From Map-and-Encap to BIER: Observations on Network Routing Scalability

TIANYUAN YU, LAN WANG, BEICHUAN ZHANG, LIXIA ZHANG

ICNRG @ IETF 126, JULY 22, 2026

Why revisiting routing scalability

- Routing scalability: a research challenge since day one of networking
- Routing on data names: seemingly magnifies the problem

This talk:

- A systematic review of history reveals the root cause
  - unicast routing
  - multicast routing
  - Observed common approach to scale
- Key observations and lessons
  - What this tells us about future routing system designs

Unicast: The root cause: IP's dual role

IDENTIFIER	LOCATOR
• Stable identifier for an endpoint • Embedded in configs, contracts, ACLs • Must not change → Independent of providers	• Assigned by provider • Encodes position in topology • Enables prefix aggregation

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

Unicast: The root cause: IP's dual role

IDENTIFIER	LOCATOR
• Stable identifier for an endpoint • Embedded in configs, contracts, ACLs • Must not change → Independent of providers	• Assigned by provider • Encodes position in topology • Enables prefix aggregation

conflict

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

An Architectural Fix: Map-and-Encap

Separate the two namespaces. Map at the boundary. Scale network core by topology (RFC1955)

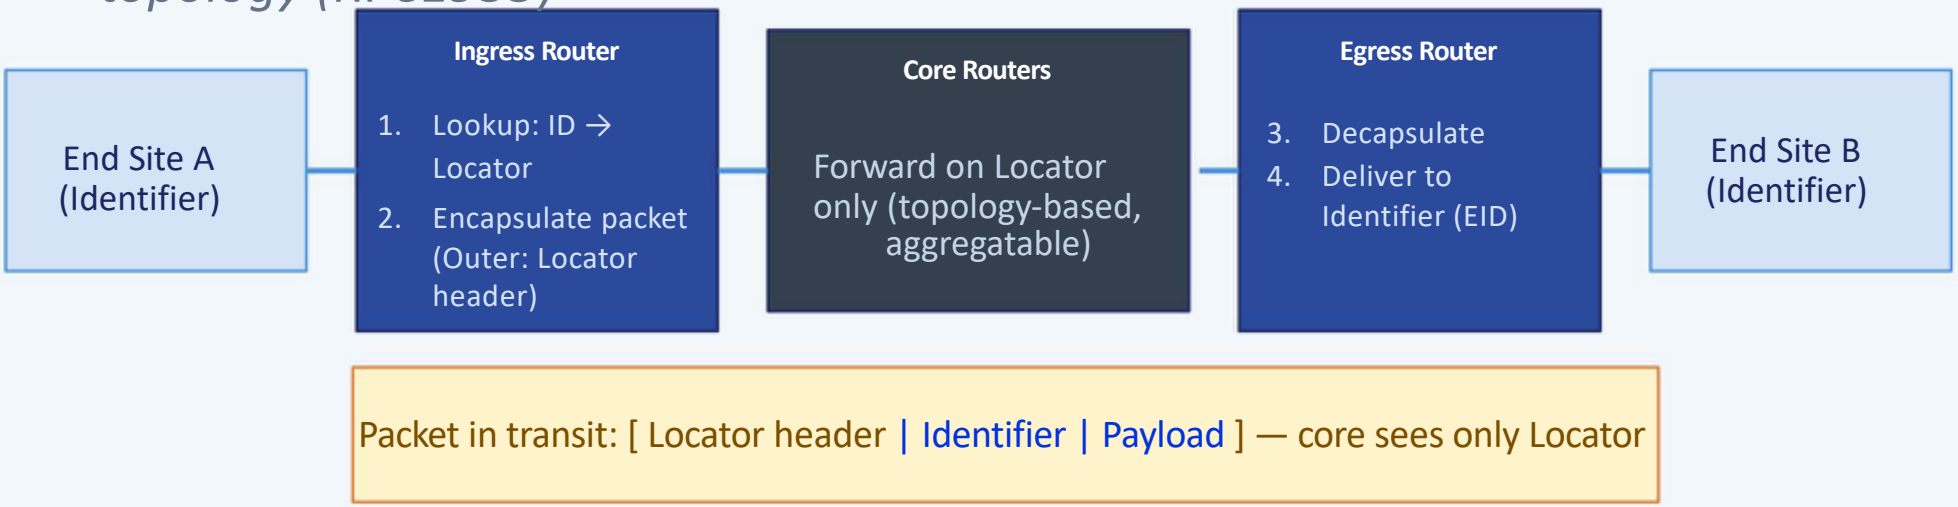

Key benefit: Core routing table bounded by topology (provider networks) — not by number of end sites.

