IETF 96 ROLL

Routing over Low-Power And Lossy Networks

Chairs:

Peter van der Stok

Ines Robles

Secretary:

Michael Richardson

Note Well

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

- The IETF plenary session
- The IESG, or any member thereof on behalf of the IESG
- Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices
- Any IETF working group or portion thereof
- Any Birds of a Feather (BOF) session
- The IAB or any member thereof on behalf of the IAB
- The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice. Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

Source: https://www.ietf.org/about/note-well.html

Meeting Materials

- 15:50-17:20 Wednesday Afternoon session II
- Remote Participation
  - Jabber Room: roll@jabber.ietf.org
  - Meetecho: http://www.meetecho.com/ietf96/roll
- Etherpad:
  - http://tools.ietf.org/wg/roll/minutes
- Audio Streaming:
- Minutes taker: Michael Richardson (in etherpad)
- Jabber Scribe:
- **Please sign blue sheets :-)**

Agenda

15:50-17:20 Wednesday Afternoon session II

Item	Time	Presenter
Status of the working group	15:50 - 15:52 (2 min.)	Peter/Ines
Use of rpl info draft - draft-ietf-roll-useofrplinfo	15:52 - 16:10 (18 min.)	Michael Richardson
DIS Modifications - draft-gundogan-roll-dis-modifications	16:10 - 16:20 (10 min.)	Cenk Gündogan
Source-Routed Multicast for RPL - draft-bergmann-bier-ccast	16:20 - 16: 35 (15 min.)	Carsten Bormann
Mpl Forwarder select - draft-vanderstok-roll-mpl-forw-select	16: 35 - 16:45 (10 min.)	Peter van der Stok
AODV-RPL - draft-satish-roll-aodv-rpl	16:45 - 16:55 (10 min.)	Charles Perkins
No-Path DAO Problem Statement - draft-jadhav-roll-no-path-dao-ps-00	16:55 - 16:05 (10 min.)	Rahul Jadhav
Charter discussion + AOB	16:05 - 17:20 (15 min.)	Peter and Ines

Milestones

Date	Milestone
May 2016	Submit draft about how to compress RFC6553, RFC6554, and IP headers in the 6LoWPAN adaptation layer context to the IESG. draft-ietf-6lo-routing-dispatch
Apr 2016	Submit draft about when to use RFC6553, RFC6554, and IPv6-in-IPv6 encapsulation to the IESG. draft-robles-roll-useofrplinfo
Jan 2016	Evaluate WG progress, recharter or close

State of Active Internet-Drafts

Draft	Status
draft-ietf-roll-applicability-ami-13 Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks	IESG Evaluation::Revised I-D Needed
draft-ietf-roll-applicability-template-09 ROLL Applicability Statement Template	No to be submitted
draft-ietf-roll-routing-dispatch-00 6LoWPAN Routing Header	Submitted to IESG for Publication
draft-ietf-roll-useofrplinfo-05 When to use RFC 6553, 6554 and IPv6-in-IPv6	New version submitted- need advice for working group

Related Internet-Drafts

Related Internet-Drafts	
draft-gundogan-roll-dis-modifications-00 **DIS Modifications**	2016-07-08 15 pages New
draft-jadhav-roll-no-path-dao-ps-01 **No-Path DAO Problem Statement**	2016-06-27 10 pages
draft-satish-roll-aodv-rpl-00 **Asymmetric AODV-P2P-RPL in Low-Power and Lossy Networks (LLNs)**	2016-06-07 14 pages
draft-vanderstok-roll-mpl-forw-select-00 **MPL Forwarder Select (MPLFS)**	2016-07-04 8 pages New
draft-vanderstok-roll-mpl-yang-01 **A YANG model for Multicast Protocol for Low power and lossy Networks (MPL)**	2016-06-27 20 pages

Expired Internet-Drafts

draft-thubert-roll-dao-projection-02	Root initiated routing state in RPL

Open Tickets

Ticket	Summary	Component	Version	Type	Owner	Status	Created
#173	Example of Flow from RPL-aware-leaf to non-RPL-aware-leaf	useofrplinfo		defect	mariainesrobles@gmail.com	new	2016-01-29
#176	Security items to consider for applicability-ami draft - IESG Evaluation-	applicability-ami		defect	draft-ietf-roll-applicability-ami@ietf.org	new	2016-05-13
#177	Editorial Comments for ami draft - IESG Evaluation-	applicability-ami		defect	draft-ietf-roll-applicability-ami@ietf.org	new	2016-05-13
#174	consider changing RPI HbH header value from 63 to 43.	routing-dispatch		enhancement	draft-ietf-roll-routing-dispatch@ietf.org	new	2016-03-29
#175	processing of multiple SRH-6LoRH	routing-dispatch		enhancement	Pascal Thubert	new	2016-04-30
#170	Use of ESC Dispatch value in new IETF header compression	draft-richardson-roll-applicability-template		task	draft-richardson-roll-applicability-template@tools.ietf.org	new	2015-05-02

RPL RPI/RH3 uses

draft-ietf-roll-useofrplinfo

Michael Richardson
Pascal Thubert
Ines Robles

Scenarios analyzed in draft-ietf-roll-useofrplinfo

{Storing,Non-Storing} X {RPL-aware-leaf,non-RPL-aware,root, Internet} X {RPL-aware-leaf,non-RPL-aware,root,Internet}

(but Internet->Internet cases removed, so 24, not 32)

STORING

1. Flow from RPL-aware-leaf to root
2. Flow from root to RPL-aware-leaf
3. Flow from non-RPL-aware-leaf to root
4. Flow from root to non-RPL-aware-leaf
5. Flow from RPL-aware-leaf to Internet
6. Flow from Internet to RPL-aware-leaf
7. Flow from non-RPL-aware-leaf to Internet
8. Flow from Internet to non-RPL-aware-leaf
9. Flow from RPL-aware-leaf to RPL-aware-leaf
10. Flow from RPL-aware-leaf to non-RPL-aware-leaf
11. Flow from non-RPL-aware-leaf to RPL-aware-leaf
12. Flow from non-RPL-aware-leaf to non-RPL-aware-leaf

NON-STORING

13. Flow from RPL-aware-leaf to root
14. Flow from root to RPL-aware-leaf
15. Flow from non-RPL-aware-leaf to root
16. Flow from root to non-RPL-aware-leaf
17. Flow from RPL-aware-leaf to Internet
18. Flow from Internet to RPL-aware-leaf
19. Flow from non-RPL-aware-leaf to Internet
20. Flow from Internet to non-RPL-aware-leaf
21. Flow from RPL-aware-leaf to RPL-aware-leaf
22. Flow from RPL-aware-leaf to non-RPL-aware-leaf
23. Flow from non-RPL-aware-leaf to RPL-aware-leaf
24. Flow from non-RPL-aware-leaf to non-RPL-aware-leaf

https://tools.ietf.org/html/draft-ietf-6man-rfc2460bis-05#section-4.3

NOTE: While [RFC2460] required that all nodes must examine and process the Hop-by-Hop Options header, it is now expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so.

https://tools.ietf.org/html/draft-ietf-6man-rfc2460bis-05#section-4

The insertion of Extension Headers by any node other than the source of the packet breaks PMTU-discovery and can result in ICMP error messages being sent to the source of the packet that did not insert the header.

The current approach to allowing a header to be inserted is to encapsulate the packet using another IPv6 header and including the additional extension header after the first IPv6 header, for example, as defined in [RFC2473].

Storing mode: **without** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

Storing mode: **without** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)
(ignore how F knows that G is nRaI)

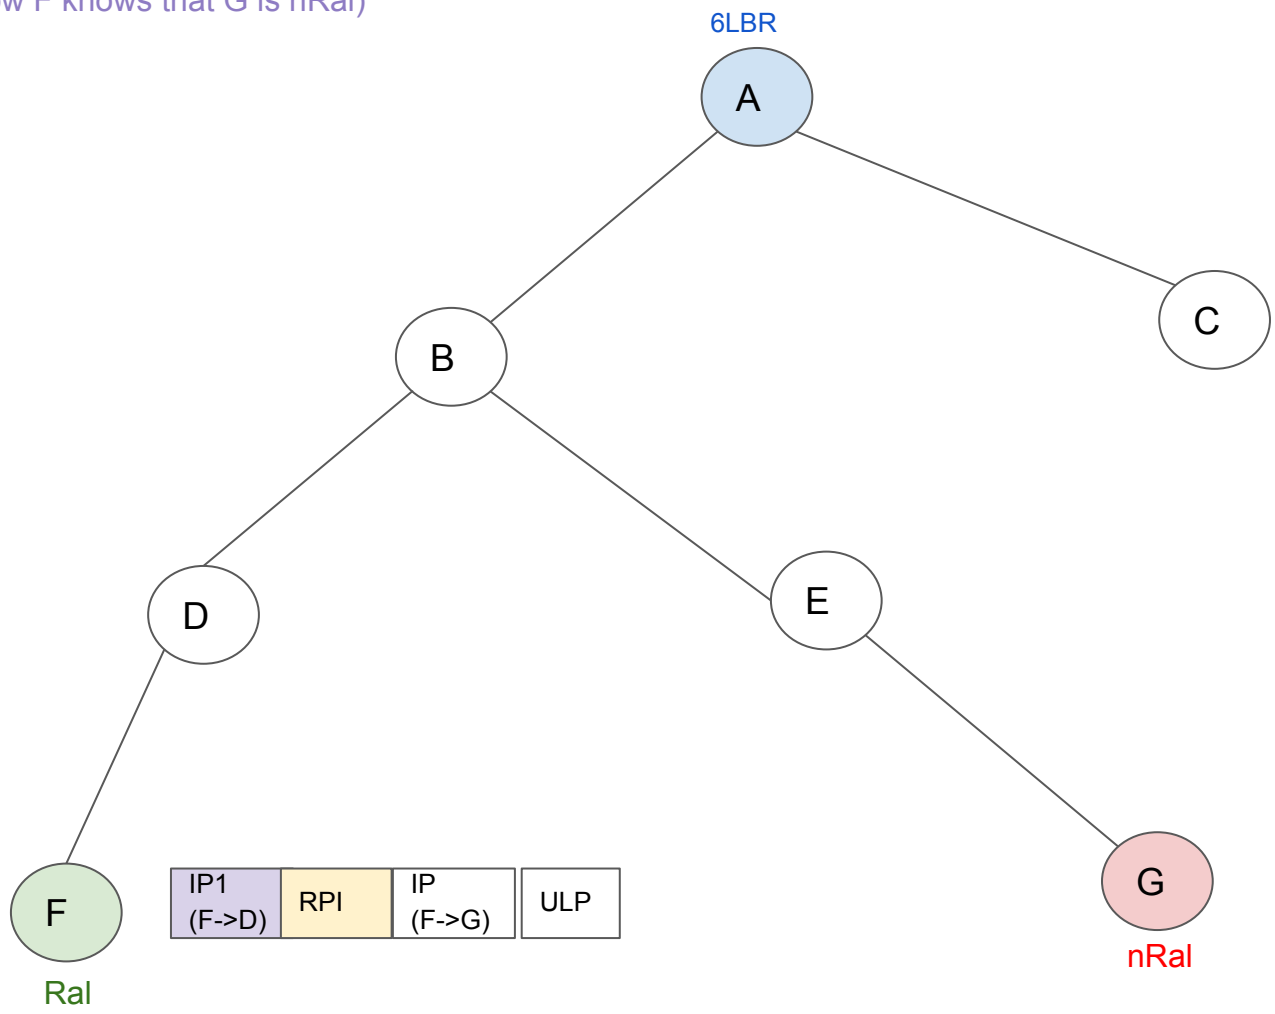

Storing mode: **without** the inclusion of 2460bis proposal

Form ral (Rpl-aware-leaf) to nRal (not-Rpl-aware-leaf)

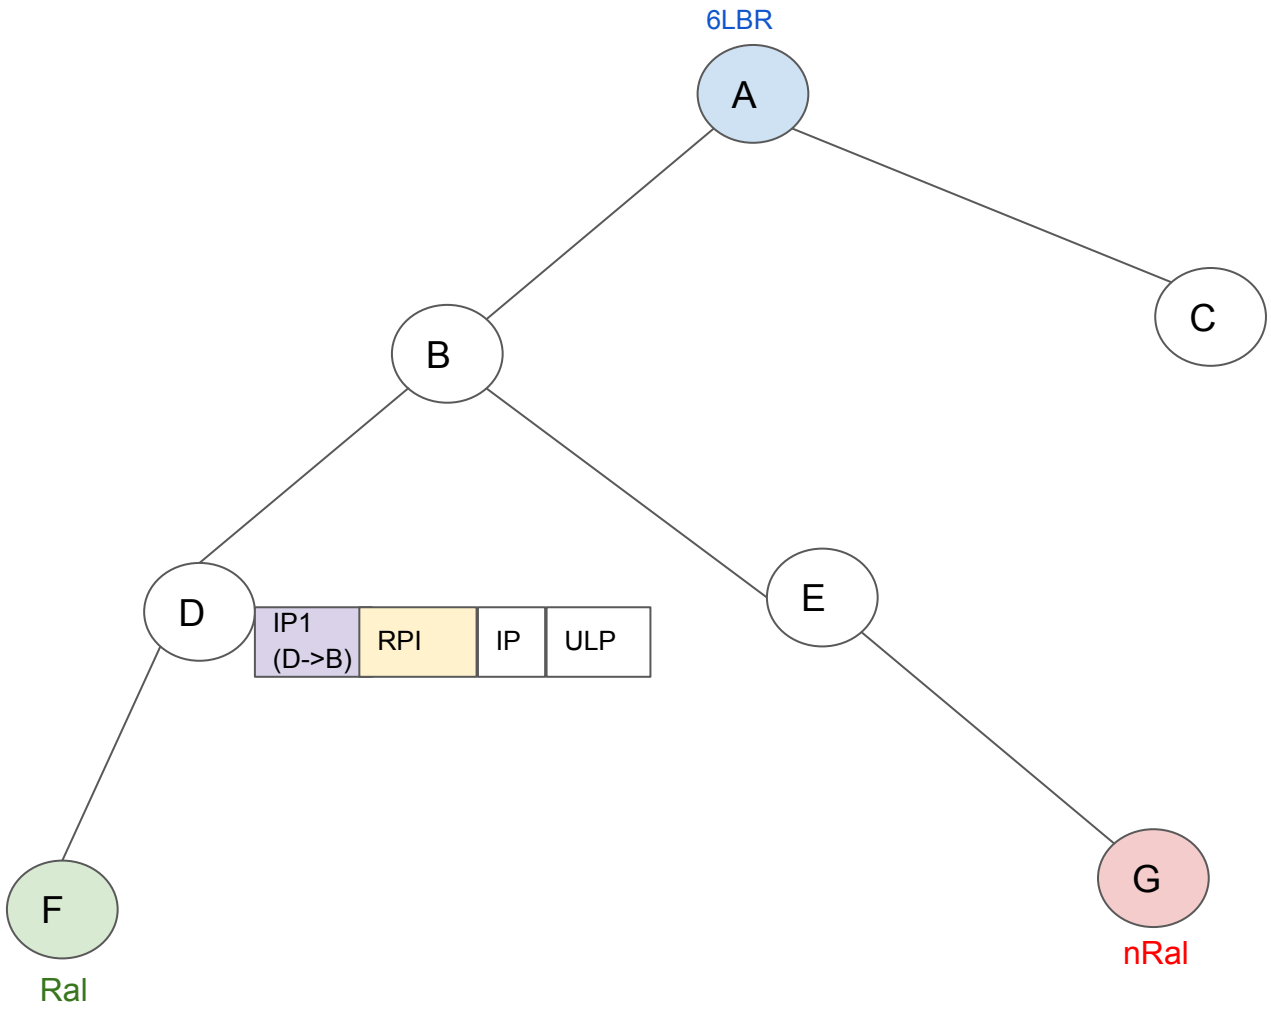

Storing mode: **without** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

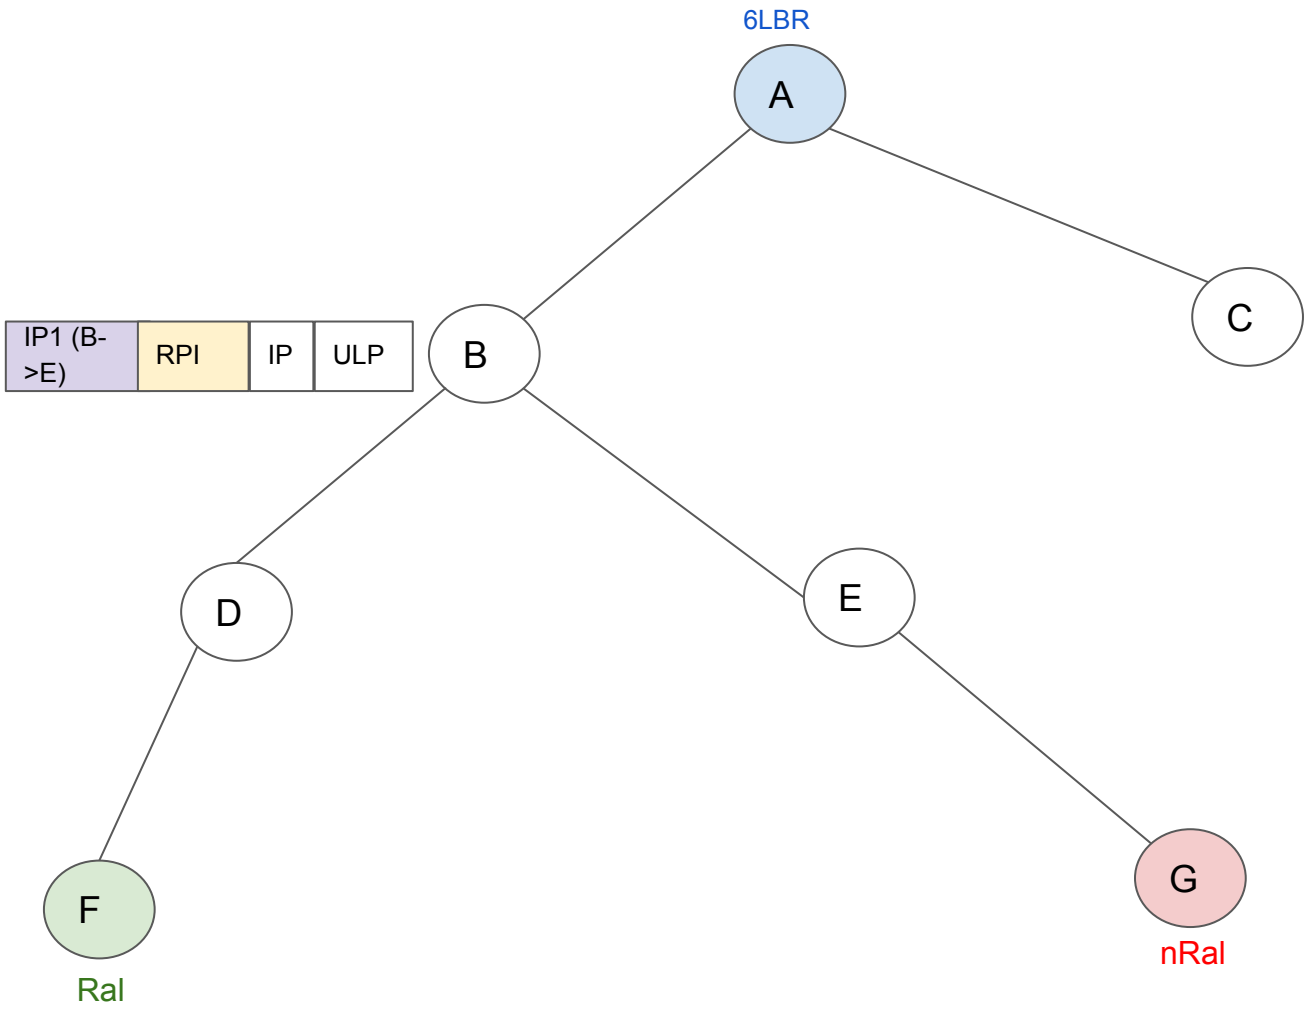

Storing mode: **without** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

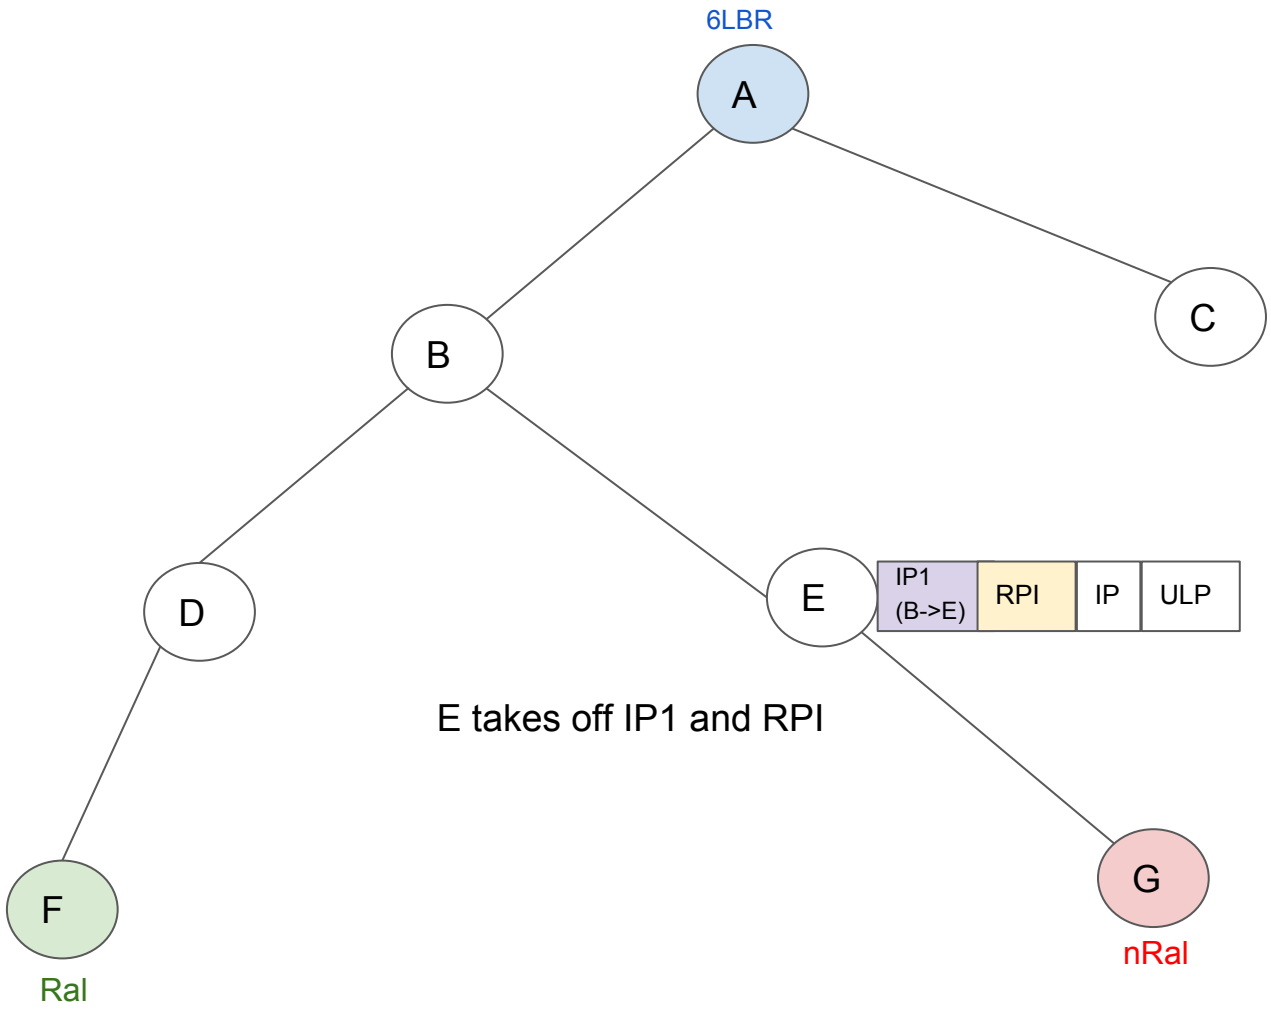

Storing mode: **without** the inclusion of 2460bis proposal

Form Ral (Rpl-aware-leaf) to nRal (not-Rpl-aware-leaf)

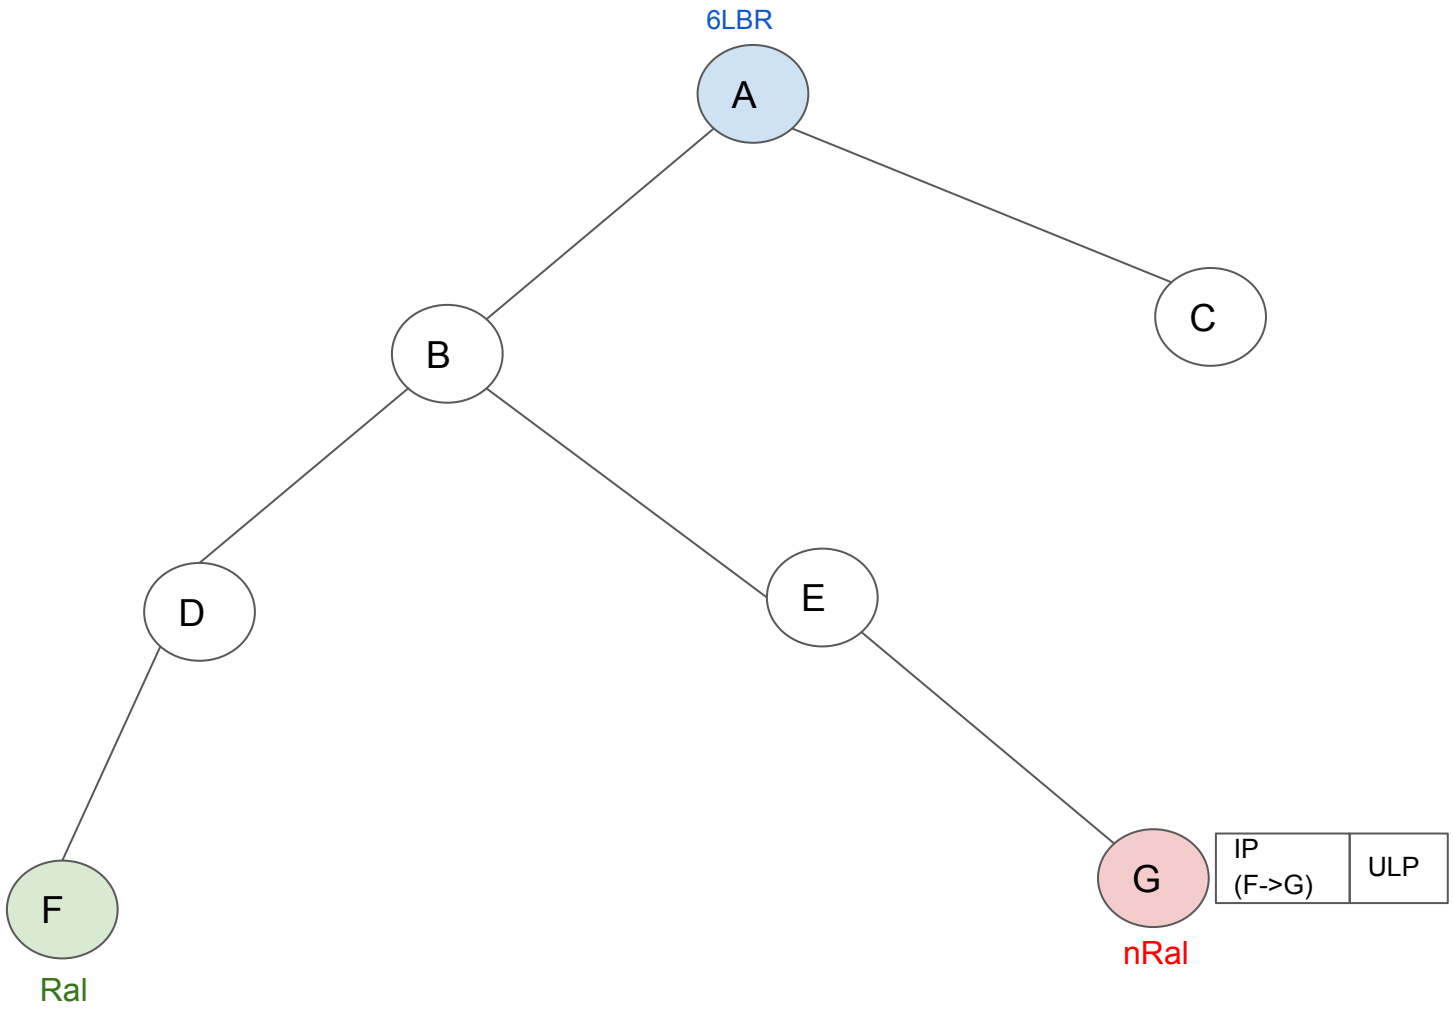

Since in general, F can not know if G is nRal or Ral, F must ALWAYS use a Hop-by-Hop IPIP header.

