IPv6 Multicast

6DEPLOY. IPv6 Deployment and Support

Intro

Multicast is inherent to the IPv6 protocol

No broadcasts

- Multicast used instead

But some parts need to be configured

- for building the multicast trees
- for topology information (routing)

Agenda

Multicast addressing

MLD & MLDv2

PIM SM/SSM

Interdomain multicast

Multicast addressing (1)

Multicast addresses format: (RFC 3513)

8 bits		4 bits	4 bits	112 bits
1111	1111	flags	scope	group ID
F	F			

- 8 high order bits set to 1 · Addresses derived from FF00::/8 prefix
- **flag** field(4 bits) : 0RPT values
  - T = 0 for permanent addresses (Defined by IANA)
  - T = 1 for transient addresses
  - Bits P and R discussed later
- **scope** field → Makes it possible to limit the scope of the multicasting
  - 0 - Reserved
  - 1 – Node-local
  - 2 – Link-local
  - 3 – Subnet-local
  - 4 - Admin-local
  - 5 - Site-local
  - 8 - Organization-local
  - E - Global (Internet)

Multicast addressing (2)

Scopes must be configured on routers!

Examples of IANA allocated addresses

- Flag bits T=P=R=0
  - Flag = 0
- Group ID 101 → NTP servers
  - FF01:0:0:0:0:0:0:101 : All the NTP servers on the sender's host
  - FF02:0:0:0:0:0:0:101 : All the NTP servers on the sender's link
  - FF05:0:0:0:0:0:0:101 : All the NTP servers on the sender's site
  - FF0E:0:0:0:0:0:0:101 : All the NTP servers on the Internet

Reserved multicast addresses: examples

Given on RFC 2375

Addresses available only for a given scope

- FF02:0:0:0:0:0:0:1 : All the nodes of the link
- FF02:0:0:0:0:0:0:2 : All the routers of the link
- FF05:0:0:0:0:0:0:2 : All the routers of the site
- FF02:0:0:0:0:0:0:D : All the PIM routers of the link
- ...

Addresses available for all scopes

- FF0X:0:0:0:0:0:0:101 : Network Time Protocol (NTP)
- FF0X:0:0:0:0:0:0:109 : MTP Multicast Transport Protocol
- ...

IPv6 multicast and Ethernet

Ethernet is multicast capable (not always implemented)

Requires 8th bit of MAC address to be set to 1

For IPv6 : @MAC = 33-33-xx-yy-zz-kk

- xx-yy-zz-kk are 32 lower bits of the IPv6 address

Example:

- IPv6@ = **FF3E:40:2001:660:3007:123**:1234:5678
- MAC@ = **33-33**-12-34-56-78

Solicited node multicast addresses (for NDP)

- Multicast address built from unicast address
- Concatenation of
  - FF02::1:FF00:0/104
  - 24 low order bits of the unicast address
- Nodes build their own IPv6 solicited node multicast address
- Nodes that know the IPv6 address of a host but not its MAC address can use the solicited node multicast address
  - NDP protocol (Neighbor Discovery Protocol)
  - Protocol for DAD management
- Avoids sending MAC broadcasts (FF-FF-FF-FF-FF-FF)
- Example:
  - 2001:0660:010a:4002:4421:21FF:FE24:87c1
  - FF02:0000:0000:0000:0000:0001:FF00:0000/104
  - FF02:0000:0000:0000:0000:0001:FF24:87c1
  - 33-33-FF-24-87-C1 -> MULTICAST MAC ADDRESS

Multicast addresses derived from unicast prefixes

Described in RFC 3306

Flag : 0RPT

11111111	flag	scop	reserved	Prefix Length	Network prefix	Group ID
8 bits	4	4	8 bits	8	64 bits	32 bits

- Flag : 0RPT
  - P=0 · Address not based on the unicast prefix
  - P=1 · Address based on the unicast prefix
  - If P=1 · T=1 · FF30::/12 prefix
    - (T=1 because not allocated by IANA)
- Reserved : 0
- Example: prefix **2001:660::/32** (RENATER)
- address FF**3**E:**20**:**2001:660:0:0**:**1234:abcd**

SSM addresses

Are also RFC3306 addresses

SSM addresses range: FF3X::/32

Only addresses in FF3X::/96 should be used now.

