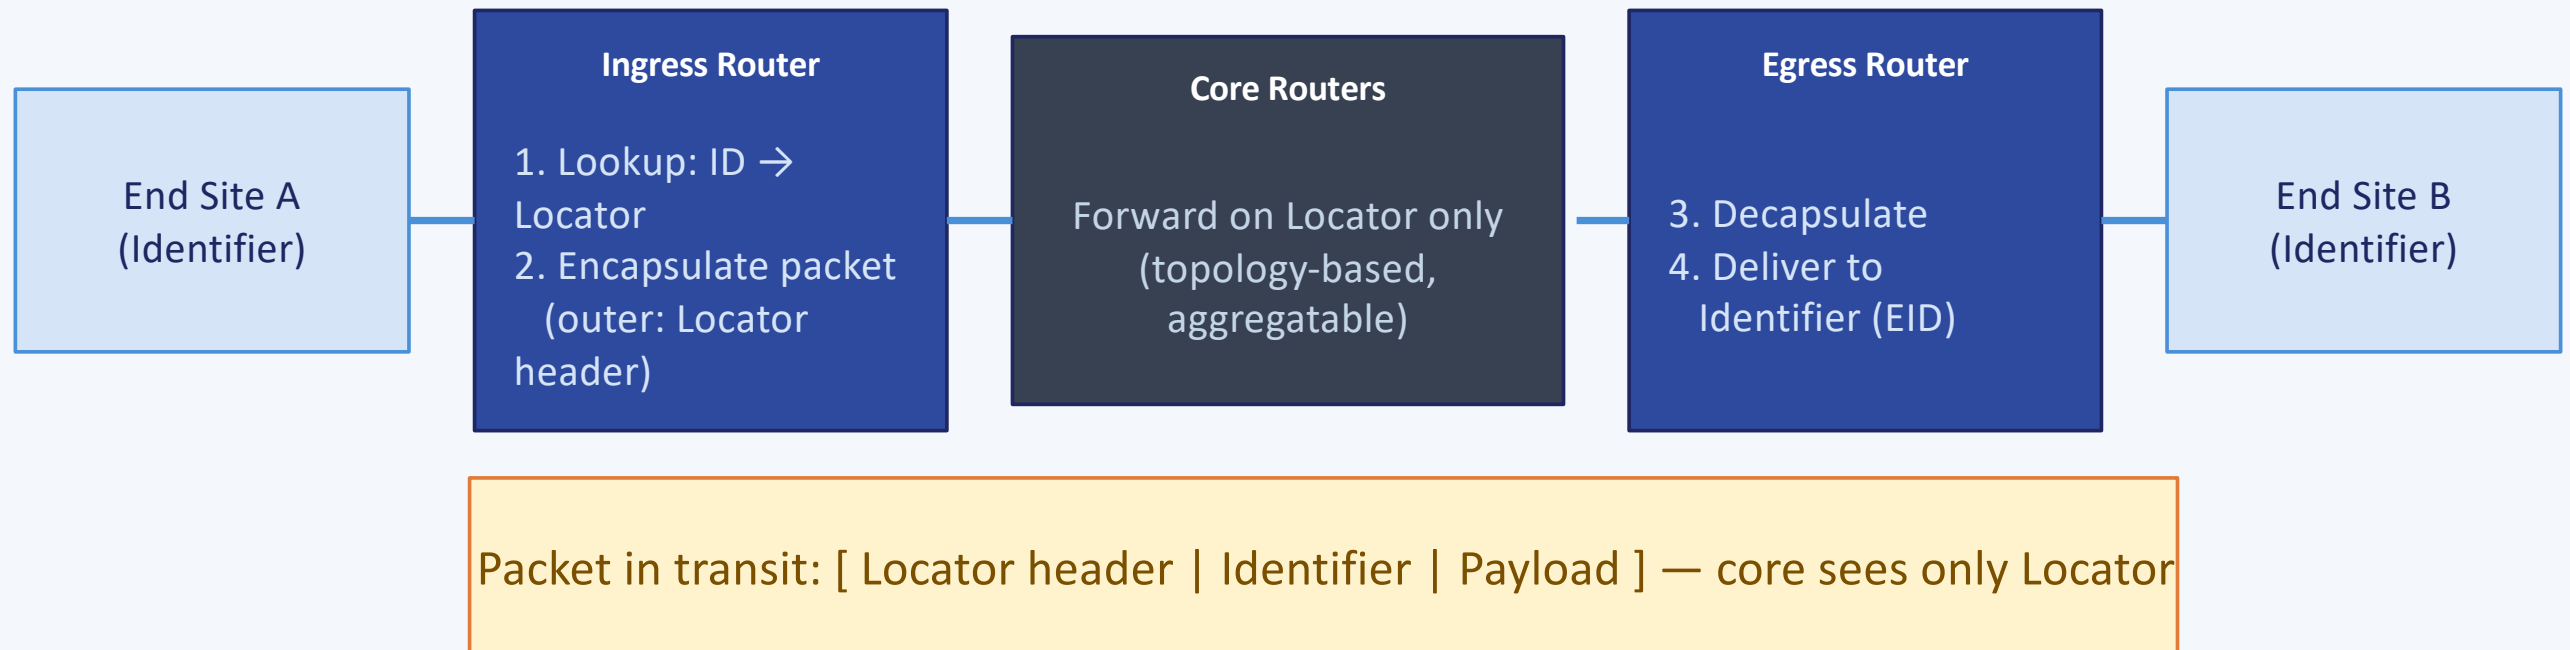
✗ NOT number of routers and links

✗ NOT network topology

Growth is driven by external administrative decisions — not by network topology.

The Architectural Fix: Map-and-Encap

Separate the two namespaces. Map at the boundary. Keep the network core clean. (Chiappa, 1992)



Key benefit: Core routing table bounded by topology (provider networks) — not by number of end sites.

LISP: The Principled Attempt

What LISP Proposed

- Endpoint Identifiers (EIDs) — provider-independent end-user addresses
- Routing Locators (RLOCs) — topological provider addresses
- Delegated Database Tree (DDT) — maps EIDs → RLOCs at network boundary
- Ingress tunnel routers (ITRs) encapsulate; egress (ETRs) decapsulate
- Core forwards on RLOCs only

The Deployment Reality

- Standardized in experimental RFCs 2013, full standard 2022
- Requires a new global mapping system (DDT) with no agreed governance
- All border routers across administrative domains must be modified
- First movers gain nothing until critical mass is reached
- Limited public Internet deployment to this day

Observation 1: A solution requiring global coordination before early adopters see benefits faces a nearly insurmountable deployment barrier.