Storing mode: **with** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

Storing mode: **with** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

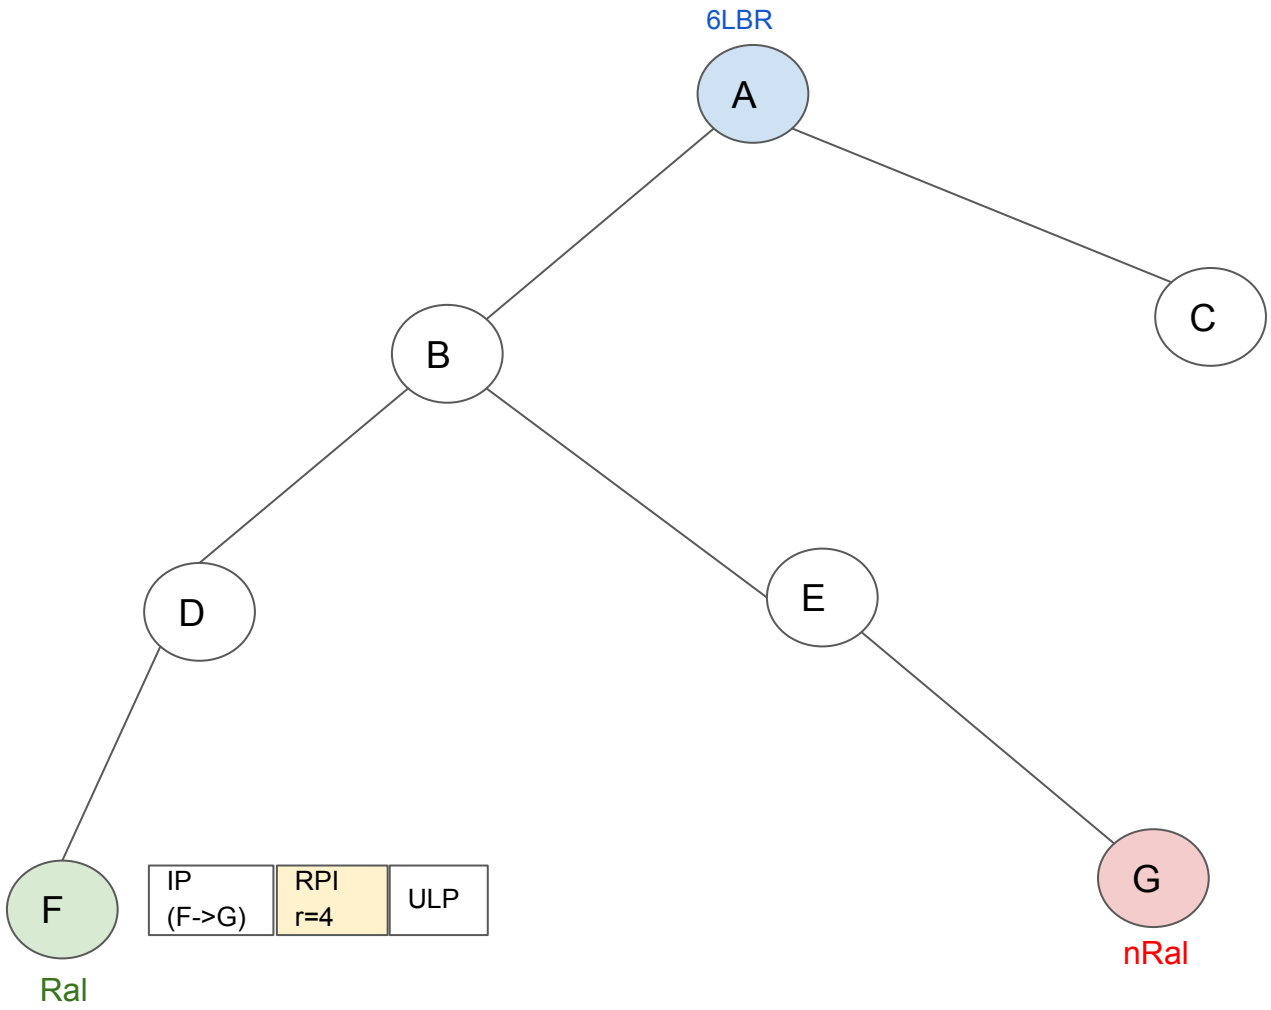

We do not need encapsulation!!

Storing mode: **with** the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

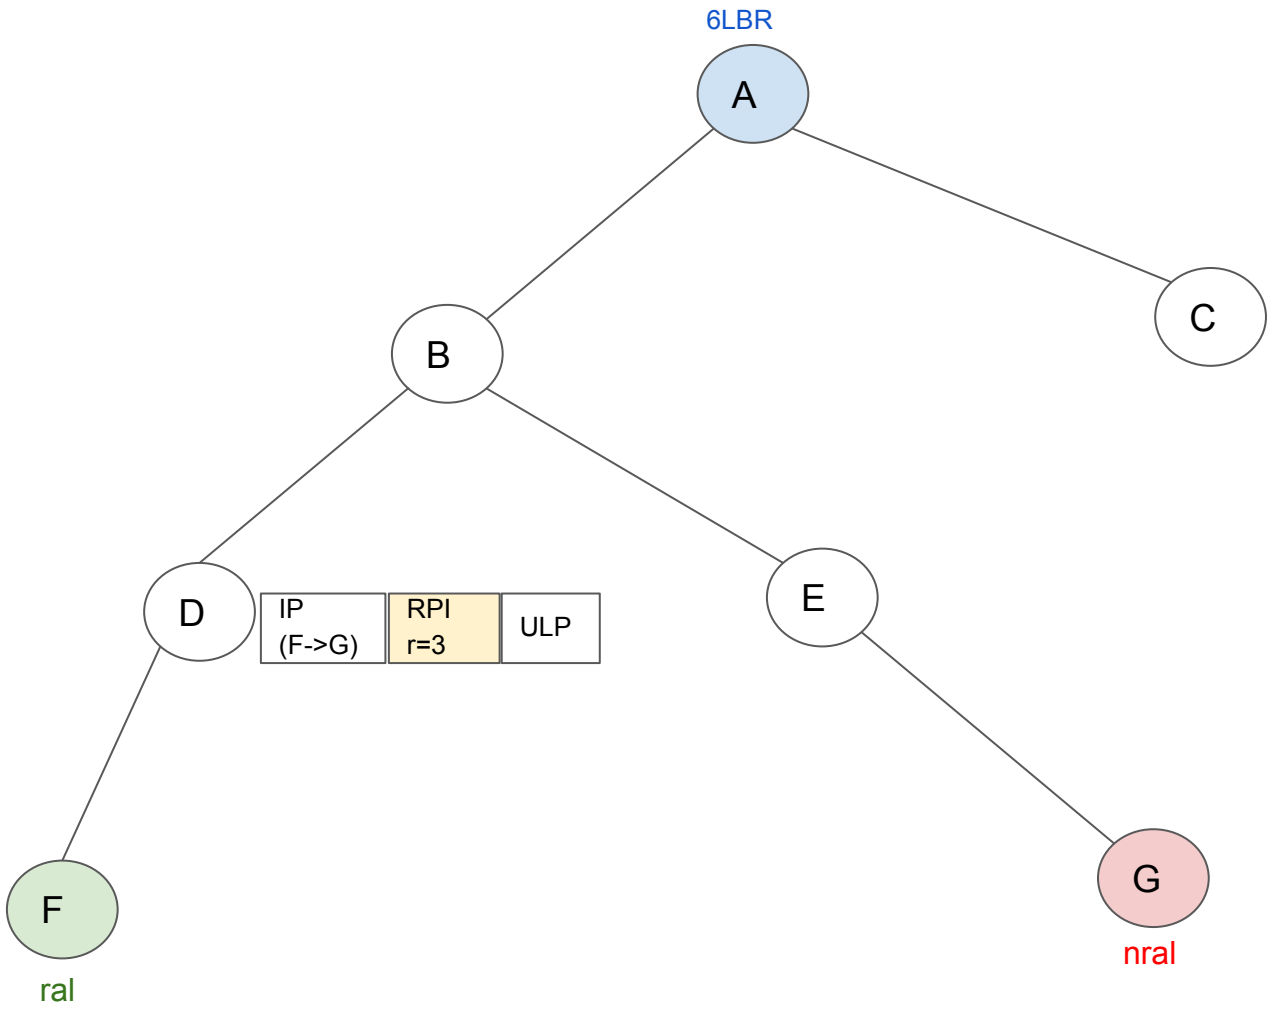

Storing mode: **with** the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

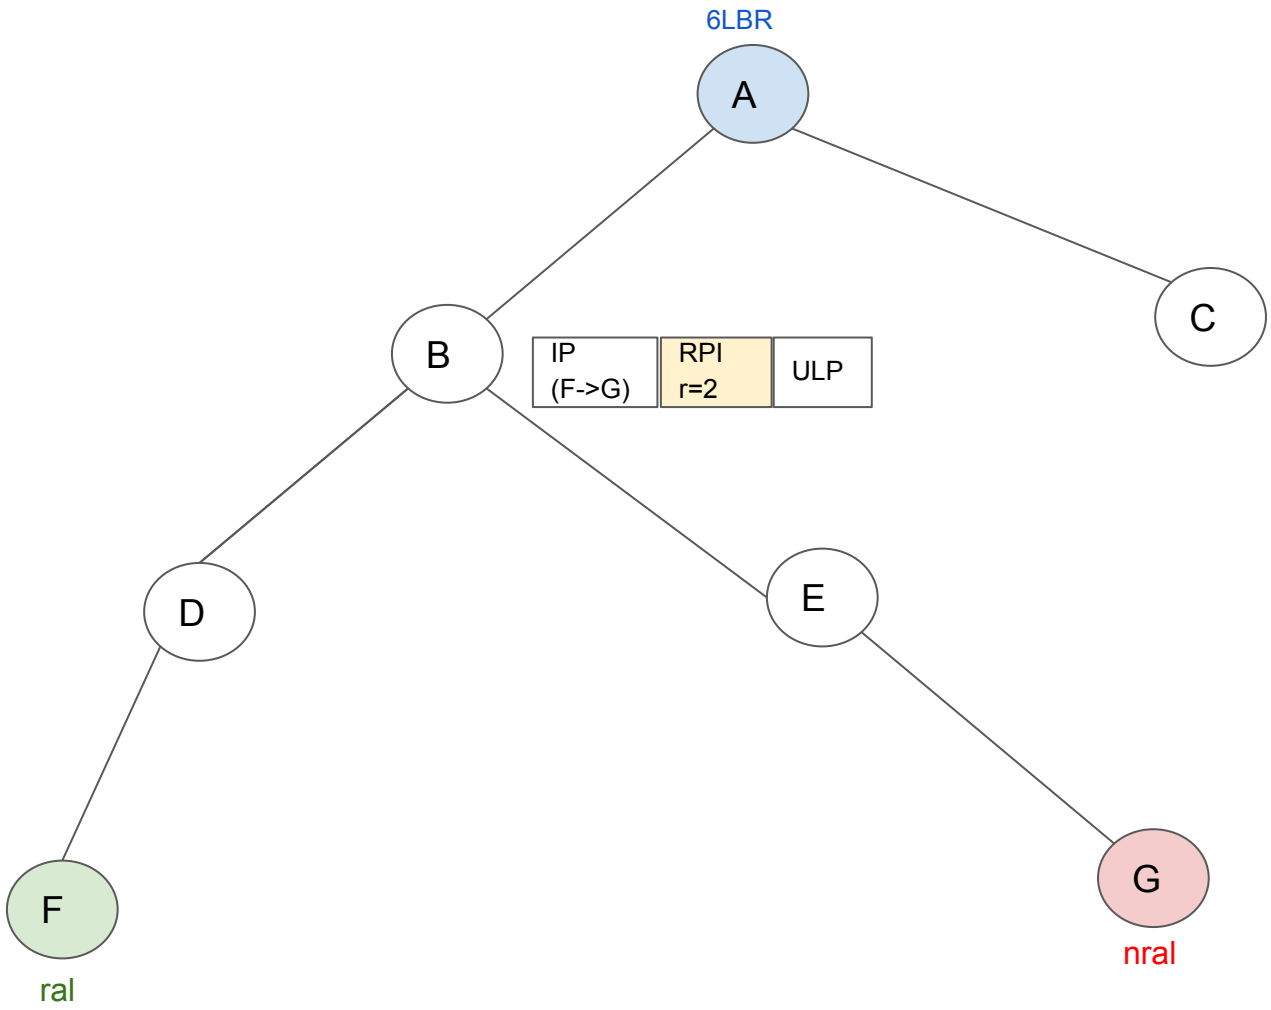

Storing mode: **with** the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

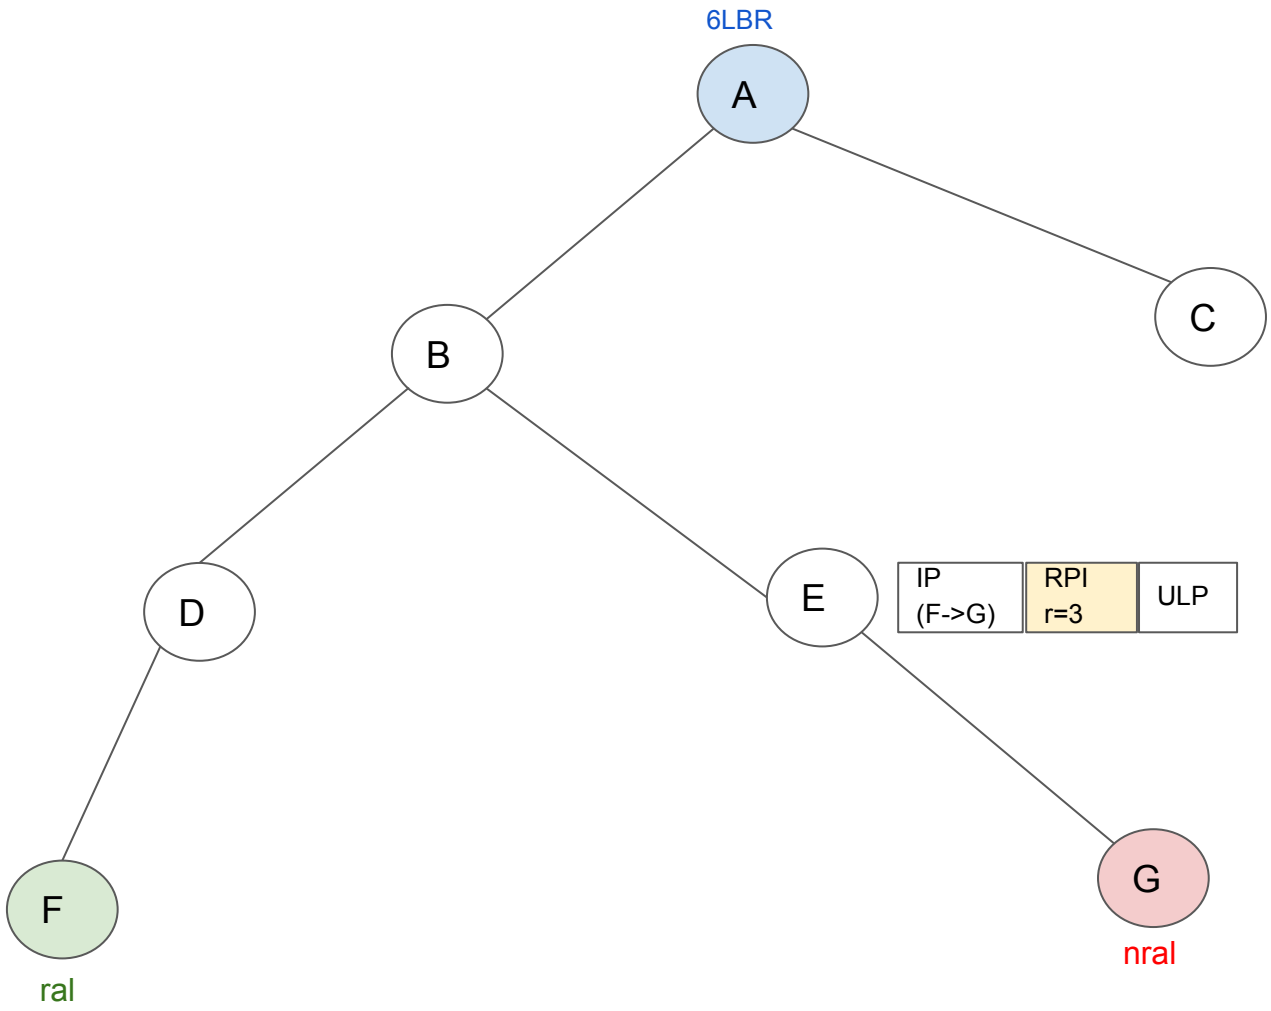

Storing mode: **with** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (non-Rpl-aware-leaf)

