Tuesday 2021/08/31

IETF ROLL Interim - online

Routing over Low-Power And Lossy Networks

Chairs:
Dominique Barthel
Ines Robles

Secretary:
Michael Richardson

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

BCP 9 (Internet Standards Process)
BCP 25 (Working Group processes)
BCP 25 (Anti-Harassment Procedures)
BCP 54 (Code of Conduct)
BCP 78 (Copyright)
BCP 79 (Patents, Participation)
https://www.ietf.org/privacy-policy/ (Privacy Policy)

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

Meeting Materials

- Session: Tuesday, 31th August 2021 - 15:00-16:30 UTC
- Remote Participation
  - Meetecho: https://meetings.conf.meetecho.com/interim/?short=e89314d3-d762-4877-8f84-1108420ad8a4
  - CodiMD: https://codimd.ietf.org/notes-ietf-interim-2021-roll-02-roll
  - Material: https://datatracker.ietf.org/meeting/interim-2021-roll-02/session/roll
  - Jabber: xmpp:roll@jabber.ietf.org?join
  - Minute takers: **Please volunteer, thank you :)**

Agenda

Time (UTC)	Duration	Topic	Presenter
15:00 - 15:05	5 min	WG Status	Ines/Dominique
15:05 - 15:35	30 min	draft-ietf-roll-enrollment-priority	Michael
15:35 - 16:05	30 min	draft-iwanicki-roll-rnfd	Konrad
16:05 - 16:25	20 min	draft-ietf-roll-dao-projection	Pascal
16:25 - 16:30	5 min	Open Floor	Everyone

State of Active Internet-Drafts

	Draft	Status
2 IPRs	**draft-ietf-roll-efficient-npdao-18**	RFC 9009
	draft-ietf-roll-turnon-rfc8138-18	RFC 9035
	draft-ietf-roll-unaware-leaves-30	RFC 9010
	draft-ietf-roll-useofrplinfo-44	RFC 9008
	draft-ietf-roll-capabilities-07	Work in Progress
1 IPR	**draft-ietf-roll-dao-projection-16**	**Discussion Today**
	draft-ietf-roll-enrollment-priority-04	**Discussion Today**
	draft-ietf-roll-mopex-02	Work in Progress
2 IPRs	**draft-ietf-roll-nsa-extension-10**	Submitted to the IESG for publication
	draft-ietf-roll-aodv-rpl-09	AD Evaluation::Revised I-D Needed
	draft-ietf-roll-rpl-observations-05	Work in progress

Apr 2021

Related Internet-Drafts

Draft	Status
draft-iwanicki-roll-rnfd-00	Discussion Today
draft-pthubert-roll-rfc6550bis-01	Open to Work

Expired Internet-Drafts

Draft	Status
draft-jadhav-roll-storing-rootack-02	Call for adoption issued Jan 26th - No replies gotten
draft-thubert-roll-eliding-dio-information	To be Continued later
draft-ietf-roll-dis-modifications-01	To be Continued later
Draft-ietf-roll-mpl-yang-02	To be Continued later**?**
Draft-ietf-roll-bier-ccast-01	To be Continued later**?**

Done Milestones

Date	Milestone
DONE	Initial submission to the IESG of mechanism to turn on RFC8138 compression feature within a RPL network draft-ietf-roll-turnon-rfc8138
DONE	Initial submission of Common Ancestor Objective Functions and Parent Set DAG Metric Container Extension to the IESG draft-ietf-roll-nsa-extension
DONE	Initial submission of routing for RPL Leaves draft to the IESG draft-ietf-roll-unaware-leaves
DONE	Initial submission of a reactive P2P route discovery mechanism based on AODV-RPL protocol to the IESG draft-ietf-roll-aodv-rpl
DONE	Initial Submission of a proposal with uses cases for RPI, RH3 and IPv6-in-IPv6 encapsulation to the IESG draft-ietf-roll-useofrplinfo
DONE	Initial submission of a solution to the problems due to the use of No-Path DAO Messages to the IESG draft-ietf-roll-efficient-npdao

Milestones

Date	Milestone	Comment/Action Point
Oct 2021	Recharter WG or close	New topic might trigger **rechartering**
Dec 2020	Initial submission of Mode of Operation extension and Capabilities for RPL to the IESG - **draft-ietf-roll-mopex-cap**	Will **split** into 2 milestones
Jul 2020	Initial submission of a root initiated routing state in RPL to the IESG - **draft-ietf-roll-dao-projection**	Work in Progress, **to update the date**
Jul 2020	Initial submission of a YANG model for MPL to the IESG -**draft-ietf-roll-mpl-yang**	**?**
Jun 2020	Initial submission of Enabling secure network enrollment in RPL networks draft to the IESG - **draft-ietf-roll-enrollment-priority**	Work in Progress, **to update the date**
Jun 2020	Initial submission of a proposal to augment DIS flags and options to the IESG - **draft-ietf-roll-dis-modifications**	**To update the date**
Jun 2020	Initial submission of a proposal for Source-Route Multicast for RPL to the IESG - **draft-ietf-roll-ccast**	**?**

Tickets

- https://trac.ietf.org/trac/roll/report/2
  - aodv-rpl (#199, #200): should be fixed with the new version to submit, new comments from IESG review
  - dao-projection (#179, #180), RPLv2 (#187, #188)
- https://github.com/roll-wg/xxx/issues
  - rpl-observations (3 Open)
  - dao-projections (5 Open, new tickets to add based on recently emails)
  - Mopex (1)

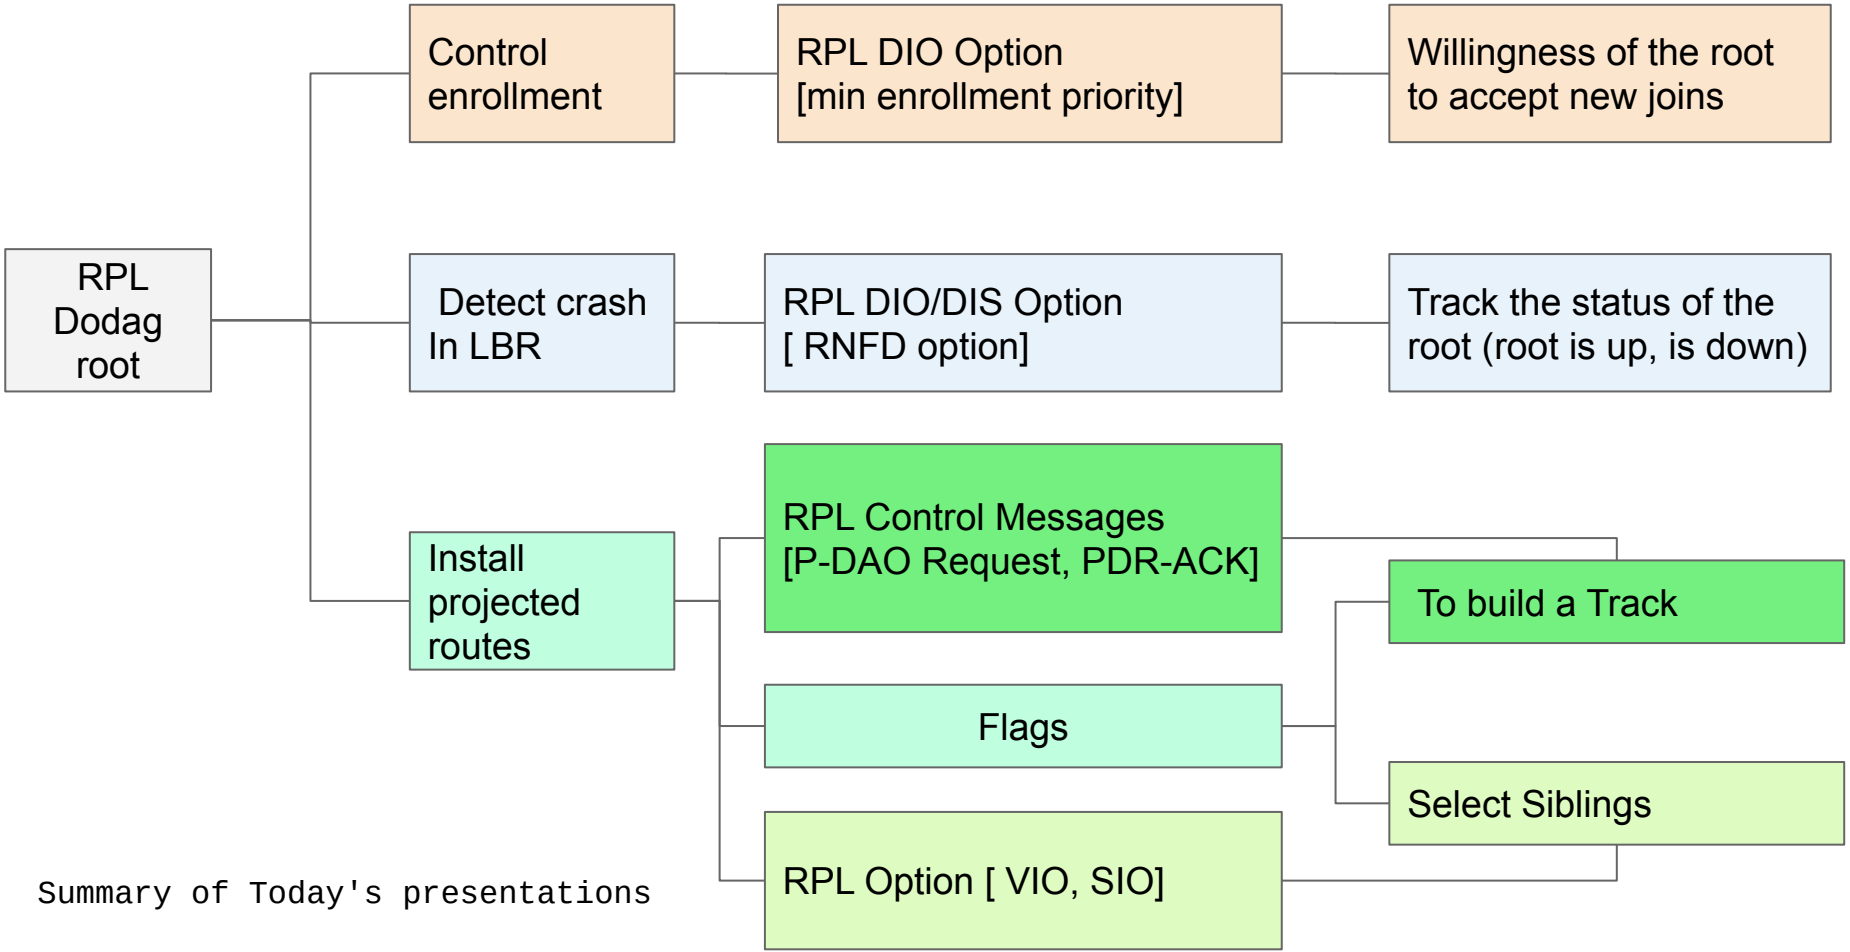

Summary of Today's presentations

Controlling Secure Network Enrollment in RPL Networks

draft-ietf-roll-enrollment-priority-04

Rahul Arvind Jadhav <rahul.ietf@gmail.com>
Pascal Thubert <pthubert@cisco.com>
Huimin She <hushe@cisco.com>
Michael Richardson mcr+ietf@sandelman.ca

post-IETF111 Virtual Interim, August 31 2021

The Story So Far

- Behaviour assumed in RFC9032
- Document Adopted March 2020
- Merged with draft-hushe-roll-dodag-metric after virtual interim meeting January 2021.
- Version -04 posted with merged document
- Reviews and Discussion Summer 2021

RFC9032 – Enhanced Beacon

Protocol Definition

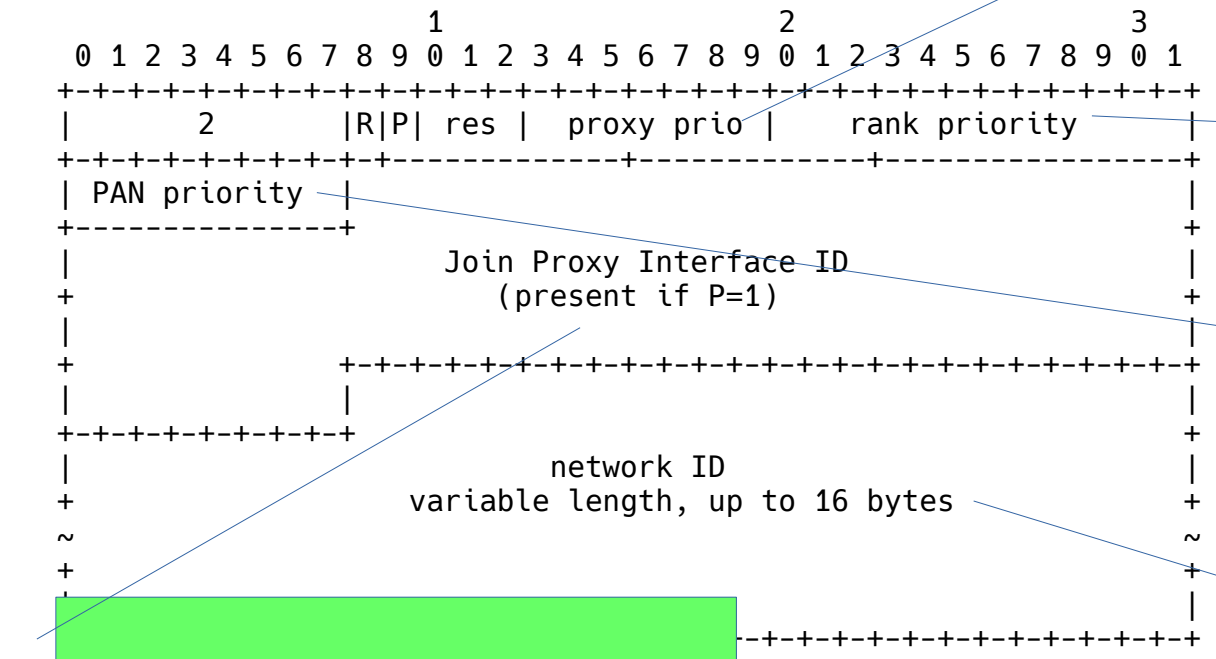

Priority of Join Proxy
(< is better)

Rank is locally known

Unclear how it is configured, could be added to RFC9031 (CoJP)