What Actually is deployed: MPLS

MPLS can be thought of map-and-encap within a single administrative domain



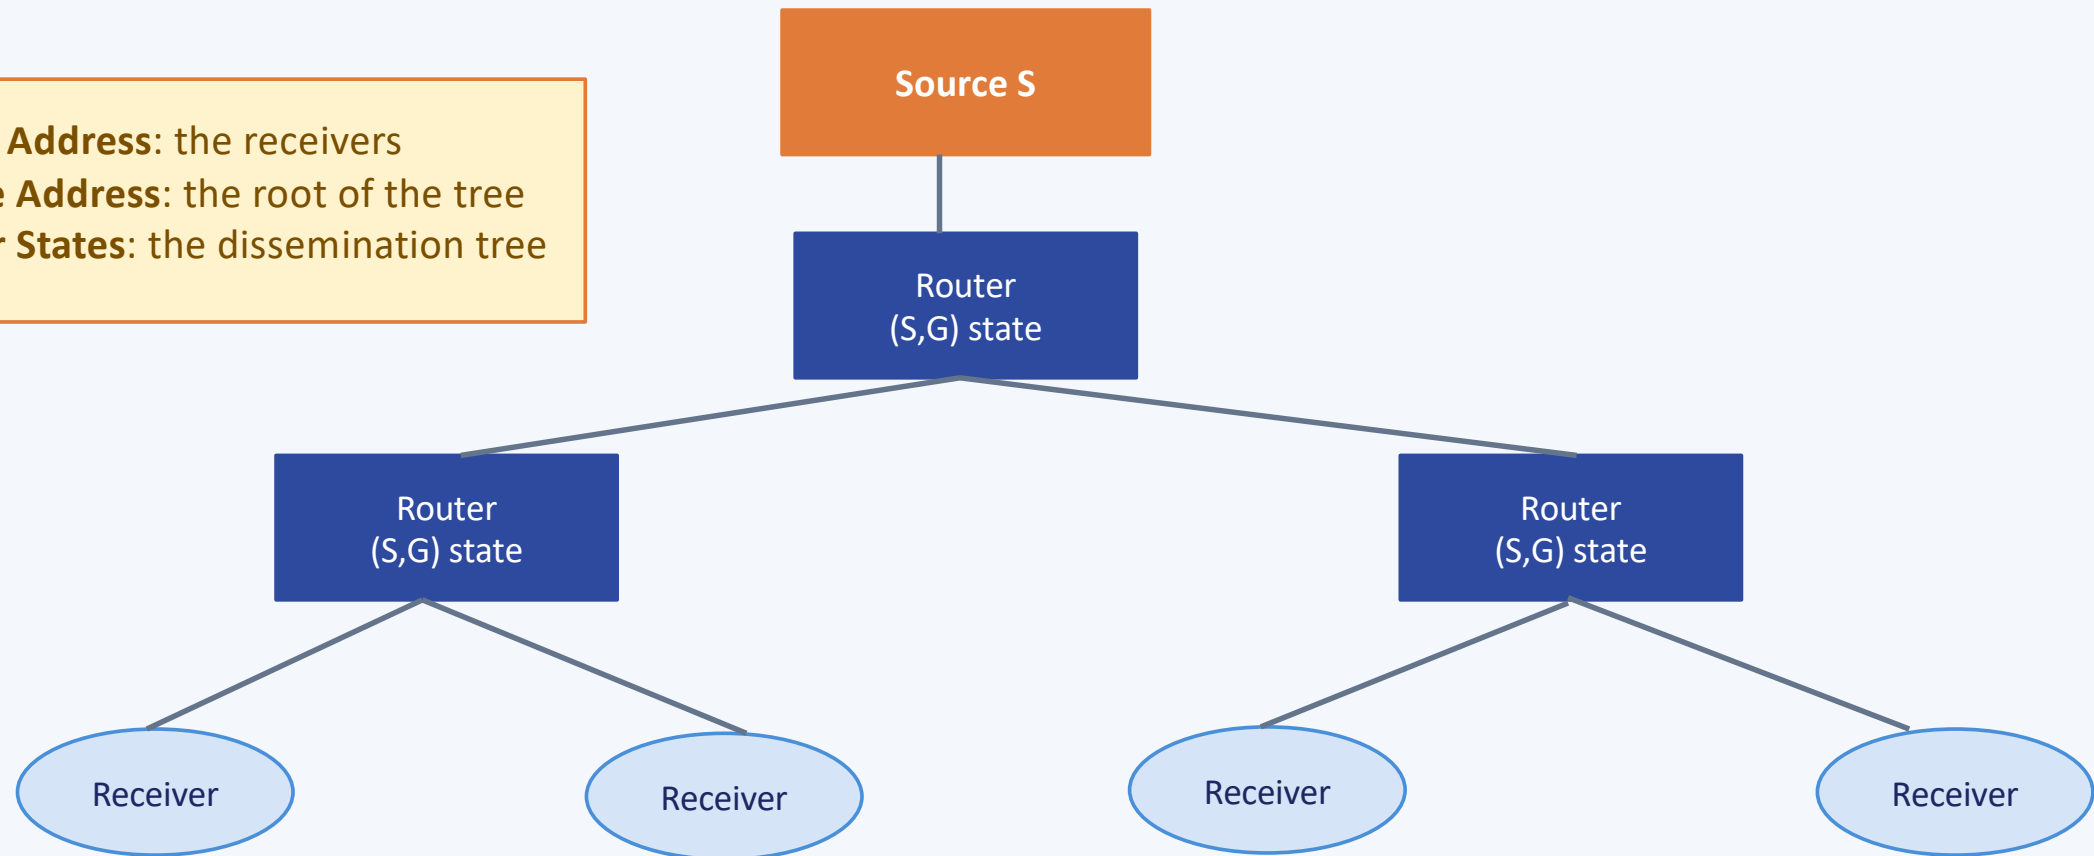
Why MPLS succeeded where LISP didn't:

- ✓ Single operator controls deployment — no cross-domain coordination needed
- ✓ Immediate local gains: core routers do label lookups, not global IP prefix lookups
- ✓ Reuses existing routing protocols (iBGP/IGP) — no new global mapping system required

Unicast ID/Locator conflict remains unresolved globally — MPLS manages symptoms within a domain.

IP Multicast: Per-Group State Explosion

Group Address: the receivers
Source Address: the root of the tree
Router States: the dissemination tree



Every router's state grows by $\# \text{ groups} \times \# \text{ sources}$ — which are outside operator control.

Multicast: The Problem becomes much worse

A group address is a purely logical identifier — it tells you nothing about where receivers are.