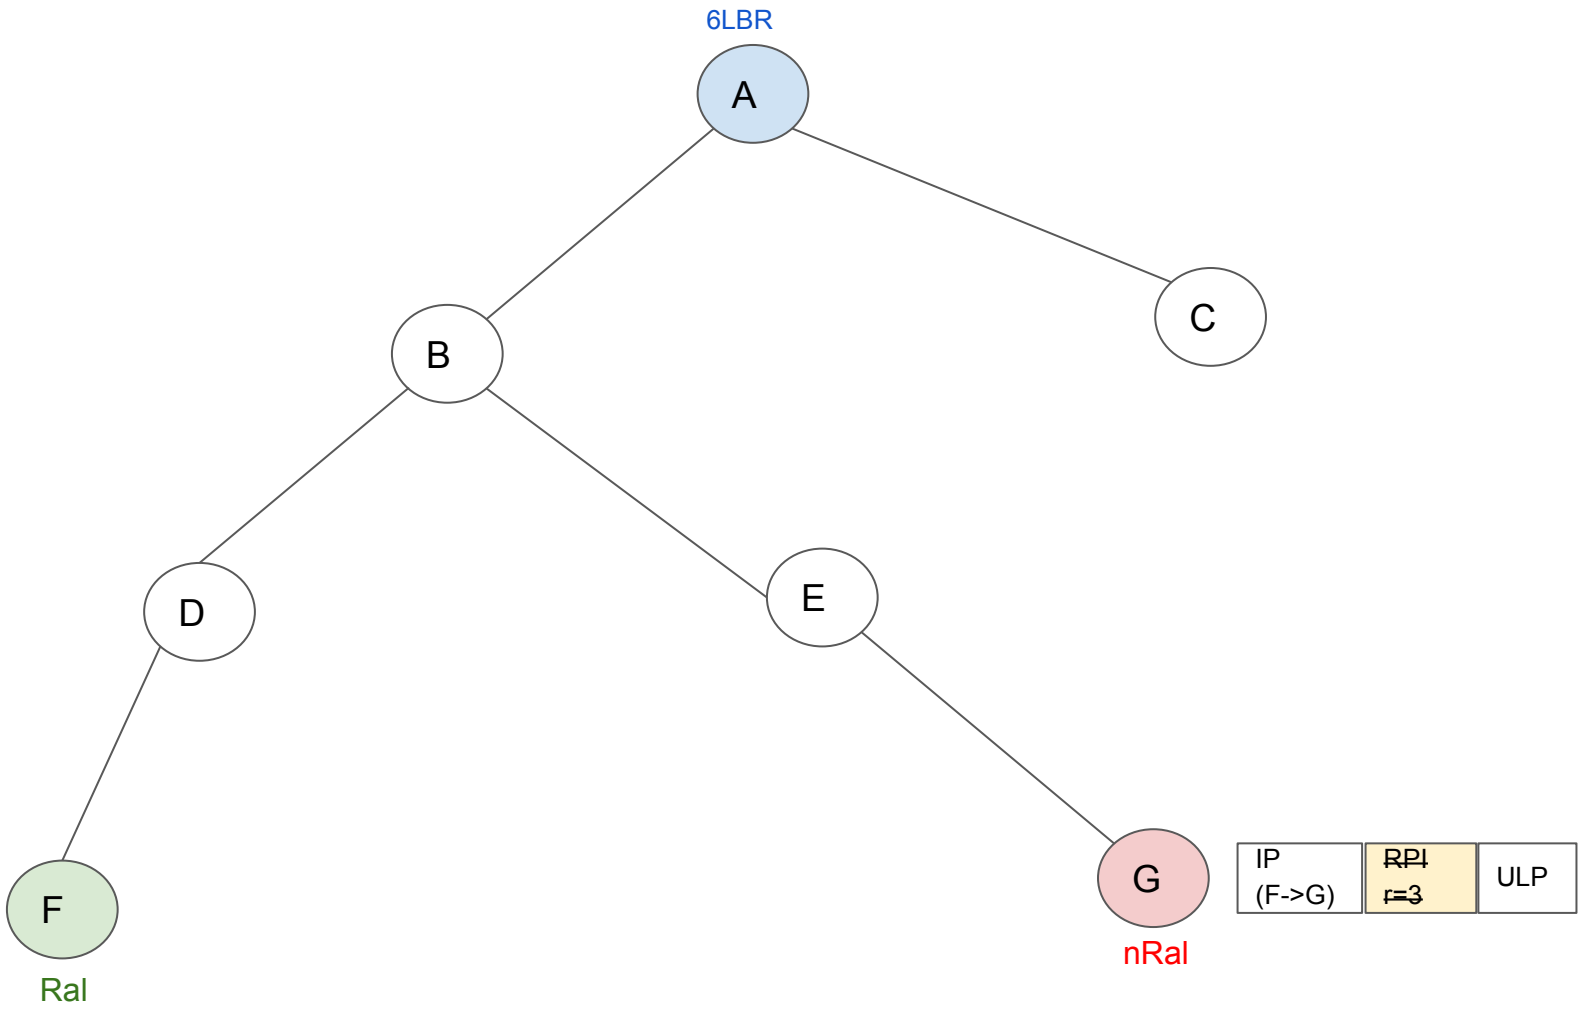

G ignores the option and processes the packet

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

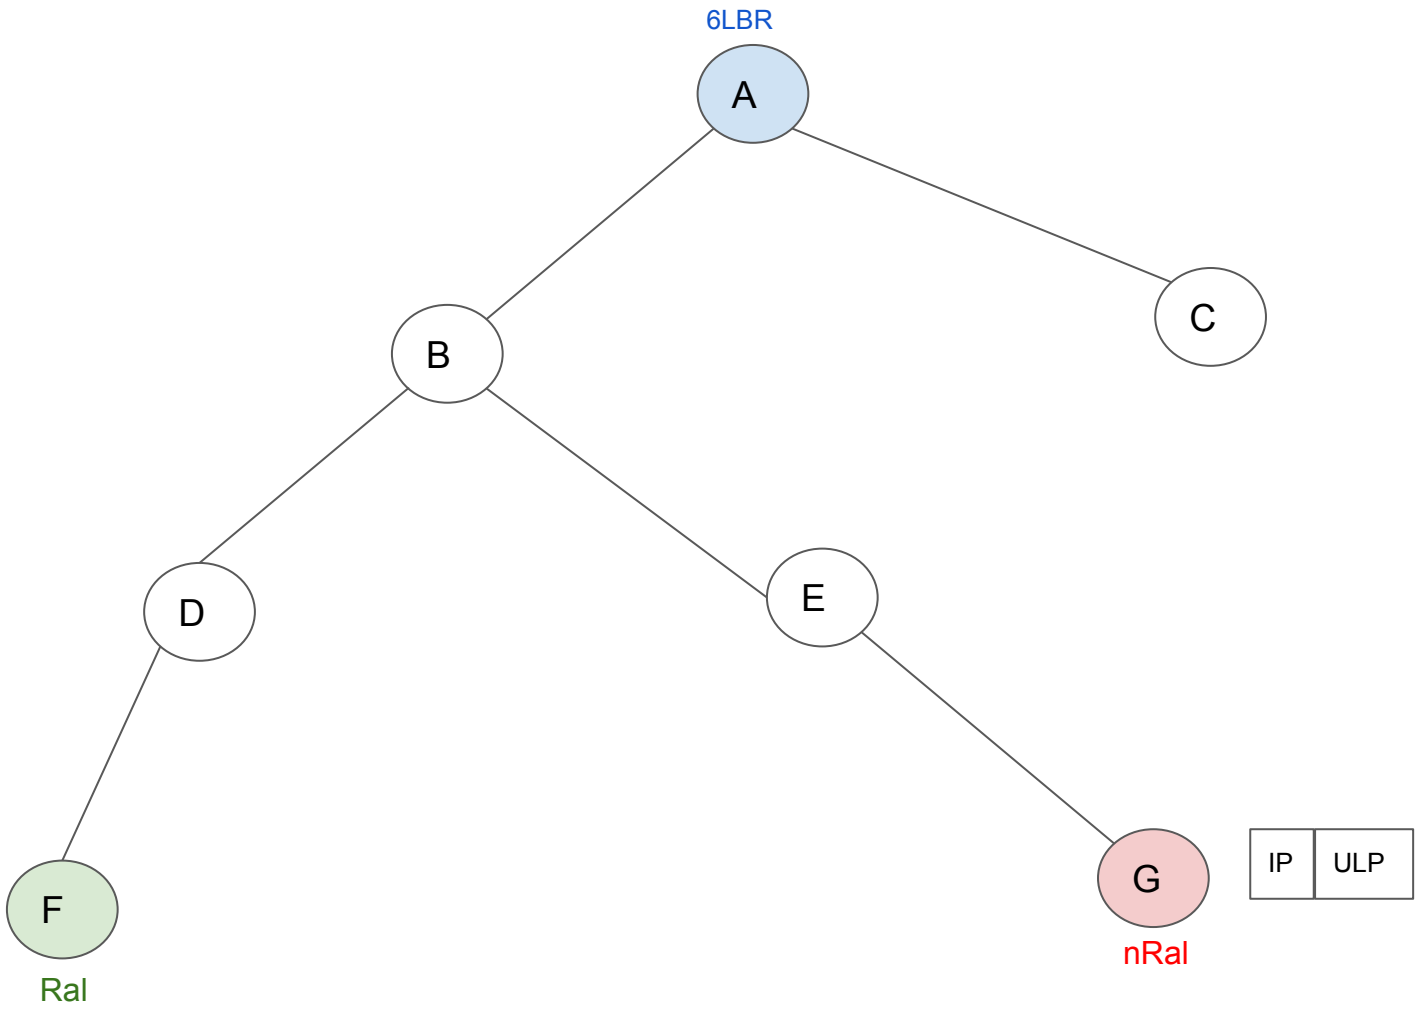

G sends a packet to E

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

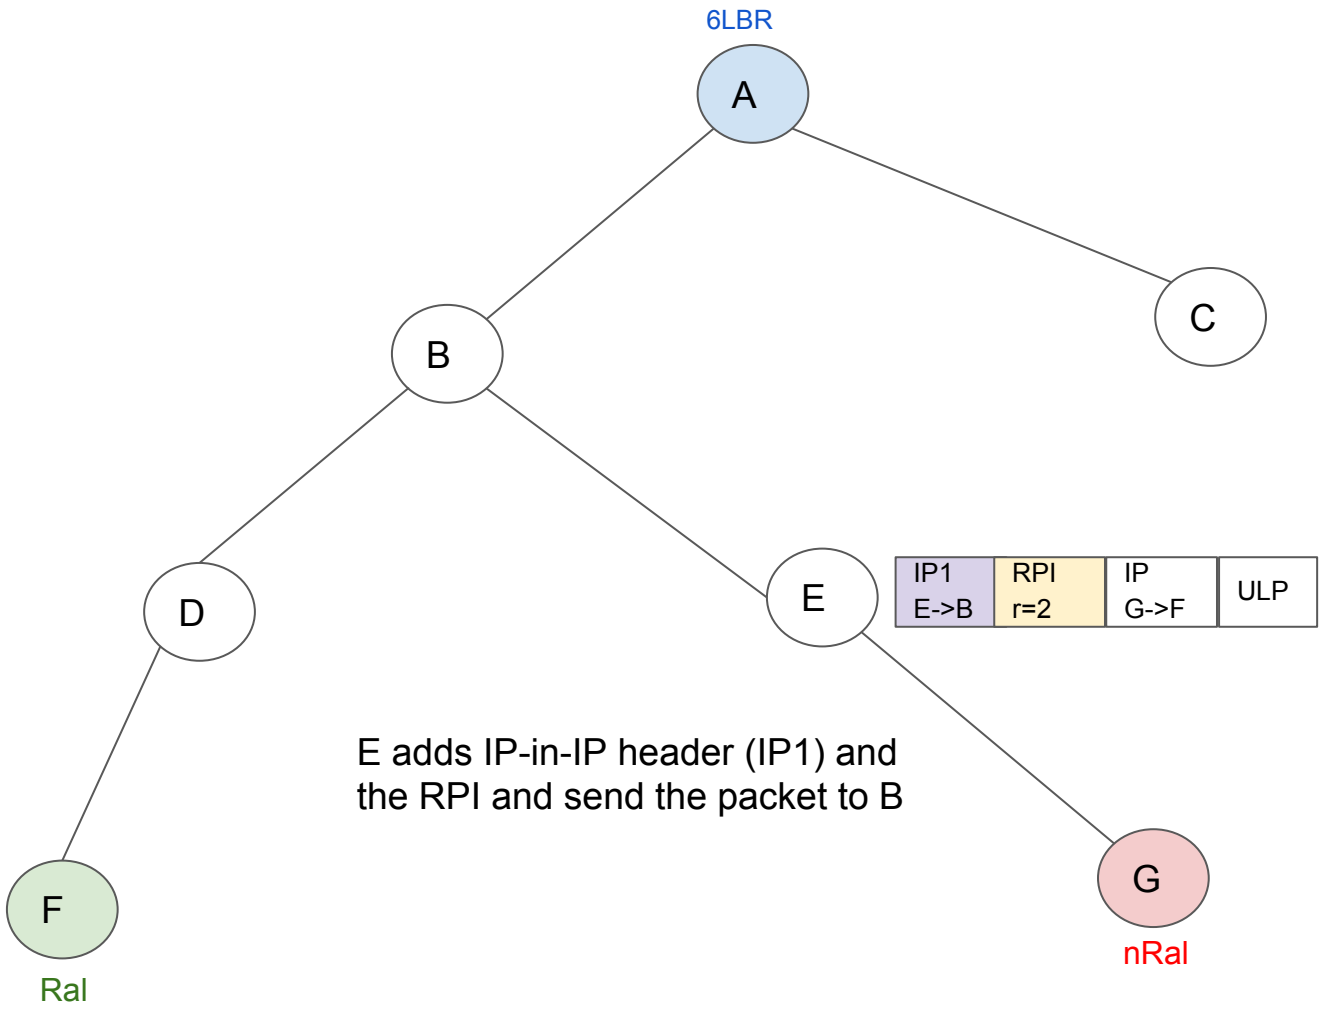

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

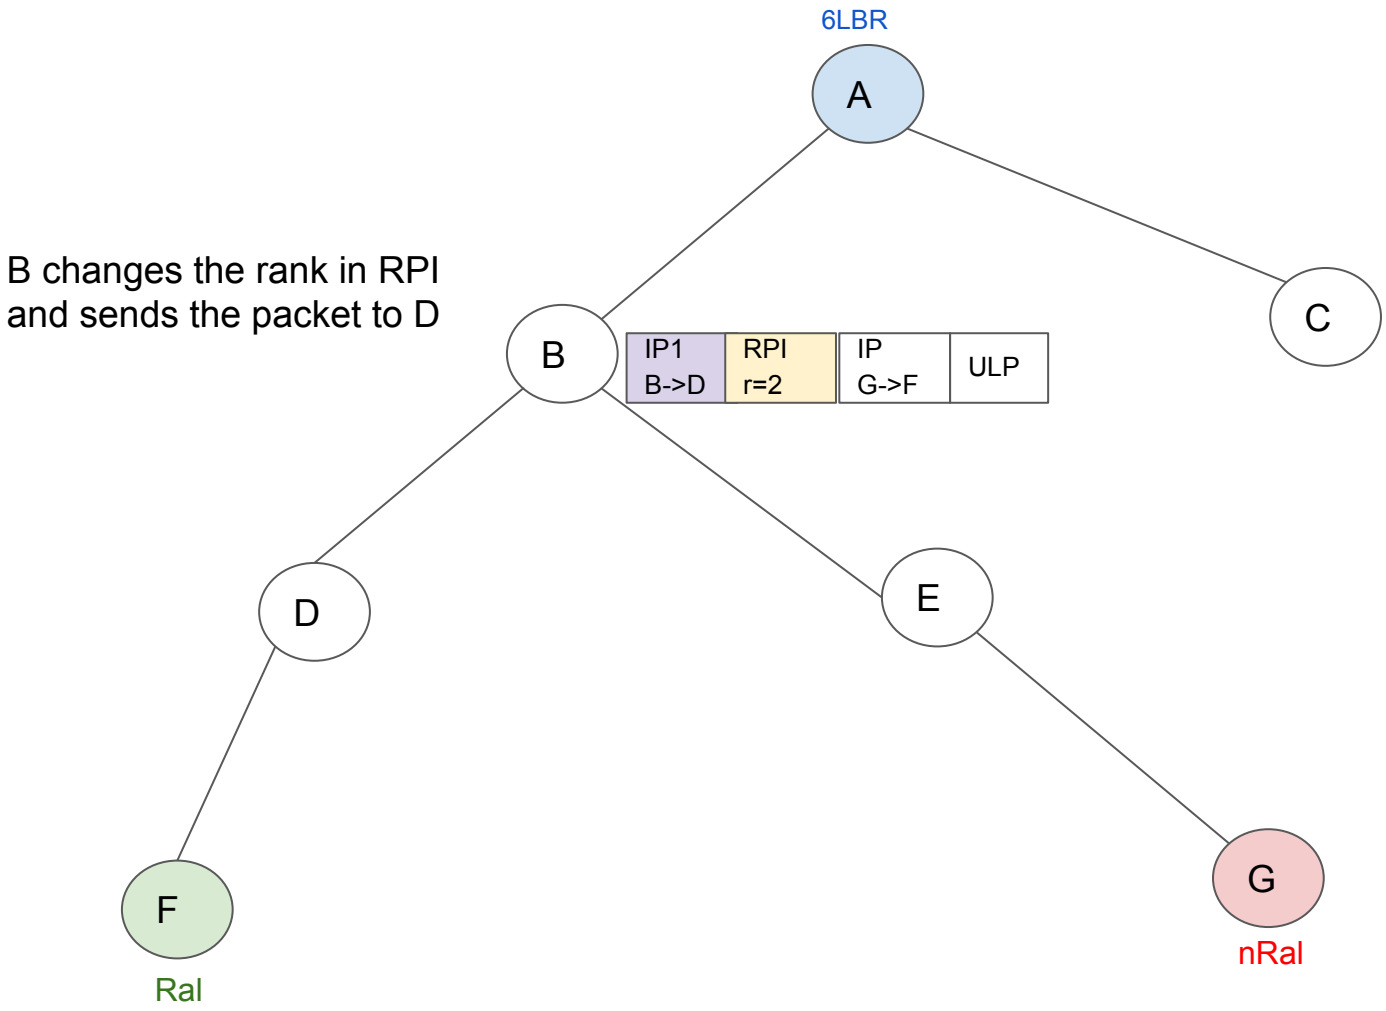

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

D removes the IP1 and RPI and sends the packet to F

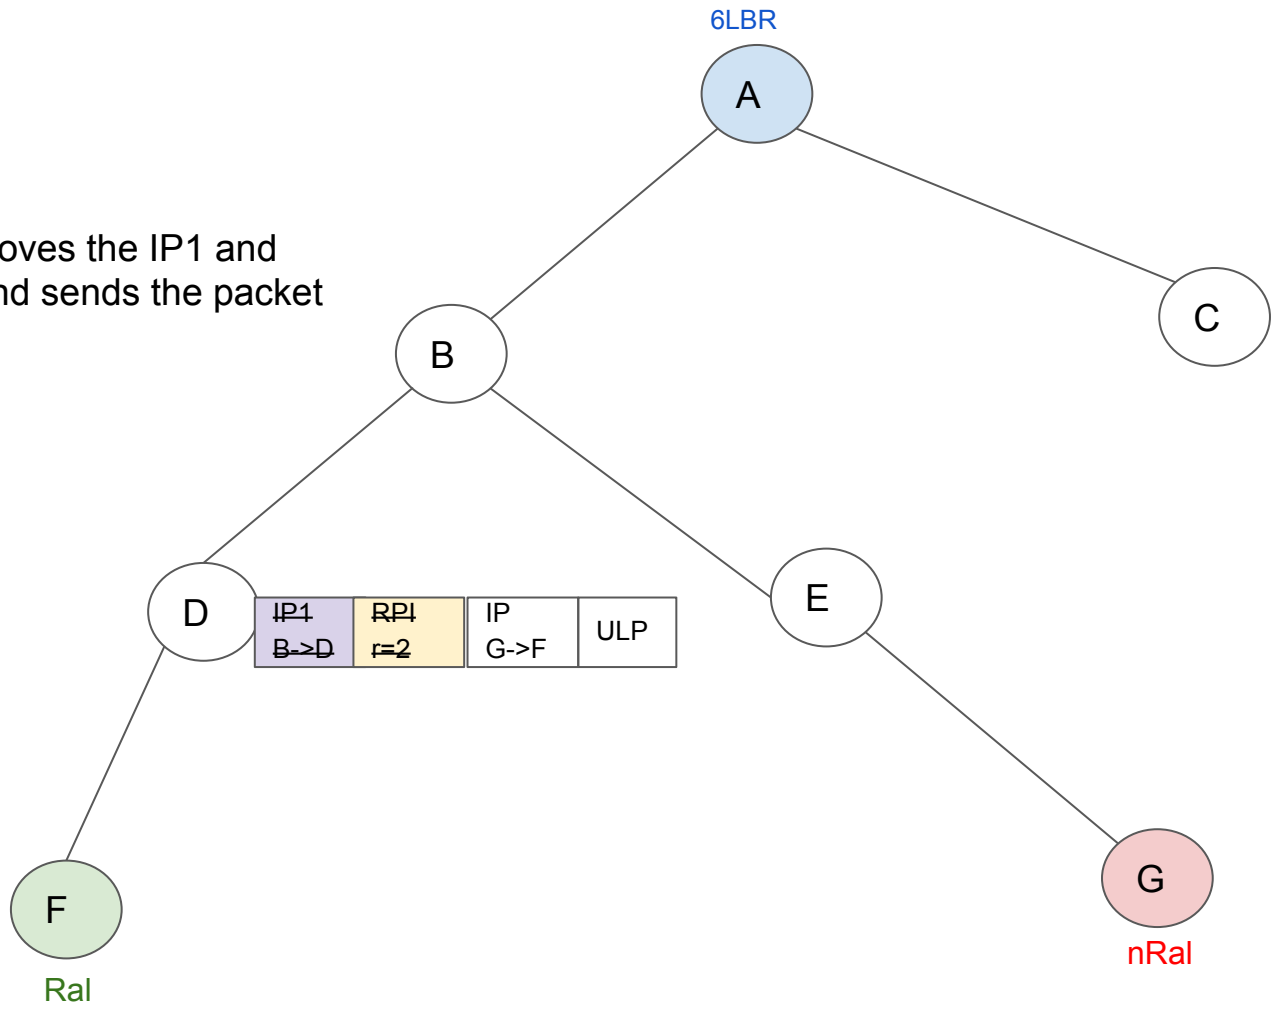

Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

F gets the packet and process it.

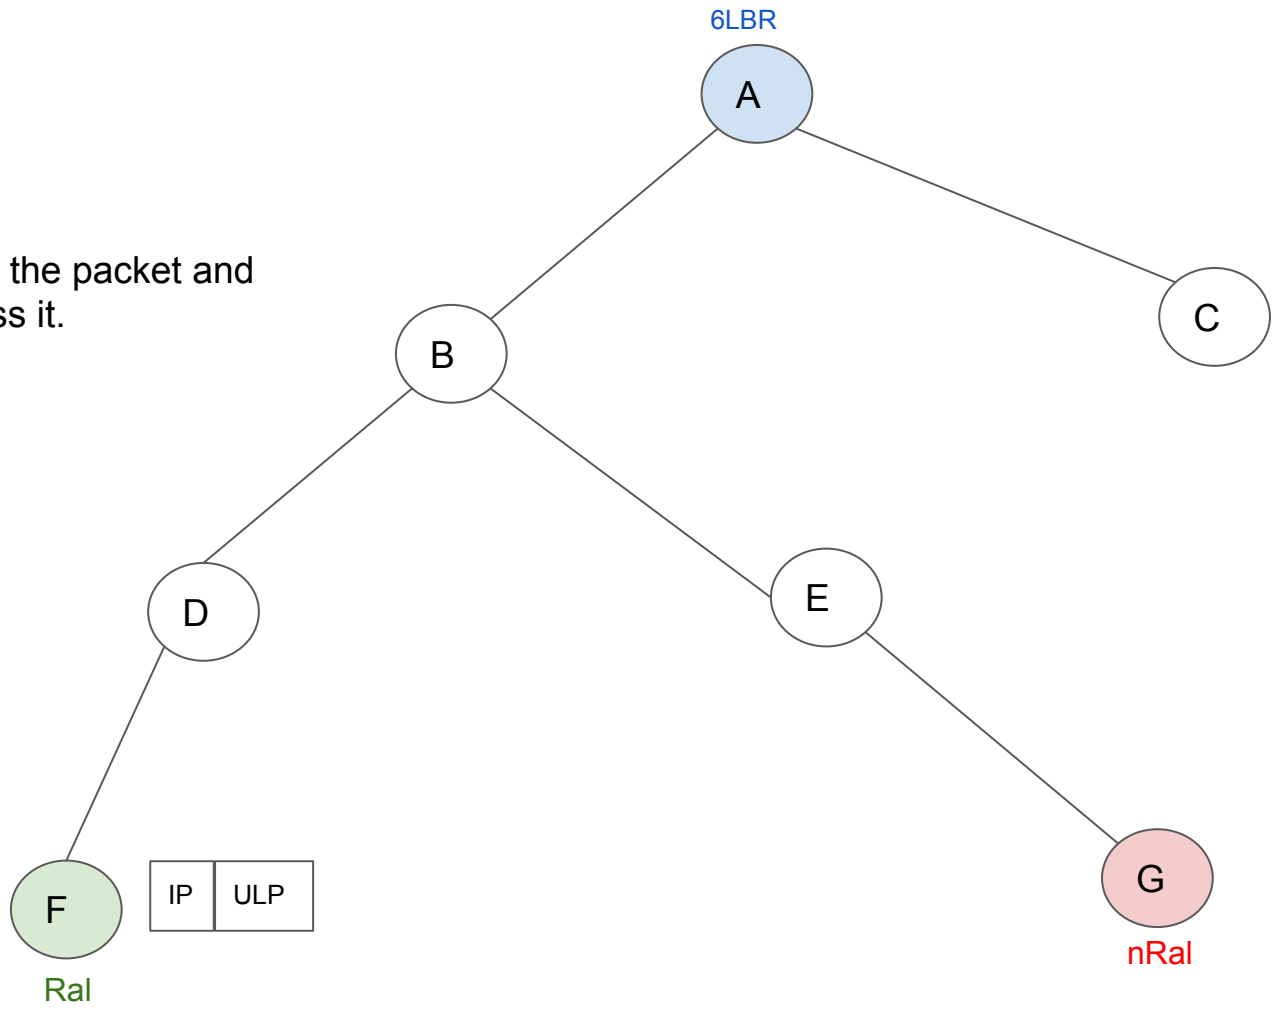

Storing mode: **with** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

Storing mode: **with** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

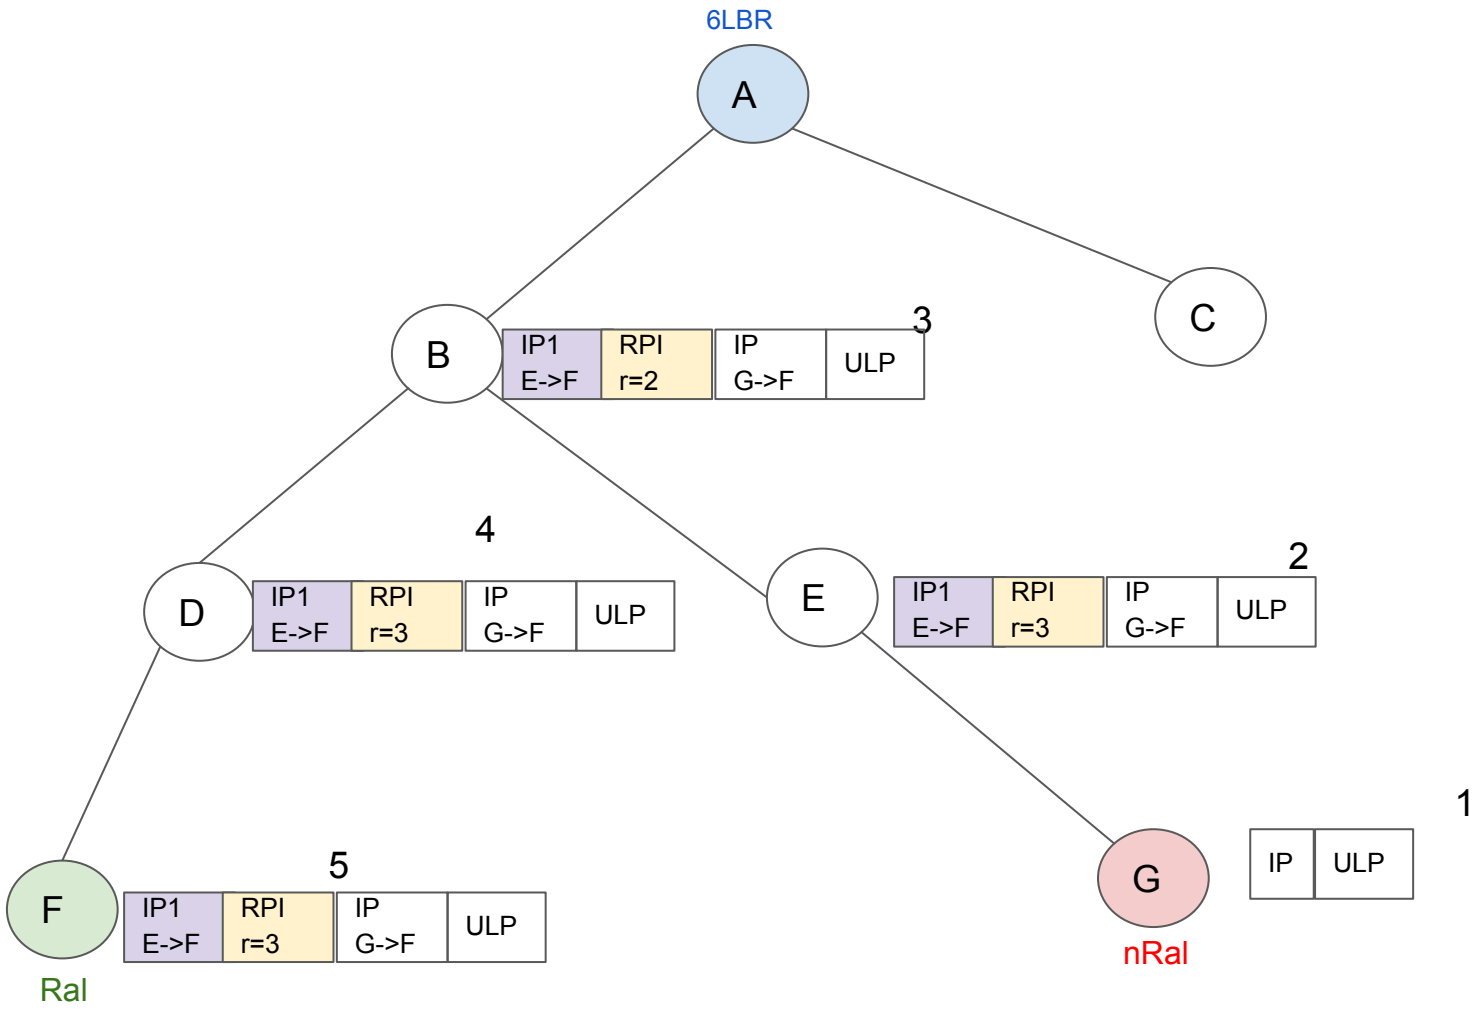

Non-Storing mode: **without** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