Calculated from DODAGID

Included if L2 addr != L3 addr

Enhanced Beacon Samples

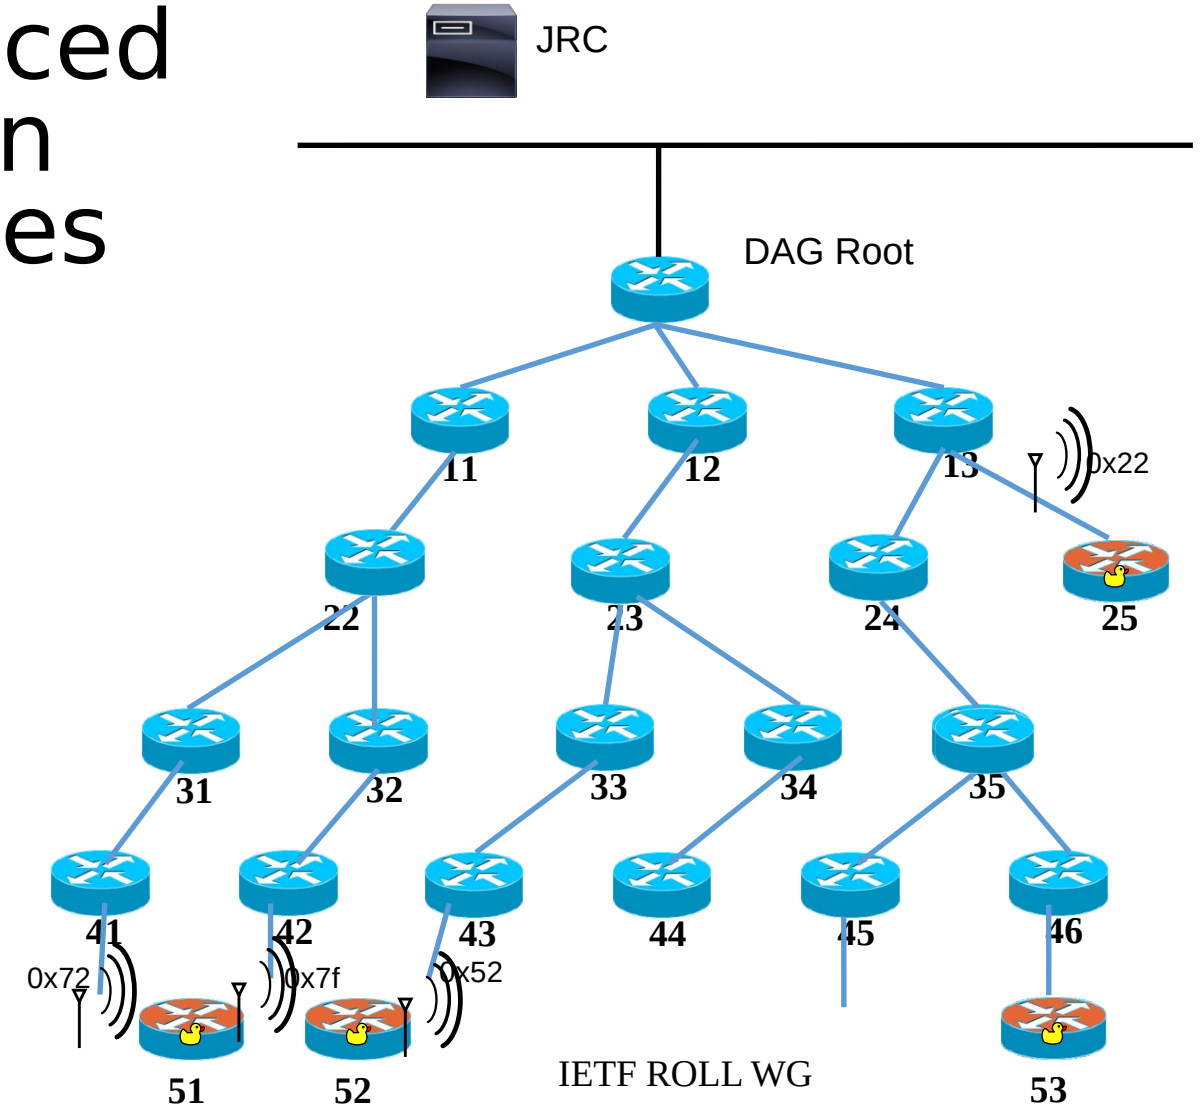

Enhanced Beacons

base diagram from PThubert

“The Minimum Enrollment Priority”

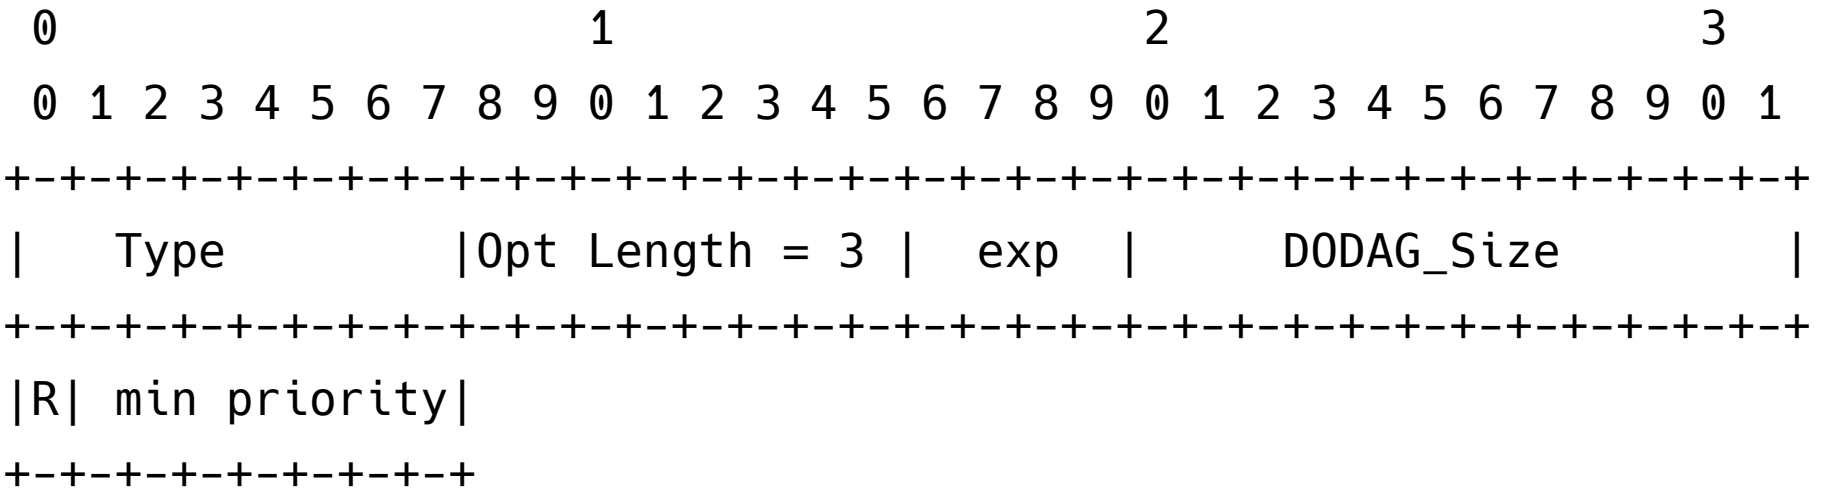

Issues from Review

- Trickle timer means that DIOs are not sent if there is no topology change.
  - So would changes to min priority be considered a change?
  - The DODAG size field could change quite often, particularly during network formation, how should it be dealt with?
- If updated min priority does not reset Trickle Timer, then this option needs to go into some new flooded control.
  - What are the desired properties of this new control, and what other things should go into it?

Discussion!