Unicast Problem

- IP address = ID + Locator
- Growth driven by # of end sites
- Administrative decisions (provider choice, multihoming)
- Hardware advances provided a workaround

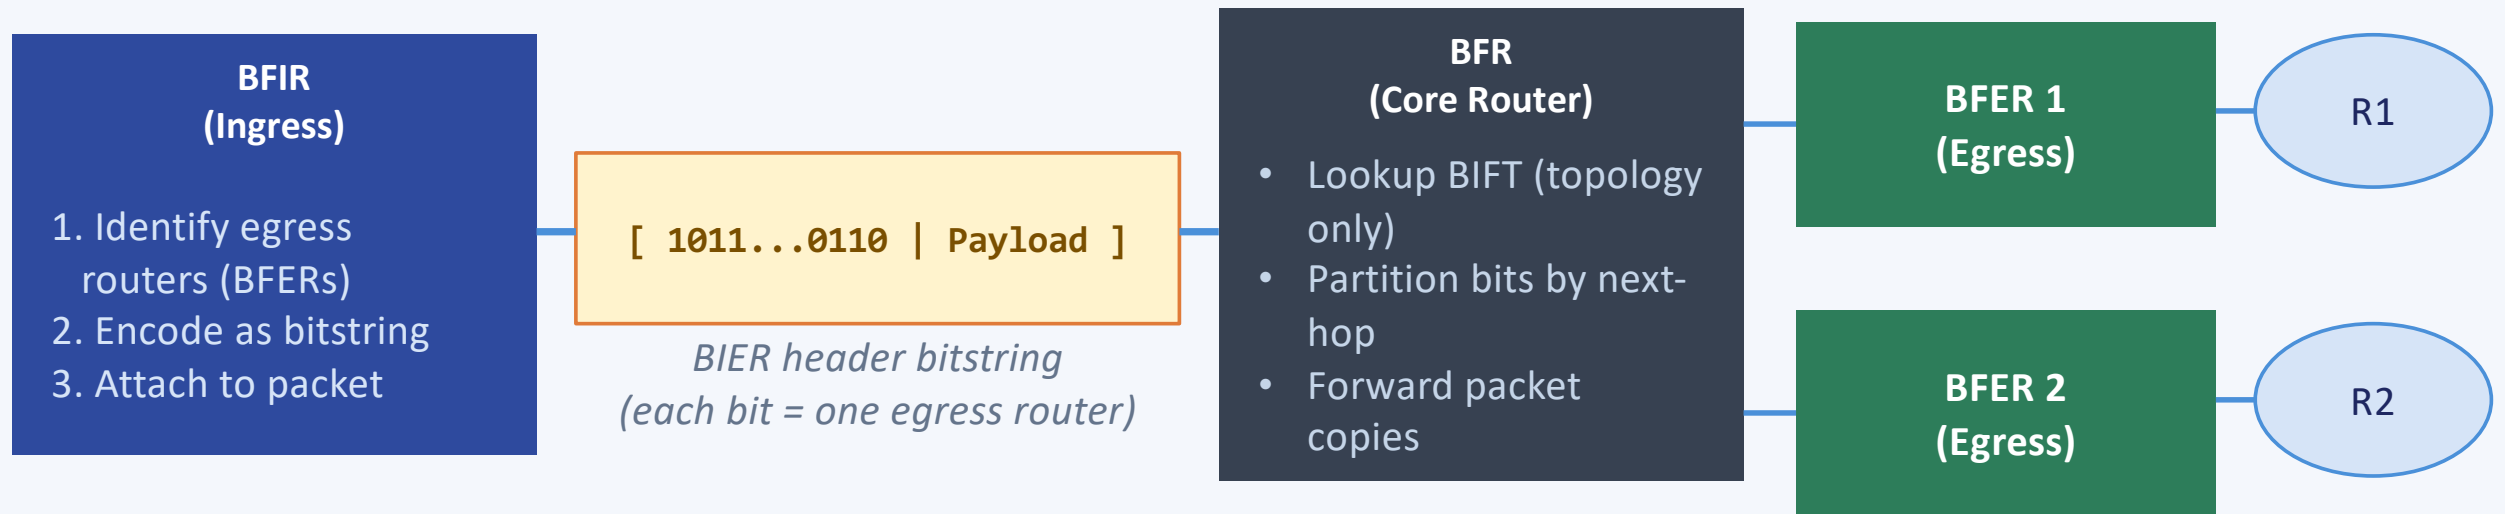
Multicast — Same Root, Harder

- Group address is purely logical — no location info
- Each (Source, Group) pair needs per-router forwarding state
- State grows with active multicast groups & members
- No hardware escape — no brute-force fix available

In practice: multicast is confined to controlled environments (enterprise, IPTV) where application behavior is predictable, and state can be bounded by policy — not by architecture.