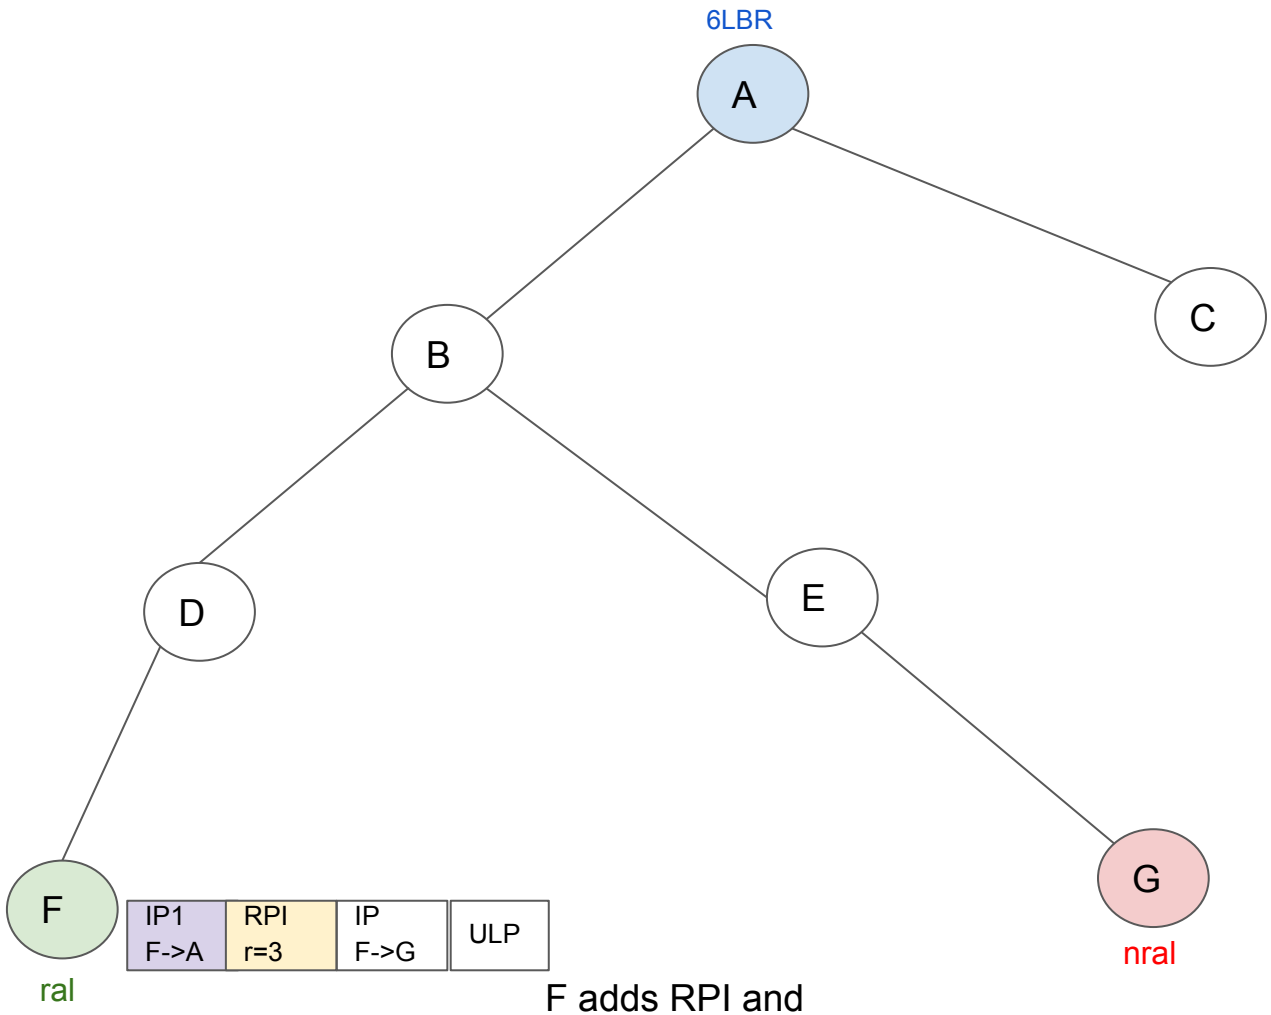

F adds RPI and
“Knowing” that G is
nRAL, sends to A

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

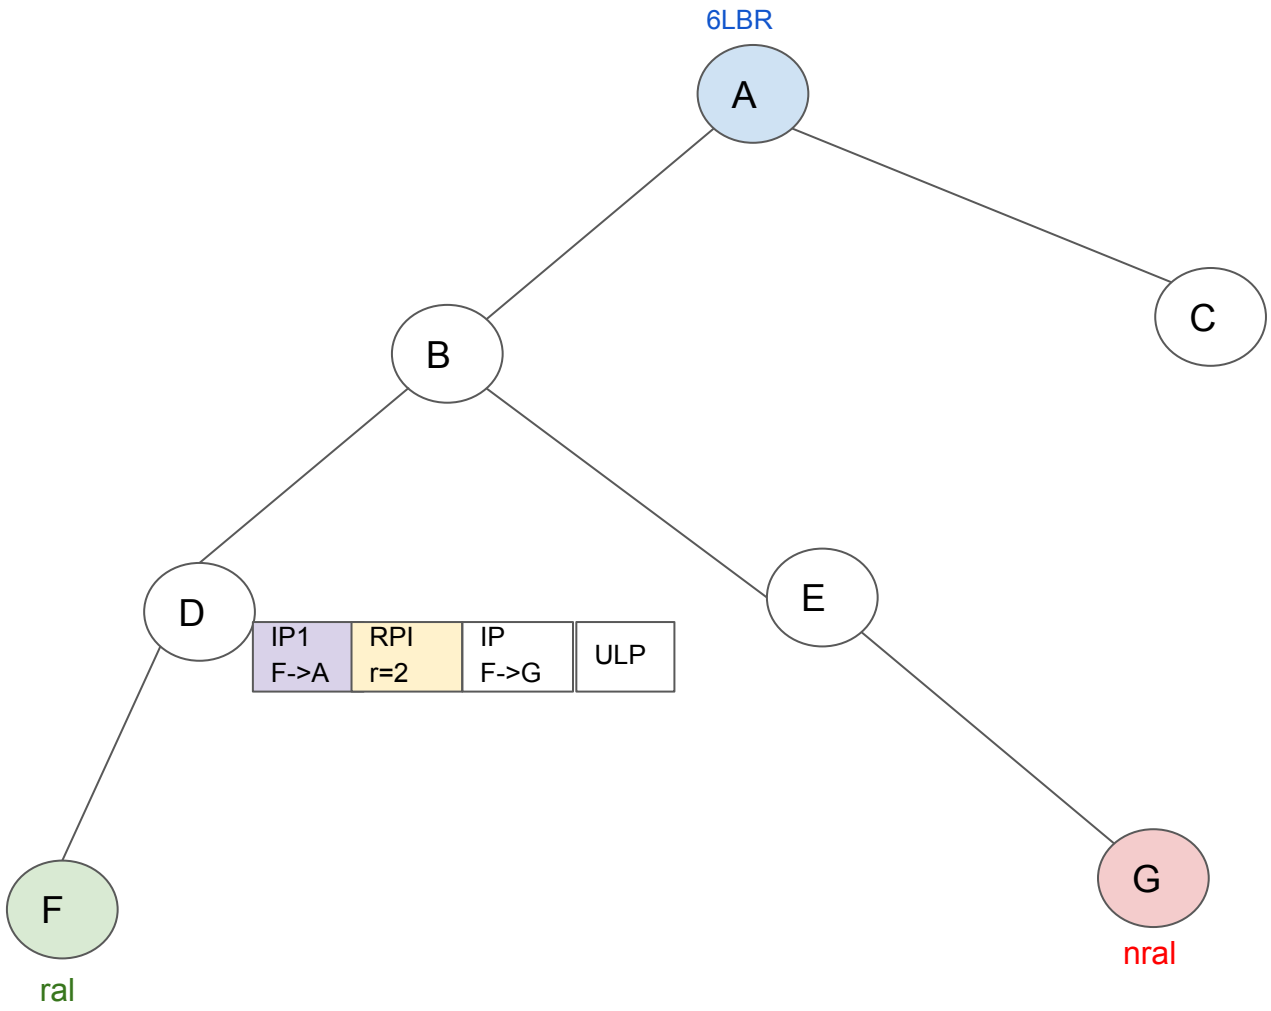

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

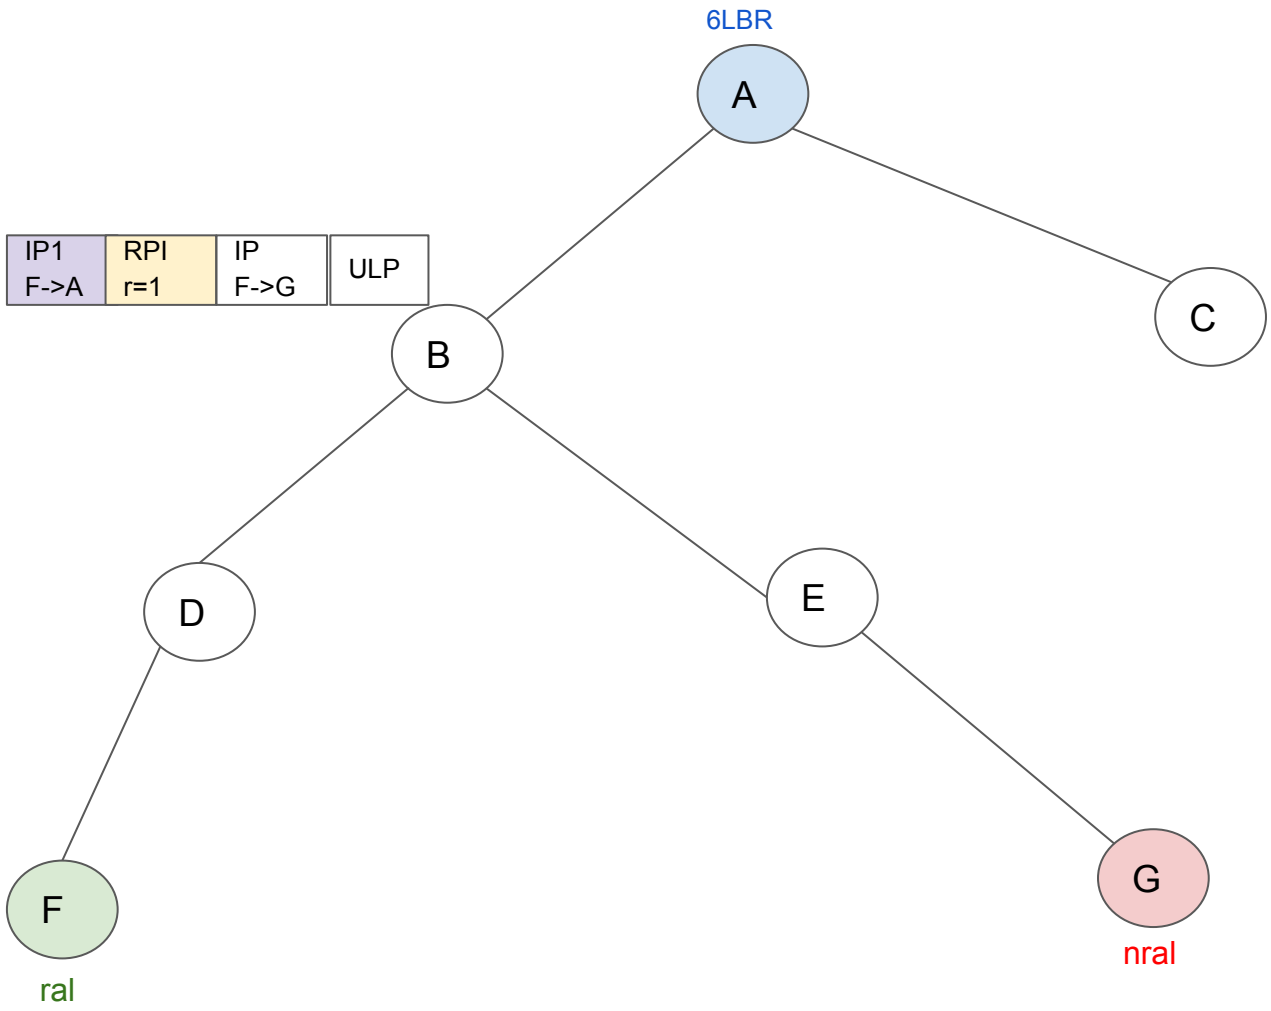

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

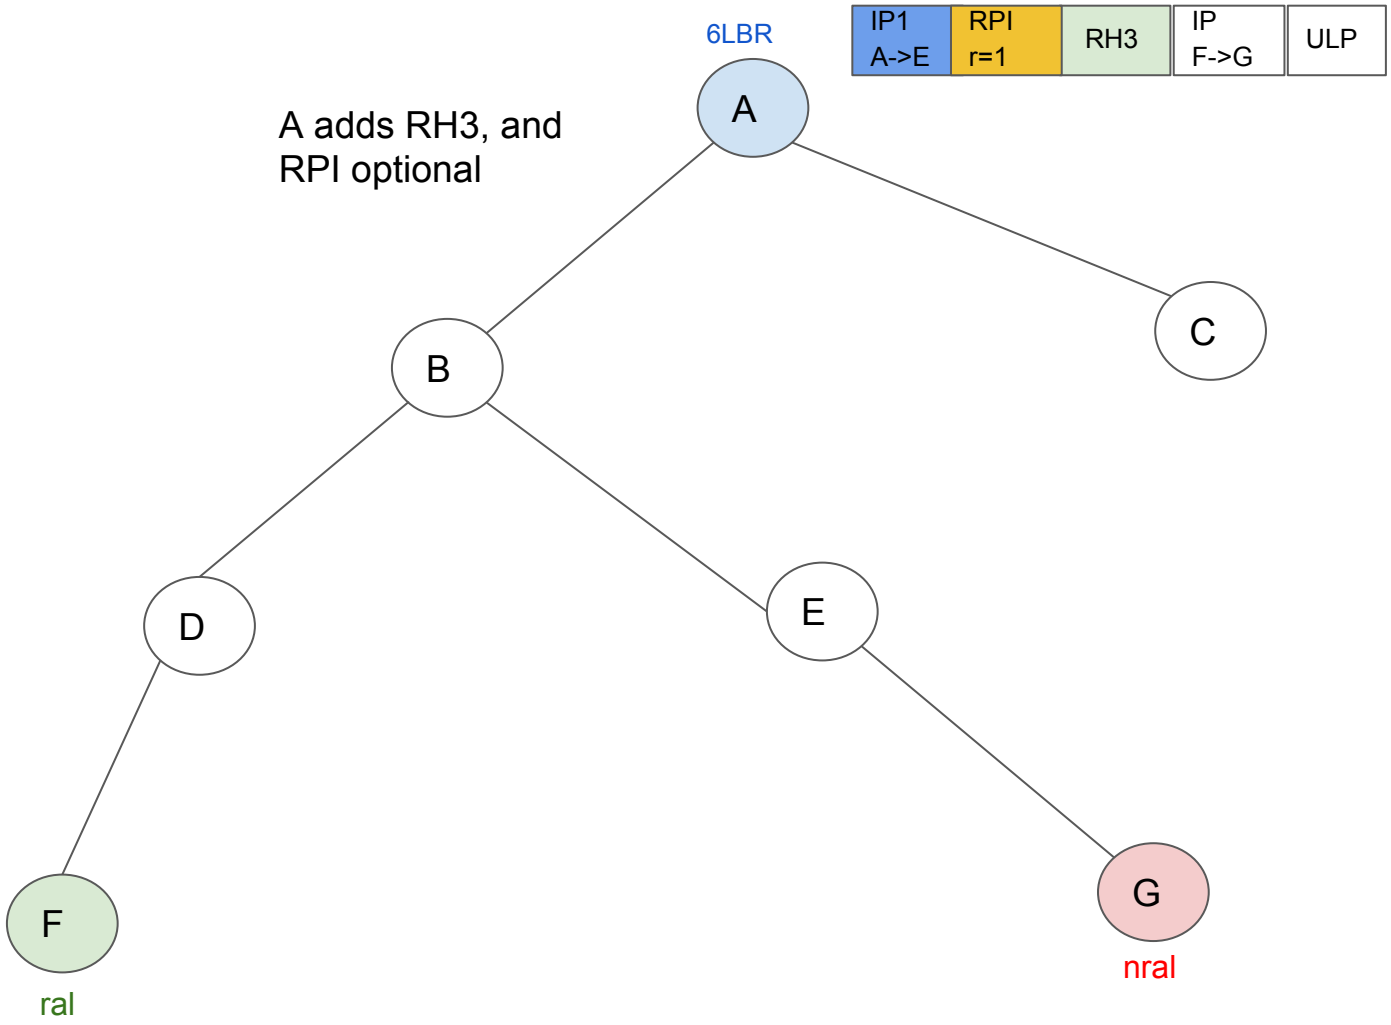

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

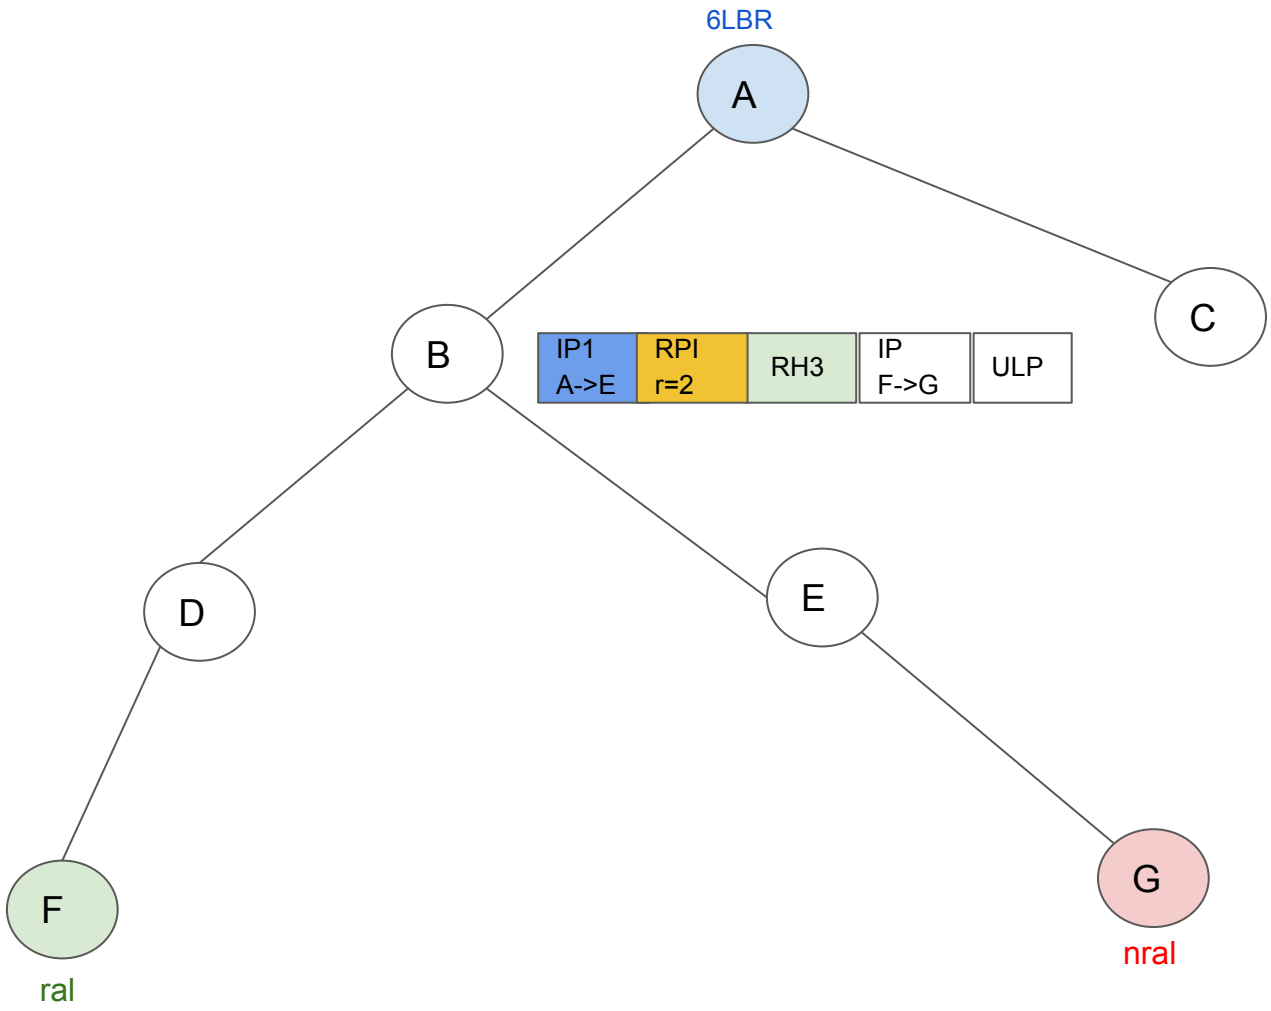

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

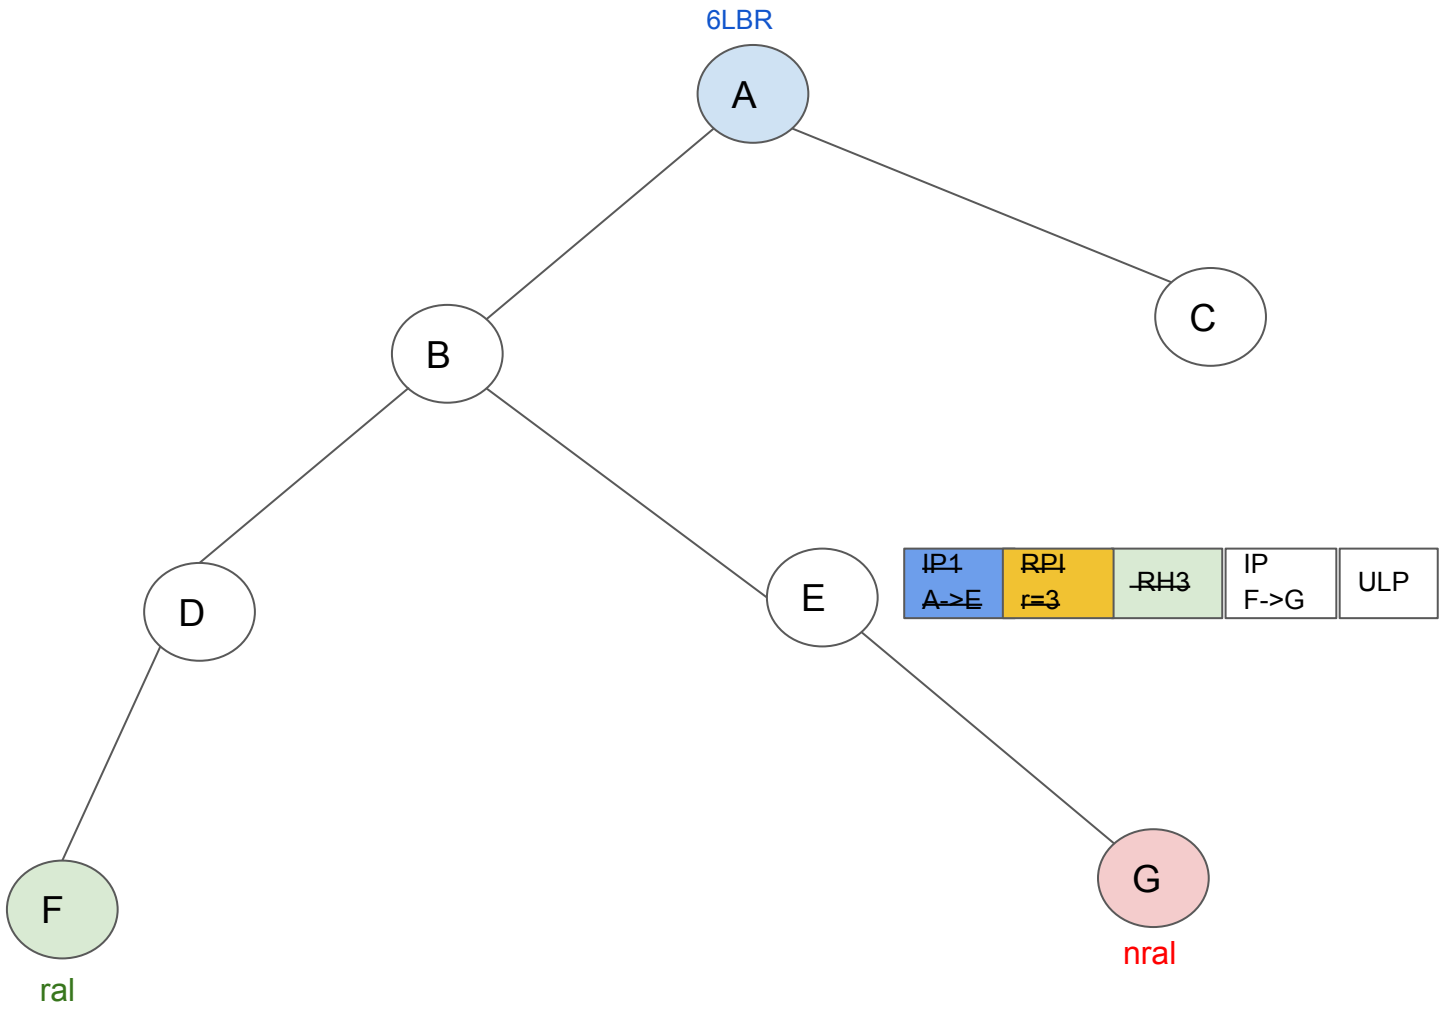

Non-Storing mode: without the inclusion of 2460bis proposal

Form ral (rpl-aware-leaf) to nral (not-rpl-aware-leaf)

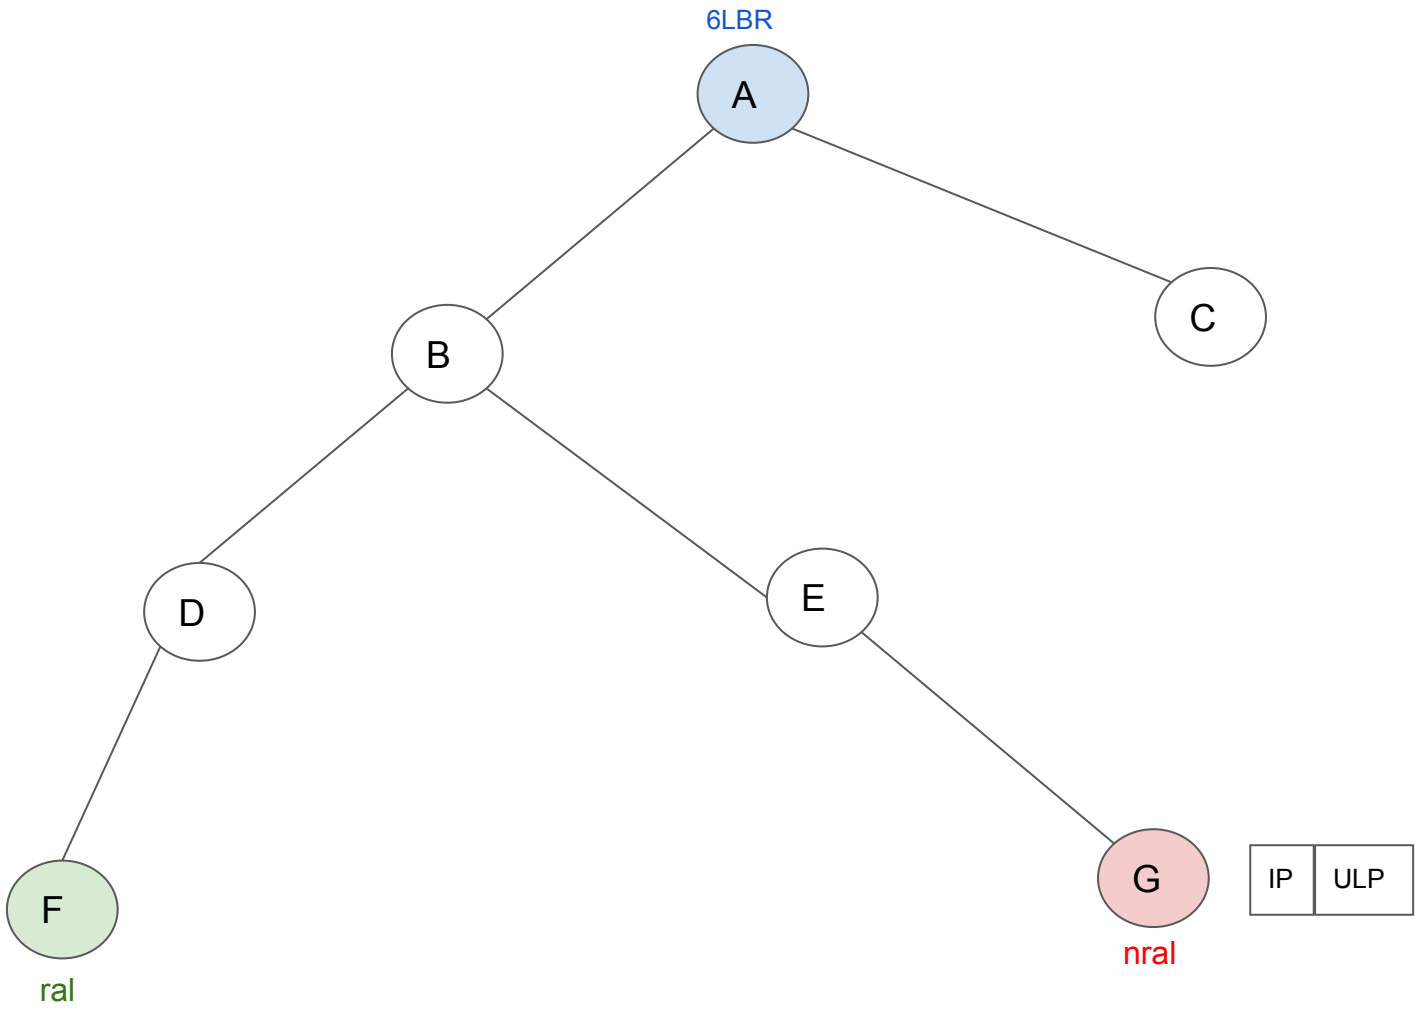

Non-Storing mode: **with** the inclusion of 2460bis proposal

Form RaI (Rpl-aware-leaf) to nRaI (not-Rpl-aware-leaf)

The Same Case as previous one

Non-Storing mode: **without** the inclusion of 2460bis proposal

Form nRal (not-Rpl-aware-leaf) to Ral (Rpl-aware-leaf)