These are RFC3306 addresses with:

- Plen = 0
- Prefix = 0

Example:

- FF3x::1234:abcd /96
- 1234:abcd being the Group ID

Multicast addresses allocation

« Manual » choice of multicast address and port

Dynamic

- Session Announcement Protocol, (SAP), ID
  - SDR implements SAP (not scalable for a global scope)
- MADCAP, RFC 2730
  - Multicast Address Dynamic Client Allocation Protocol (too much complex, very few implementations and no deployment)
- GLOP, RFC 2770
  - Useless as we have RFC 3306

Multicast addresses derived from unicast prefixes (RFC 3306)

- Any host can derive a multicast address from the network prefix where it is connected
- Makes allocation easier
- How to assign addresses to end user remains a problem

Agenda

Multicast addressing

MLD & MLDv2

PIM SM/SSM

Interdomain multicast

MLD

Interaction protocol between

- Multicast router on the link-local
- Multicast hosts on the link-local

Host can say: « I want to join group FF0E::1234 and receive the related flow »

MLDv1 (RFC 2710)

- MLD <-> IGMPv2 <-> ASM only

MLDv2 (RFC 3810)

- MLDv2 <-> IGMPv3 <-> SSM + ASM

MLD messages are sent in ICMPv6 packets

MLD Packet

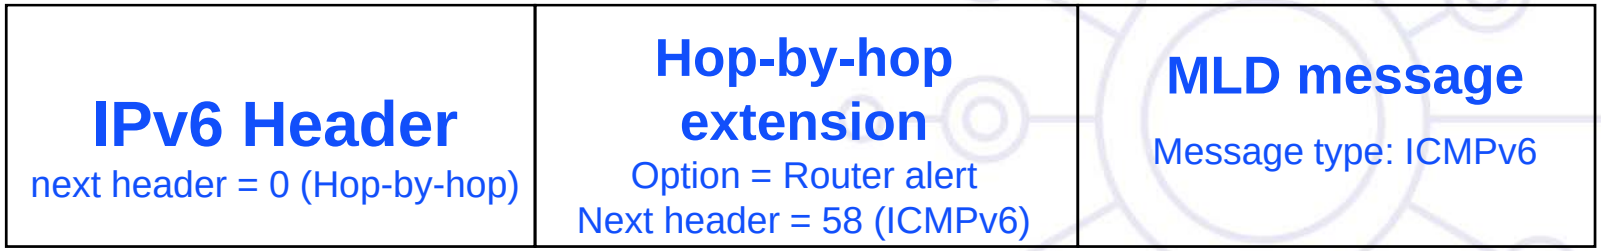
IPv6 Header
next header = 0 (Hop-by-hop)
Hop-by-hop extension
Option = Router alert
Next header = 58 (ICMPv6)
MLD message
Message type: ICMPv6

MLDv1 Message

- **Type** : Messages types
  - General Query and Multicast-Address-Specific Query (130)
  - Multicast Listener Report (131)
  - Multicast Listener Done (132)
- **Code** : Set to 0 by sender and ignored then
- **Checksum** : for the complete packet (headers+MLD message)
- **Maximum Response Delay** : For query messages, time by which hosts must respond
- **Reserved** : Not used: set to 0 and ignored then
- **Multicast Address**: IPv6 multicast address or 0 according to message type