LISP: A 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

- Published as experimental RFCs 2013, proposed standard 2022
- Requires a new global mapping system (DDT), but with no agreed root authority or governance rules
- 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 insurmountable deployment barrier.

What actually got deployed: MPLS

MPLS can be viewed as map-and-encap within a single administrative domain

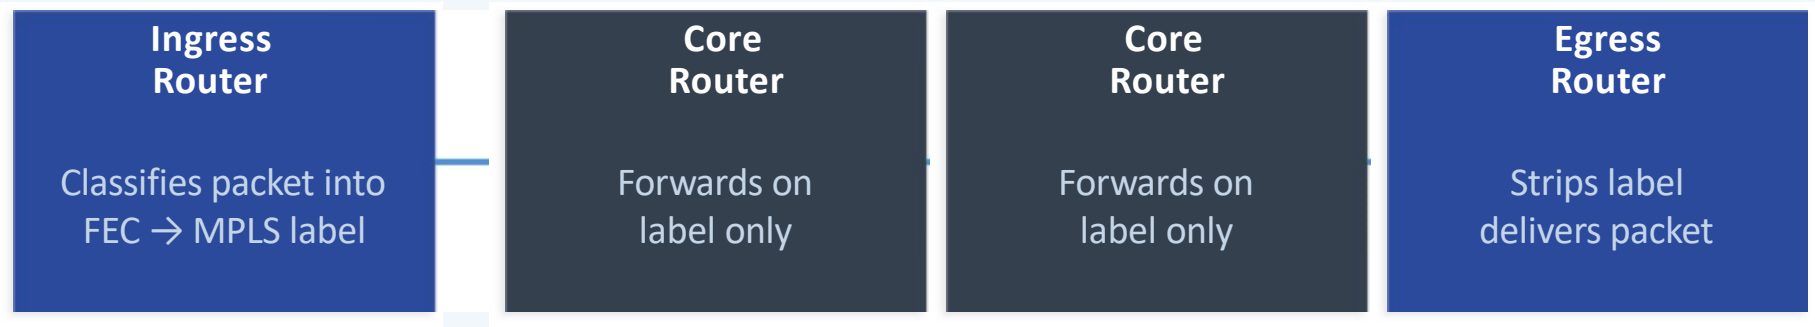

Why MPLS succeeded where LISP didn't:

- ✓ 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.