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

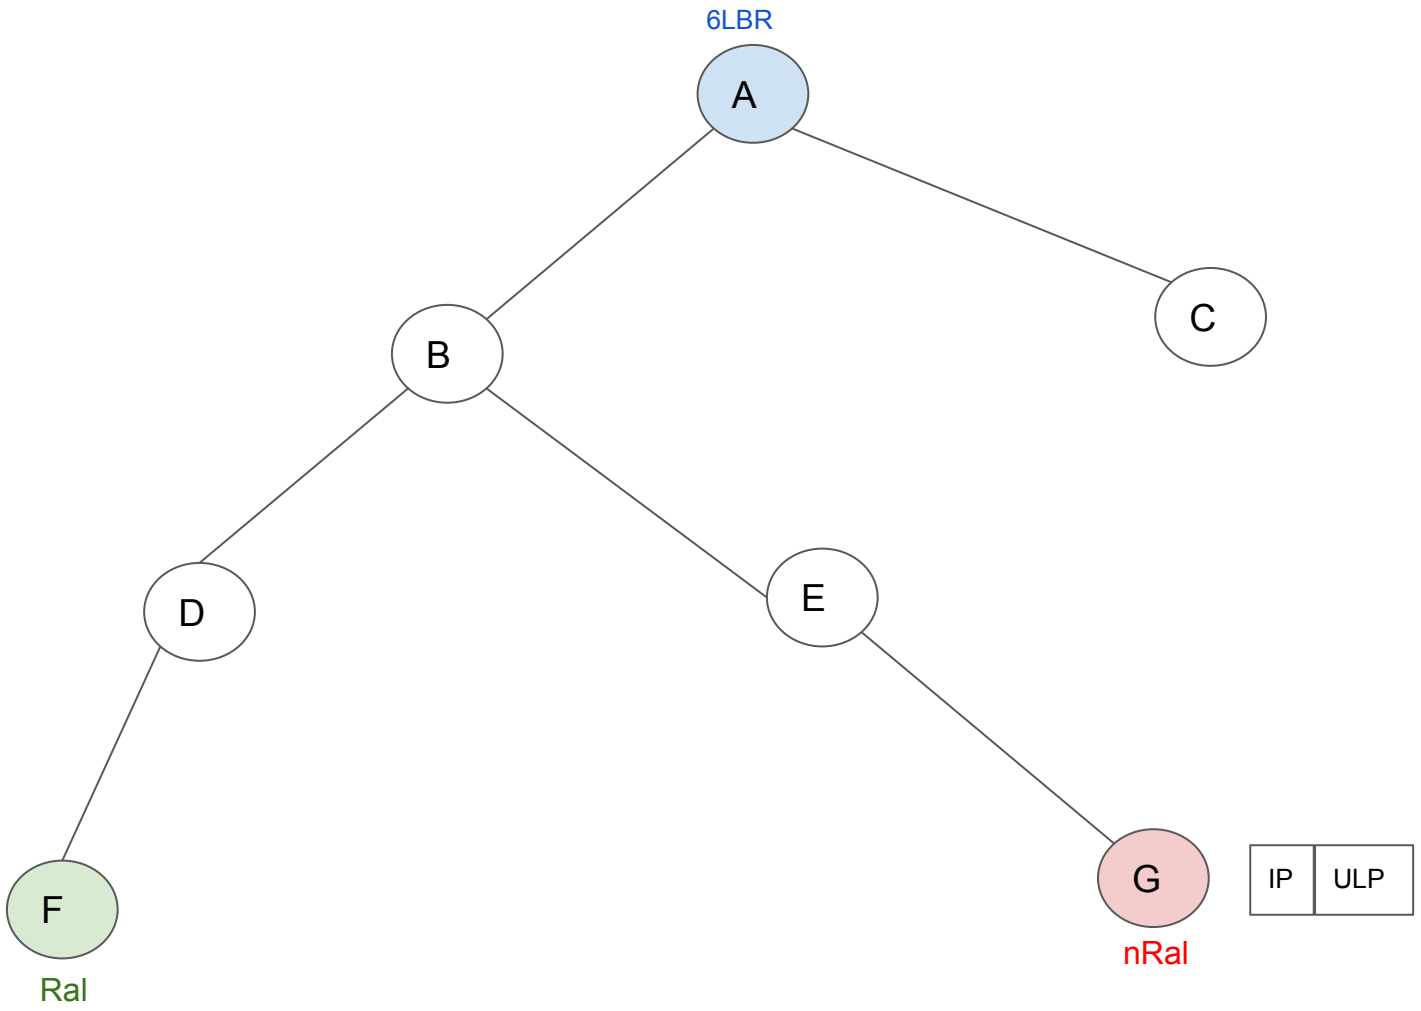

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

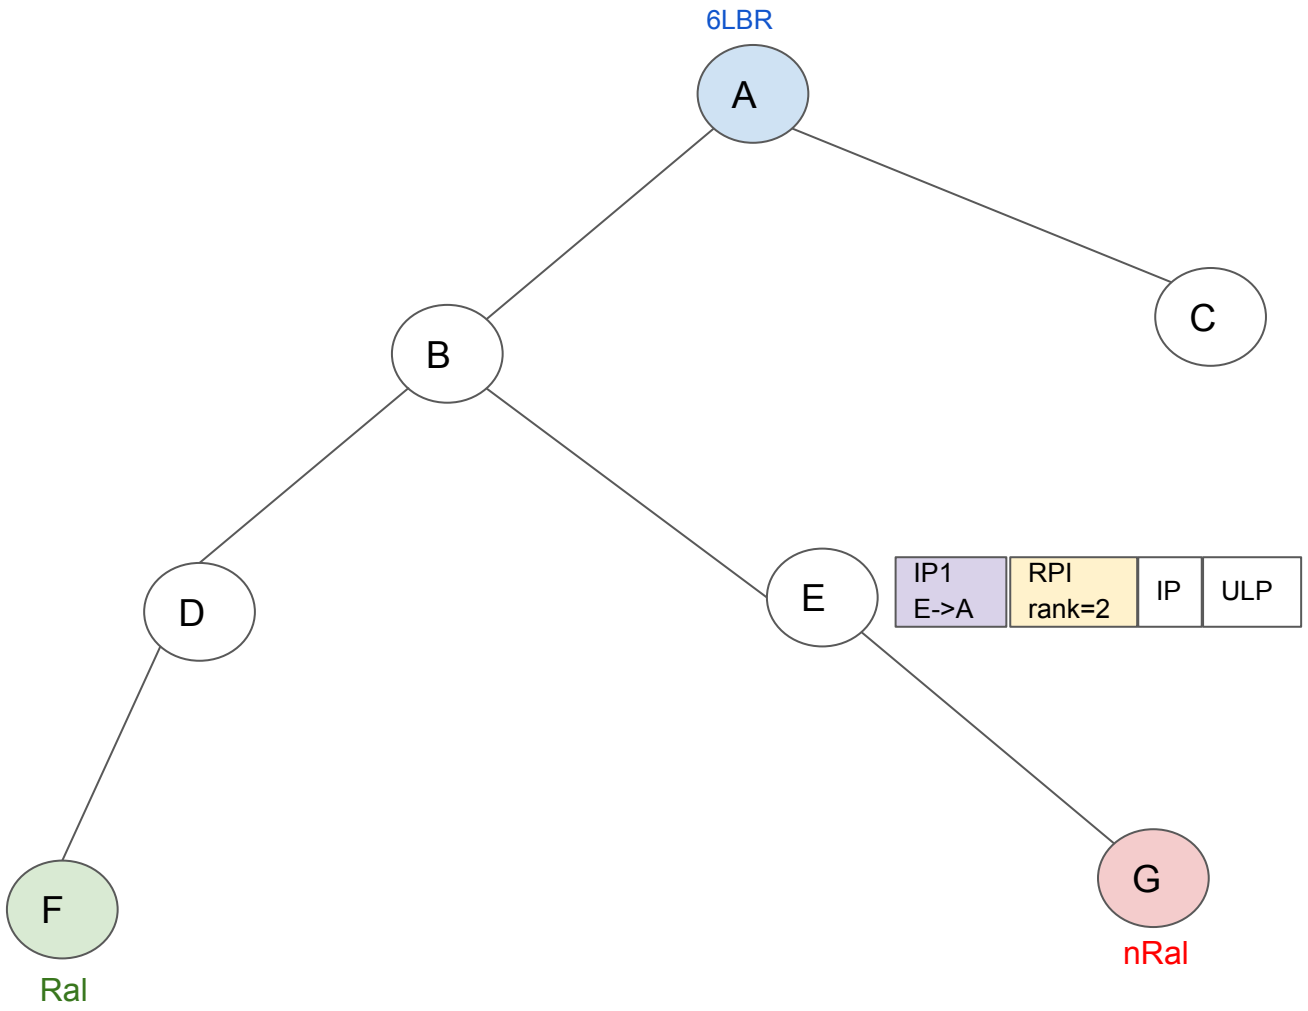

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

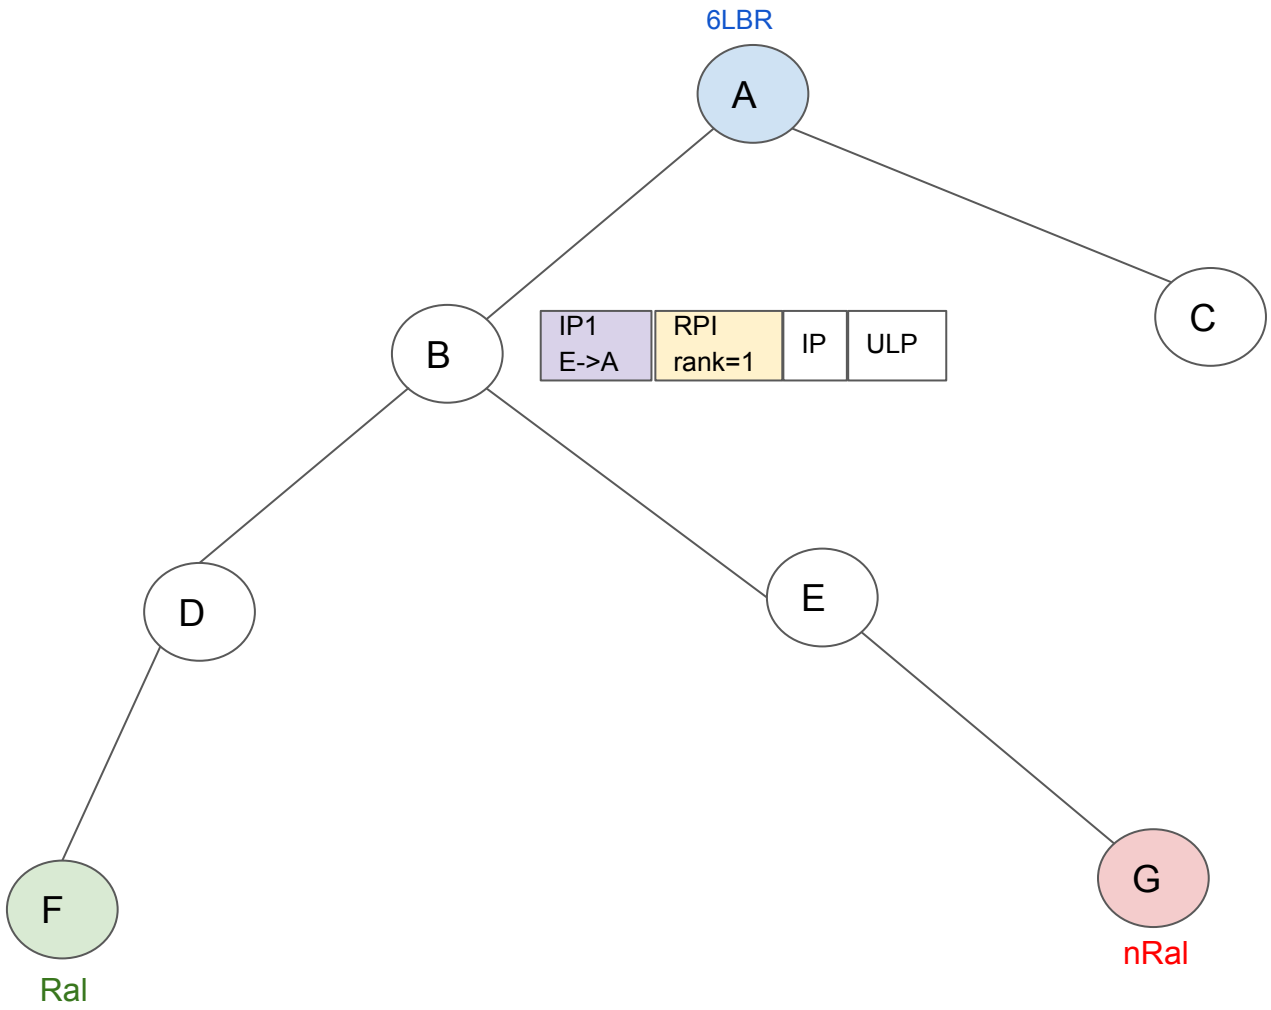

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

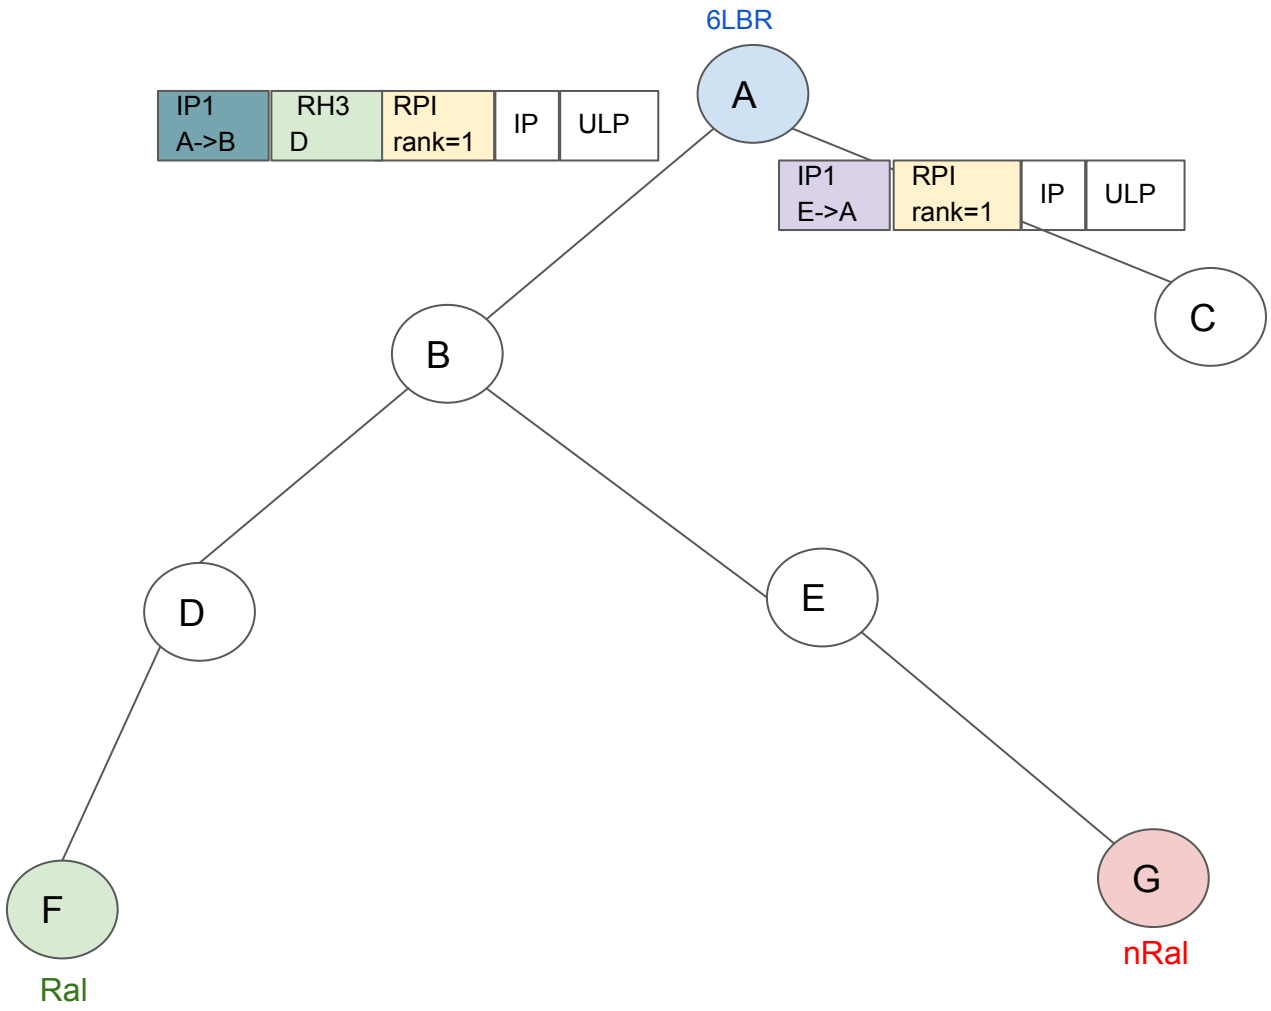

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

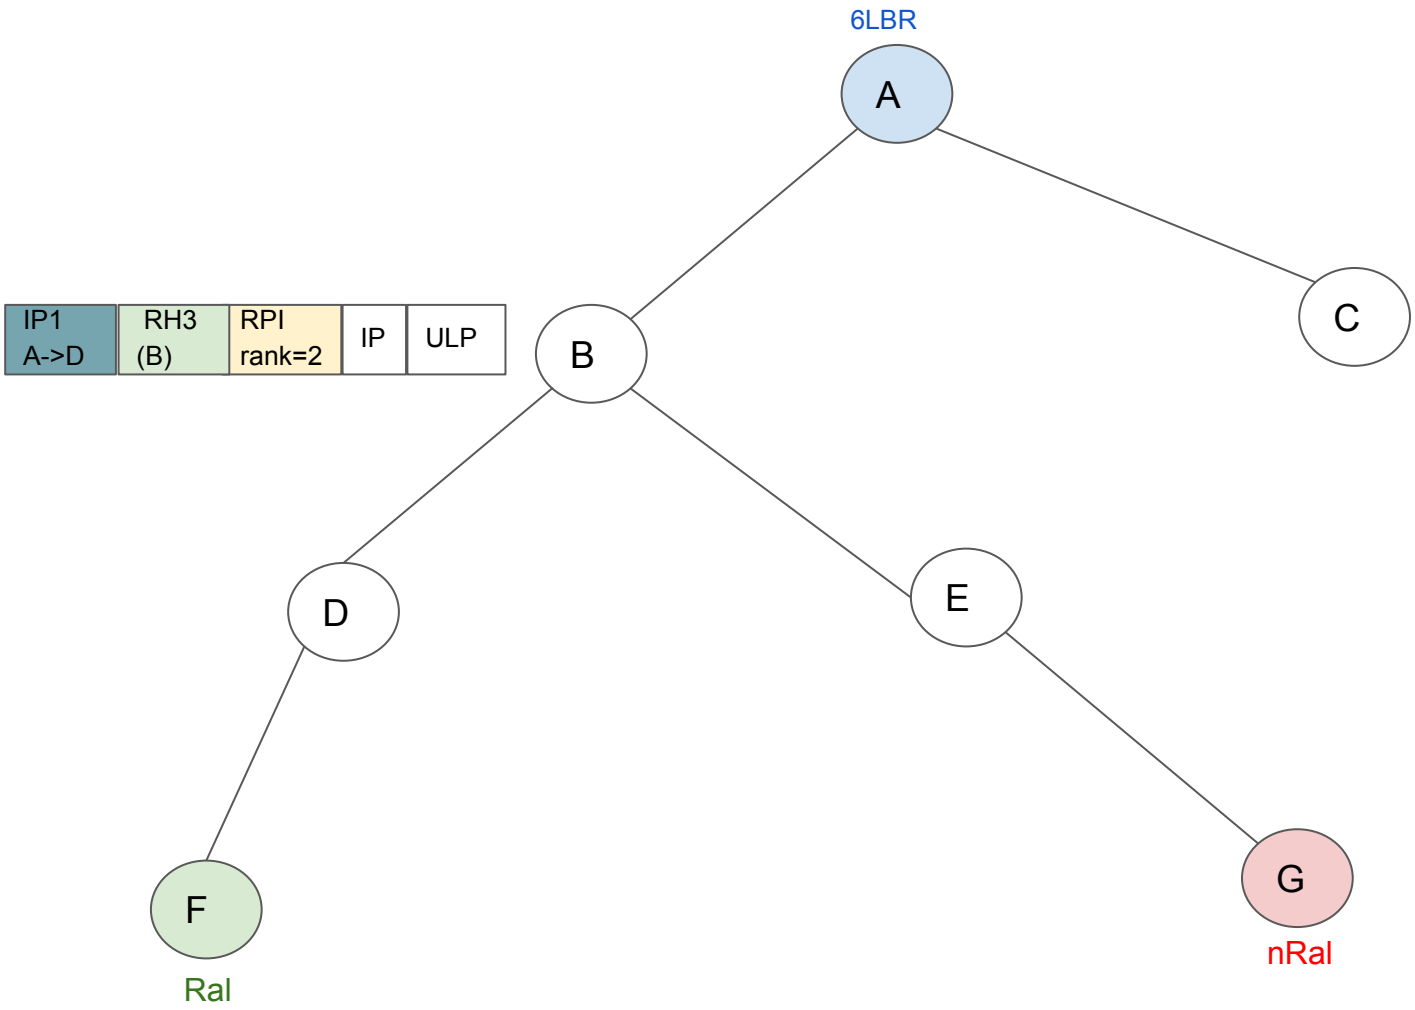

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

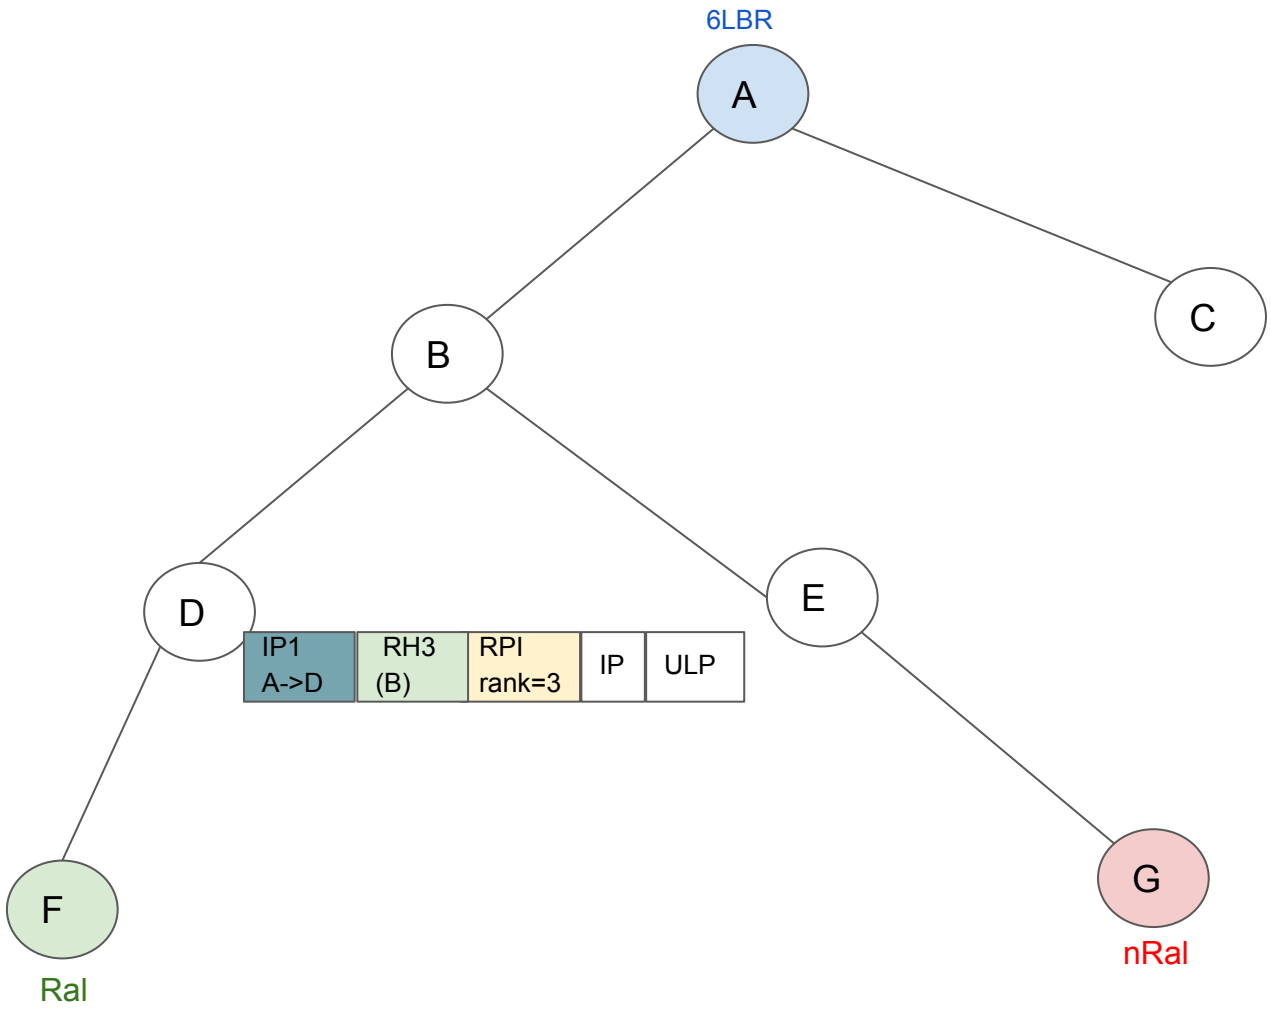

Non-Storing mode: without the inclusion of 2460bis proposal

Form nRal (not-rpl-aware-leaf) to Ral (Rpl-aware-leaf)

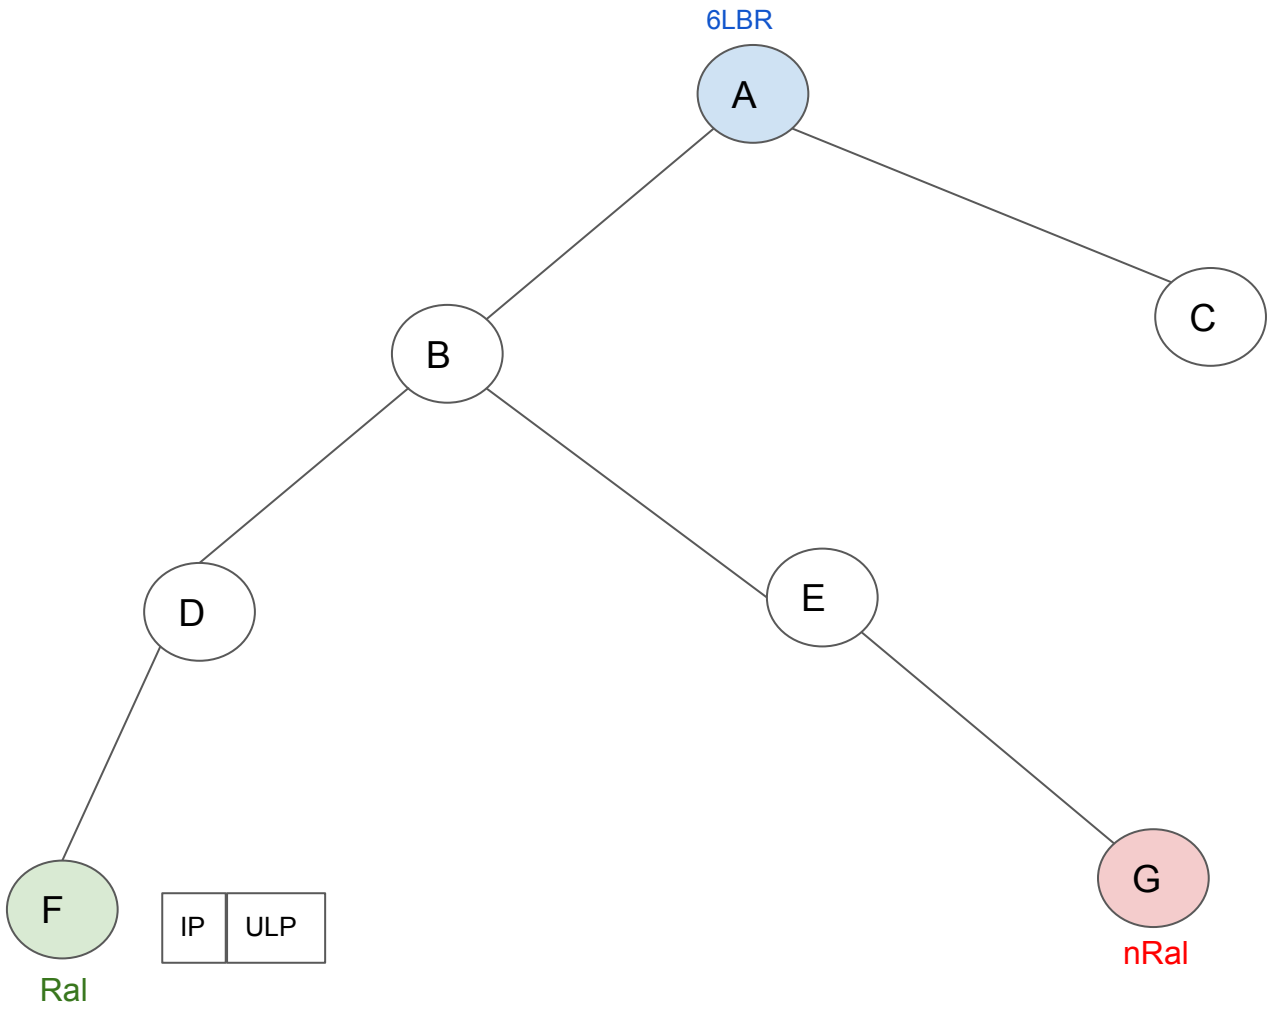

Can anyone write this clearer?

<section title="hop-by-hop IPv6-in-IPv6 headers">
<t>
The term "hop-by-hop IPv6-in-IPv6" header refers
to: adding a header
that originates from a node to an adjacent node,
using the addresses (usually the GUA or ULA, but could use
the link-local addresses)
of each node. If the packet must traverse multiple
hops, then it must be decapsulated at each hop, and then
re-encapsulated again in a similar fashion.

https://www.ietf.org/mail-archive/web/roll/current/msg09845.html

https://www.ietf.org/mail-archive/web/roll/current/msg09856.html

---

```
peter van der Stok <stokcons at xs4all.nl> wrote:
    >> (unless we mandate that the E flag is
    >> propogated all the way upwards), then the root has to always use
    >> hop-by-hop IPIP for both 5.2 and 5.3, because the root can't tell
    >> which situation is in play.

    > I thought that this was a good occasion to point that out in the
    > rplinfo document.

Noting that this issue goes away with rfc2460bis... but assuming that wasn't
the case...

if we need to clarify that the E flag is sent upwards, then we *can* say that
in useofrplinfo, because it was designed to be standards track *UPDATE* to
RFC6550.

--
Michael Richardson <mcr+IETF at sandelman.ca>, Sandelman Software Works
 -= IPv6 IoT consulting =-
```

Thank you!

DIS Modifications
draft-gundogan-roll-dis-modifications-00

Cenk Gündoğan, Ed.
Dominique Barthel
Emmanuel Baccelli

ROLL – IETF 96

July 20, 2016

Recap

- DODAG Information Solicitation (DIS)
- Used to solicit DODAG Information Object (DIO)

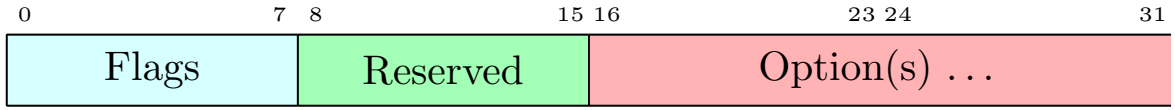

DIS Base Object

Unicast DIS	Multicast DIS
Don't Reset Trickle	Reset Trickle
One-Shot Unicast DIO	Regular Multicast DIO
Conf. Option	No Conf. Option

DIS Operation

Objectives of Draft

Identify Use Cases that prompt for `DIS` enhancements

- Node Joining DODAG
- Identifying Defunct DODAG

Three behaviors that can be enhanced

- Multicast `DIS` and Trickle behavior
- Selectivity of multicast `DIS` messages
- Information carried by `DIOs`

Multicast DIS & Trickle behavior

- Consistent Network → large Trickle intervals for DIOs
- Appearing Node requests DIOs with multicast DIS

Current Behavior

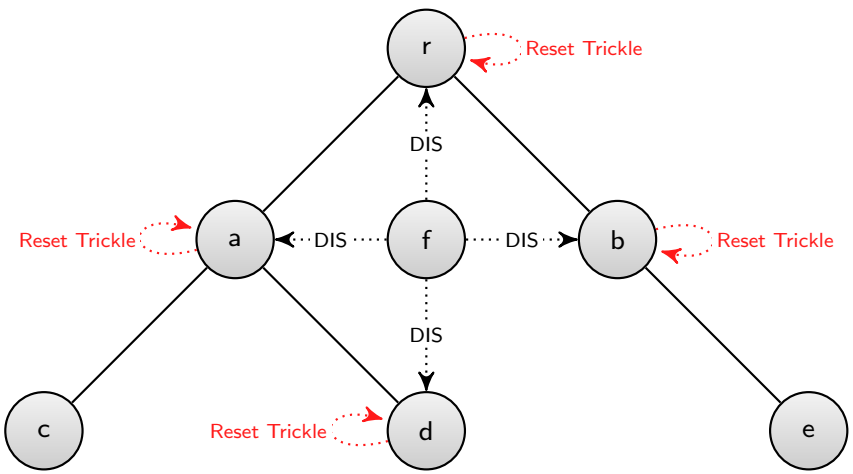

Mcast DIS – Reset Trickle

Desired Behavior

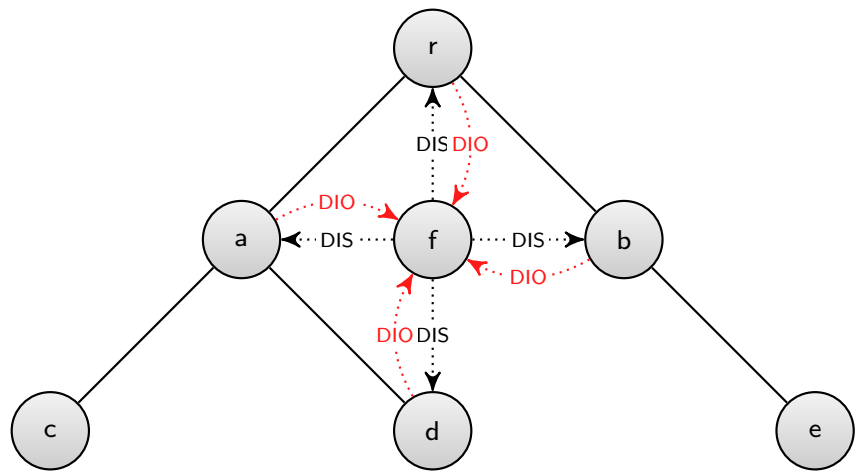

Mcast DIS – One-Shot Ucast DIOs

Multicast DIS & Trickle behavior

Proposed Protocol Improvements

- New DIS Flags: N (No Inconsistency) + T (DIO Type)

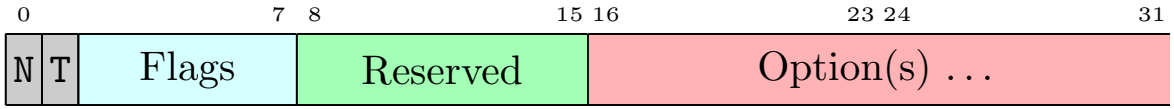

Unicast DIS	Multicast DIS		
	N = 0	N = 1, T = 0	N = 1, T = 1
Don't Reset Trickle One-Shot Ucast DIO Conf. Option	Reset Trickle Regular Mcast DIO No Conf. Option	Don't Reset Trickle One-Shot Mcast DIO No Conf. Option	Don't Reset Trickle One-Shot Ucast DIO No Conf. Option

- Delay DIOs by $\left[0 \ldots 2^{SpreadingInterval}\right]$ ms (randomly chosen)

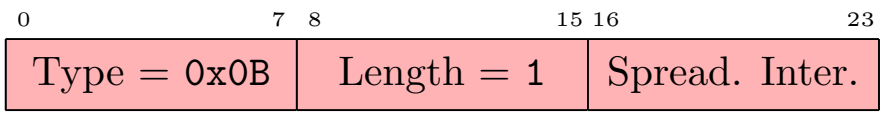

Response Spreading Option for less DIO collisions

Selectivity of multicast DIS messages

Current Behavior

- Selectivity with Solicited Information Option
  - Instance ID
  - DODAG ID
  - DODAG Version

Desired Behavior

- More selectivity

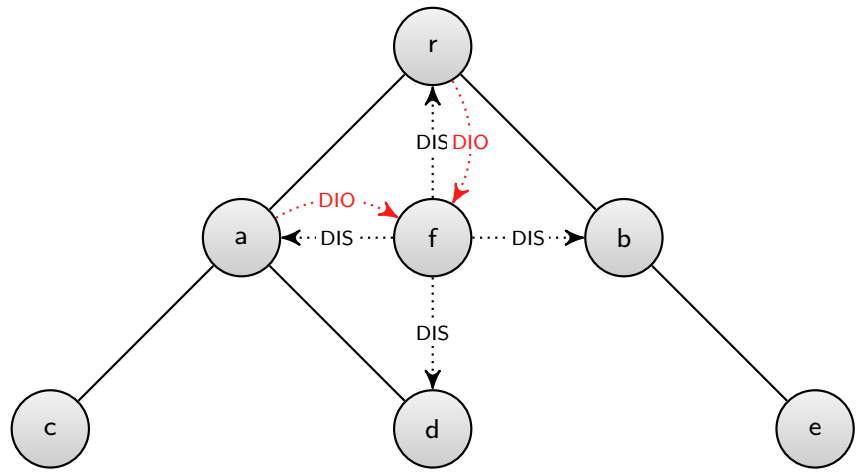

Proposed Protocol Improvements

- Allow Metric Container Options in DIS messages
  → Less neighbors respond w.r.t. specified routing constraints

Information carried by DIOs

Current Behavior

- Conf. Opt. in ucast DIO (**MUST**)
- Inflation by 16 bytes

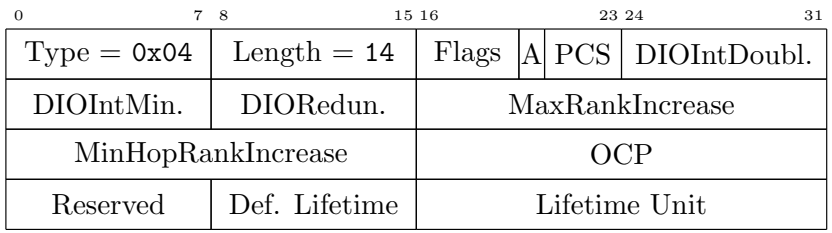

0 – 7	8 – 15	16	17	18 – 23	24 – 31
Type = 0x04	Length = 14	Flags	A	PCS	DIOIntDoubl.
DIOIntMin.	DIORedun.	MaxRankIncrease			
MinHopRankIncrease		OCP			
Reserved	Def. Lifetime	Lifetime Unit			

DODAG Configuration Option

Desired Behavior

- More control over inclusion of Conf. Option
- More flexibility: request inclusion of other options
  - PIO, 6CO, ...