MLDv1 : Join a group

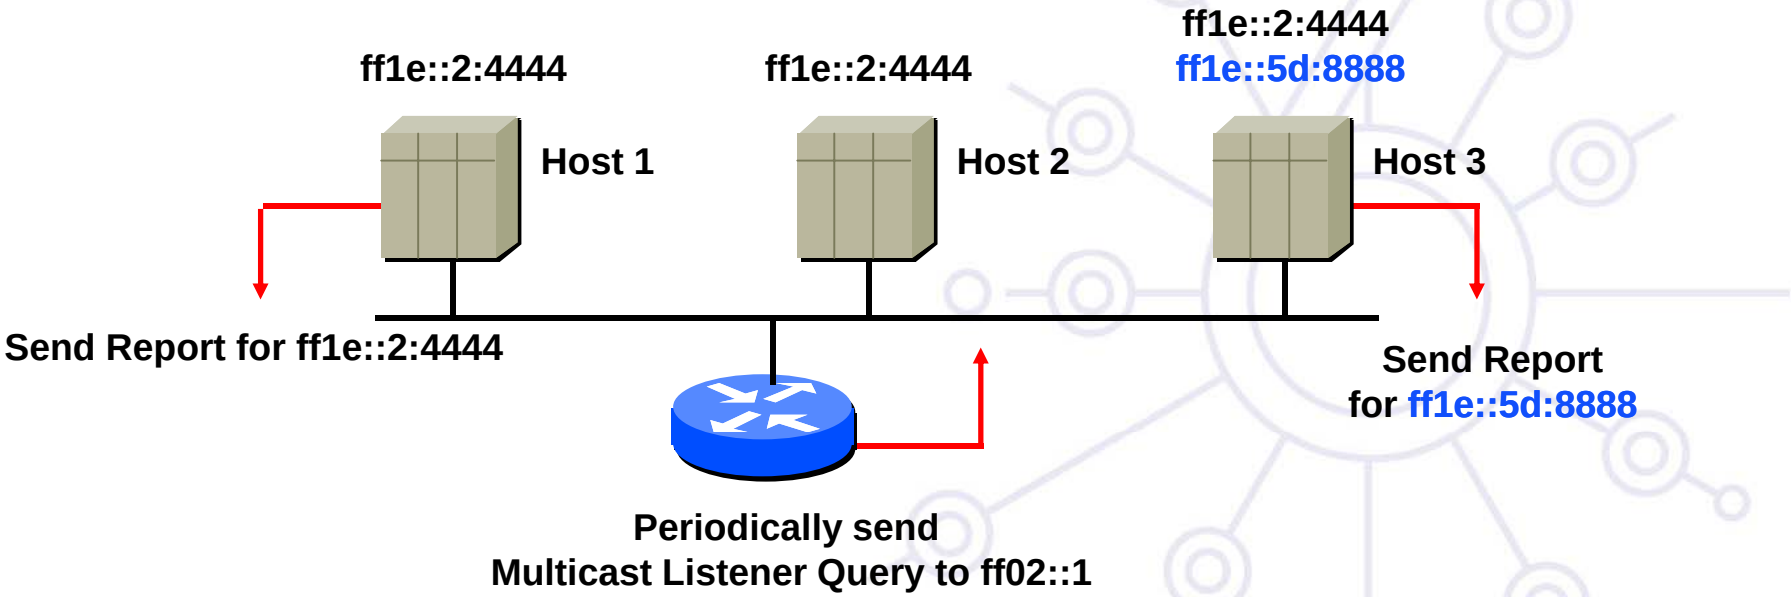

MLDv1 : Leave a group (1)