draft-ietf-roll-enrollment-priority-04

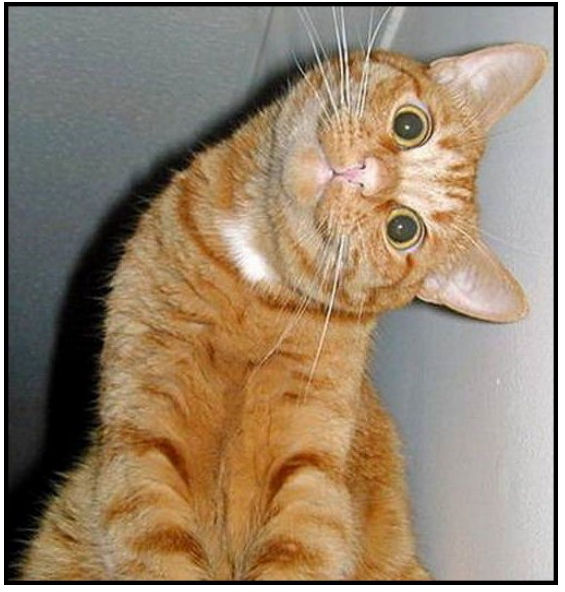

Auxiliary Slides Follow

Example enrollment priority

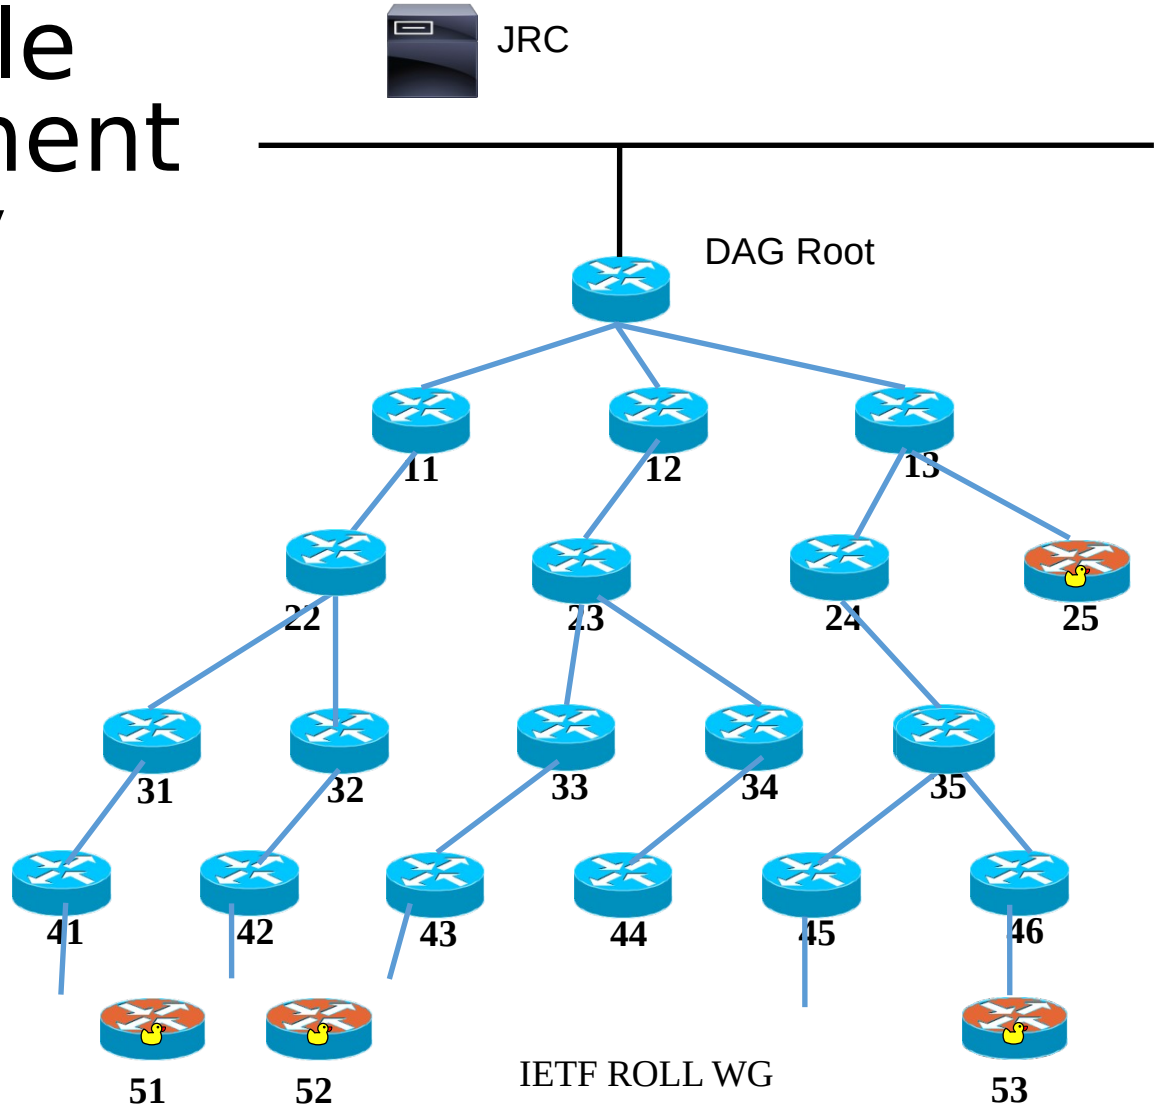

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

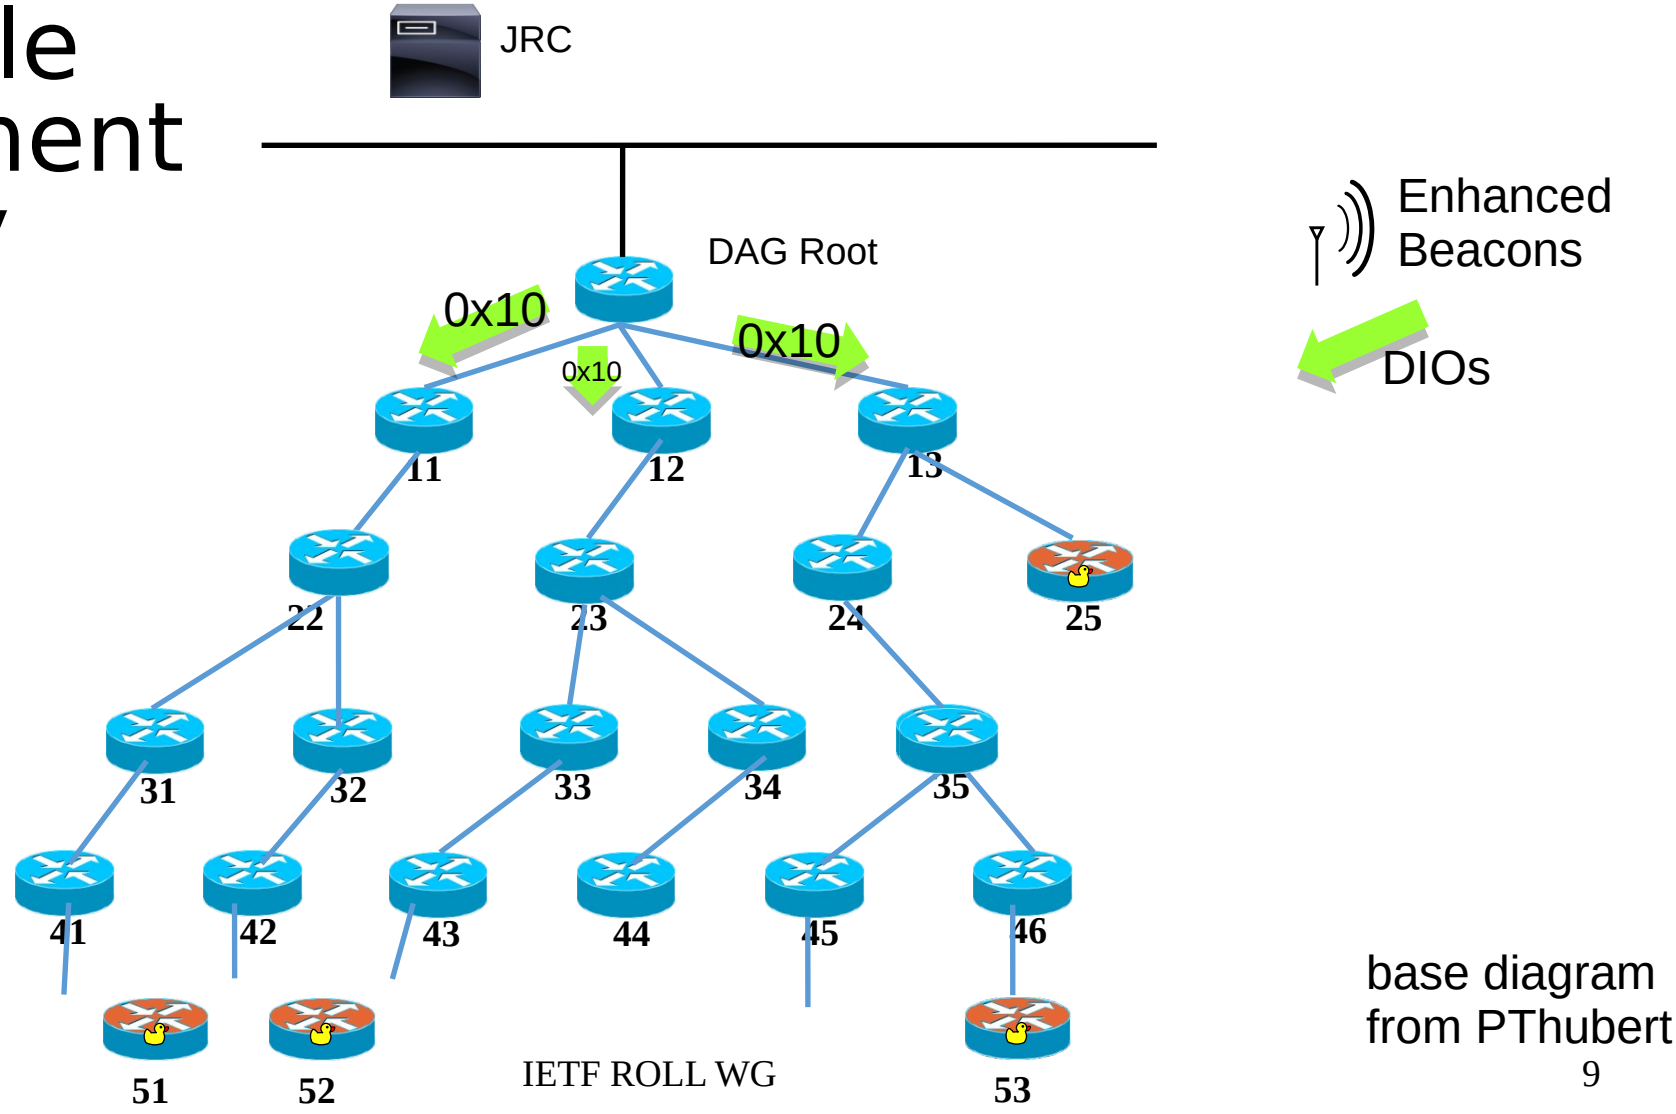

base diagram from PThubert

Example enrollment priority

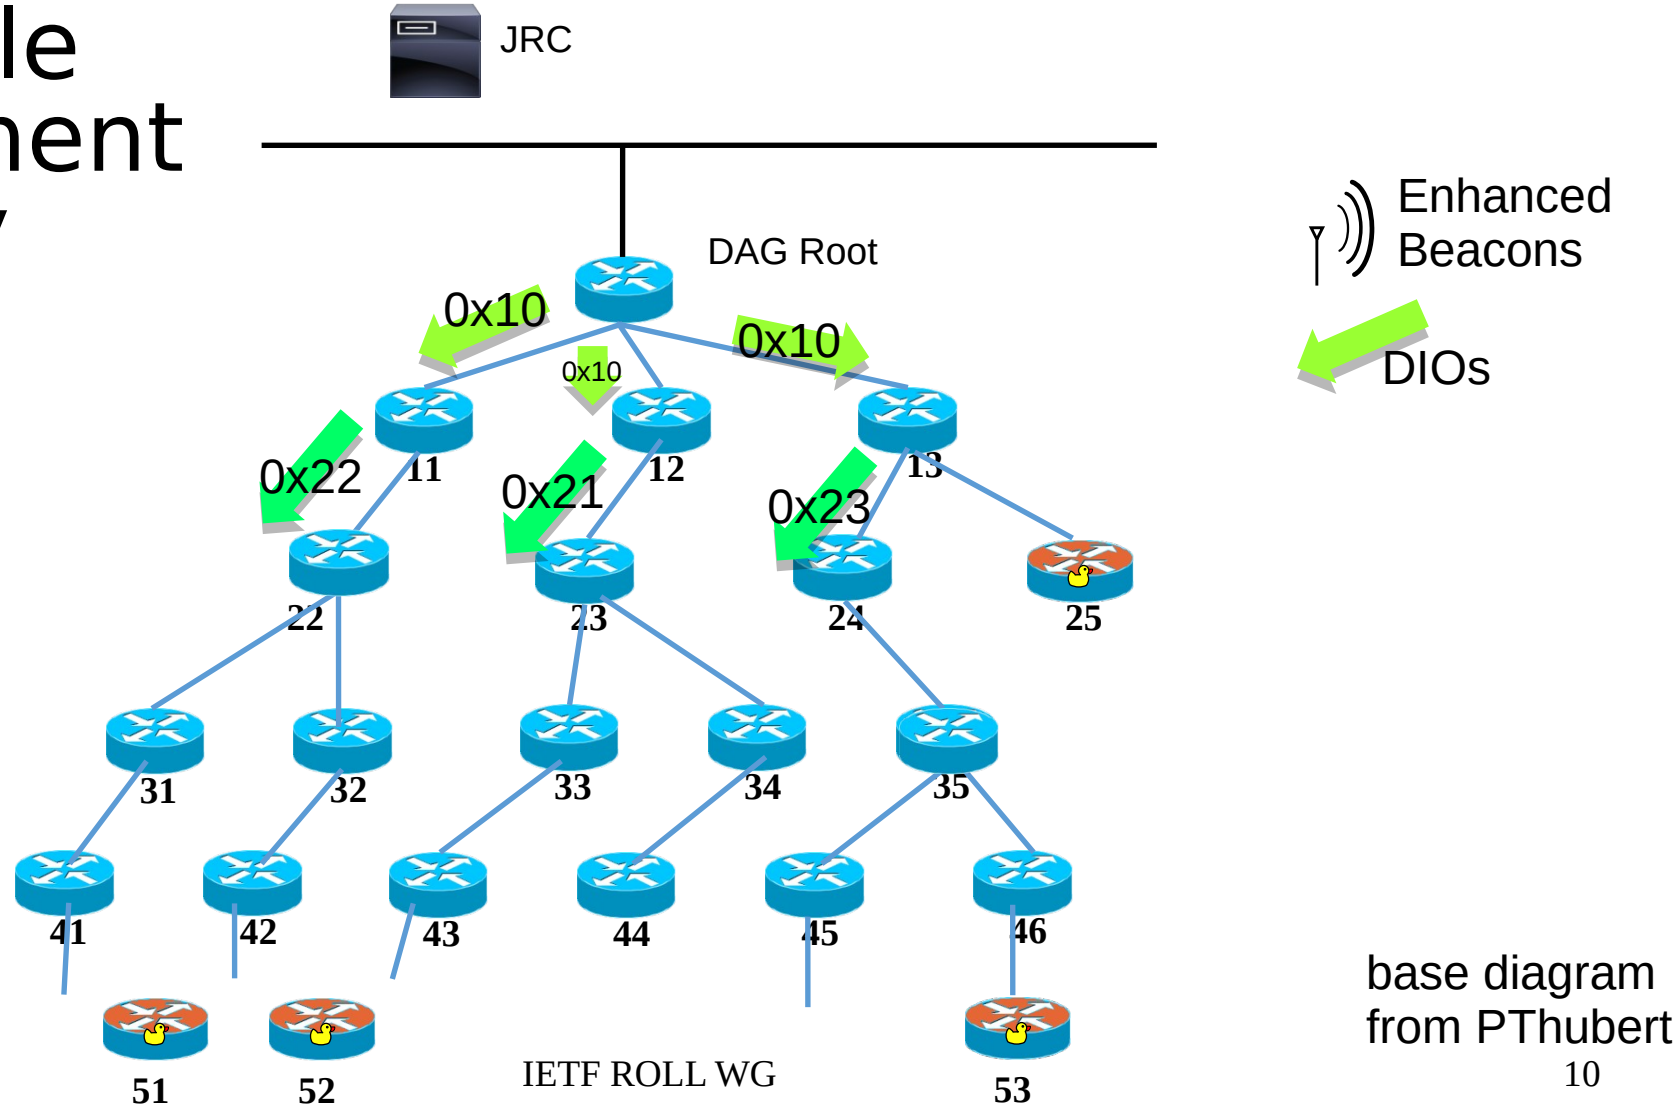