Information carried by DIOs

Proposed Protocol Improvements

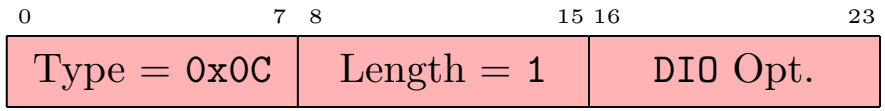

DIO Option Request Option

- DIS may include one or more DIO Opt. Request Options

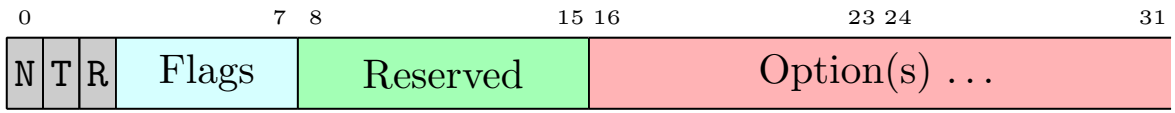

- if R=0: old behavior (Conf. Options in solicited unicast DIOs)
- if R=1: transmit explicitly requested DIO options only

Summary

Proposed DIS enhancements so far (Work In Progress!)

- 3 new flags (N + T + R) for DIS
- DIS Options
  - Response Spreading Option to reduce DIO collisions
  - Allow Metric Container Options in DIS messages
  - DIO Opt. Request Option for explicitly requested DIO options
- Backwards compatible if flags are unset

We would like to engage WG in a discussion about

- currently proposed enhancements
- further DIS enhancements

Backup – Use Case: Node Joining DODAG

- multicast `DIS` (N + T flags set) for unicast one-shot `DIOs`
- `DIS` may include Response Spreading Option
- include Metric Container with strict routing constraints
- unsuccessful `DIS` $\rightarrow$ relax routing constraints with new `DIS`
- stop when joined or routing constraints relaxed to lowest level

Proven in the field to be energy efficient

Constrained-Cast: Source-Routed Multicast for RPL

Olaf Bergmann, Carsten Bormann, Stefanie Gerdes

Supported by Siemens Corporate Technology

IETF91, Honolulu, November 2014, updated for IETF96, Berlin, July 2016

RPL Multicast

- Assumes **Storing Mode**
- Multicast DAOs indicate subscription to multicast address
  - percolate up and create bread-crumbs

- Specification in RFC 6550:
  - incomplete, untested

How to do this for Non-Storing Mode?

- Non-Storing is based on source-routing by root
- How to source-route multicast packets?
  - List all outgoing interfaces?

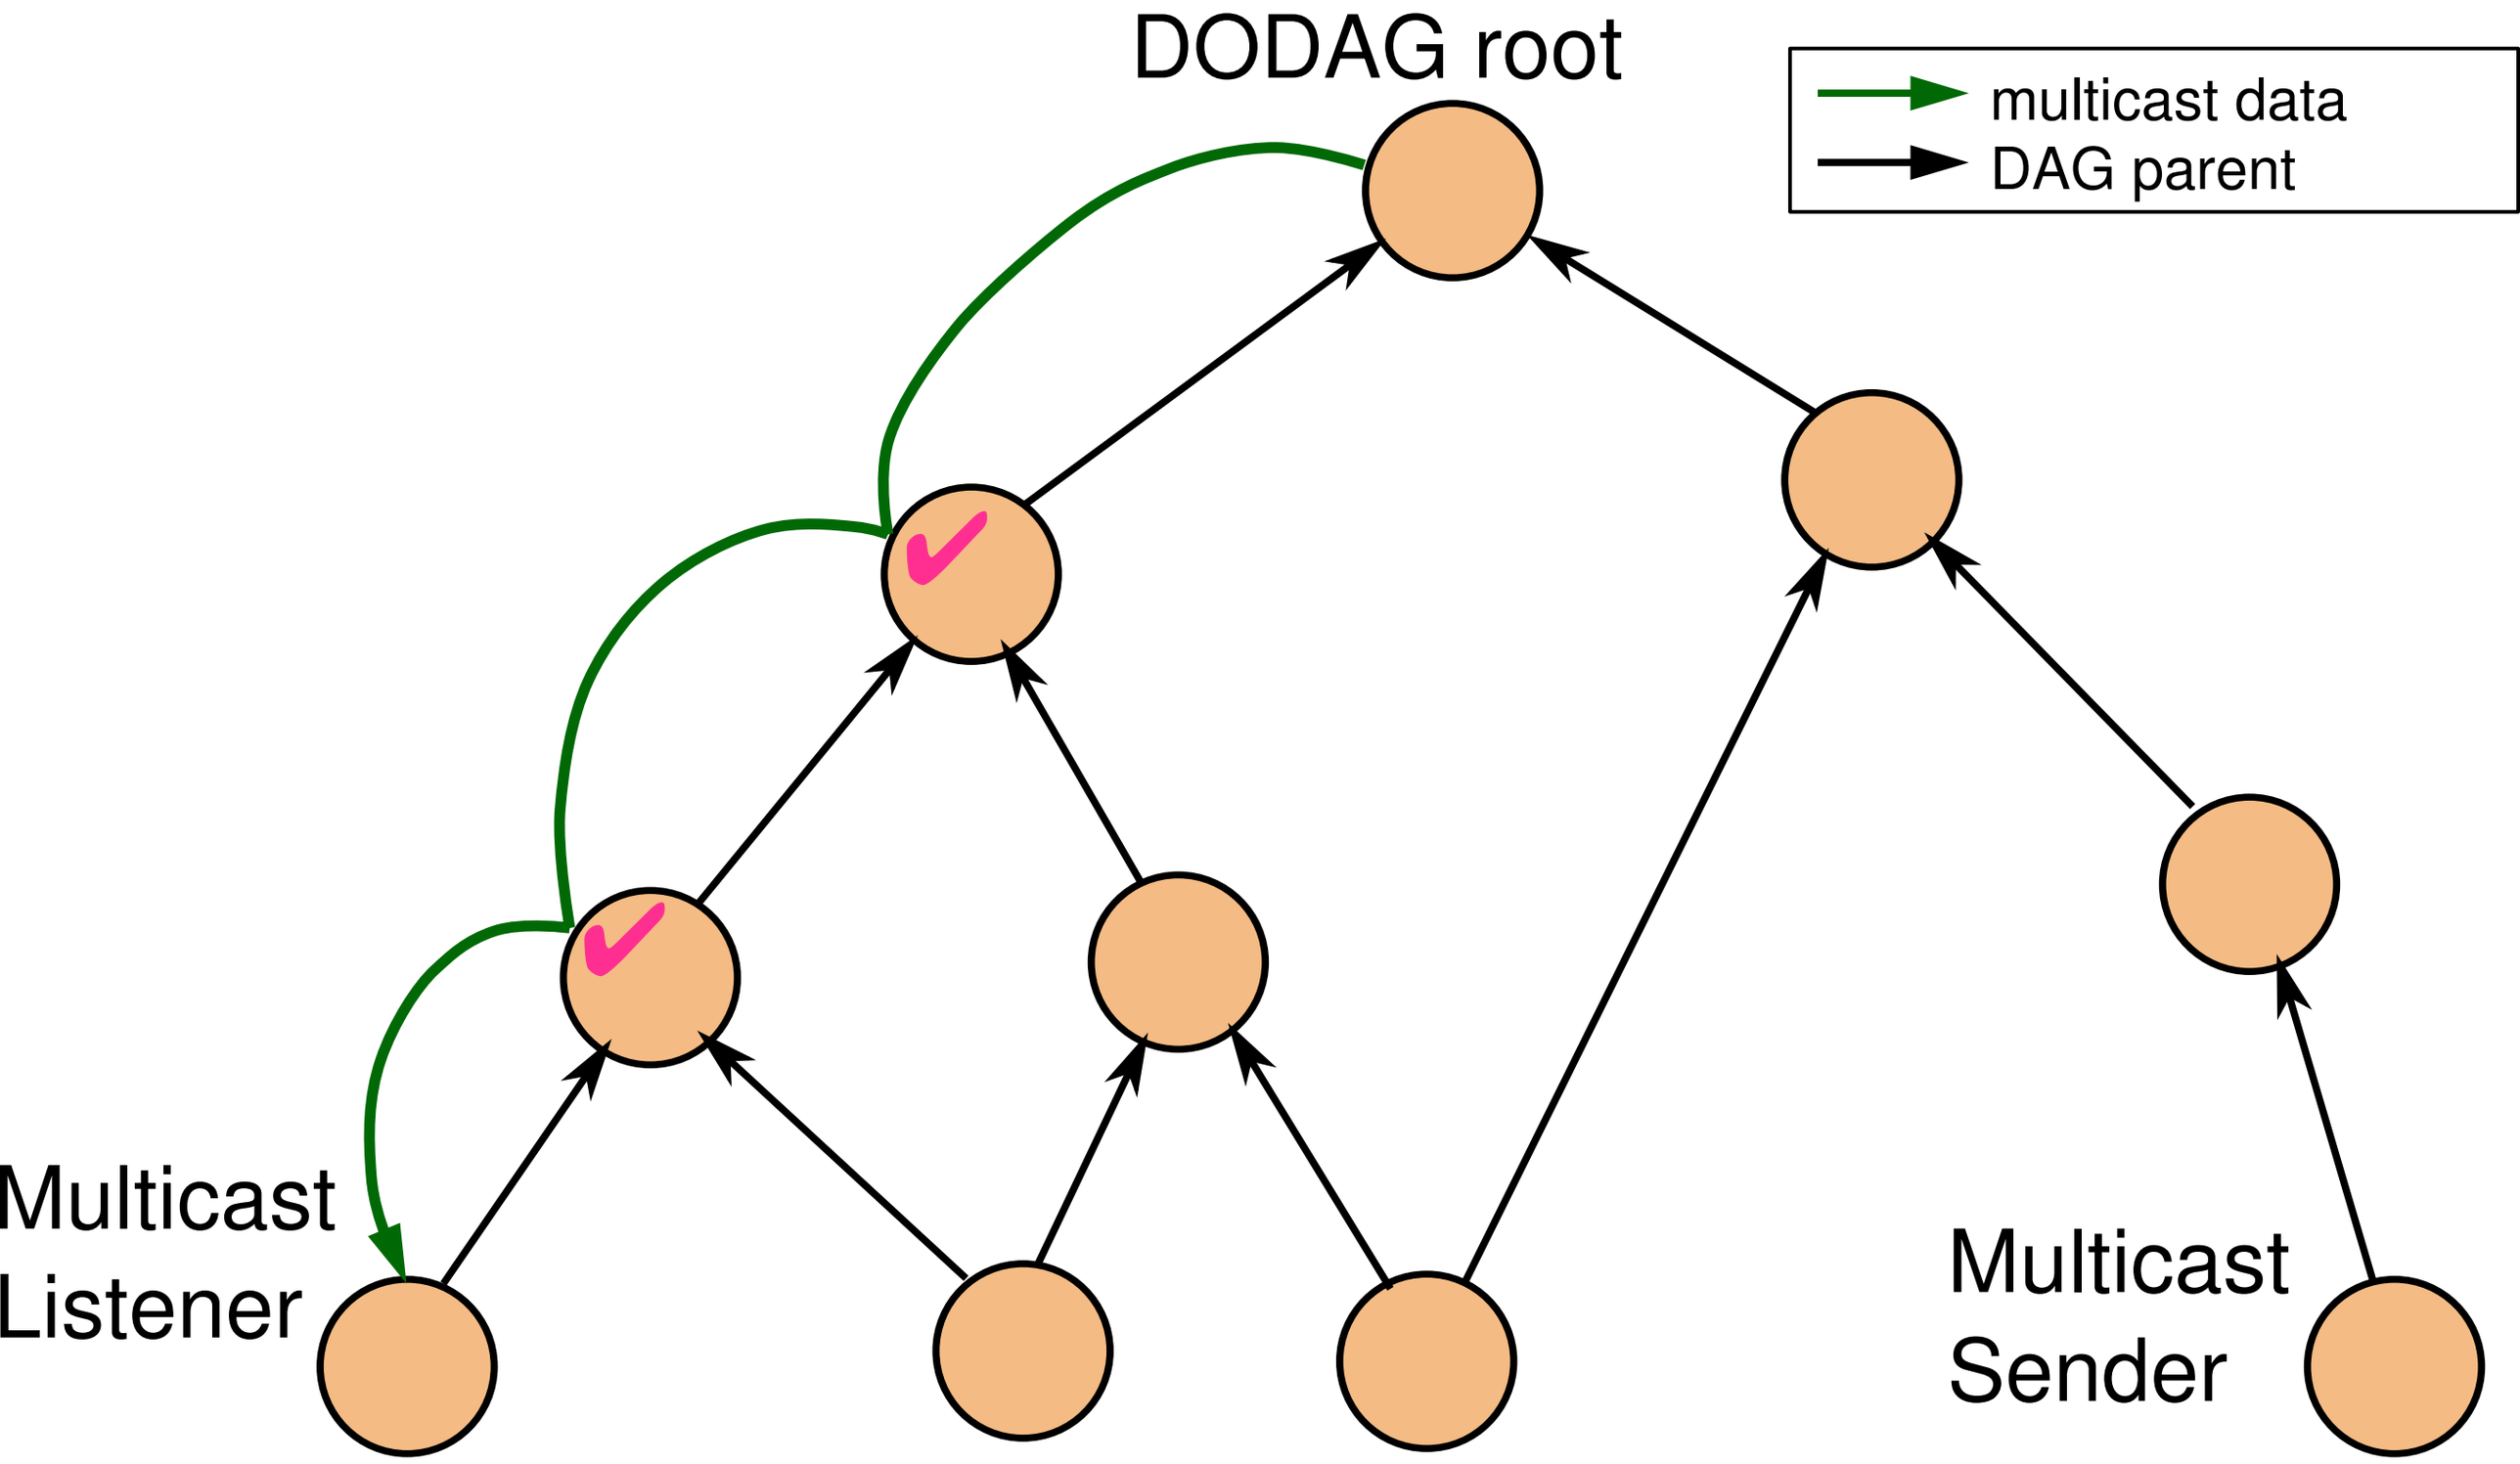
DODAG root
multicast data
DAG parent
Multicast
Listener
Multicast
Sender

Embed an Outgoing Interface List?

- Non-starters:
  - List of 16-byte IPv6 Addresses
  - Even with RFC 6554 style compression

BIER

- Bit-Indexed Explicit Replication (BOF@IETF91)
  - lists all egress routers in a bitmap
    - requires **numbering** all egress routers
    - requires network **map** in forwarders

- So use the bitmap for something different

Bloom Filters

- Compact representation of a match/non-match for elements in a large set
  - Hash the element with multiple hash functions, result → index bitmap
  - If **all** corresponding bits are set: match
- [B. H. Bloom, CACM July 1970]

Send Bloom Filter with packet, match OIF

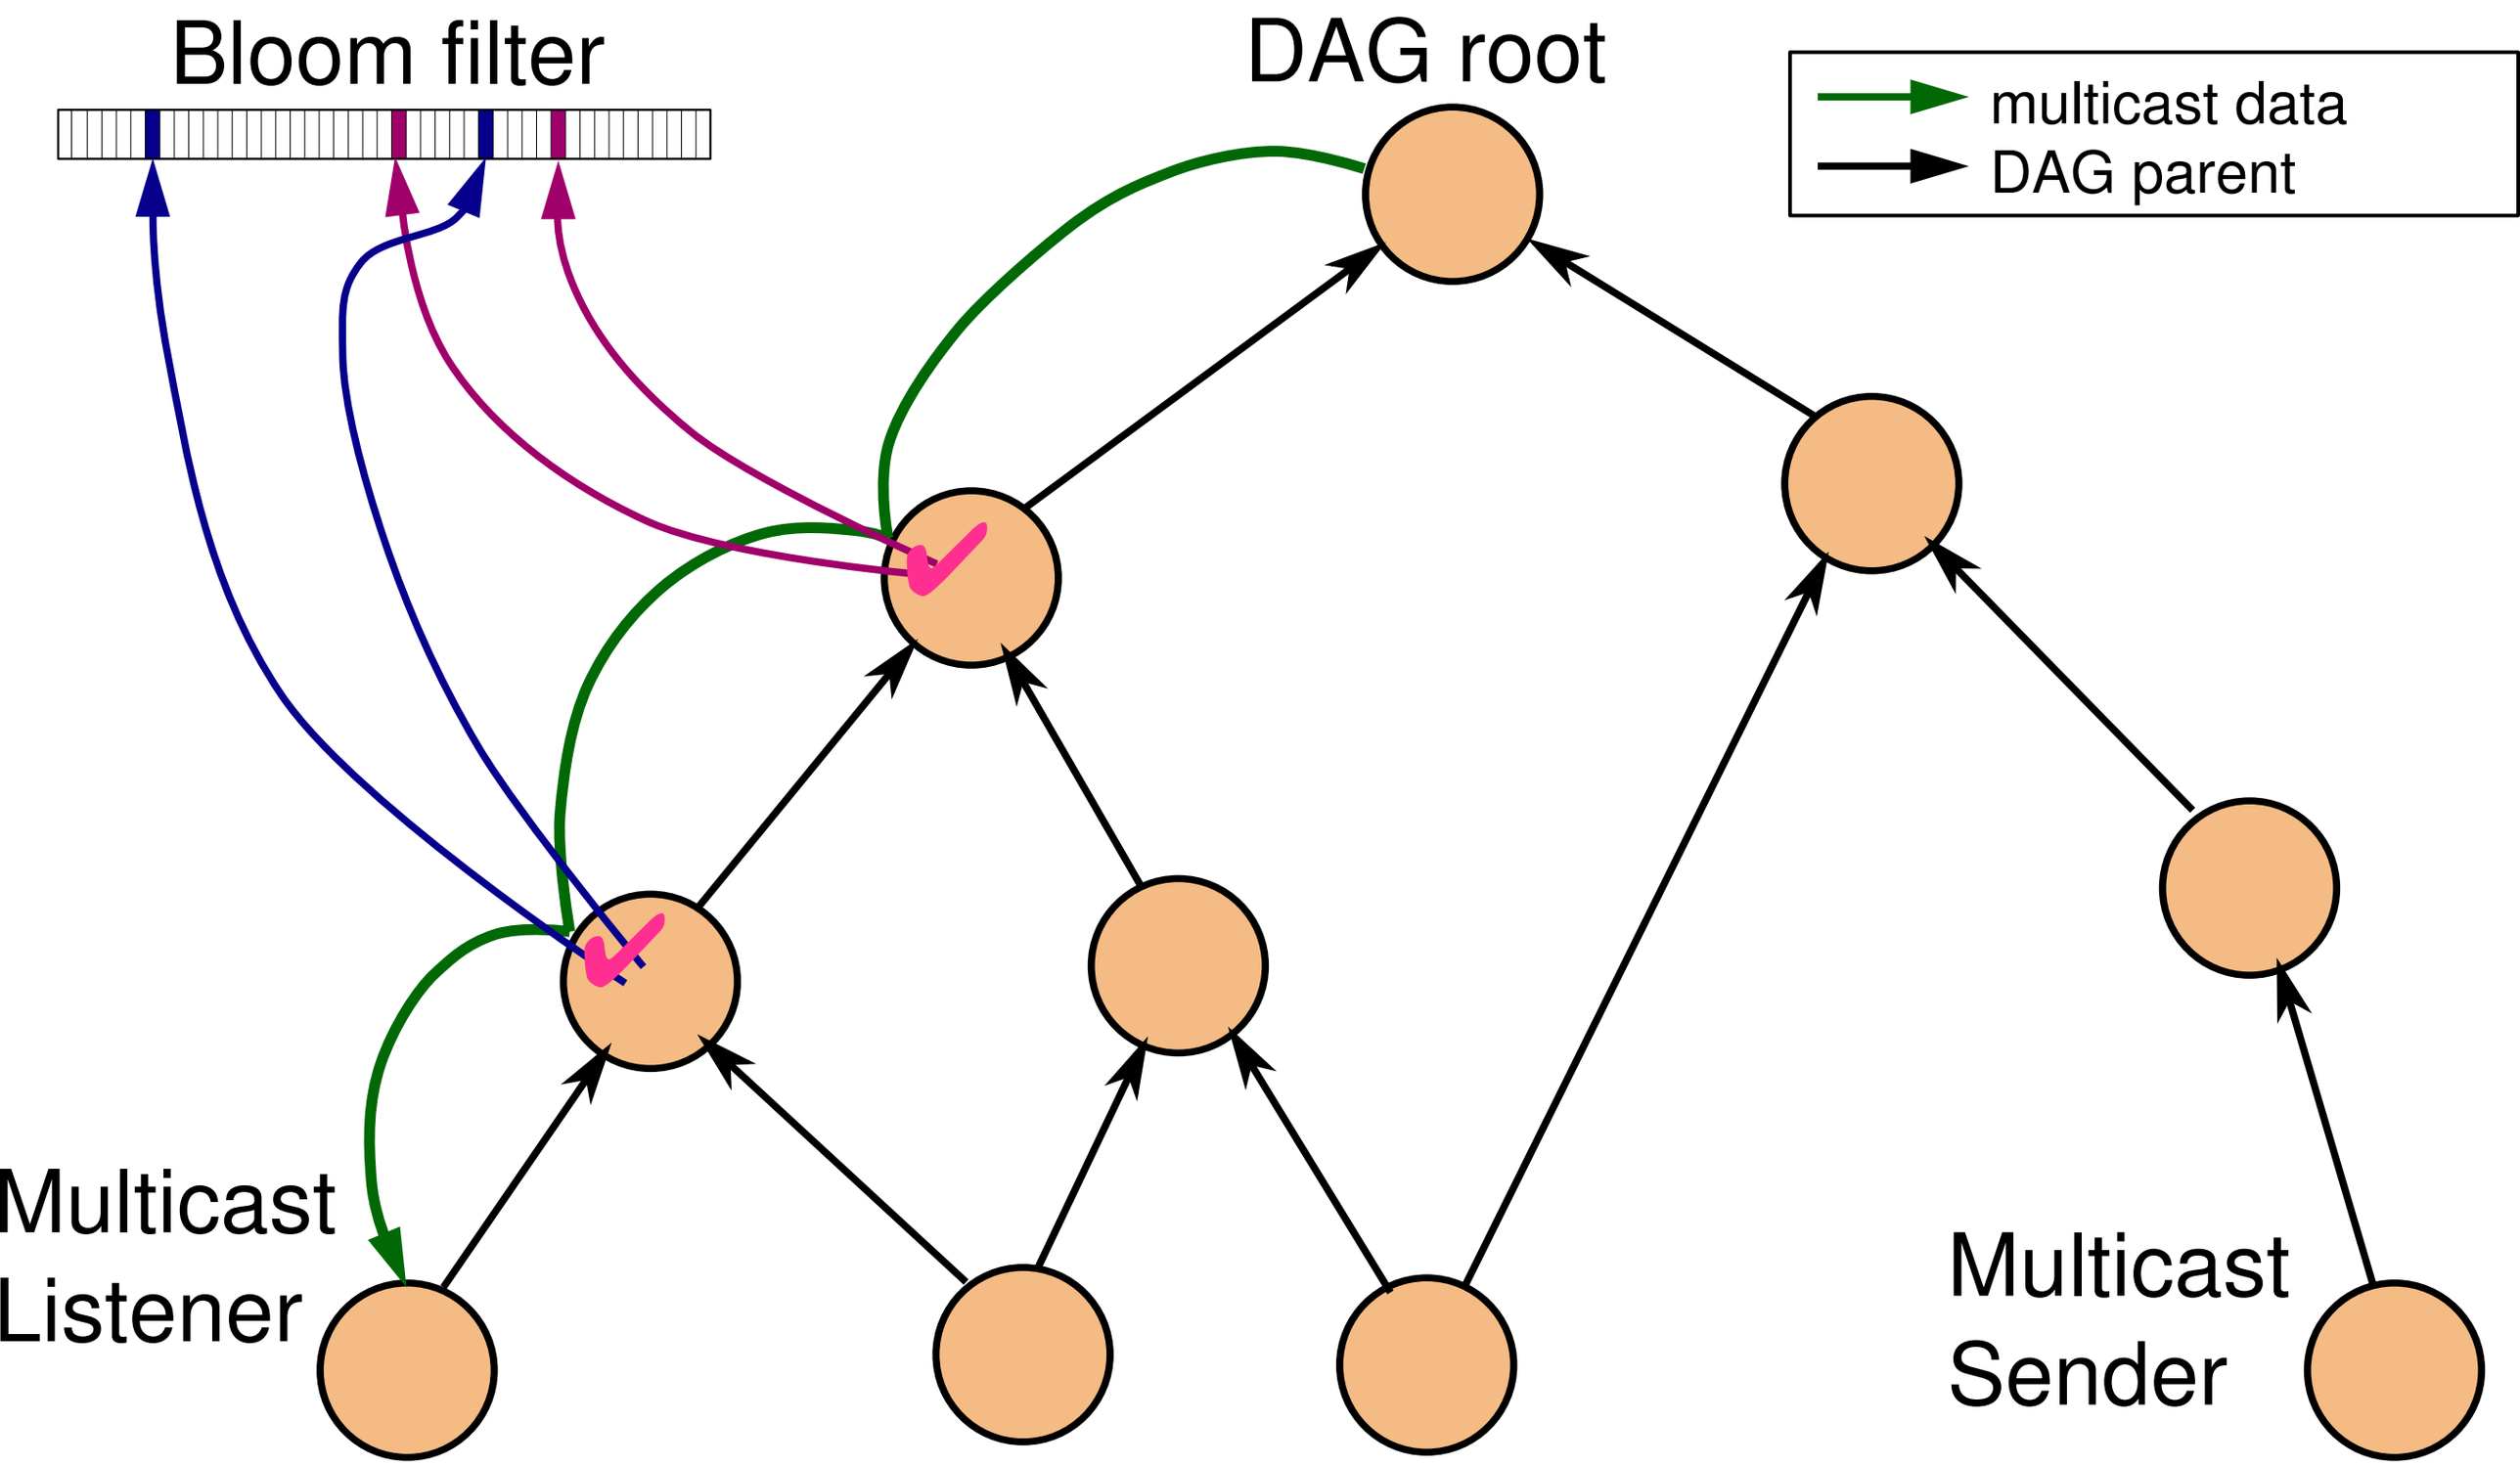

False positives?