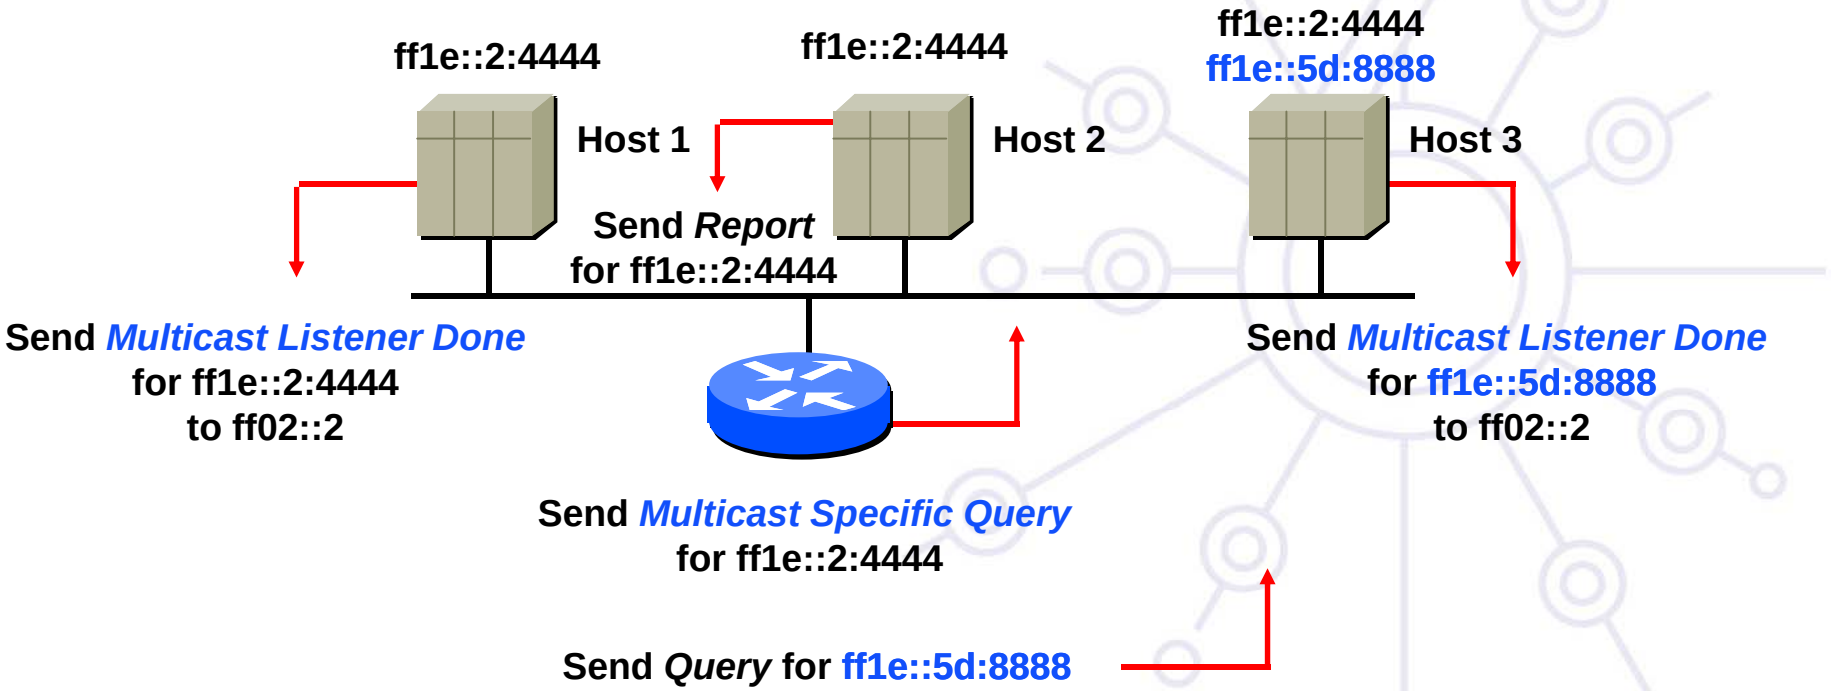

MLDv1 : Leave a group (2)

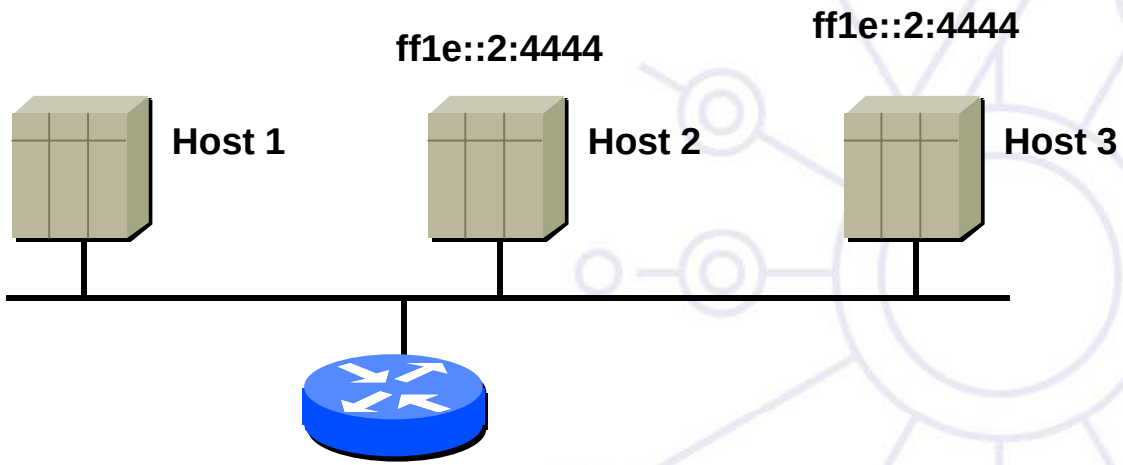

MLDv2 (RFC 3810)