(traditional) IP Multicast forwarding: Per-Group Forwarding State

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

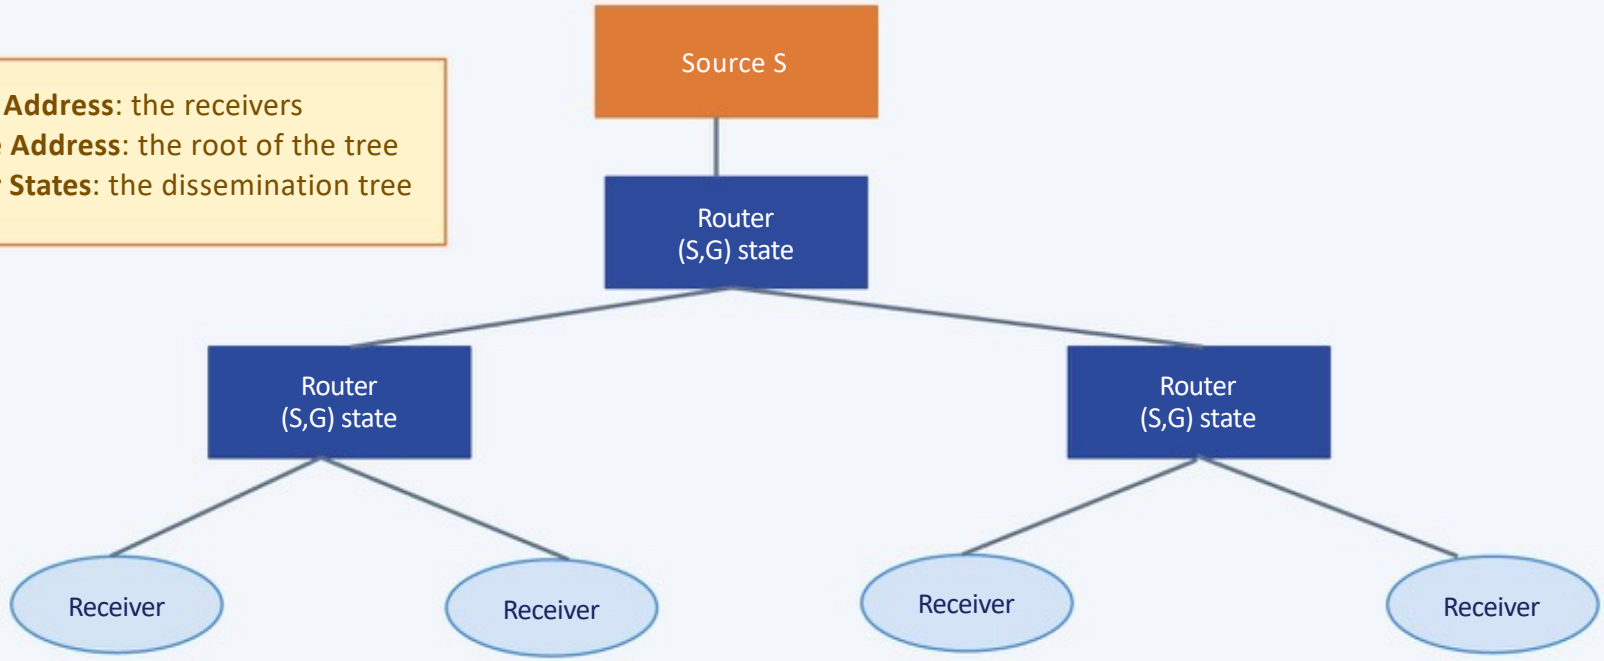

Every router's state grows by # groups × # sources — which are outside network control

Multicast forwarding: even worse than unicast scaling

A group address is purely a 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 configuration — not by architecture.

New multicast forwarding design: BIER=Map-and-Encap for Multicast

Bit Index Explicit Replication — RFC 8279 (2017)

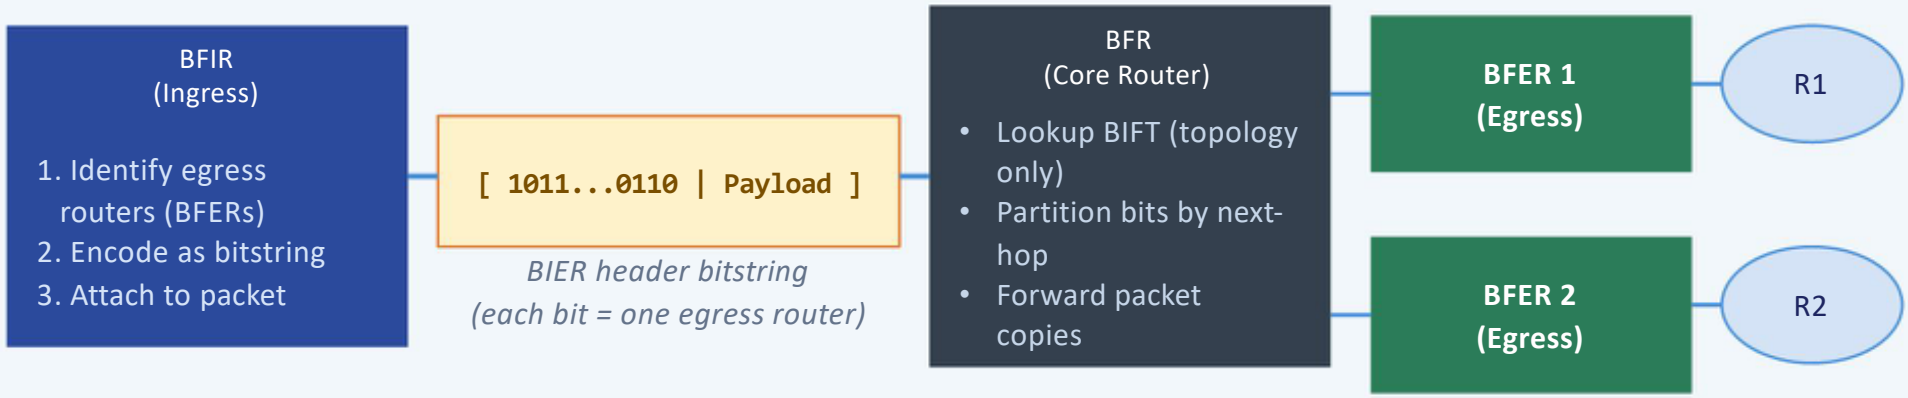

BIFT (Bit Index Forwarding Table): derived from unicast topology. Maps each bit (= BFER ID) to a next hop and outgoing interface. No per-group state — adding new multicast groups does not change the BIFT.