base diagram from PThubert

Example enrollment priority

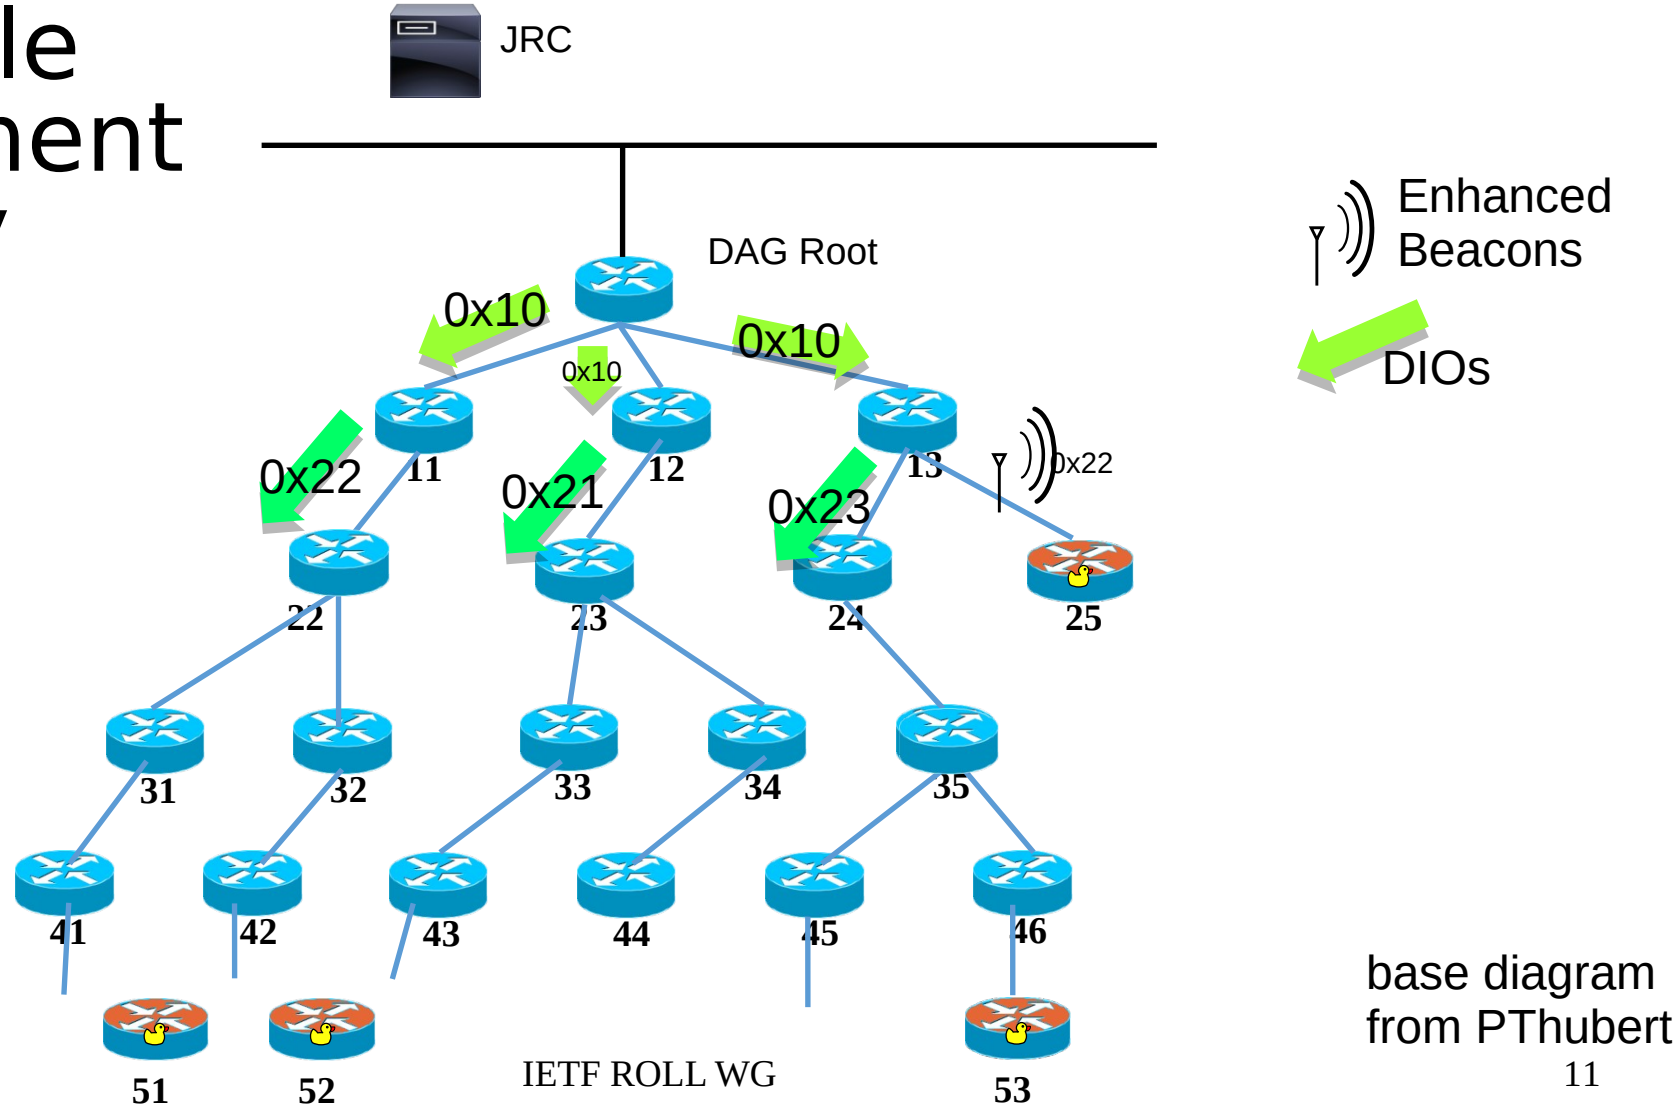

base diagram from PThubert

Example enrollment priority

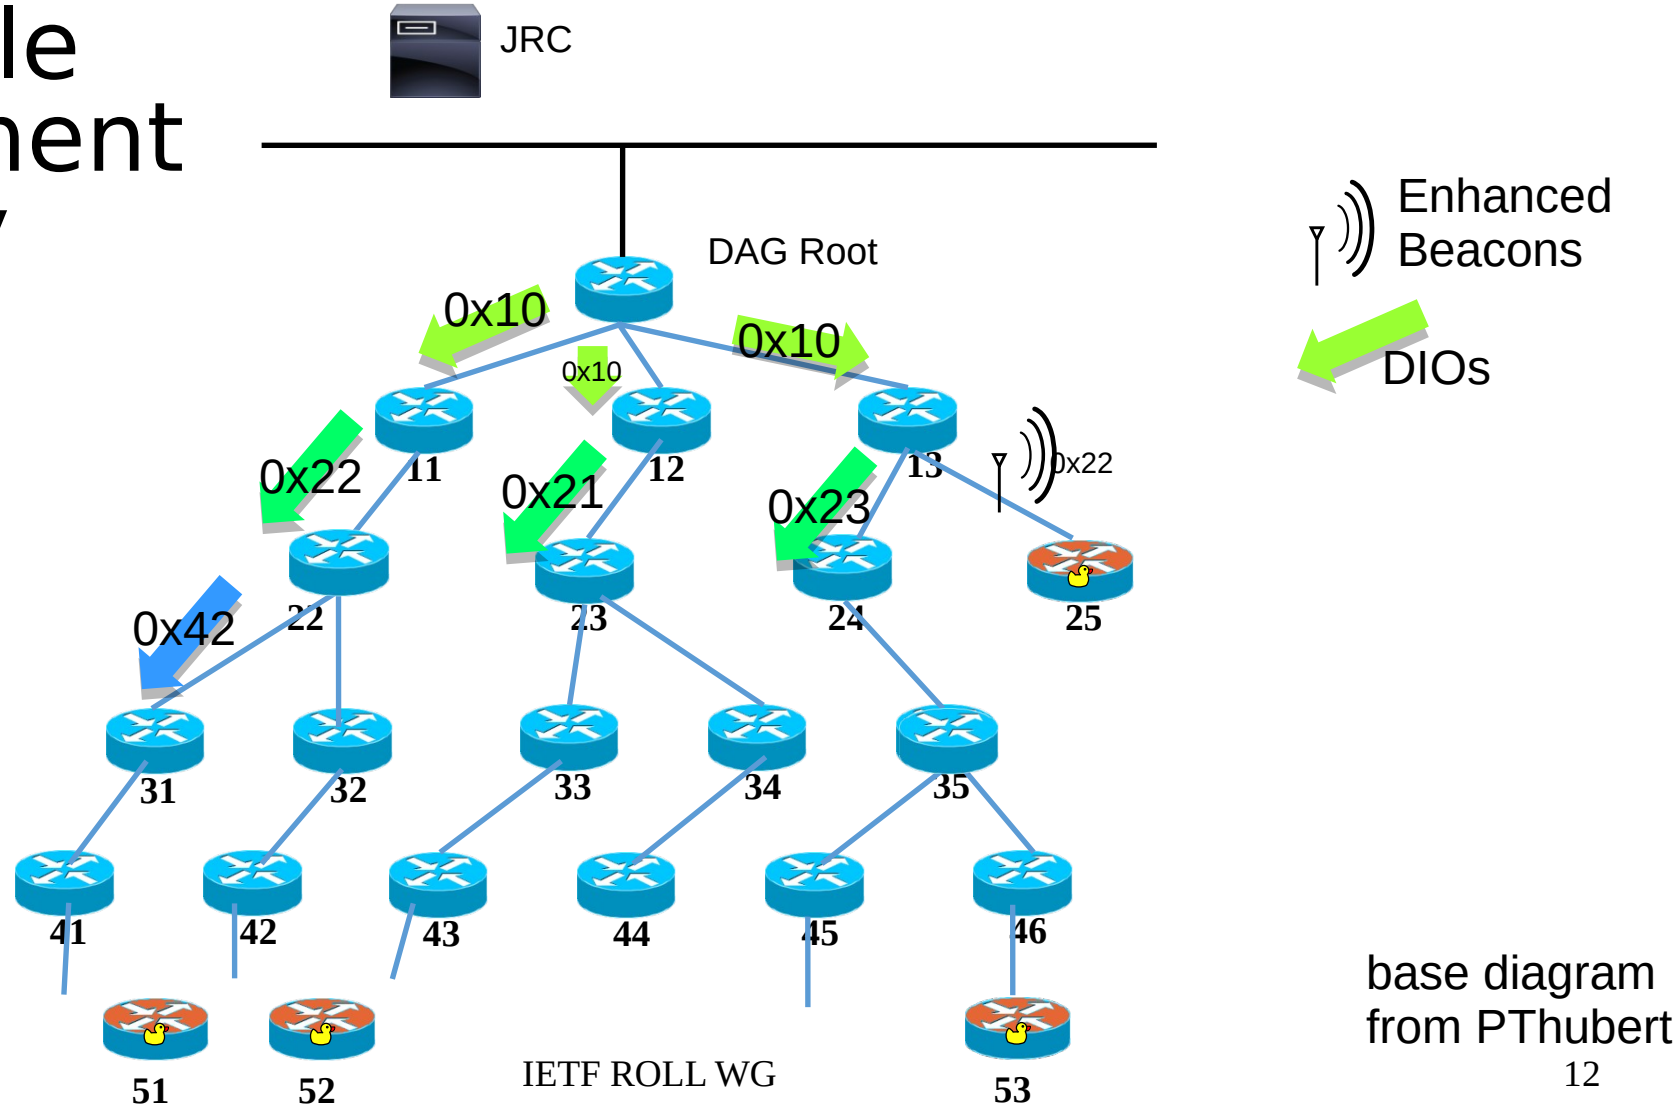

base diagram from PThubert

Example enrollment priority

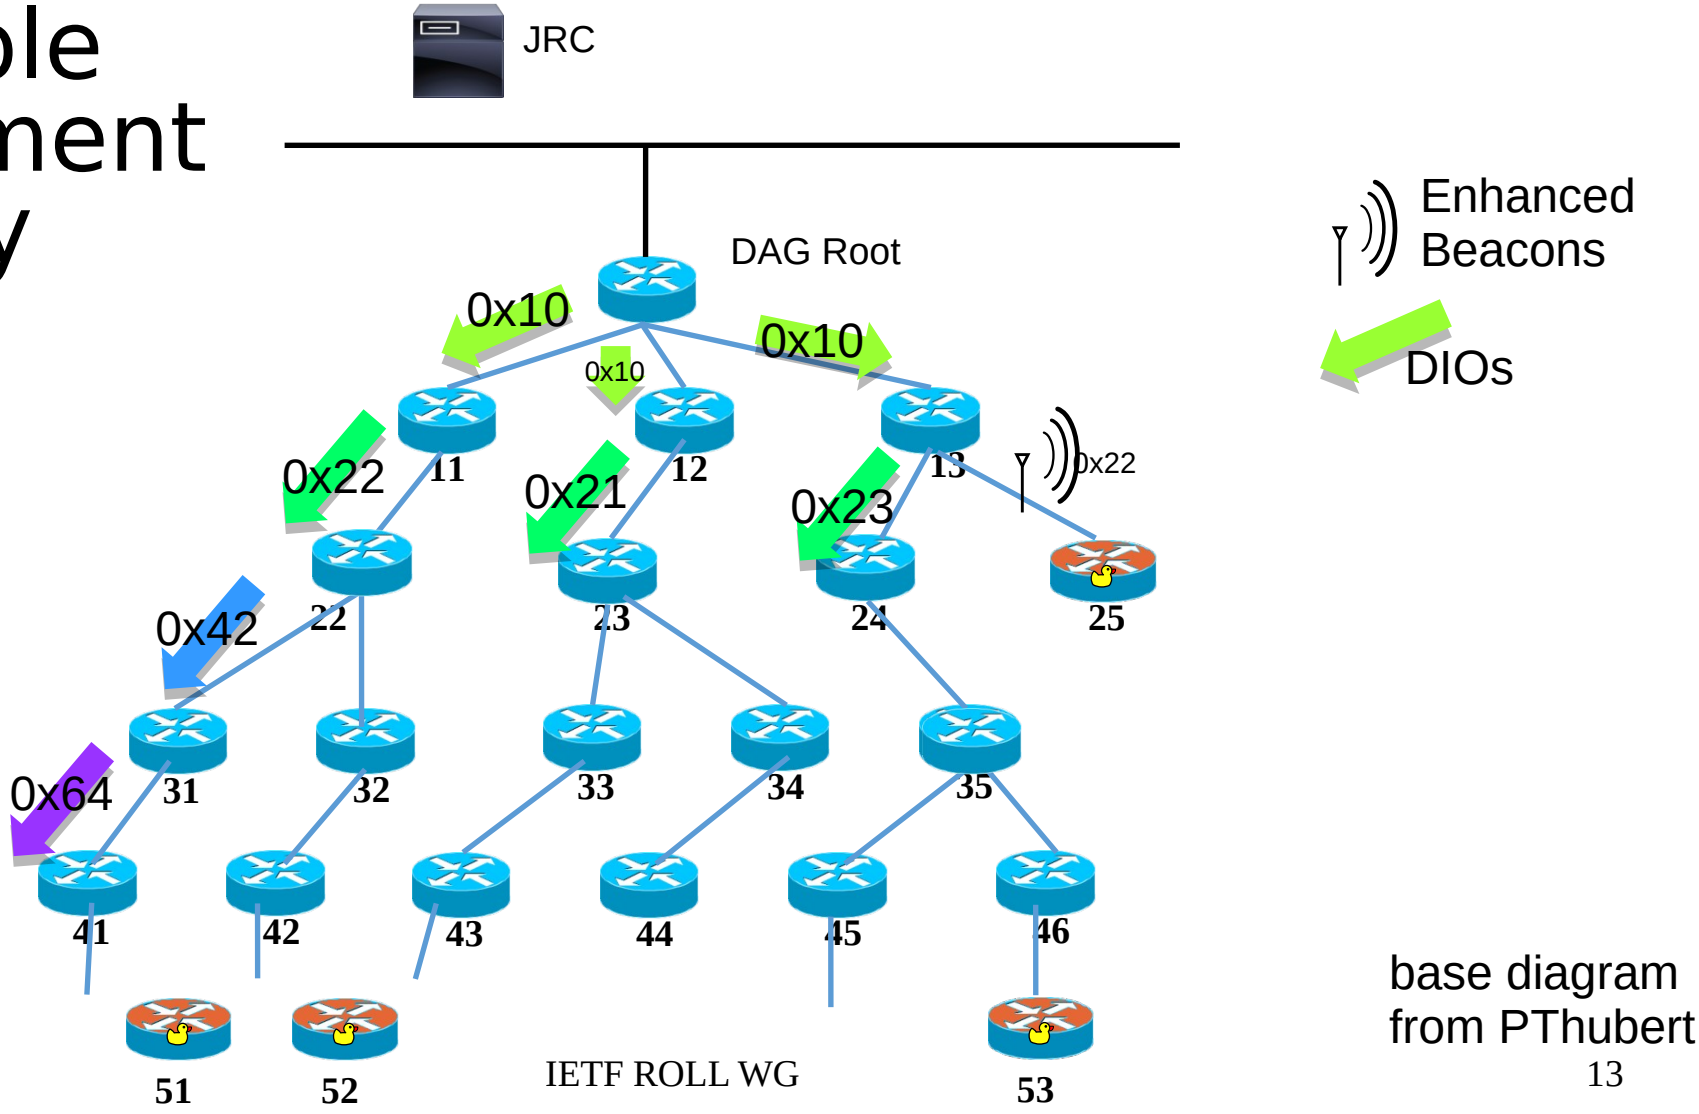

base diagram from PThubert

Example enrollment priority

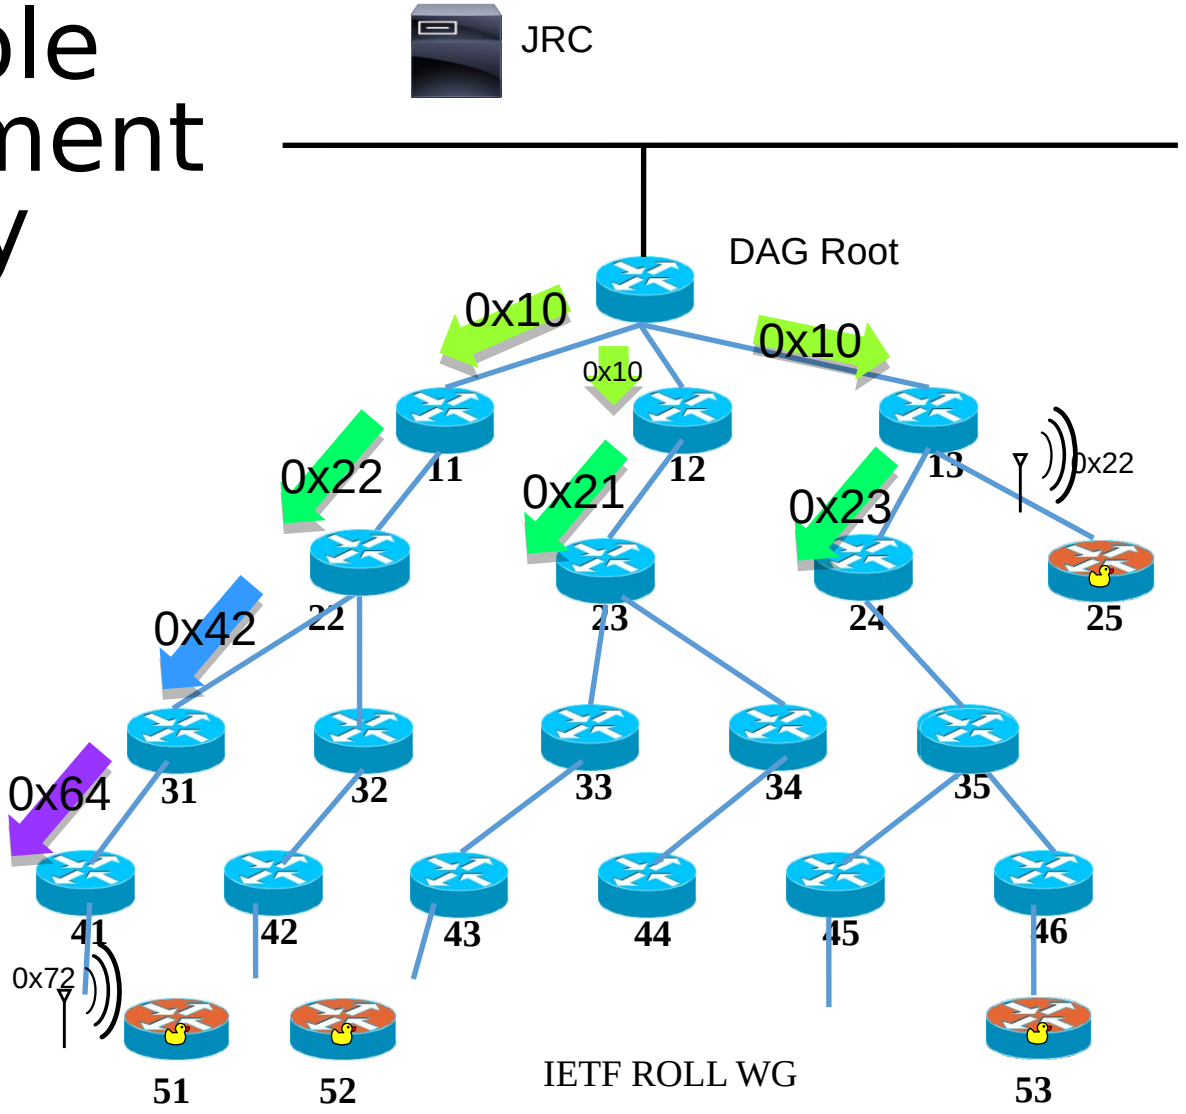