Management of group & sources

- INCLUDE : to receive packets from sources specified in the MLDv2 message
- EXCLUDE : to receive packets from all sources except the ones specified in the MLDv2 message

2 types of messages

- Multicast listener query messages
- Multicast listener report messages

Interoperable with MLDv1

Agenda

Multicast addressing

MLD & MLDv2

PIM SM/SSM

Interdomain multicast

PIM SM/SSM

Protocol Independant Multicast

No difference with PIM for IPv4

- Except PIM messages are sent with link-local IPv6 address

Creates multicast trees between senders and receivers (distribution trees)

Not a routing protocol

Relies on other routing protocols (MBGP, static…)

Agenda

Multicast addressing

MLD & MLDv2

PIM SM/SSM

Interdomain multicast

Interdomain multicast (1)

No problem for SSM. Source specific trees created from senders to receivers accross domains

ASM problem: was solved in the IPv4 world with MSDP (Multicast Source Discovery Protocol)

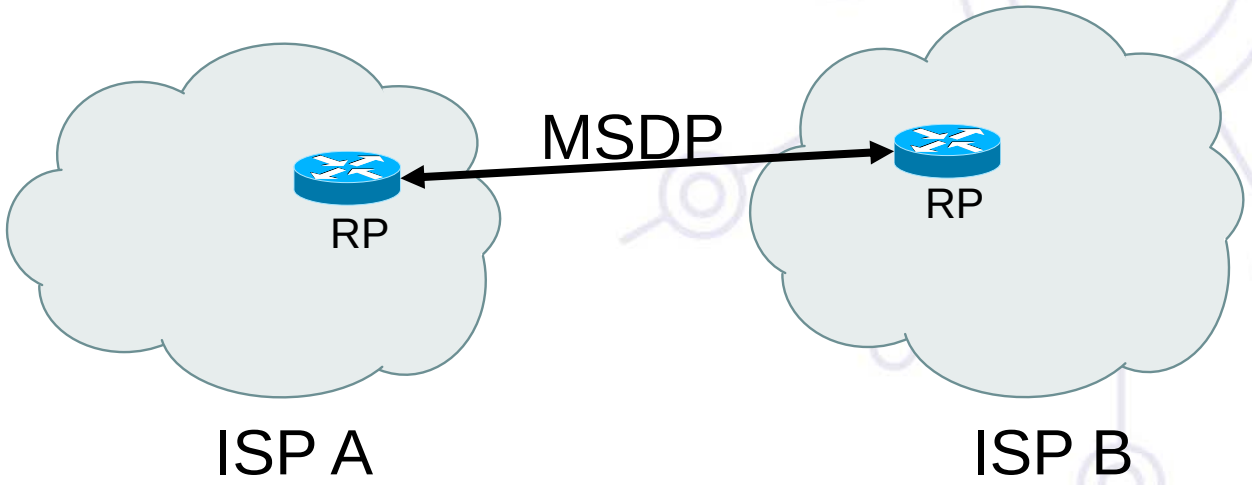

Interdomain multicast (2)

No one wants MSDP for IPv6, not manageable/scalable

SSM IETF lobby

- Some SSM applications already available

How to solve N -> M multicast ?

- Application / Middleware ?
- Not there yet (ongoing work)

Embedded-RP – RFC 3956

- For each group, everyone uses the same RP
- Embedded is a solution for group-to-RP mapping
- Requires support in all PIM routers (that are part of the tree)

Embedded-RP

Flag : 0RPT

11111111	flag	scop	res	rpad	Prefix Length	Network prefix	Group ID
8 bits	4	4	4	4	8	64 bits	32 bits

Flag : 0RPT

- R=1 → Embedded-RP address
- If R=1 → P=1 → T=1
- FF7x::/16 addresses

Res : 0

Rpad : last 4 bits of the RP address

E.g. RP address 2001:660:3001:104::8

- Multicast address FF7E:0820:2001:660:3001:104:1234:abcd

IPv6 multicast