

Core Use-Case Requirements for MSR6

IETF114 Philadelphia
v1.0 - 07/13/2022

*how to fit 30 years of IP Multicast protocol design and deployment experience
into too few slides & time*

Toerless Eckert (Futurewei USA), tte@cs.fau.de

Core MSR6 use-case requirements

There is really only one requirement:

**“Simple”⁽¹⁾ , IPv6 integrated⁽²⁾, “End-to-End”⁽³⁾, stateless⁽⁴⁾,
IPv6 multicast⁽⁵⁾ for IPv6-only networks⁽⁶⁾**

(6) What IPv6 only networks ?

A: All IPv6 networks that require IPv6 multicast!

Explained by other drafts/presentations: SP/WAN/Metro native IPv6 (with/without SRv6) (e.g.:IPTV, MVPN), DCN, OTT/Overlays

IMHO also: Any enterprise, transportation, IoT network (small..large)

Not considered by current MSR6 drafts – but should be

Refresh: BIER and IPv6-only networks⁽⁶⁾

BIER RFC8279 (arch) + RFC8296 (header) is new layer ~L2/L2.5

BIER router (BFR) forwarding is not IPv6 forwarding (RFC8200)

BIER packets are not IPv6 packets (RFC8200)

Goal: “One additional multicast forwarding plane for all unicast networks”

Initially built for MPLS: header: label field (BIFT-ID), TC, OAM, signaling, ...

Also BIER over L2 only, but little/no operator interest / stalled drafts

BIER-WG solution for IPv6 only networks / IPv6 Multicast

Do not build IPv6-only networks !

Run separate BIER hop-by-hop forwarding plane
in parallel to IPv6 unicast forwarding plane

End-to-end tunnel for IPv6 Multicast over BIER : **2 layers**


BIER header + IPv6 (multicast) Header (so-called BIER flow overlay)

Transit over IPv6-unicast only routers : **3 layers**

BIER hop-by-hop tunnel BIER over IPv6 (bier-bierin6 draft)

IPv6 unicast header (lower) + BIER header + IPv6 (multicast) header

no endorsement,
i have not read the book
i just like the cover



Simplicity⁽¹⁾, End-to-End⁽²⁾

Operational / Architectural alignment/integration⁽²⁾ with IPv6 (unicast)

30 year experience: IP Multicast solutions are only successful when they align and integrate with network unicast designs as much as possible

Do not build or afford parallel additional or unnecessary different multicast technology

1989: IP/IPv6 Multicast is re-using / extending IP (RFC1112)⁽²⁾

Allowed to re-use/extend IP/IPv6 ecosystem:

SDKs/ sockets-API(*), QoS Diffserv/IntServ(RSVP)(*), ACLs(*), any IP L2 encaps(*), IPFIX(*), IPsec(*), ..

(*) *Most of this would all have to be reinvented / duplicated for BIER but not for MSR6.*

Would have put bitstrings into IPv6 addresses if they were long enough (e.g. Cisco drafts in BIER)!

1990th: We tried novel multicast routing (MOSPF, DMVPN, ...)

Replaced by PIM + unicast routing (OSPF, ISIS, RIP,...) – do not re-invent routing for multicast!

200x: IPv4 multicast MVPN solution for MPLS/VPN SP networks – additional forward/control!

IPv4 Multicast replaced by native MPLS multicast (mLDP/RSVP-TE/P2P)

PIM/MVPN signaling replaced by BGP/MVPN signaling

IMHO: BIER driven by MPLS SP use-cases (MPLS/MVPN). Well aligned/integrated there!

But not for the wide range of IPv6 networks – end-to-end – into IPv6 applications

Non-MPLS networks do not want/need an additional BIER ecosystem

Stateless⁽⁴⁾, (1) .. and End-to-End⁽³⁾

Operational simplicity (troubleshooting), safety, reliability

All stateful multicast (IP or MPLS): Applications create (tree) state on routers in the network.

No IETF standard multicast circuit breaker / state congestion/control solutions for multicast state.

Any bad or attack multicast application can bring down stateful multicast routers

Unicast state: routing tables - do not grow with traffic, only with topology

Global MPLS SPs where deploying ingress-replication to avoid Multicast state on P nodes (RFC7988)

after we invested 10 years in IETF to specify MPLS multicast – Core reason for BIER

Scale and Convergence

Finance / Telemetry / Content distribution / adaptive streaming would require

hundreds of thousands of multicast states. Can not create, re-converge, operate!

Additional new multicast paradigm for applications

IP Multicast (and SSM): application signaling: flow based – receiver join/leave group/channels

Only with stateless multicast:

Sender can **DIRECT EVERY PACKET SEPARATELY** across different paths to different receivers

Only way to enable e.g.: adaptive streaming at scale via multicast

BIER always wanted/wants to explore this. But IMHO **NO WAY** to get a ubiquitous BIER socket API

Prior multicast socket extensions took almost 20 years (e.g.: SSM)

IPv6 extension header API already defined since 2003 (RFC3542) (no BIER API work)

Summary / Conclusion

Stateless bitstring replication is the best new multicast direction in 40 years

BIER-WG - great stateless intra-SP solution for MPLS networks

But inferior fit for IPv6 networks and end-to-end applications

IETF/BIER can-not / should not replicate the whole IPv6 ecosystem

MSR6: Keep it simple & Make IPv6 Multicast great again!

Native stateless multicast - for all IPv6 networks

***Stateless IPv6 multicast into applications
(DC, industrial, IoT, SP-edge, ...)***

Re-use / share all of BIER that fits!

Stuffed Agenda, No questions now ?!

But welcoming questions any time after the meeting!