BIER: Map-and-Encap for Multicast

Bit Index Explicit Replication — RFC 8279 (2017)



BIFT (Bit Index Forwarding Table): derived from unicast topology only. Maps each bit (= BFER ID) to a next hop and outgoing interface. No per-group state — adding new multicast groups does not change the BIFT.

BIER: Map multicast dissemination tree to unicast egress points and carry the mapping inside packets.

BIER: Core router's state scales with # egress routers (operator-controlled topology).

The Scaling Principle

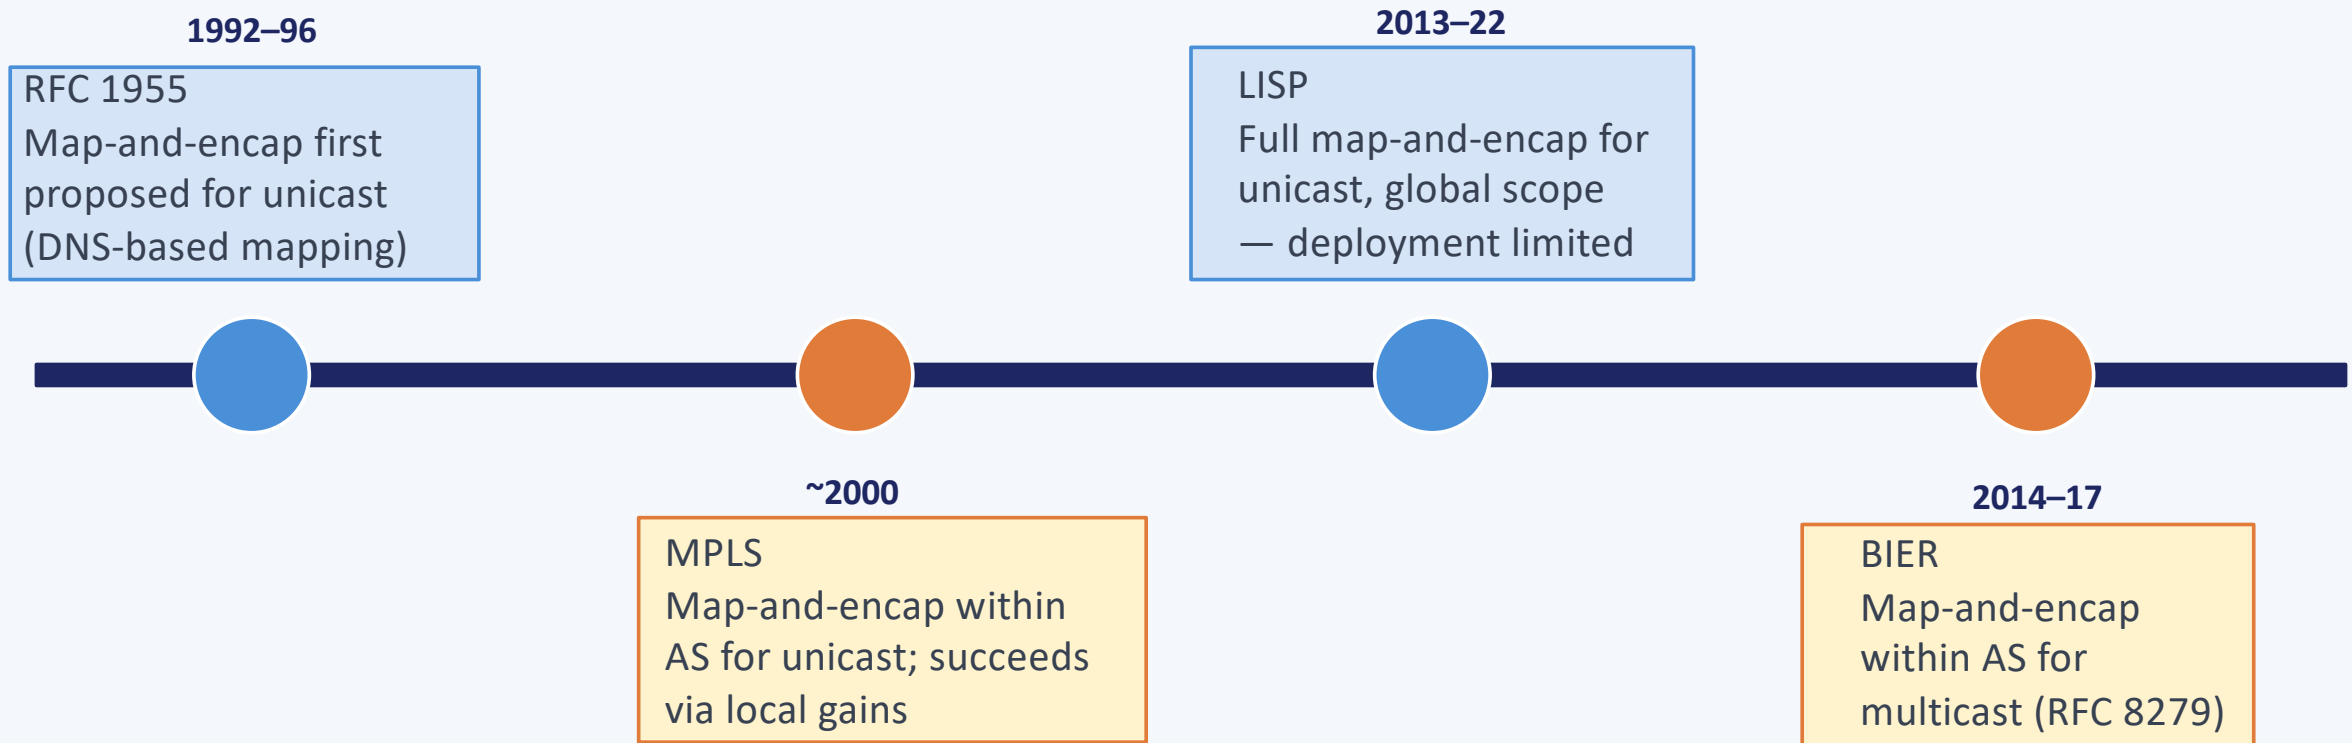
What variable should a routing design's resource use grow with?