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

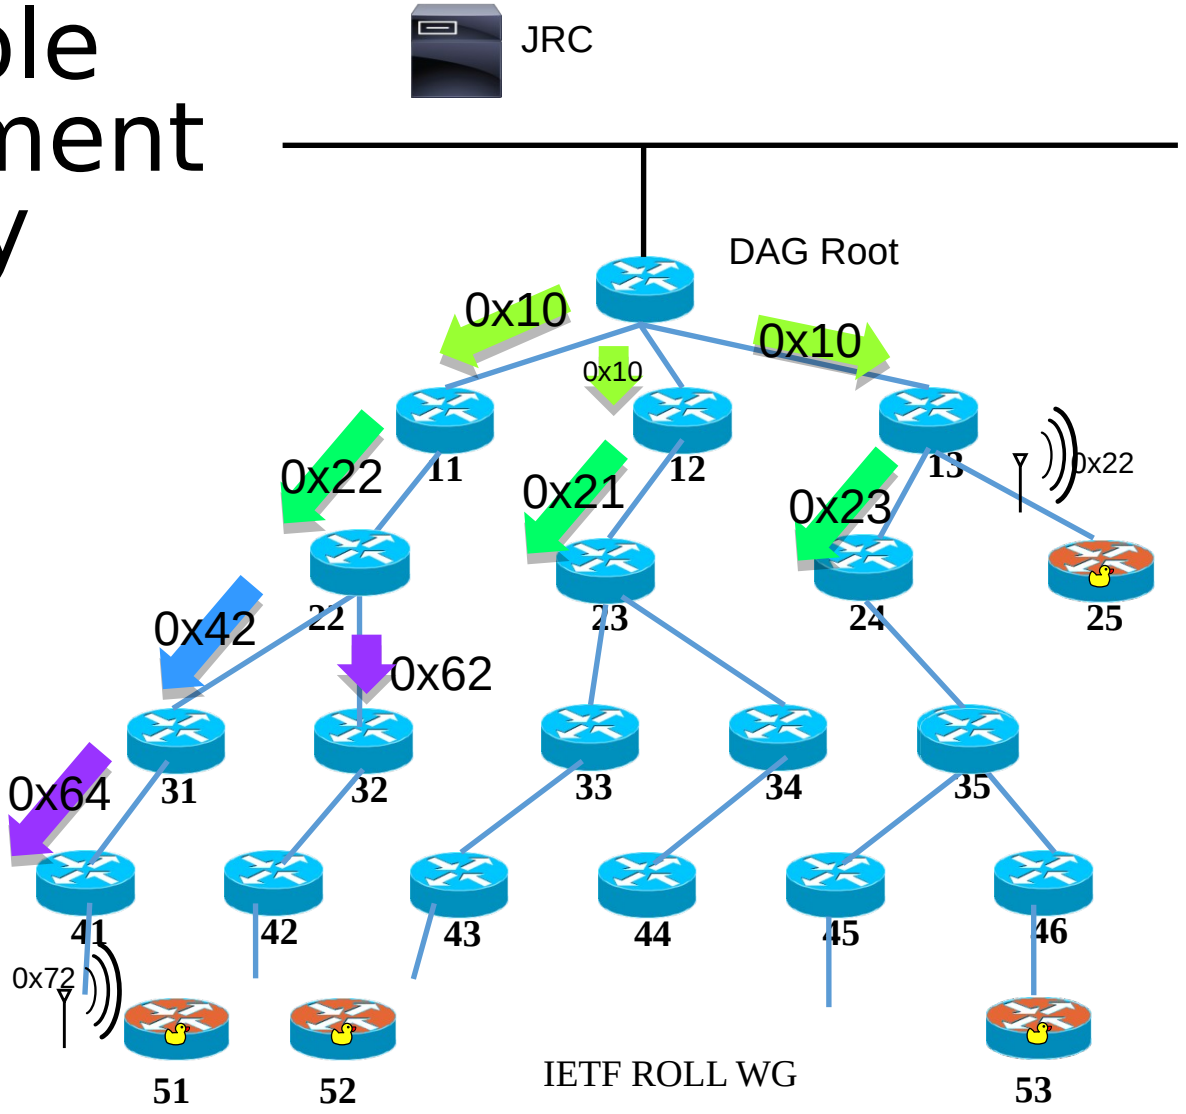

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

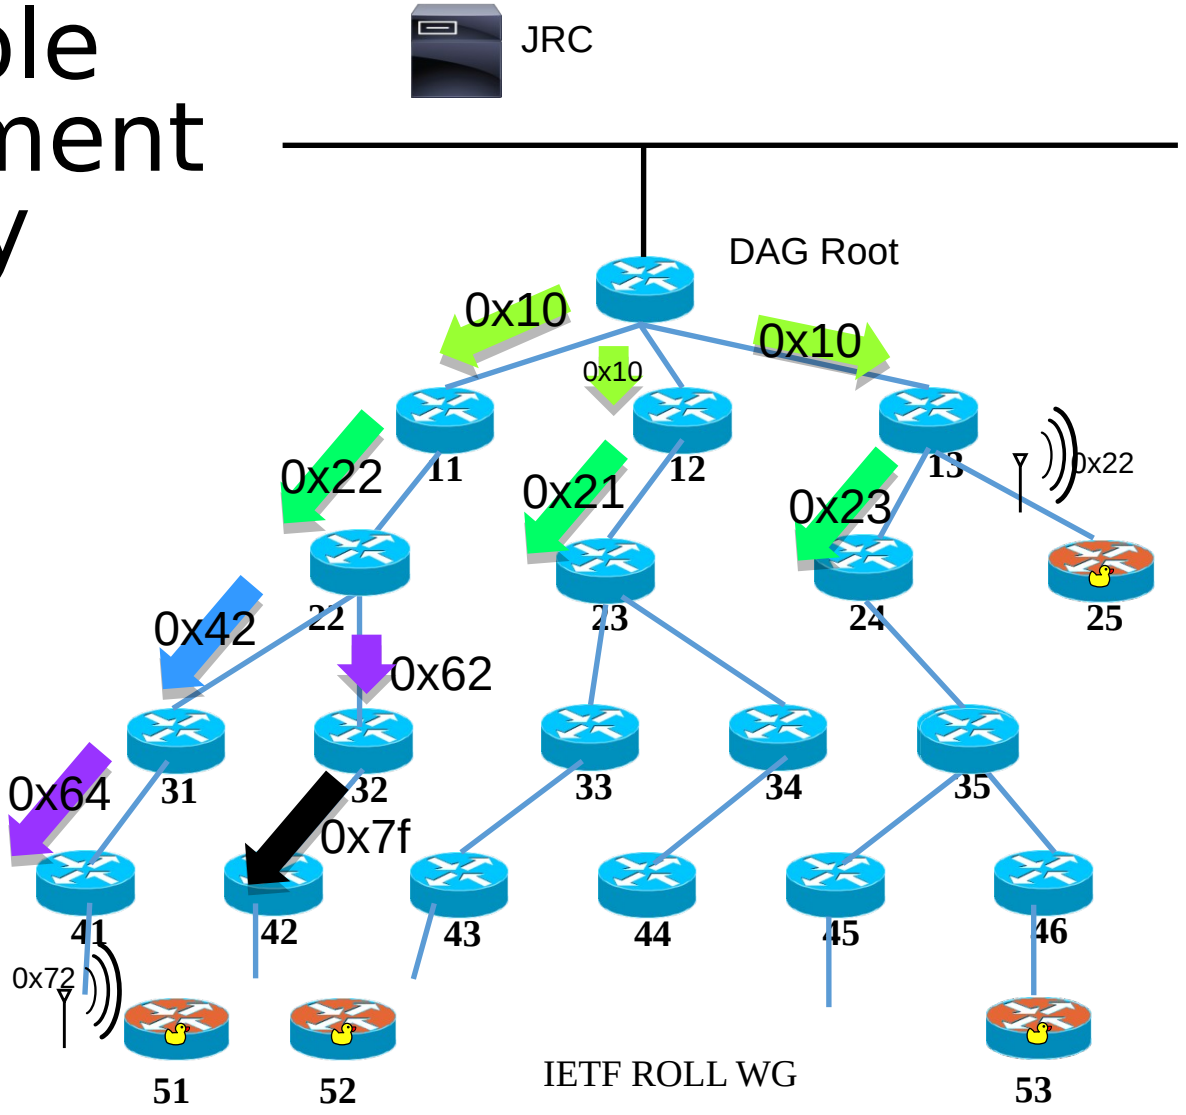

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

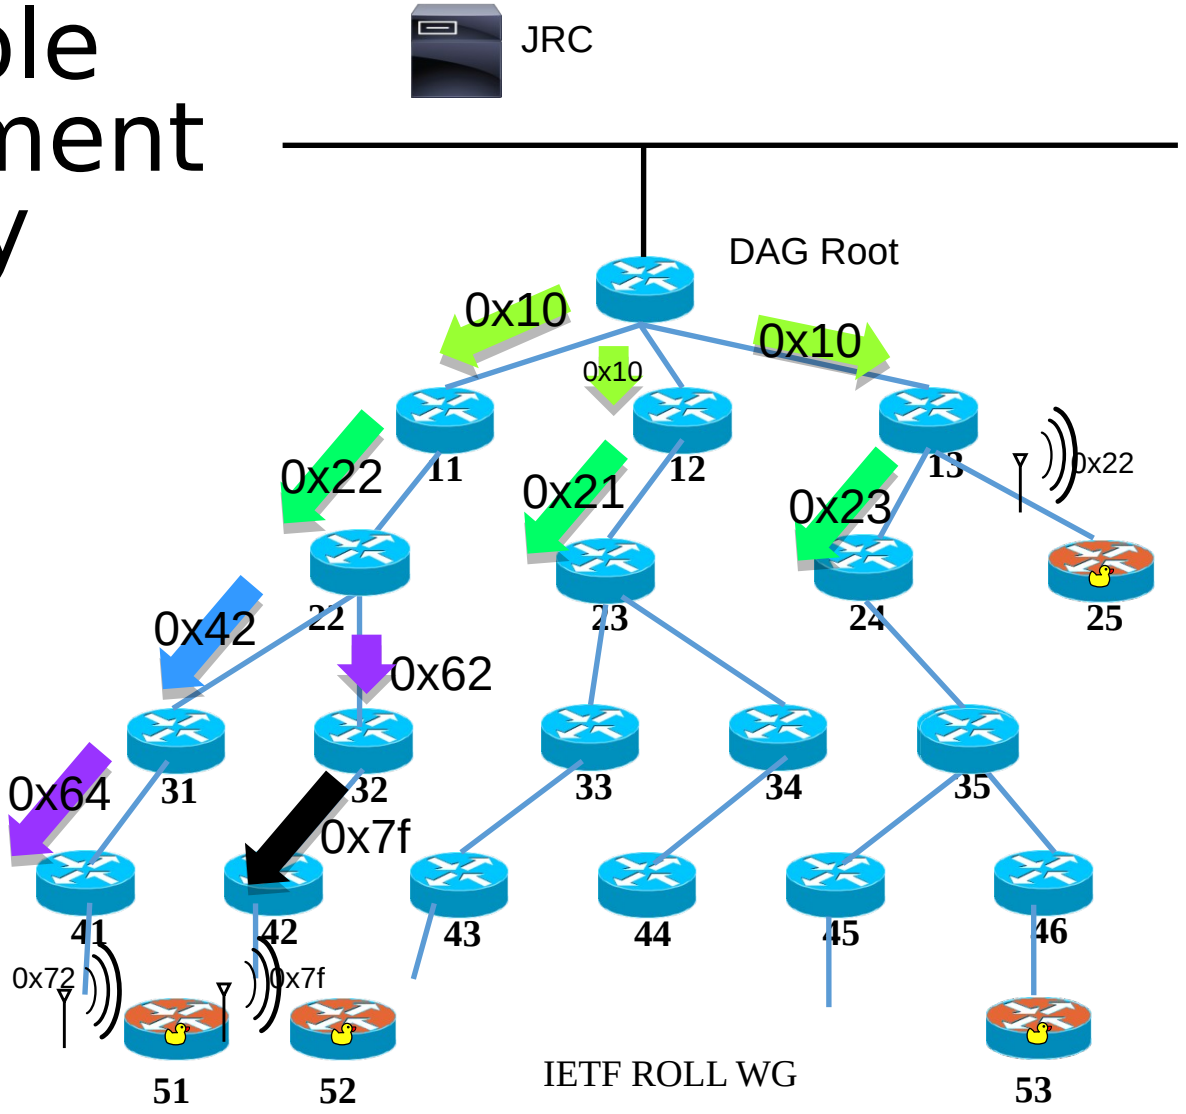

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

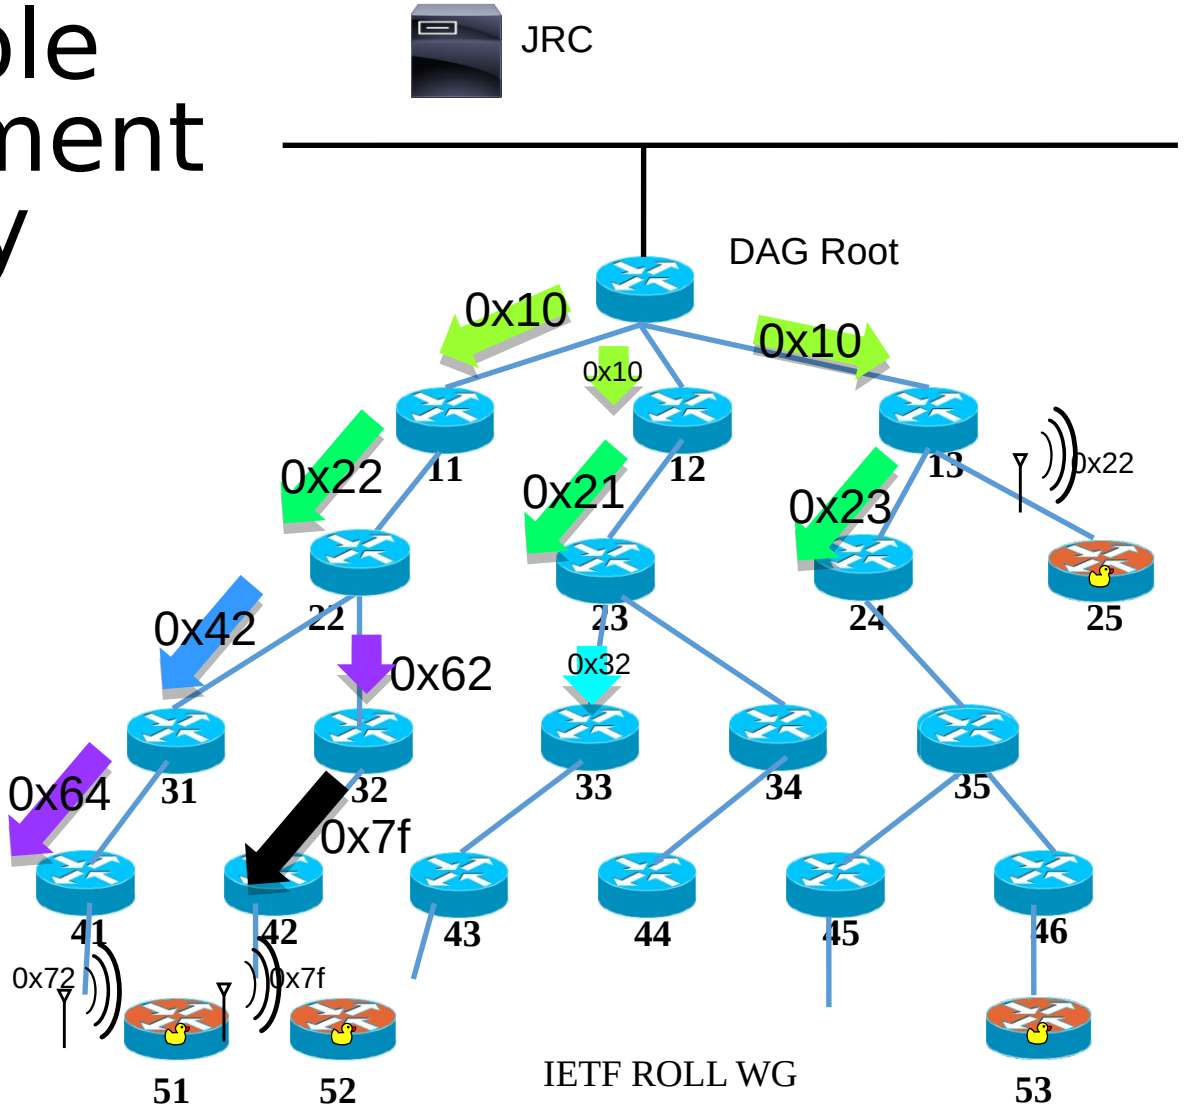