- Bloom filters are **probabilistic**
- False positive: match indicated by aliasing of hash values
- Cause spurious transmission

False positive causes spurious transmission

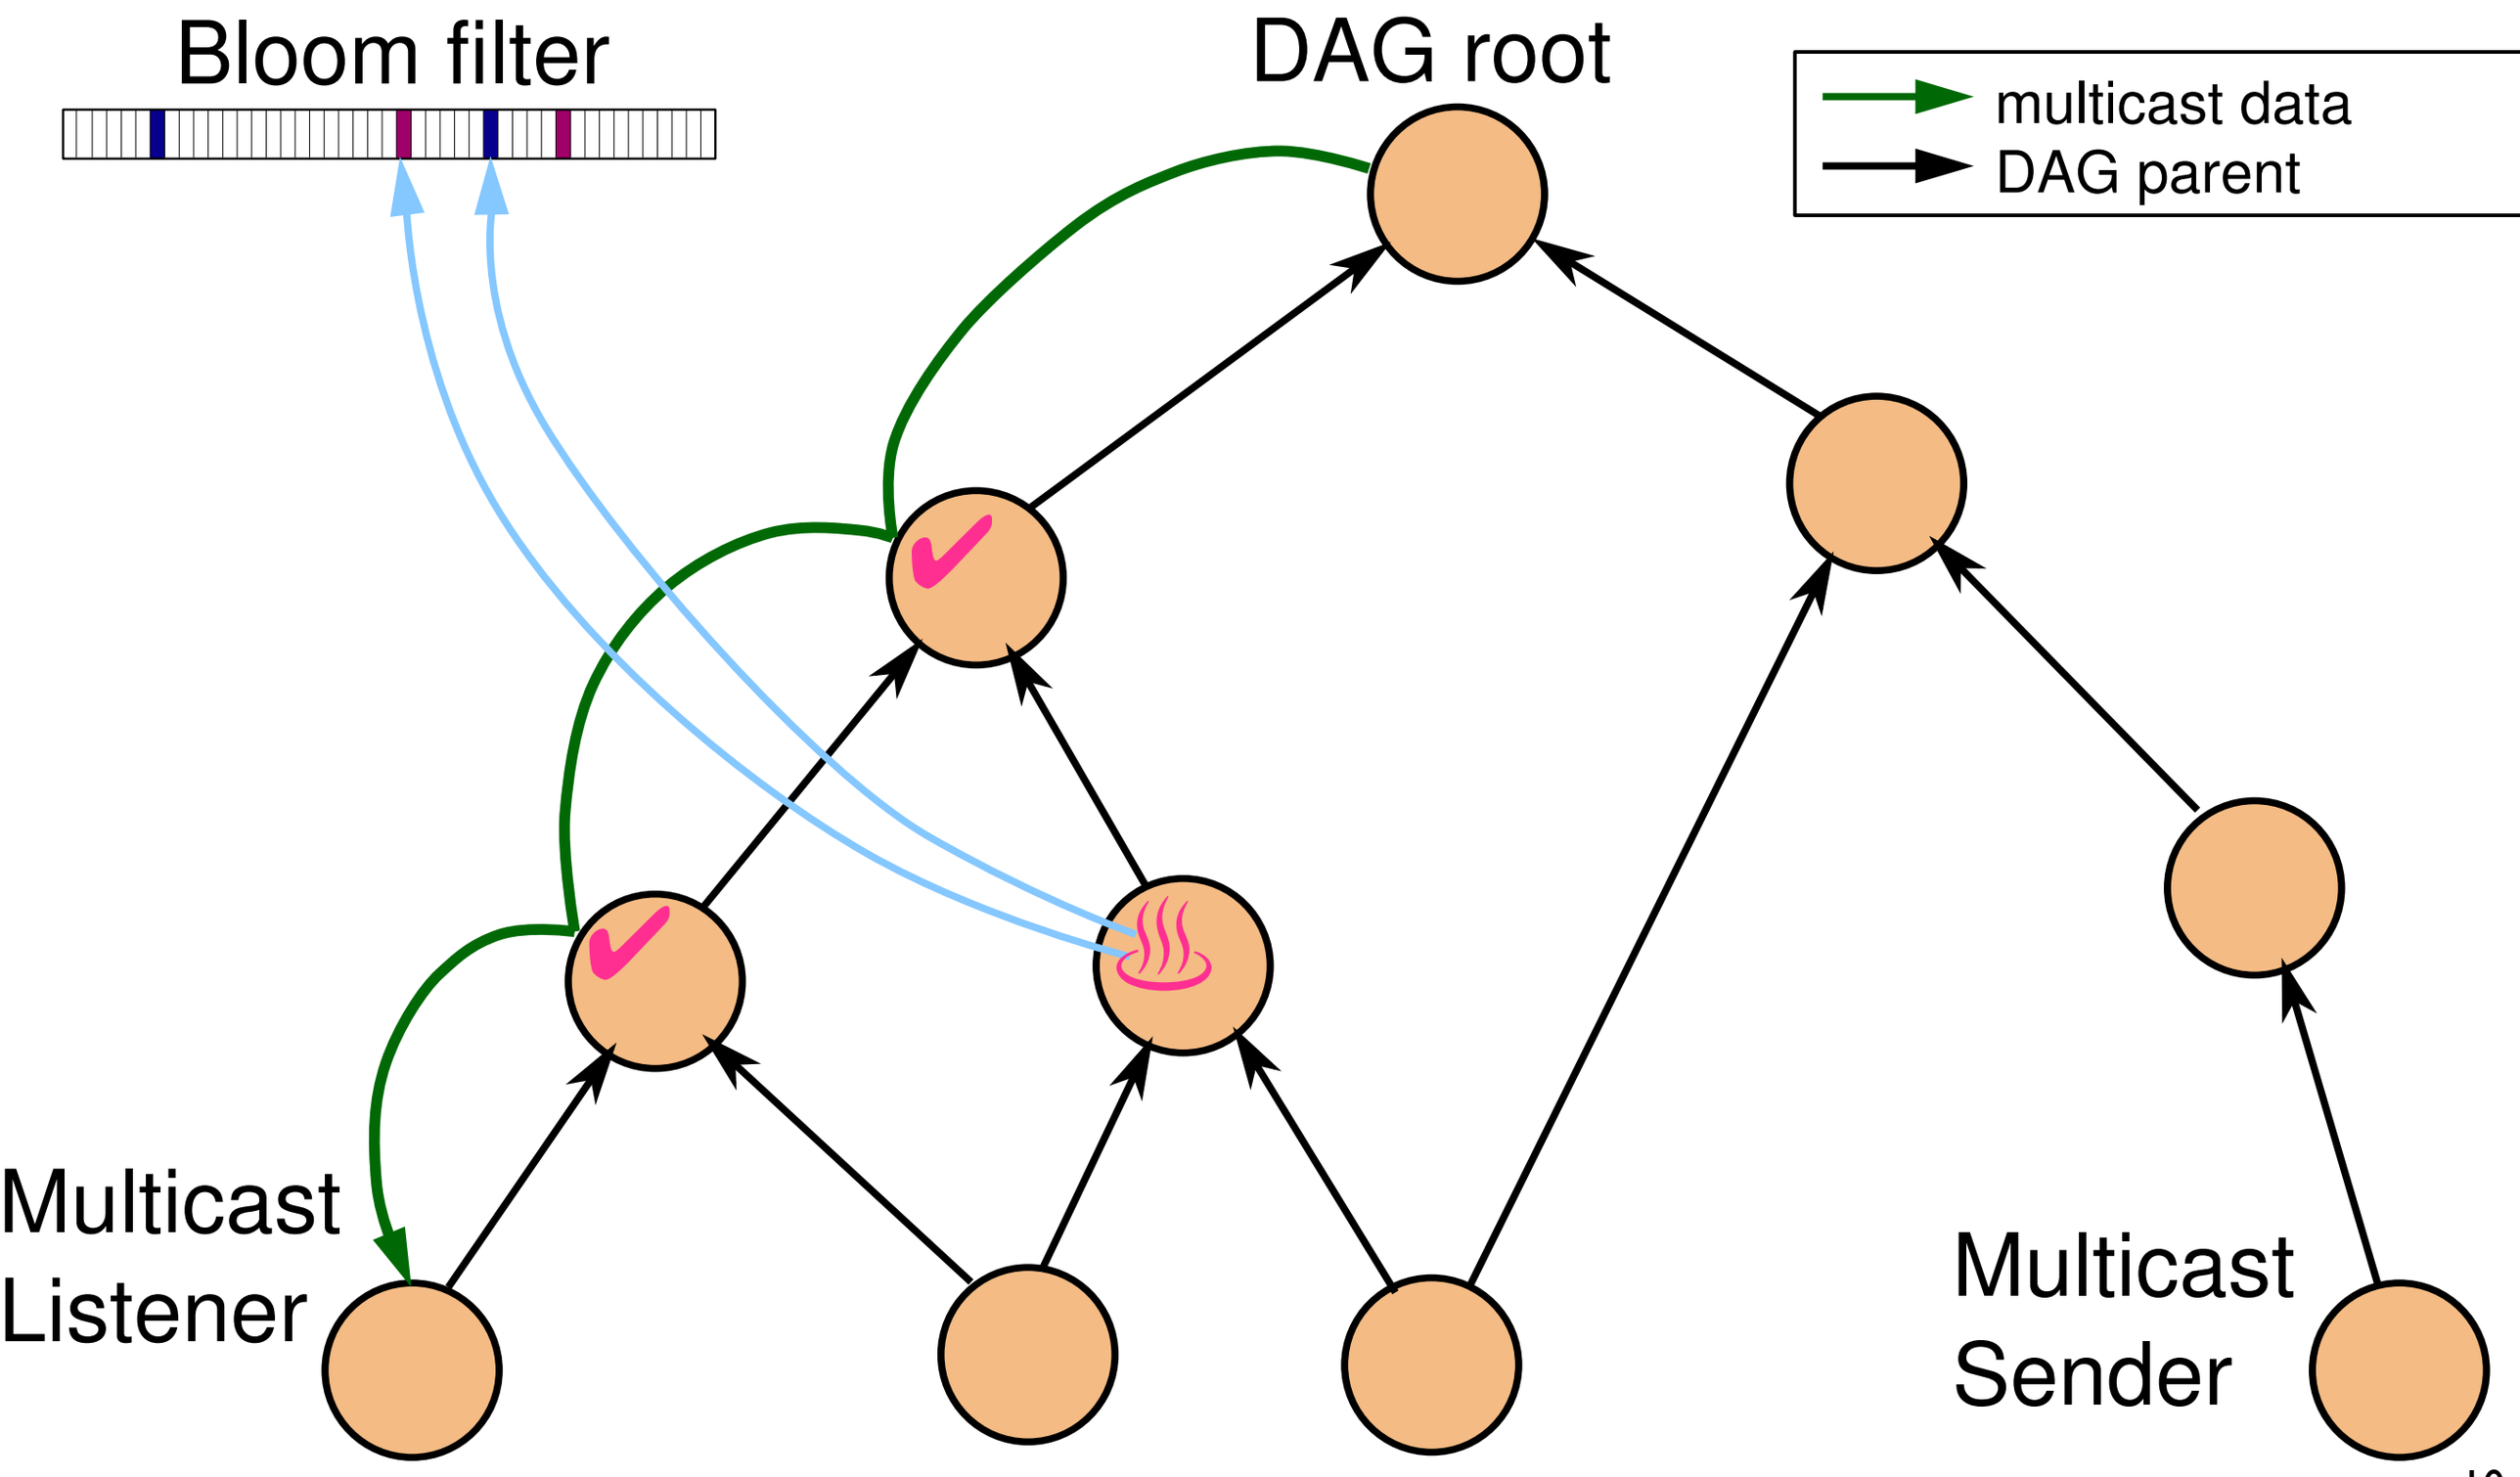

How bad are false positives?

- False positives cause spurious transmission
- No semantic damage (hosts still filter out)
- Waste in energy and spectrum:
  - ~ false-positive-rate × density
- Can easily live with significant percentage

Number of **forwarders**, filter size, f.p.r.:

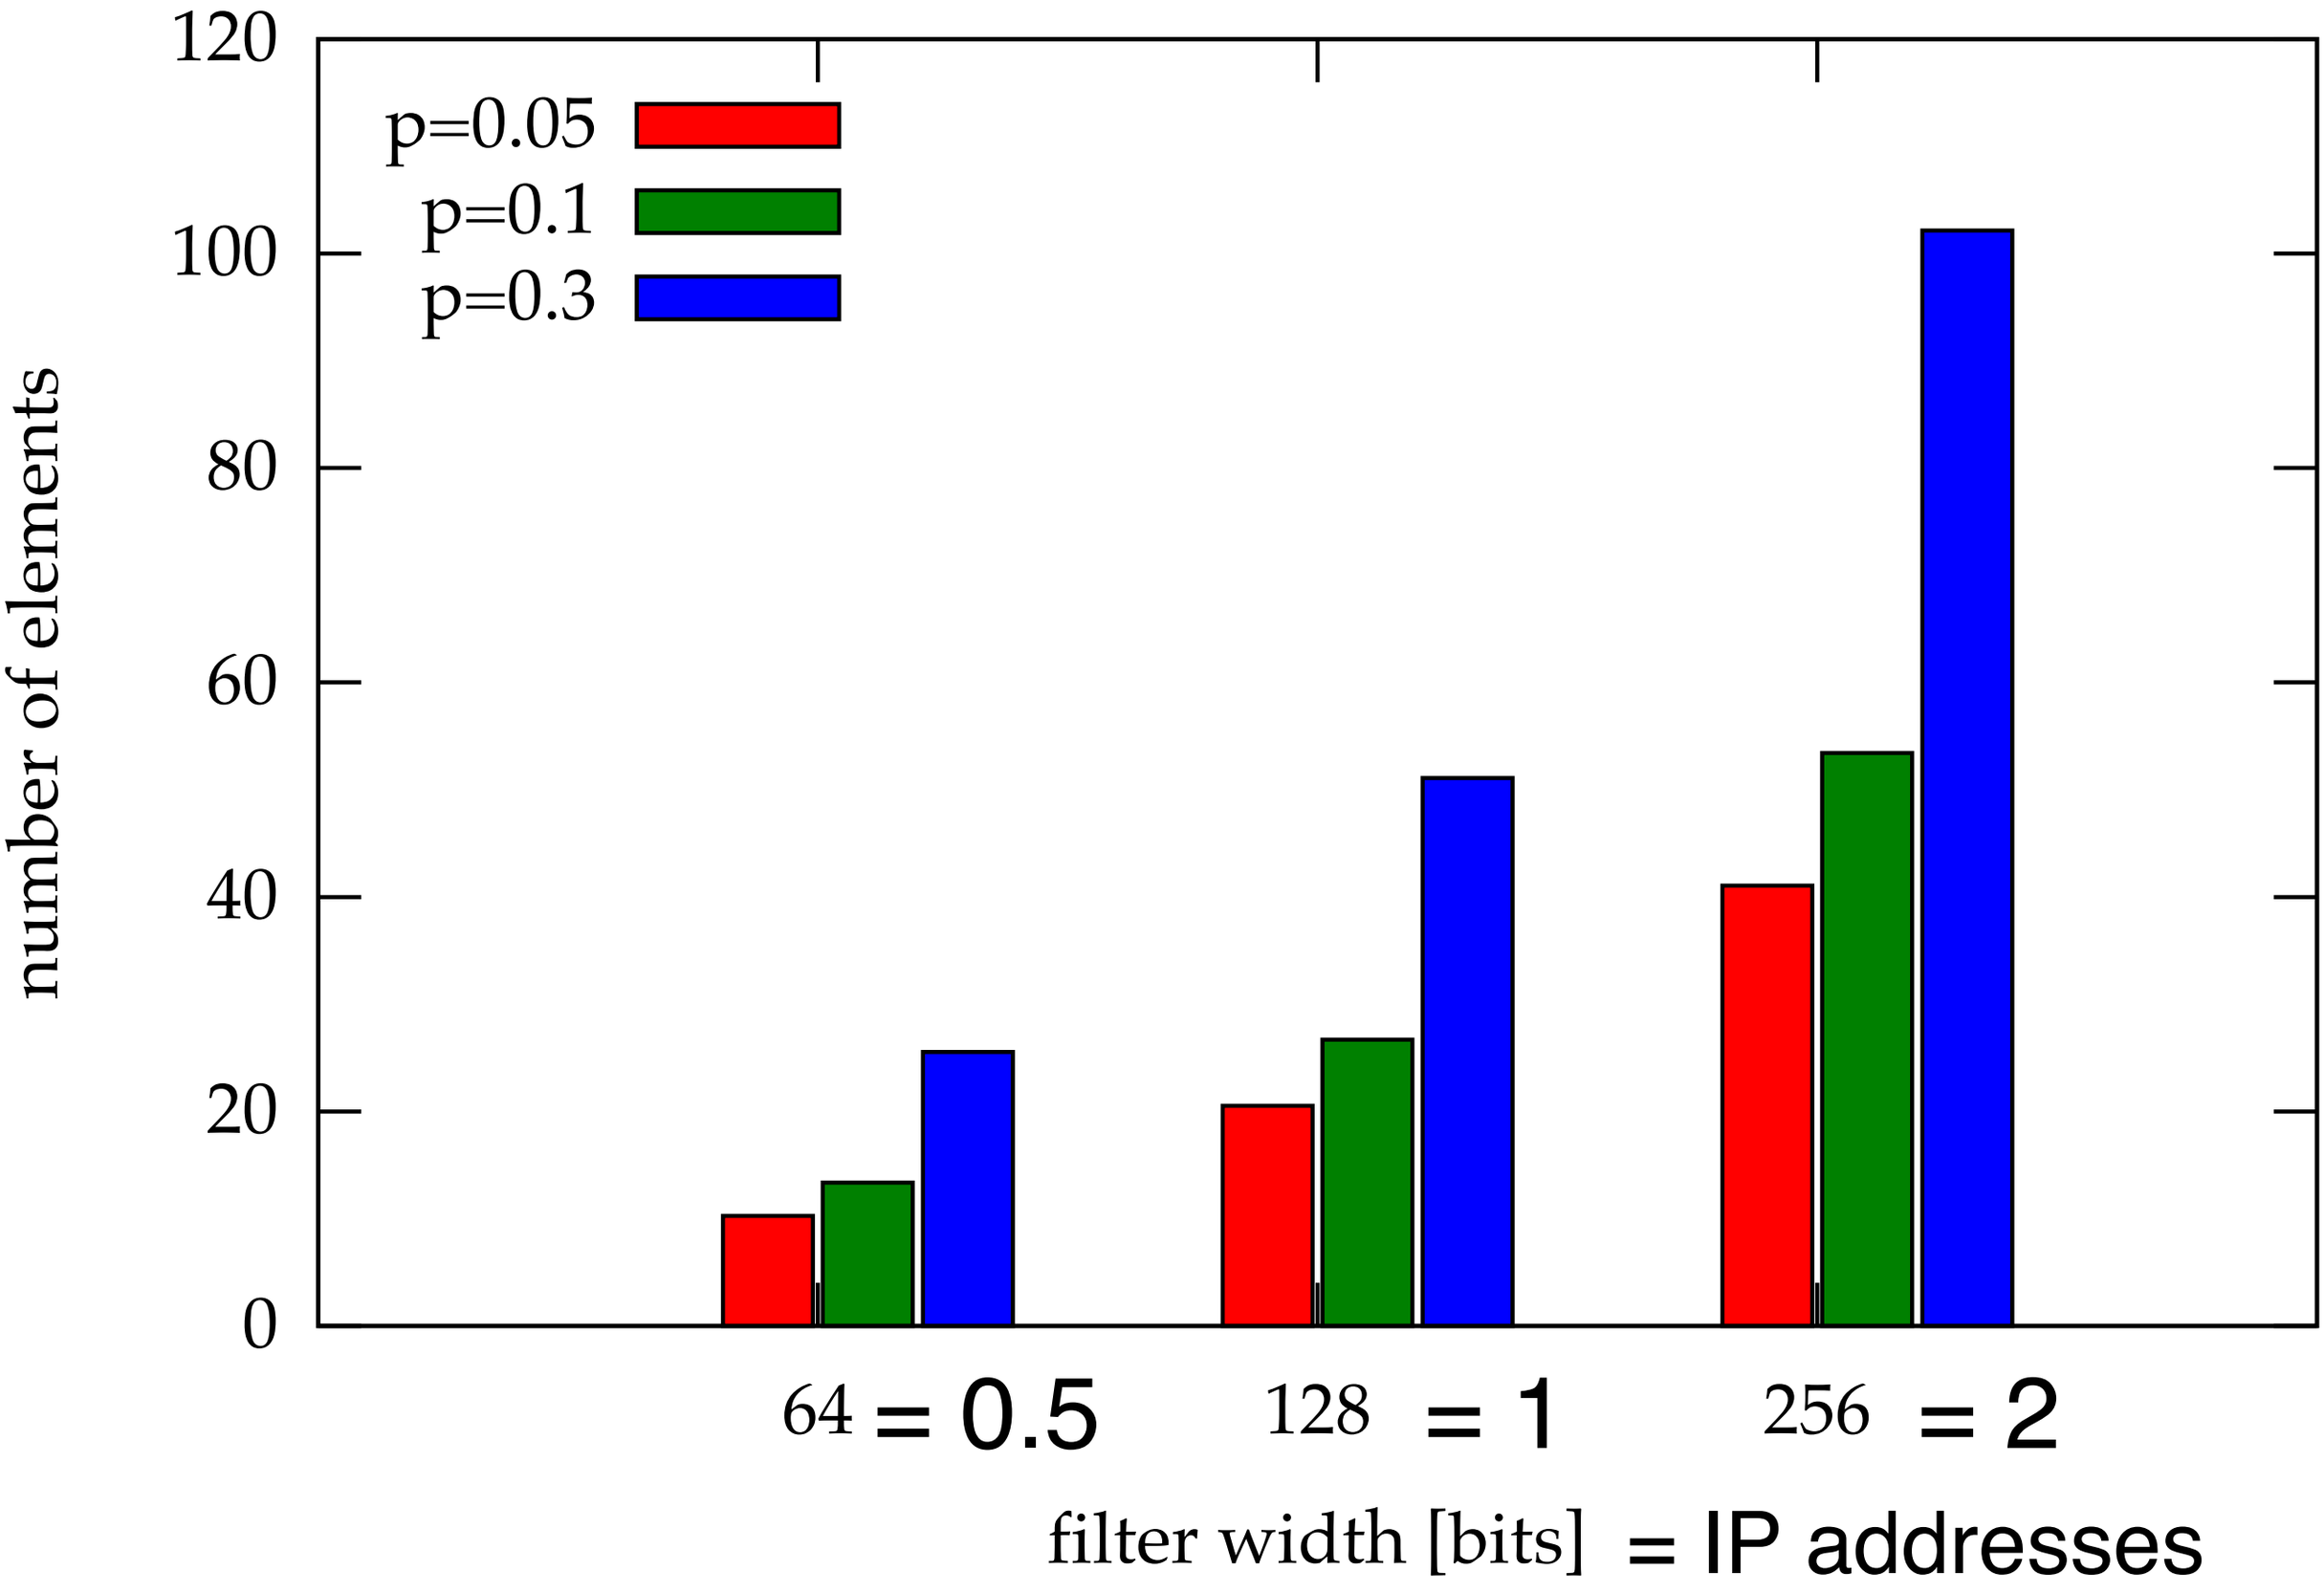

Protocol

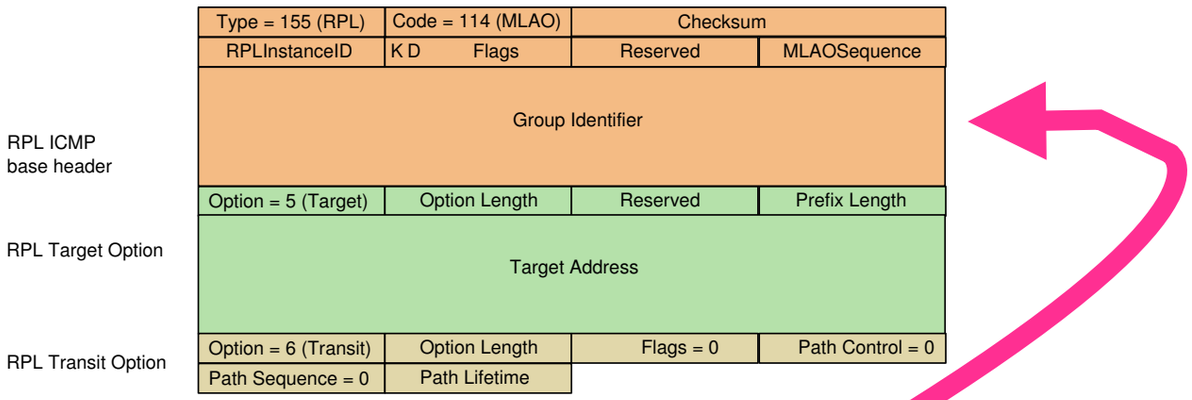

- Membership: MLAOs (go right up to root)
- Submit: tunnel up to root
- Deliver: root inserts bloom filter

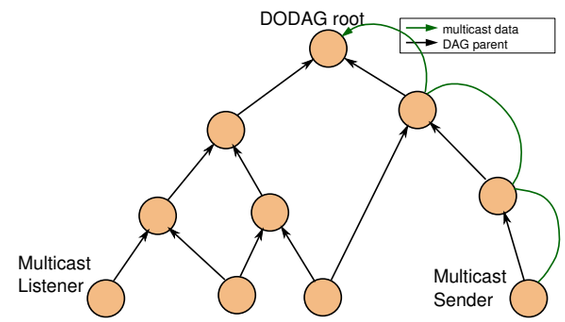

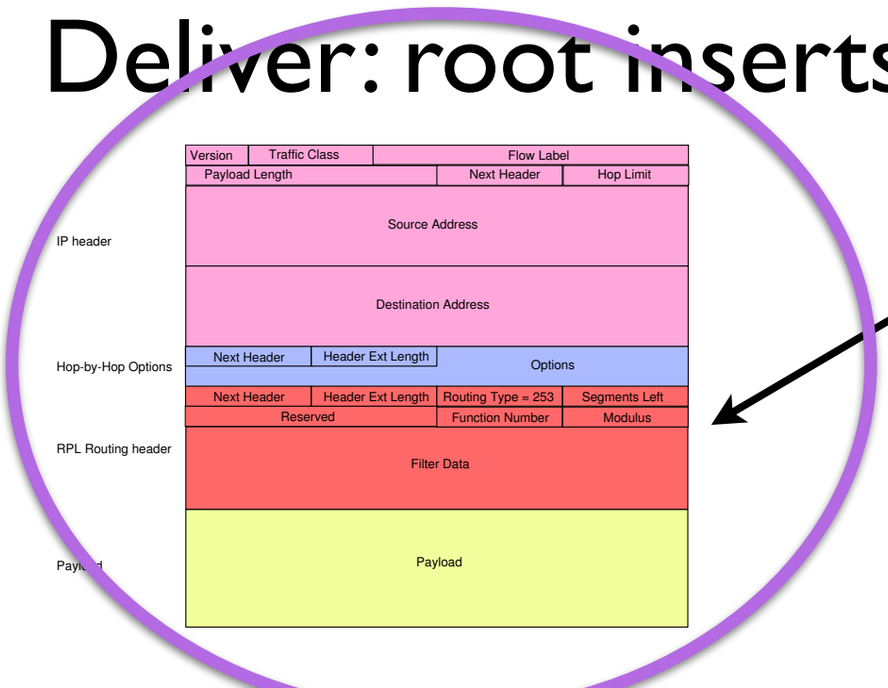

This might go into a 6LoRH Header

Implementation

- Implemented 2013 in Contiki
- Forwarders compute the hash functions **once** and then simply bit-match their OIF hash bits against bloom filter
- Root:
  - easy insertion of path for new member
  - recompute bloom filter on leave

ROLL working group

MPL Forwarder Select (MPLFS)
draft-vanderstok-roll-mpl-forw-select-00

P. van der Stok, AR. Sangi

PURPOSE

All nodes are connected to a minimum number, N_DUPLICATE, of MPL forwarders.

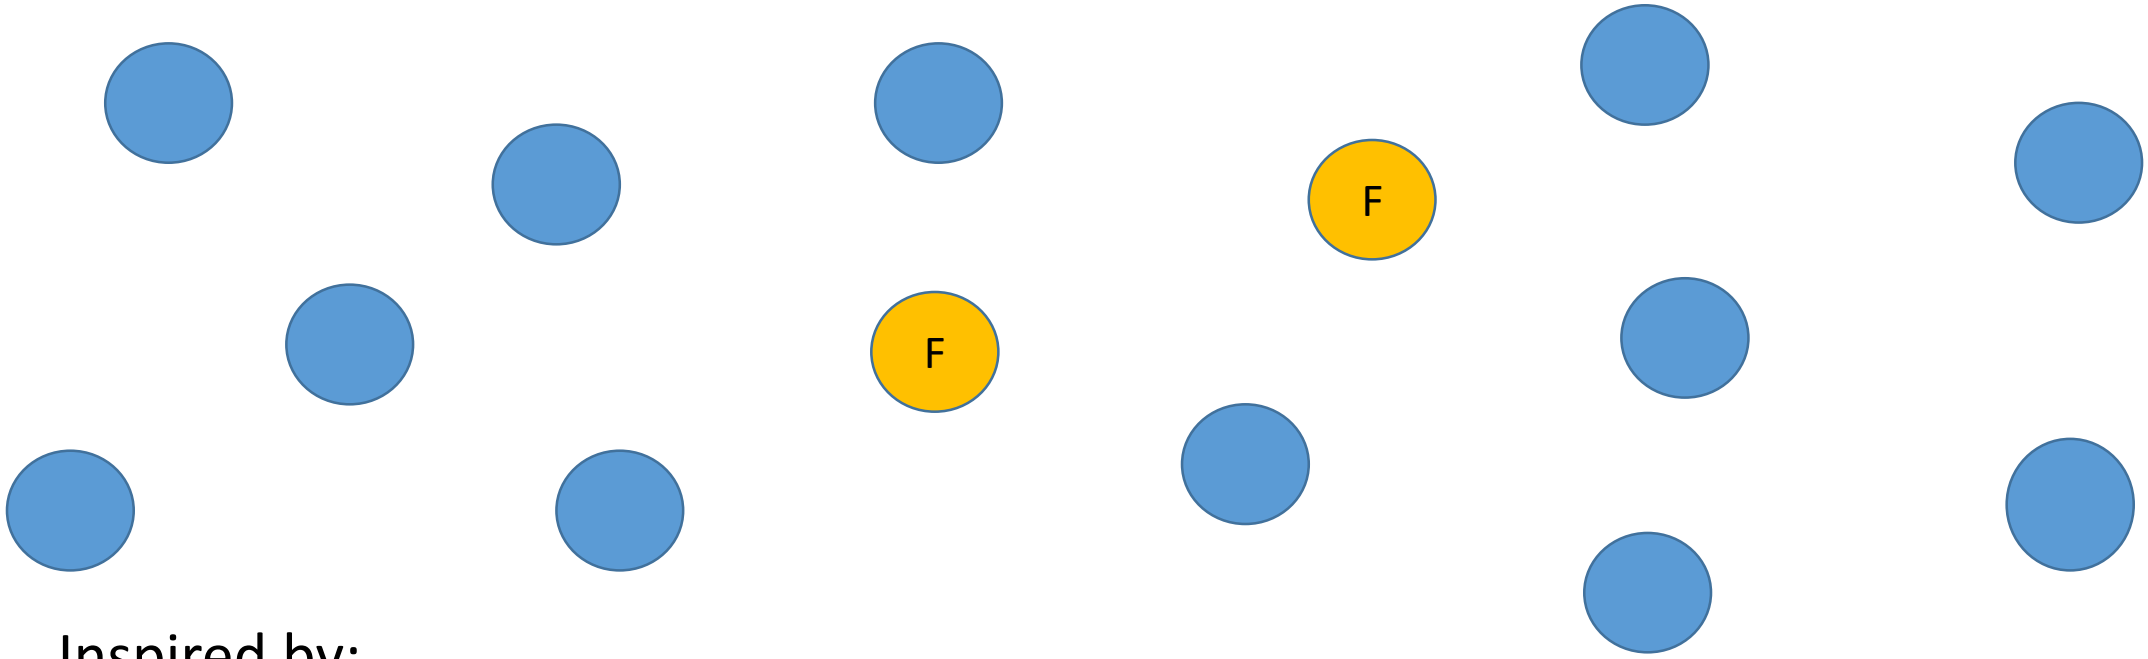

Inspired by:
NeighbourHood Discovery (NHDP) [RFC6130],
and Simple Multicast Forwarding (SMF) [RFC6621].