Design	State Grows With	Operator Controls It?
Unicast DFZ (BGP)	# end-site address blocks (multihoming, provider portability)	X No
IP Multicast (PIM, DVMRP...)	# active (S,G) pairs, group size, membership dynamics	X No
MPLS within AS	# label-switched paths to egress routers	✓ Yes
BIER within AS	# BFER egress routers in domain	✓ Yes

Observation 2: Routing resource requirements should scale with network topology — an engineered, controllable quantity — not with application dynamics outside the operator's control.

Map-and-Encap: A Recurring Discovery

Developed independently, across different eras and problem contexts — yet always the same answer.



Observation 3: Map-and-encap is the recurring architectural answer — rediscovered independently for unicast (MPLS) and multicast (BIER), always arriving at the same principle.

The Wall That Remains: Inter-Domain

Map-and-encap works, but the deployment breaks down across AS boundaries — for the same underlying reason in both unicast and multicast.

Why it works within an AS

- Single administrative authority
- Egress routers are well-defined topological entities
- IGP provides full topology → BIFT derivable
- iBGP distributes EID-to-RLOC or BFER mappings
- Bounded scope → tractable mapping system

Why it fails across ASes

- AS = policy/administrative boundary, not topological unit
- No universally defined "exit AS" abstraction
- BGP distributes reachability, not topological egress points
- Inter-domain BIER proposals reintroduce per-group state at boundaries
- No agreed global mapping governance (LISP DDT problem)

Observation 4: Inter-domain routing lacks a topological abstraction equivalent to an egress router within a domain. Both BGP and proposed inter-domain BIER have to fall back to finer-grained state at domain boundaries.

Design Lessons for Future Routing Architectures

1 Demand local gains first

Solutions requiring global coordination before anyone sees benefit face near-insurmountable deployment barriers. Design for incremental adoption.

2 Scale with topology, not applications

Forwarding state must be bounded by quantities the operator controls. Designs that scale with application dynamics accumulate architectural debt that workarounds only defer.

3 Map-and-encap is the right pattern

Resolve logical identifiers to topological locators at the domain boundary, carry the result in the packet, and let the core forward on topology alone.

4 The mapping system is as critical as forwarding

A new global mapping system cannot be retrofitted into deployed infrastructure — it must be designed in from the start, with clear governance.