Enhanced Beacons

DIOs

base diagram from PThubert

Example enrollment priority

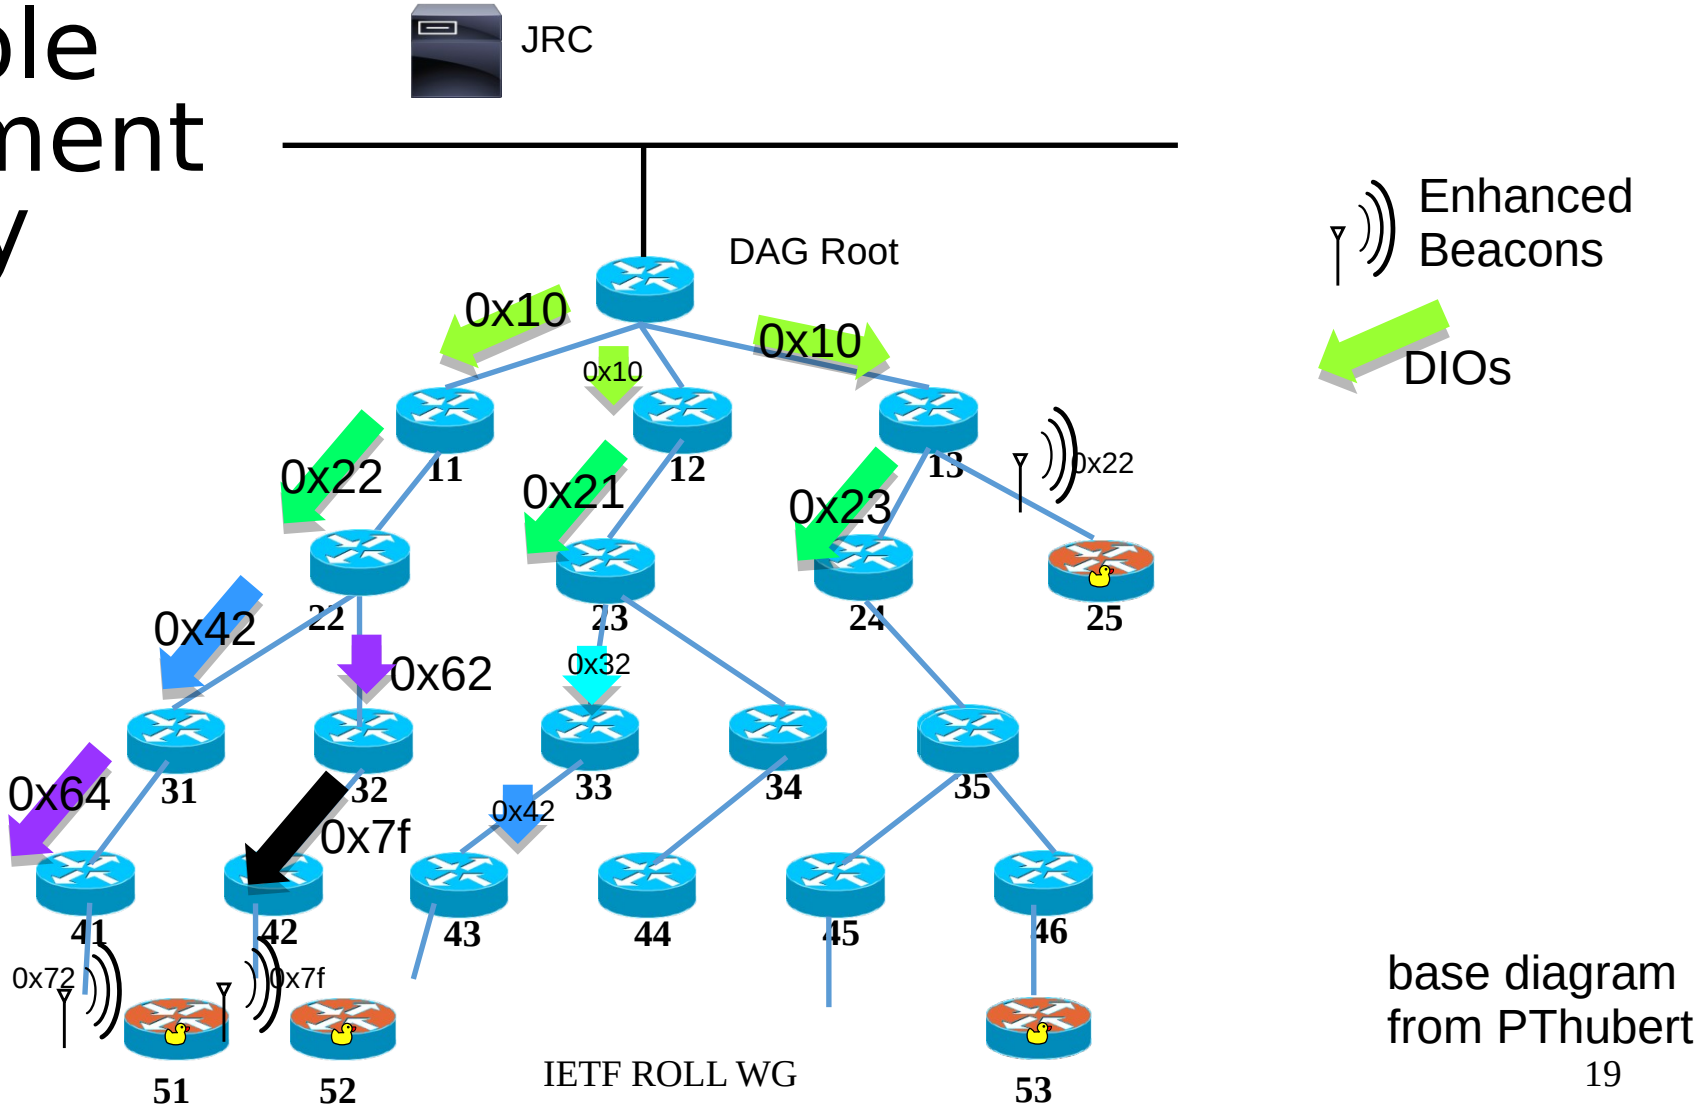

base diagram from PThubert

Example enrollment priority

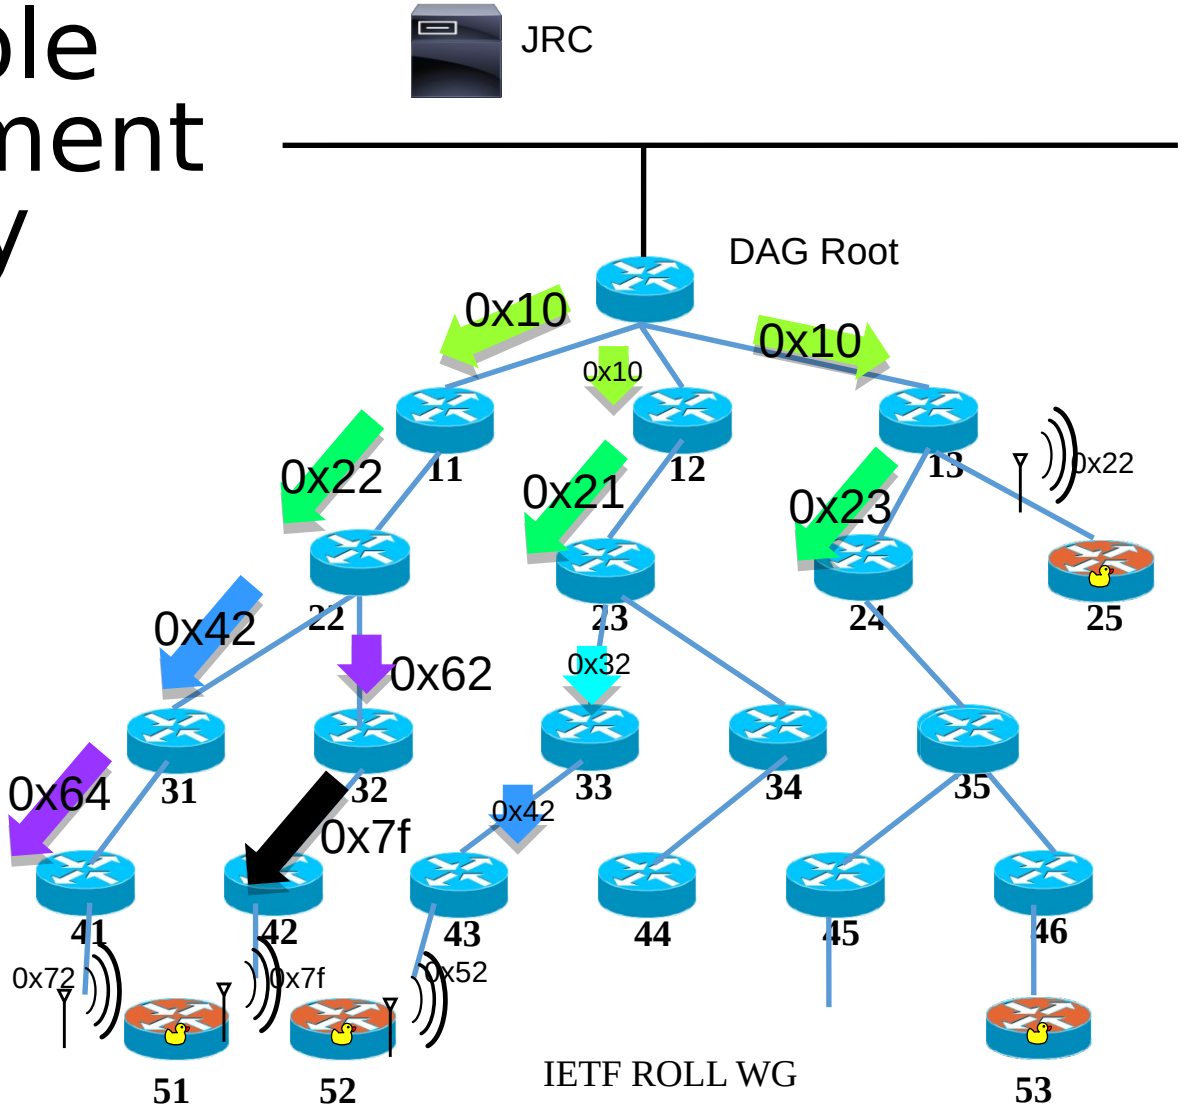

Enhanced Beacons

DIOs

base diagram from PThubert

With impaired node (24)

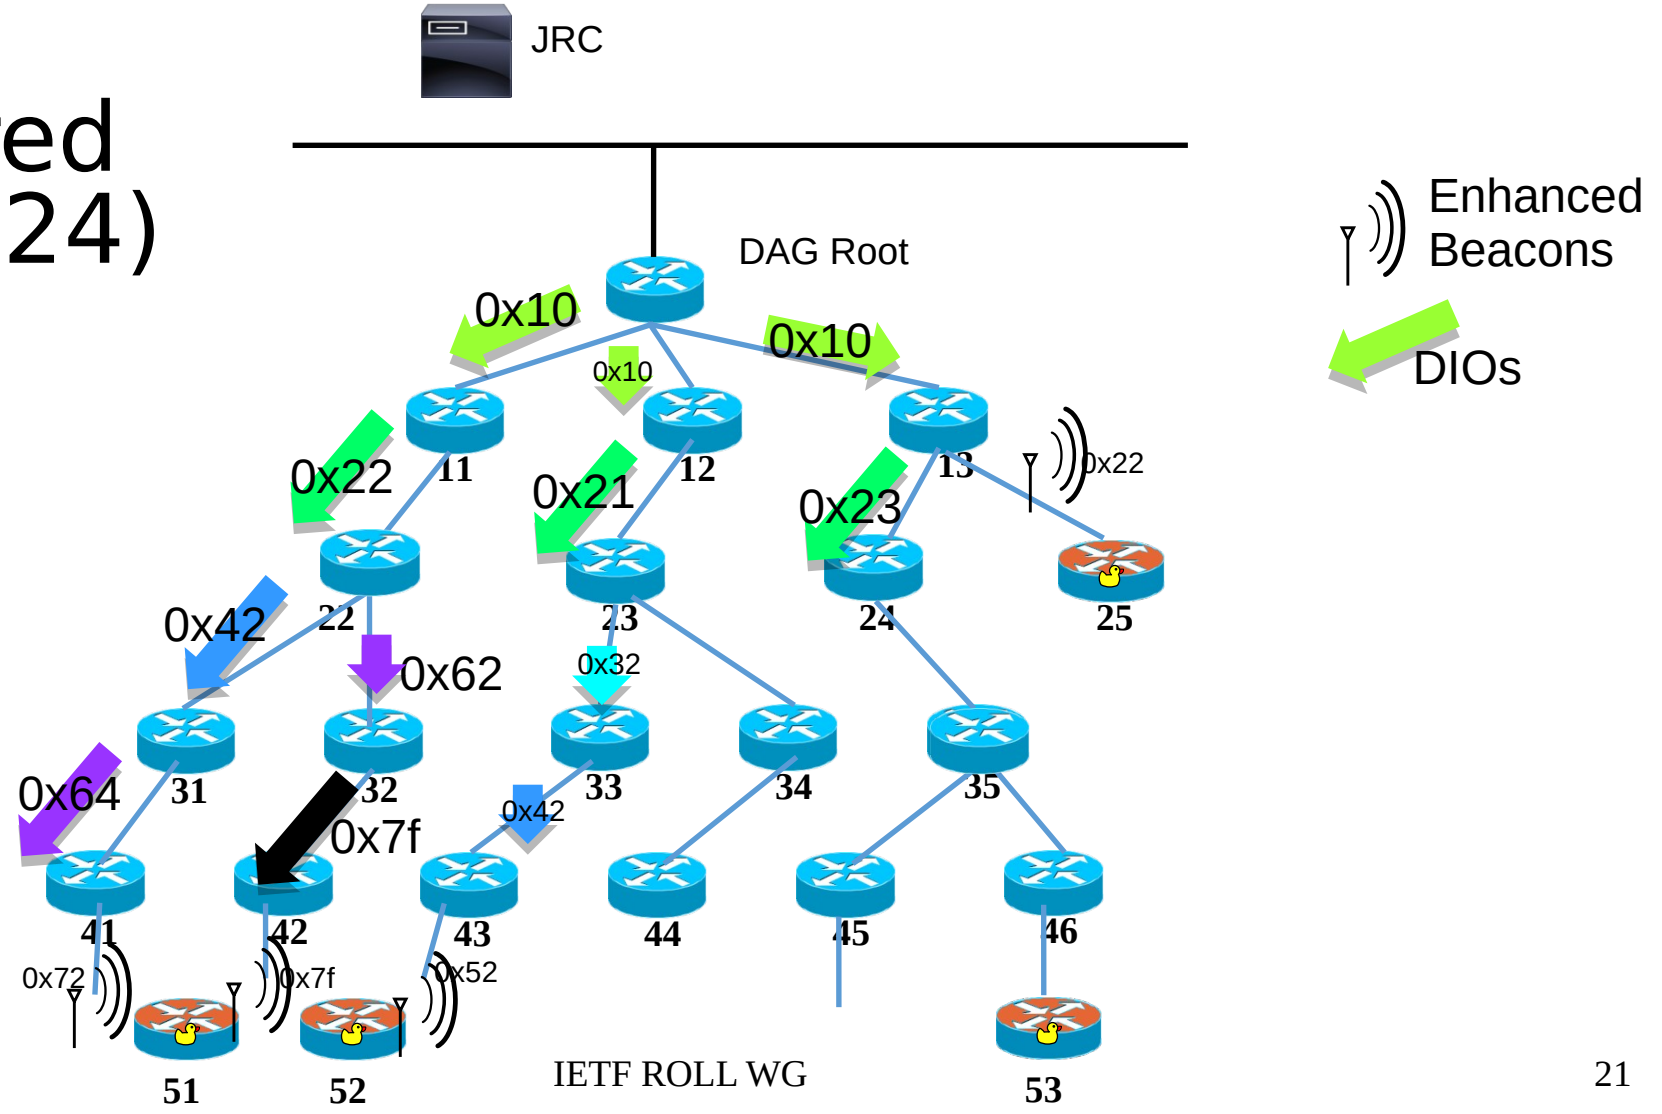

With impaired node (24)