Node data

- Nodes maintain:
  - State (FF, PF, NF)
  - Set of 1-hop neighbours
  - RSSI (in and out) to neighbour
  - Set of 2-hop neighbours

Protocol

- Nodes link-local multicast (Trickle) info about their 1-hop neighbour
  - Address of neighbour
  - RSSI of link to neighbour
  - nr_FF: number of Forwarder neighbours
  - nr_Under: Number of neighbours with nr_FF < N_DUPLICATE
  - State
- CBOR format used in message
- Only messages from valid neighbours with rssi_in and rssi_out < 4
- On message reception:
  - Update node data and execute selection algorithm

Selection algorithm

Intended for stable networks (e.g. fixed installations with rare failures)

- Order neighbours on nr-Under and address
- First neighbour has values max_under and max-address
  - If nr_Under > max_under OR
    - nr_Under == max_under and address > max_address
  - Then set State to FF

Draft state

- Errors in simulation: no reliable algorithm results yet
- Find stability criterion for nr_Under values
- How to handle node departures ?
- Investigate stability of solutions

Option: In Node, calculate nr_Under for all neighbours
Currently: nr_Under calculated in every neighbour

Asymmetric AODV-P2P-RPL in Low-Power and Lossy Networks (LLNs)

draft-satish-roll-aodv-rpl-00

Satish Anamalamudi
Mingui Zhang
Charlie Perkins
S.V.R Anand
Dongxin Liu
satish.anamalamudi@huawei.com

Overview

- Extension of P2P-RPL [RFC6997]
  - Support for purely storing mode (hop-by-hop routing).
  - Support both Symmetrical and Asymmetrical bi-directional links.
  - Avoid address vector in “P2P-RDO” and “P2P-DRO” messages.

- AODV-RPL Mode of Operation (MoP)
  - RREQ Message.
  - RREP Message.

P2P-RPL Control messages

- Paired DODAGs.
  - DODAG RREQ-Instance
  - DODAG RREP-Instance

- DODAG RREQ-Instance
  - DIO + RREQ Option
  - Control transmission from OrigNode to TargNode.
  - Data transmission from TargNode to OrigNode.

- DODAG RREP-Instance
  - DIO + RREP Option
  - Control transmission from TargNode to OrigNode.
  - Data transmission from OrigNode to TargNode.

DODAG RREQ-Instance

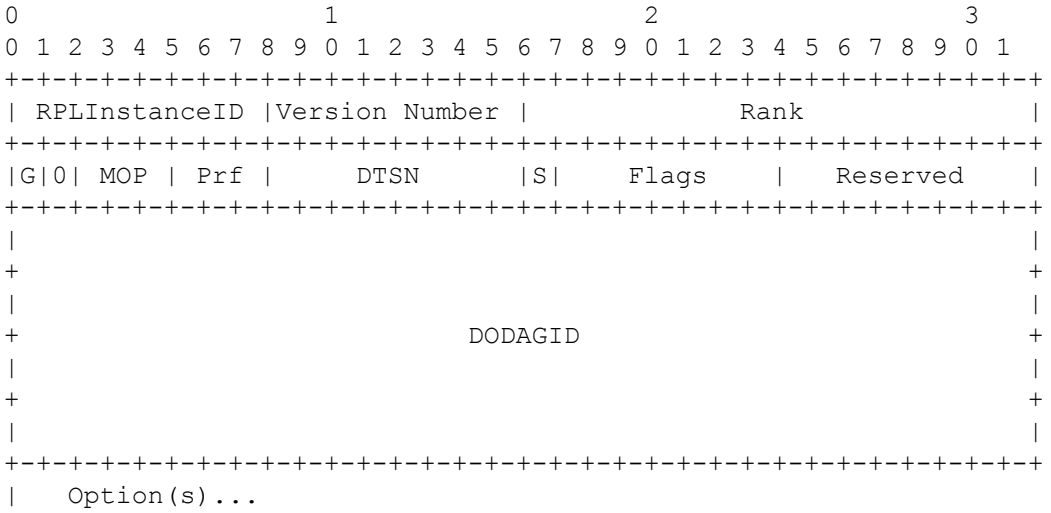

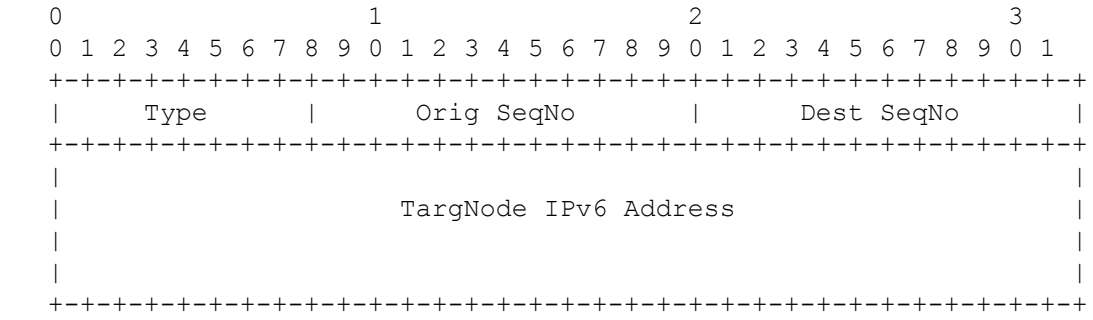

- Link Nature
  - “S" bit is added in DIO message.
  - Describe the link nature
  - Source IP address DODAGID
  - Destination IP address RREQ option

DODAG RREQ-Instance

- Symmetric Links
  - OrgNode set “S” bit to “1” during RREQ-Instance multicast.
  - ‘S’ bit is set to mean that the route is symmetric.
  - Intermediate node remain “S” bit as “1” (Symmetric links).
  - RREP-Instance is unicast for “S=1”.

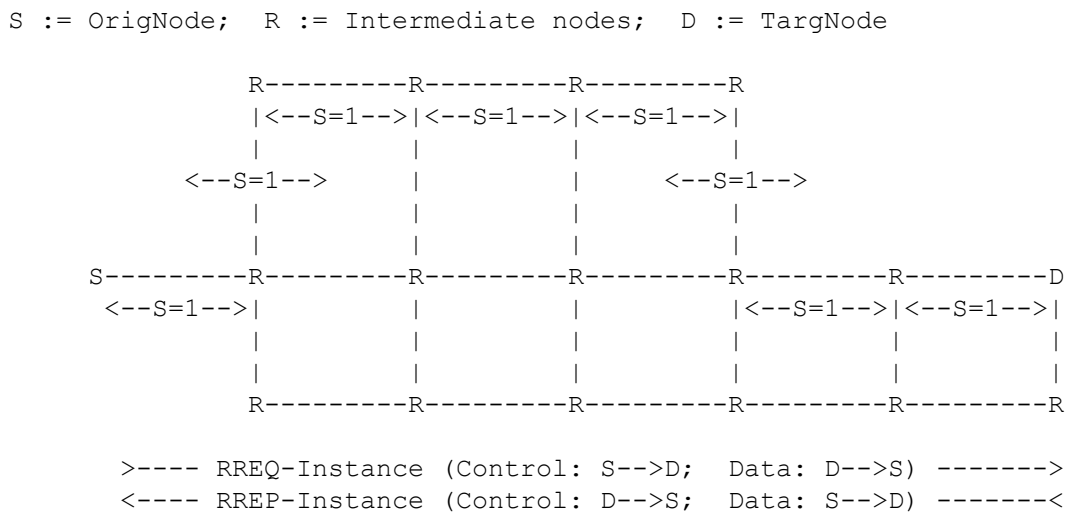

DODAG RREQ-Instance

- Asymmetric Links
  - OrgNode set “S” bit to “1” during RREQ-Instance multicast.
  - Intermediate node identify the bi-directional link reliability.
  - Update the “S” bit to “0” (for bi-directional Asymmetric links).
  - When “S” bit is already “0” then it should remain to “0”.

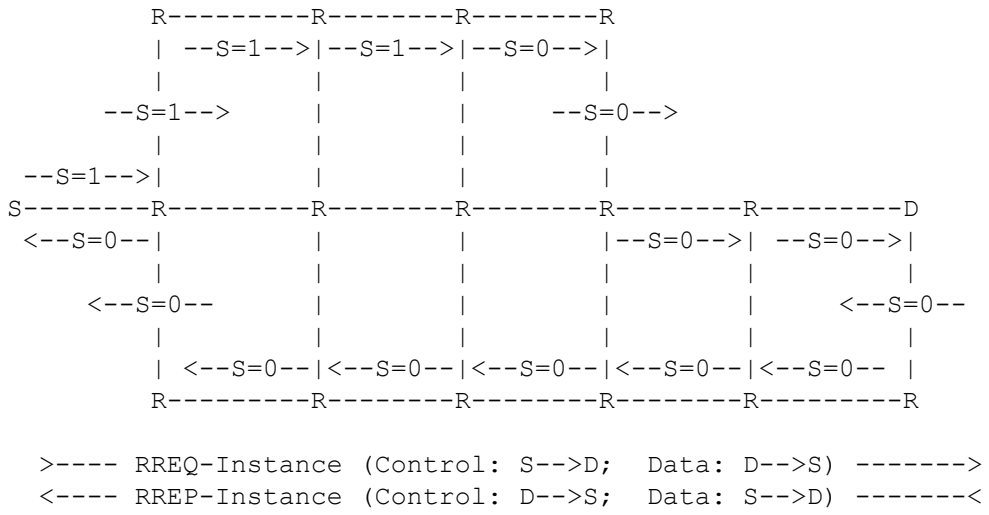

DODAG RREP-Instance

- Symmetric Links
- Destination:
  - If “S” bit is “1” then RREP-Instance is unicast.
  - Same links for RREQ-Instance and RREP-Instance.
  - Same links for upstream and downstream data transmission.

- Asymmetric Links
- Destination :
  - If “S” bit is “0” then RREP-Instance is multicast.
  - Different links for RREQ-Instance and RREP-Instance.
  - Different links for upstream and downstream data transmission.

RREP

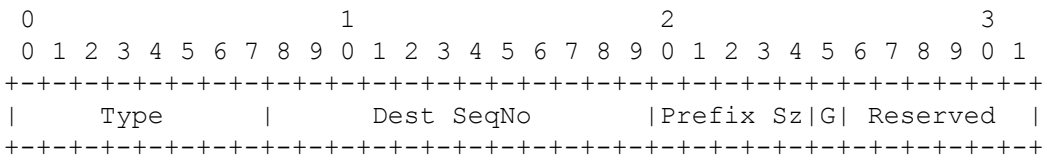

- RREP can be from TargNode or from Intermediate node.
- “G=1” means that the RREP is from Intermediate node.
  - Only if Intermediate node has a viable path towards destination.
- G-RREP can reduce the multicast control overhead by using unicast
  - Unicast G-RREP towards OrigNode.
  - Unicast RREQ-Instance towards TargNode.

Next Steps

- Comments and Questions

Thanks!

No-Path DAO of RPL

Problem Statement

https://tools.ietf.org/html/draft-jadhav-roll-no-path-dao-ps-01

Rahul, Rabi, Zhen@ Huawei

Hui Deng @China Mobile

IETF96

No-Path DAO is important to handle network dynamics

- **NPDAO Recap**
  - NPDAO = DAO (lifetime:=0)
  - Used for route invalidation
    - Release resources (for e.g. routing entries) along the previous path
  - Traverses upwards along the path from previous best parent towards the sink

- **Why NPDAO is important?**
  - Routing entries are the biggest memory-hogging component (especially in bigger storing-mode RPL networks)
    - In case of contention, its better to know which entries are non-active.
  - When a node switches parent, the sub-tree rooted at that node switches. Thus a high possibility of invalid route entries.
  - Impacts P2P traffic

Problems with NPDAO

1. NPDAO messaging depends on previous link which the node decided to no longer use
2. No route invalidation for dependent nodes
3. Possibility of Route downtime caused by async operation of NPDAO and DAO.
   - If NPDAO reaches before DAO, then the route will be unavailable till the time DAO reaches the all common parent nodes (A & BR in the example below).
4. Impact on P2P traffic because of NPDAO inefficiency

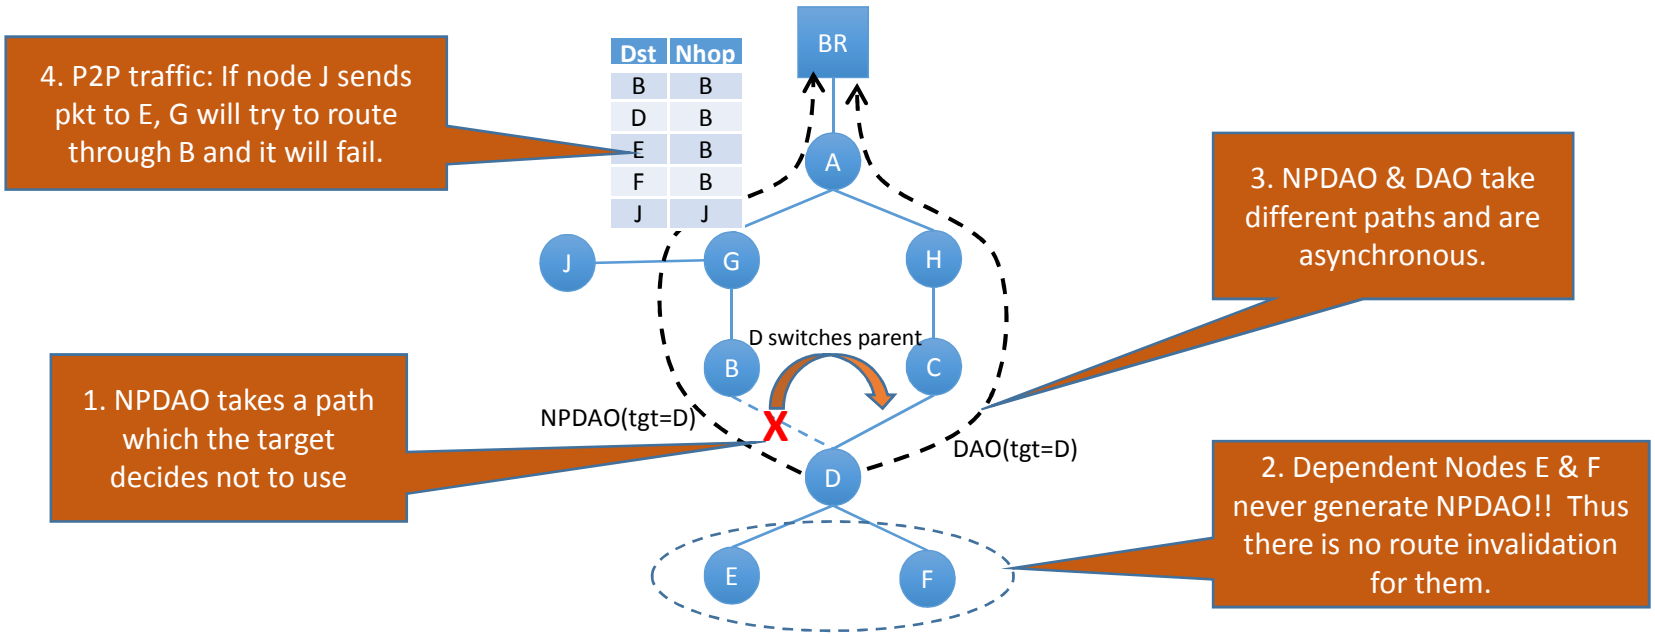

Possible existing solutions and corresponding problems…

- It may be possible for the parent to detect child unavailability*
  - Problem is parent cannot act unilaterally based on this info
  - On error detection, RFC6550 section 11.2.2.3 mentions parent can send "a packet" to clear the RPL states*… The provisions are vaguely stated…

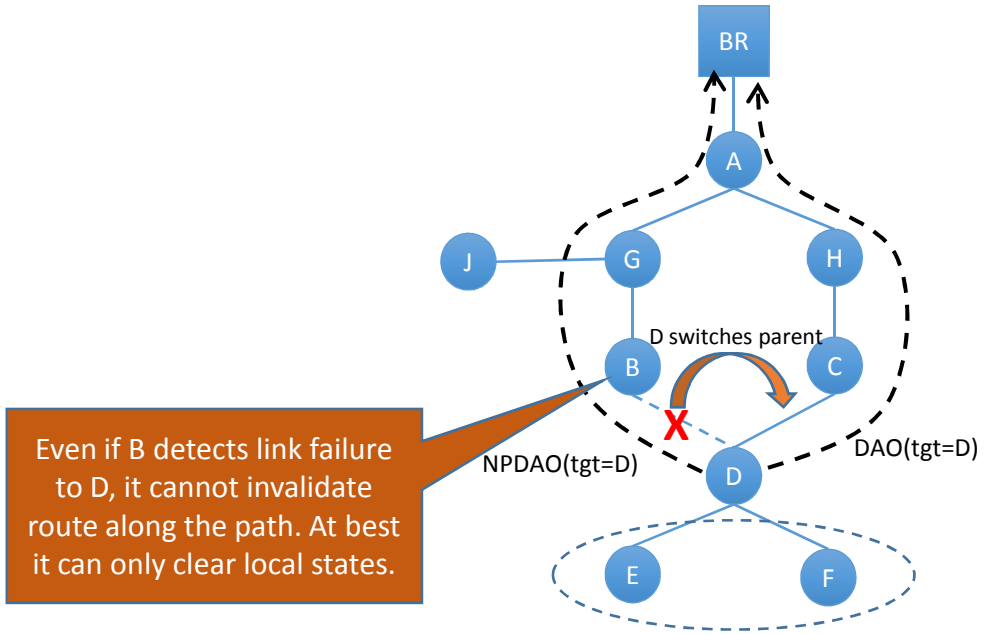

Section 11.2.2.3 states:
"With DAO inconsistency loop recovery, a packet can be used to recursively explore and clean up the obsolete DAO states along a sub-DODAG."

[*] Thanks to Cenk for pointing this out

Async NPDAO generation by parent node - Scenarios

- We received several comments describing possibility that the parent node can detect that child node is no more available and can initiate route cleanup along the previous path
  - Such detection can work only if there is any unicast P2P traffic to the child node originated along the previous path!
  - Or if the parent has some explicit detection for child node reachability which is seldom used in LLNs.
  - In case of sleepy leaf nodes such detection may not be feasible.

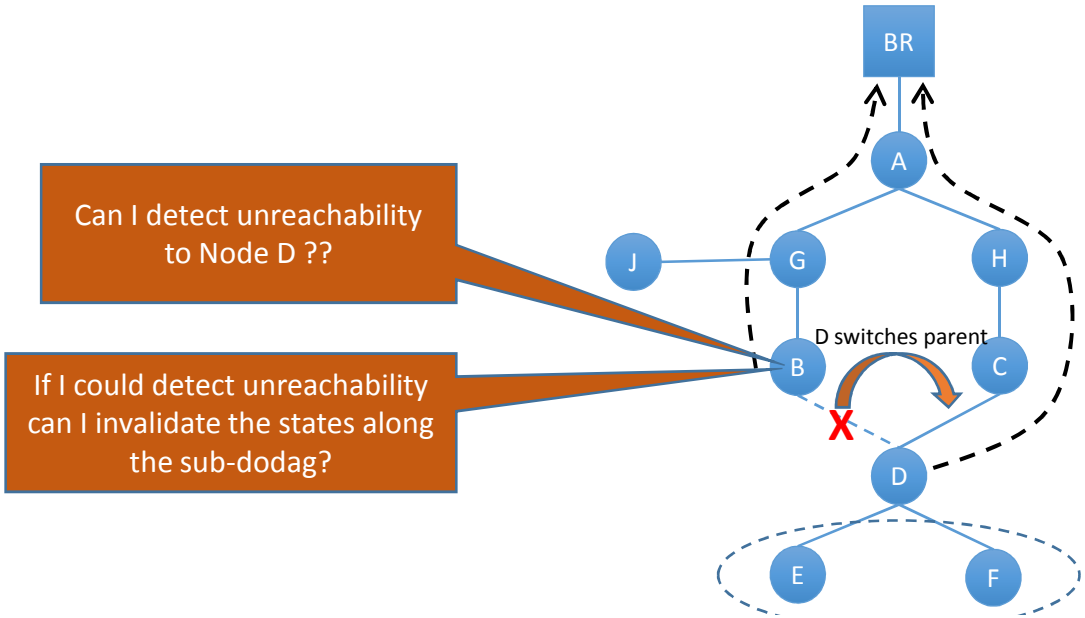

Impact of DAO state information - PathSequence

- *Assuming a parent node detects child unreachability and can generate an NPDAO on behalf of child node… What PathSequence can be used in the NPDAO ?*
- Every target is associated with a PathSequence number which relates to the latest state of the target. Every router en-route stores this sequence number to identify the freshness of the DAO.
- Consider two scenarios, Node B has stored PathSeq=X from previous DAO received from Node D.

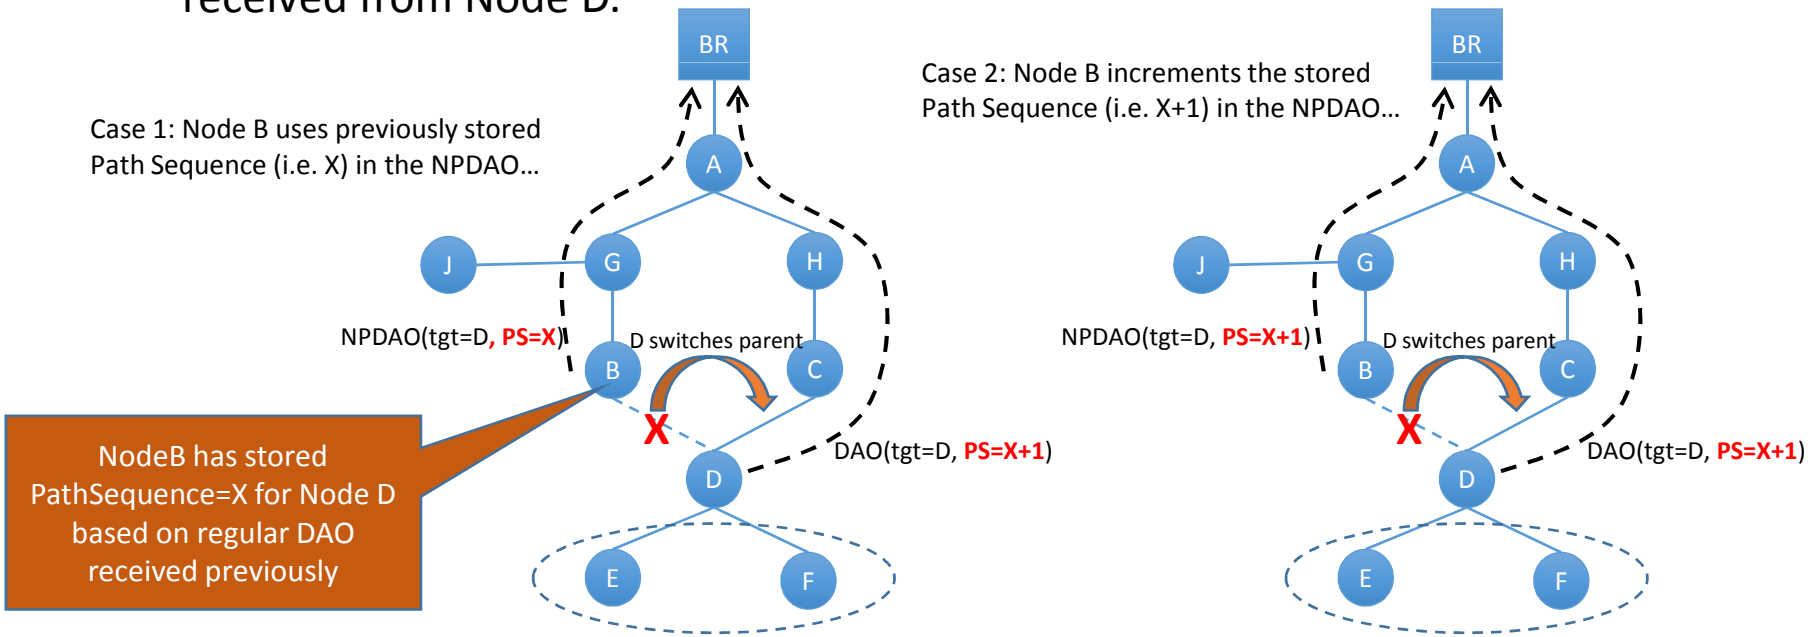

PS = PathSequence

Requirements for NPDAO improvements

- Should be tolerant to link failures to previous parent
- Should be possible to invalidate routes for dependent nodes as well
- Avoid route downtime because of NPDAO, DAO operation
- Should not introduce new memory requirement to handle route invalidation

Next Step

- Shall we work on this problem within ROLL WG?
- WG Adoption?

Thank you

ROLL working group

ROLL charter discussion

CHARTER PROPOSAL (Part I)

Charter for Working Group

Low power and Lossy Networks (LLNs) [RFC7102] [RFC7228] are made up of many embedded devices with limited power, memory, and processing resources. They are interconnected by a variety of links, such as IEEE 802.15.4, Bluetooth, Low Power WiFi, wired or other low power PLC (Powerline Communication) links. LLNs are transitioning to an end-to-end IP-based solution to avoid the problem of non-interoperable networks interconnected by protocol translation gateways and proxies.

Generally speaking, LLNs are characterized as follows, but not limited to:

- LLNs operate with a hard, very small bound on state.
- In most cases, LLN optimize for saving energy by using small packet headers and reduce amount of control packets.
- Typical traffic patterns are not simply unicast flows (e.g. in some cases most if not all traffic can be point to multipoint).
- In most cases, LLNs will be employed over link layers with restricted frame-sizes and low bit rates, thus a routing protocol for LLNs should be specifically adapted for such link layers.
- LLN routing protocols have to be very careful when trading off efficiency for generality; since LLN nodes do not have resources to waste.
- These specific properties cause LLNs to have specific routing requirements.

CHARTER PROPOSAL (Part II)

RFC 5548, 5673, 5826, and 5876 describe the requirements for LLNs from several application perspectives.

The Working Group has focused on routing solutions for the areas: connected home, building and urban sensor networks. It has developed a protocol set that takes into consideration various aspects including high reliability in the presence of time varying loss characteristics and connectivity while permitting low-power operation with very modest memory and CPU pressure in networks potentially comprising a very large number (several thousands) of nodes.

The Working Group continues to focus on routing issues for LLN and to maintain, improve and streamline the protocols developed by the working group, including RPL and MPL.

ROLL will coordinate closely with the working groups in other areas that focus on constrained node networks, such as 6lo (Internet) and CoRE (APP). behavior and the other protocols defined by the working group. The Working group will align with the 6man and BIER WGs when needed.

Work Items are:

- Guidance in using RFC6553, RFC6554, and IPv6-in-IPv6 encapsulation. The WGLC on this work will be shared with 6lo.
- Compression of RFC6553, RFC6554, and IP headers in the 6LoWPAN adaptation layer context. (coordinated with 6lo WG).
- Automatic selection of MPL forwarders to reduce message replication
- Data models for RPL and MPL management (coordinated with netmod wg)
- Alternative Multicast algorithm such as BIER forwarding.
- Methods to improve or correct the current RPL control messages behaviour, e.g. DIS and No-Path DAO.

CHARTER PROPOSAL (PART III)

Milestones

DATE	MILESTONE
May 2016	Initial submission of the draft about how to compress RFC6553, RFC6554, and IP headers in the 6LoWPAN adaptation layer context. to the IESG. draft-ietf-roll-routing-dispatch
August 2016	Initial Submission of the draft about when to use RFC6553, RFC6554, and IPv6-in-IPv6 encapsulation Draft-ietf-roll-useofrplinfo to the IESG.
October 2016	Submit draft about YANG MPL model to IESG
January 2017	Initial submission of draft about MPL selection to IESG
March 2017	Initial submission of draft about YANG RPL model to IESG
April 2017	Initial submission of draft about Bier Multicast to IESG
April 2017	Initial Submission of the No-Path DAO Problem Statement to the IESG
April 2017	Initial Submission of the DIS Modifications Document to the IESG
September 2017	Recharter WG or close

Open Floor

Thank you!