Map multicast dissemination tree to unicast egress points in packets; the bit-string ≈ encap
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)	✗ **No**
IP Multicast (PIM, DVMRP…)	# active (S,G) pairs, group size, membership dynamics	✗ **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, which can be controlled by topological engineer, not with external application dynamics

Map-and-Encap for scalability: A Recurring Discovery

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

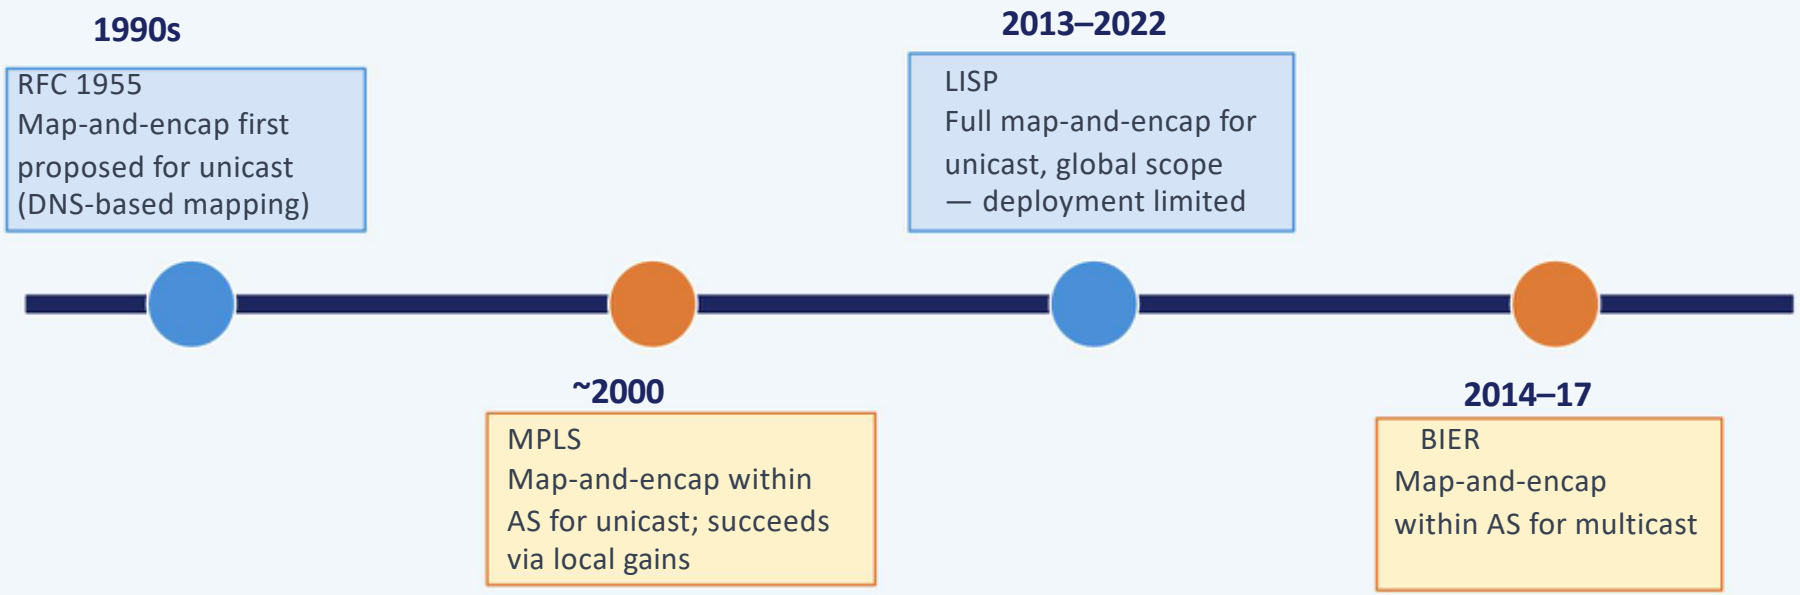

Observation 3: Map-and-encap is the recurring architectural answer — rediscovered independently for unicast (MPLS) and multicast (BIER), always arriving at the same principle.

The remaining scaling wall: Inter-Domain

Map-and-encap works, but not across AS boundaries so far — for the same underlying reason in both unicast and multicast.

Why it works within an AS	Why it fails across ASes
• 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	• 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 has no topological abstraction equivalent to an egress router within a domain. Both BGP and ongoing inter-domain BIER have to work with finer-grained state at domain boundaries

Design Lessons for Future Routing Architectures

- Scale routing/forwarding with topology
  - Bounding routing/forwarding state by network *topology*.
  - routing/forwarding designs that scale with *logical* application-level identifiers, incur architectural debt, and workarounds only defer the bill.
- Use map-and-encap to resolve the conflict between *logical* and *topological* identifiers
  - May logical identifiers to topological locators at the domain boundary, carry the result in the packet, so the core forwards on topology alone.
- The mapping system: *the critical ingredient*
  - a mapping system must be designed into the system from the start
  - Retrofitting one into the deployed infrastructure has not succeeded so far
    - because it requires coordinated adoption across independently operated networks.