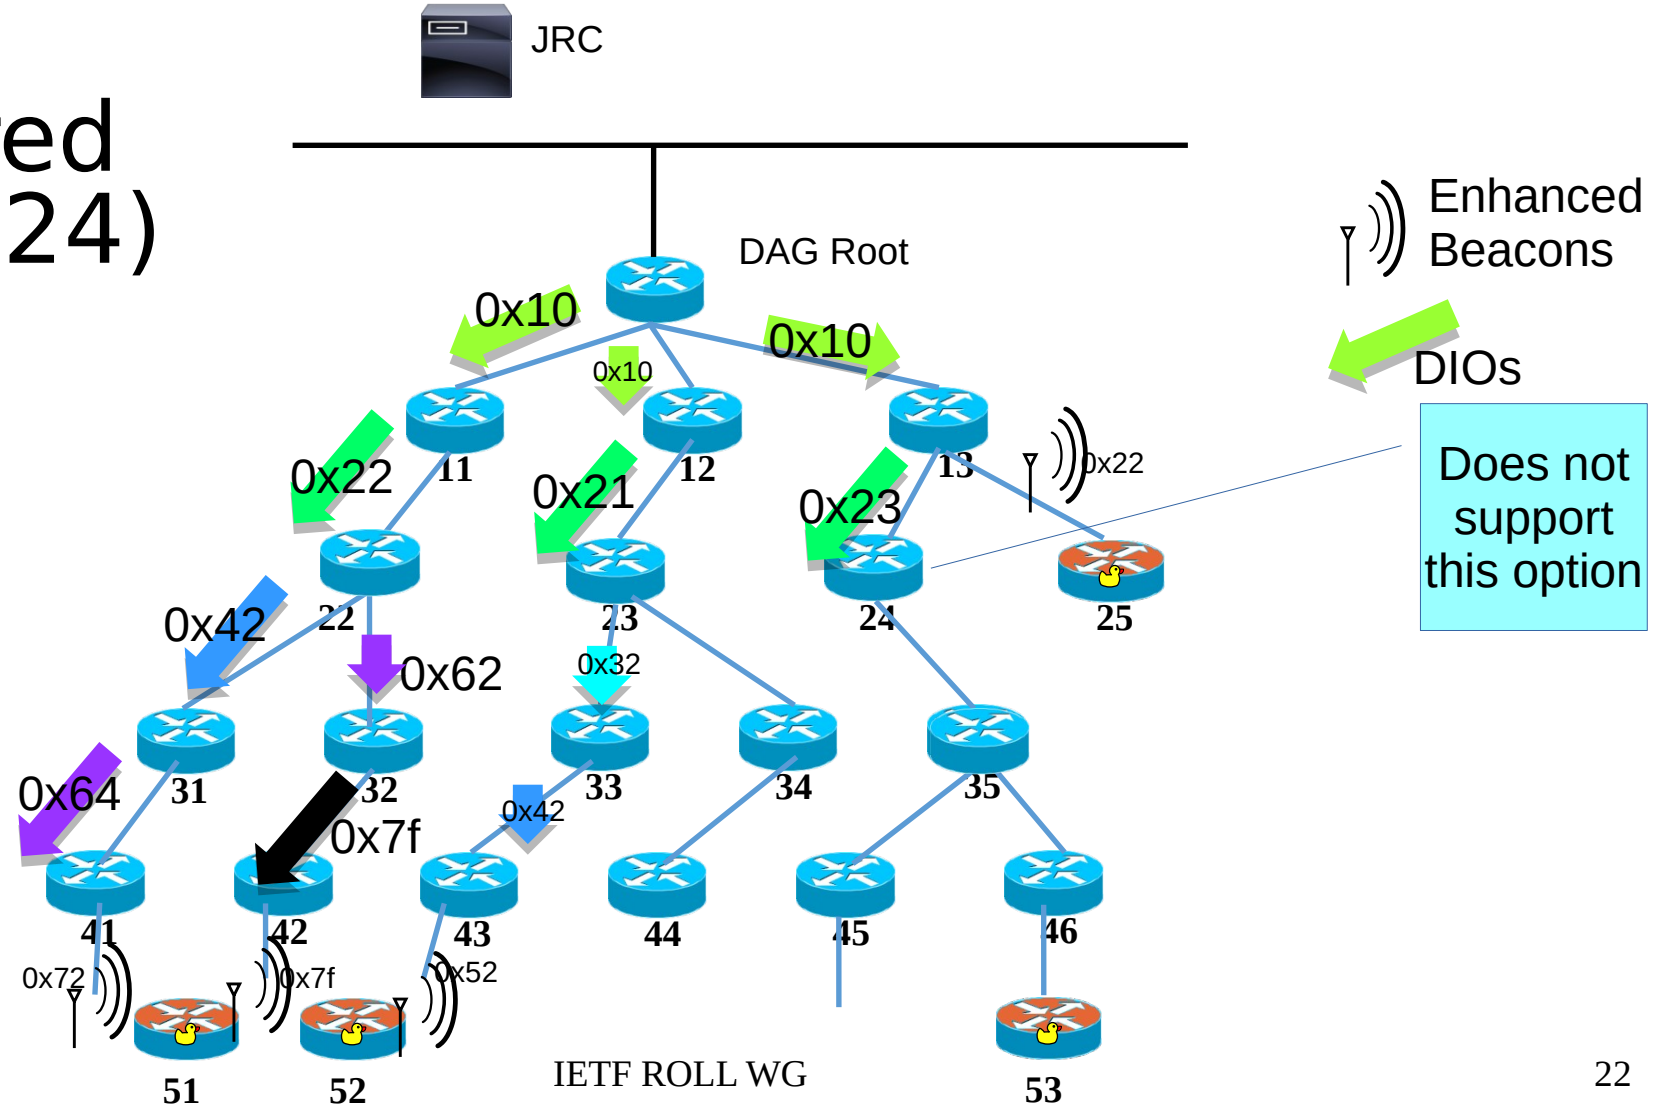

With impaired node (24)

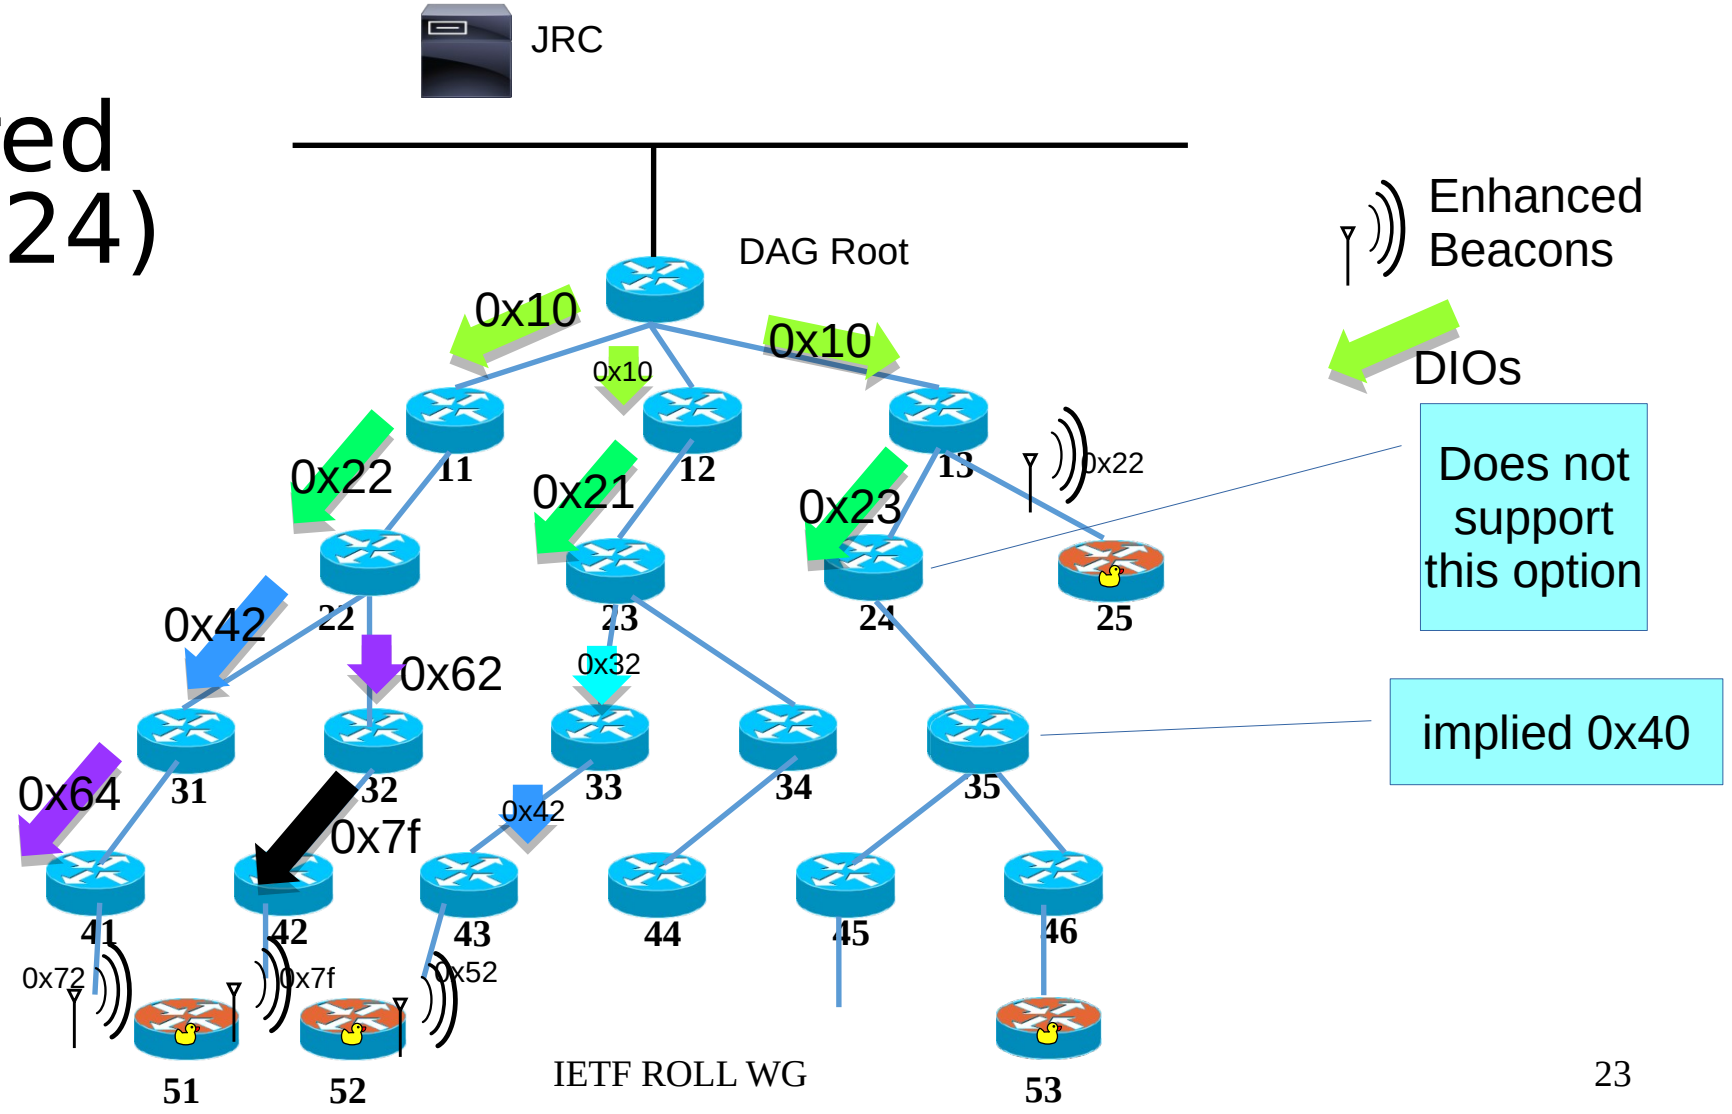

With impaired node (24)

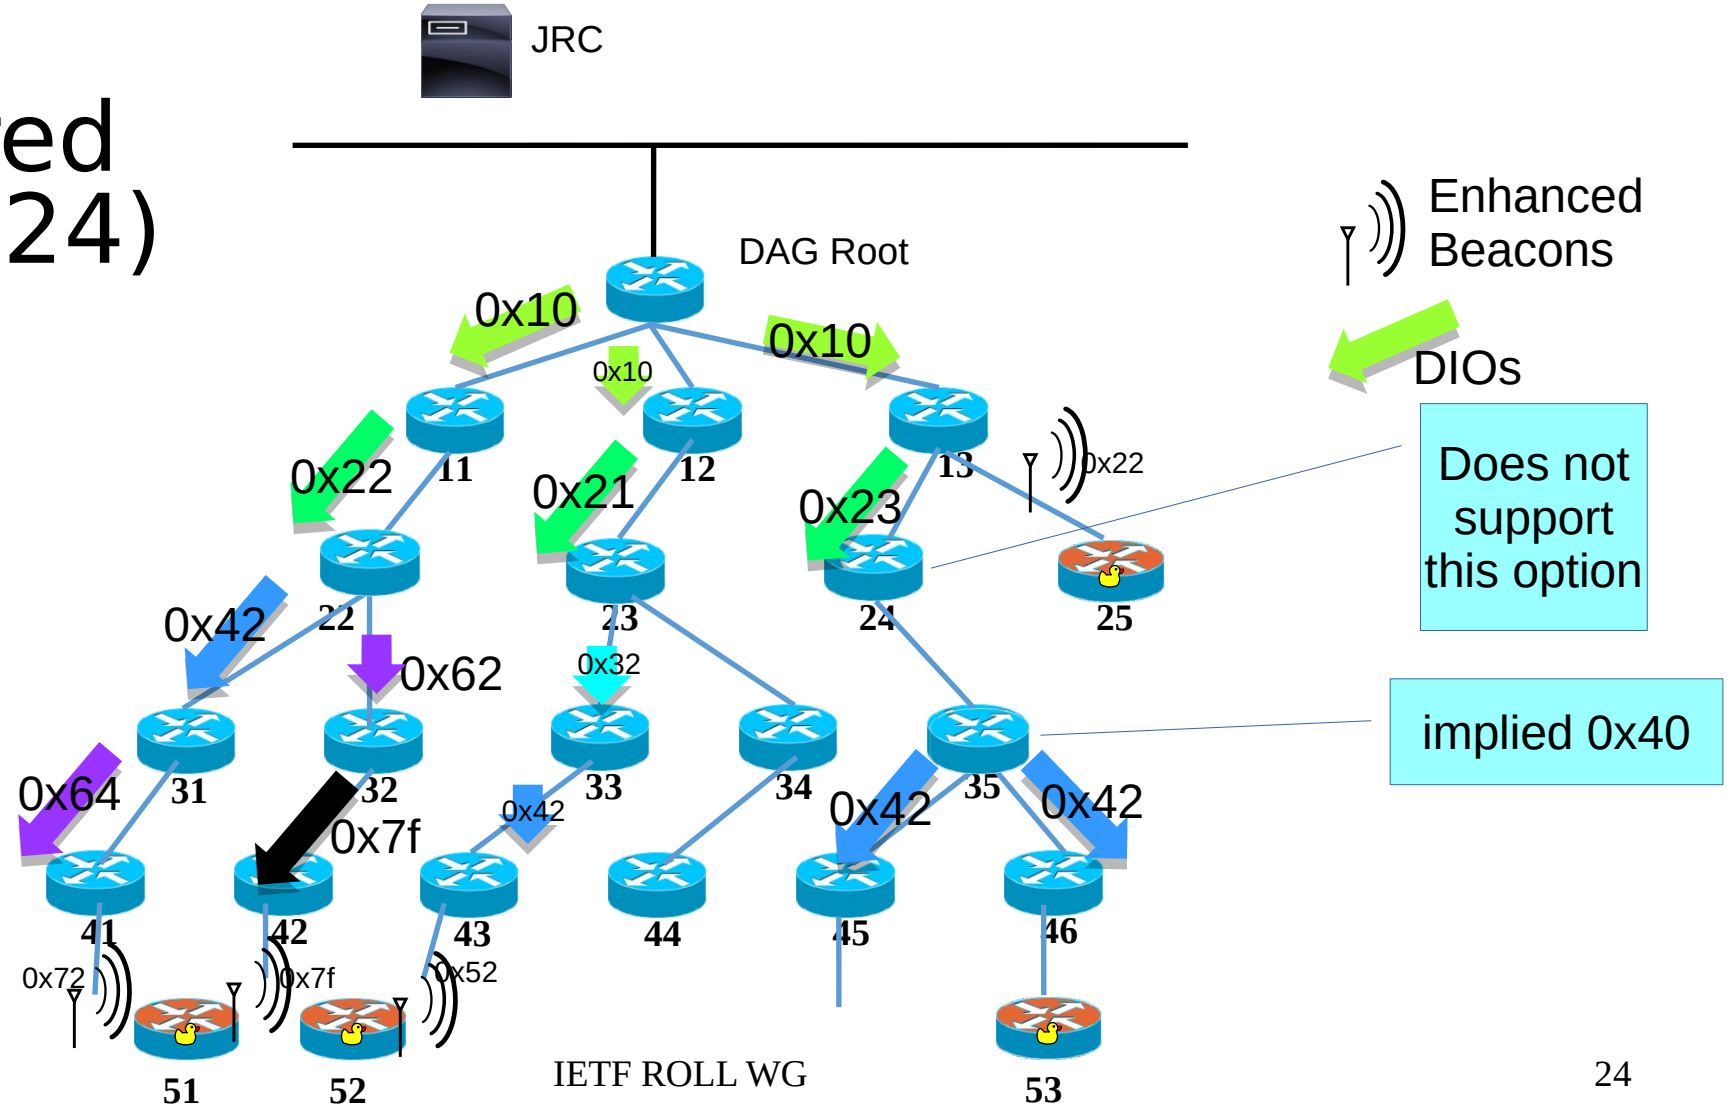

With impaired node (24)

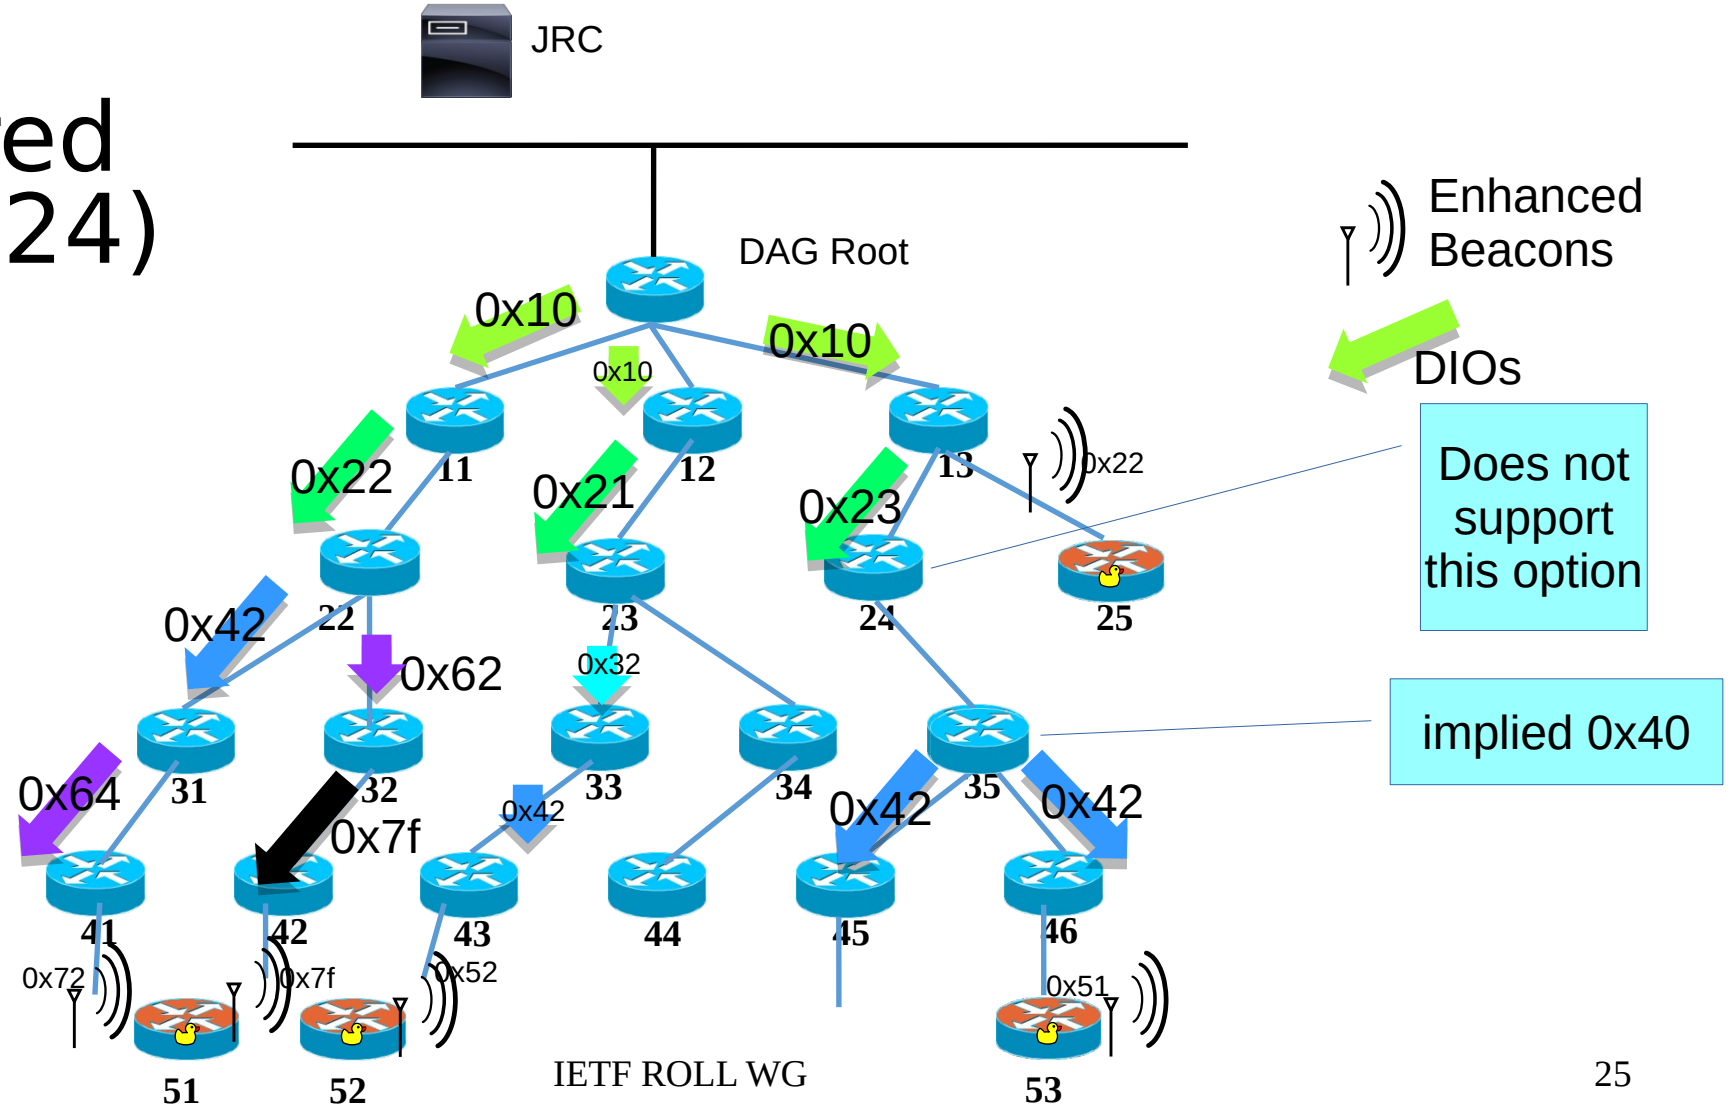

With impaired node (24)

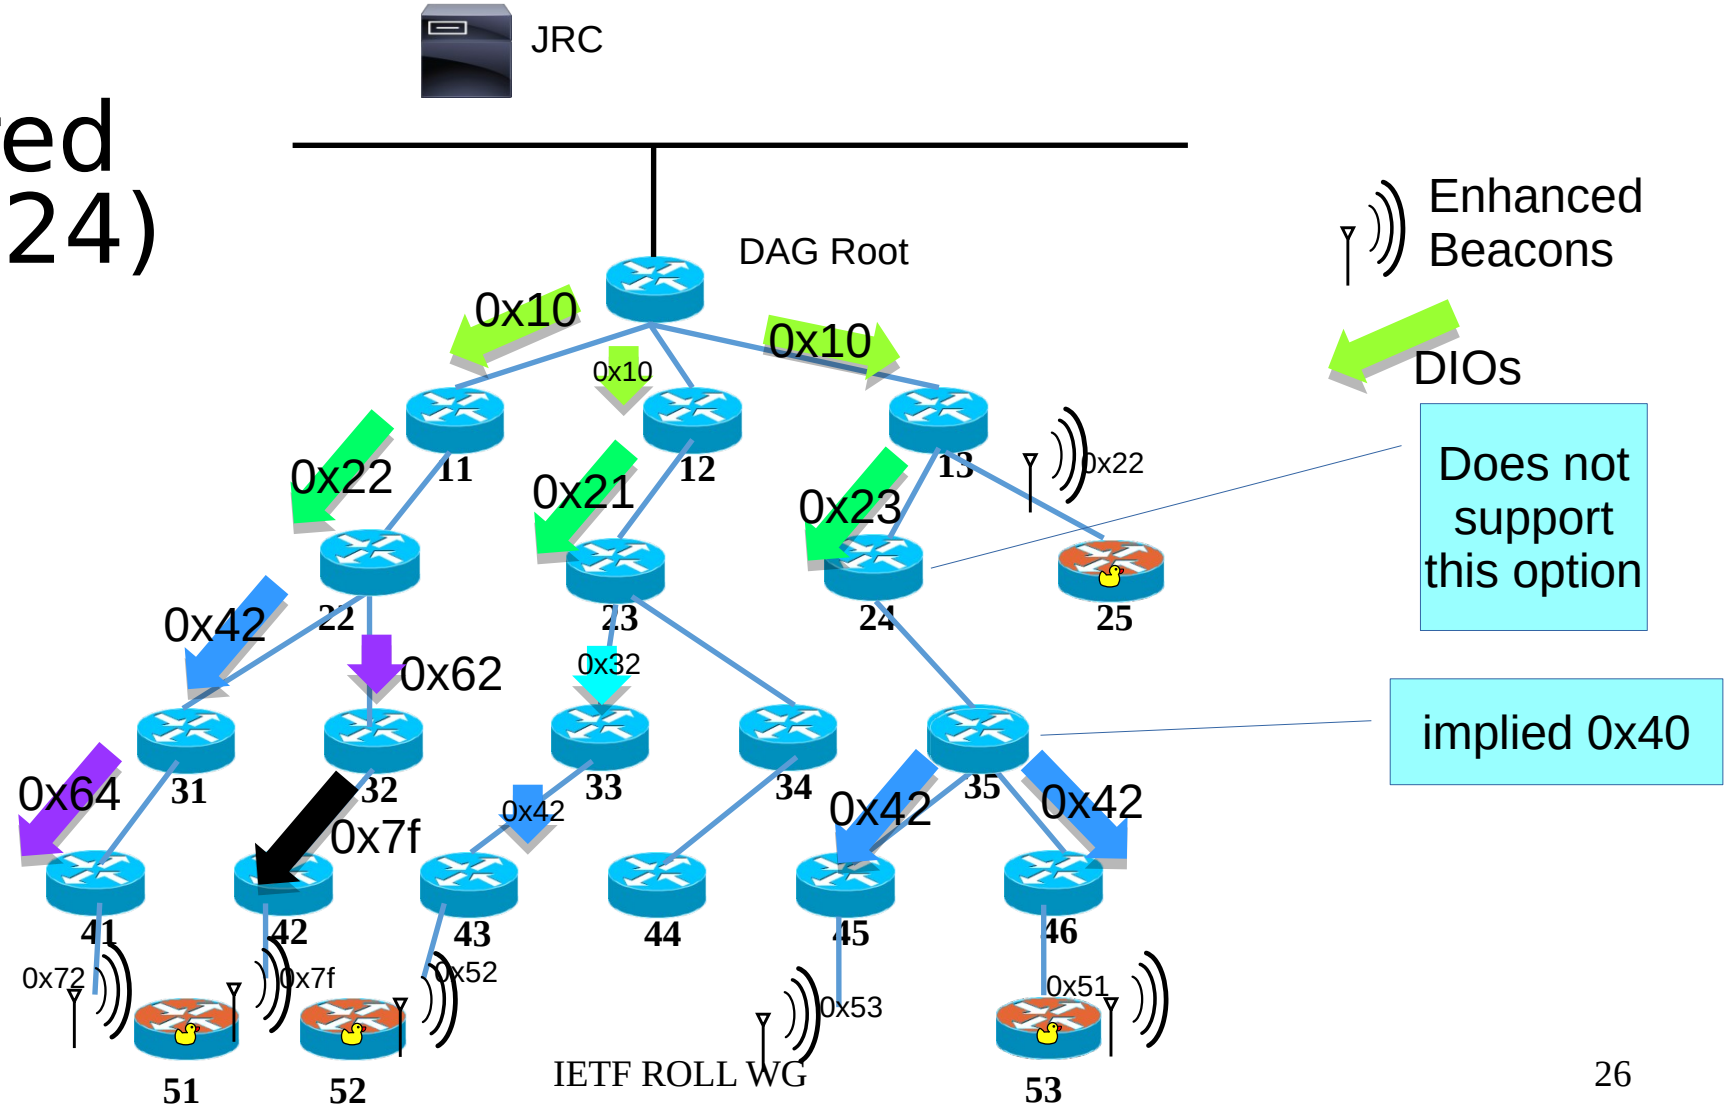

version -04

This document describes an RPL DIO option that can be used to announce a minimum enrollment priority. Each potential _Join Proxy_ would this value as a base on which to add values relating to local conditions. As explained in [I-D.ietf-6tisch-enrollment-enhanced-beacon], higher values decrease the likelyhood of an unenrolled node sending enrollment traffic via this path.

This document describes a RPL DIO option that can be used to announce a minimum enrollment priority. The minimum priority expresses the (lack of) willingness by the RPL DODAG globally to accept new joins. It may derive from multiple constaining factors, e.g., the size of the DODAG, the occupancy of the bandwidth at the Root, the memory capacity at the DODAG Root, or an administrative decision.

Each potential _Join Proxy_ would this value as a base on which to add values relating to local conditions such as its Rank and number of pending joins, which would degrade even further the willingness to take more joins.

When a RPL domain is composed of multiple DODAGs, nodes at the edge of 2 DODAGs may not only join either DODAG but also move from one to the other in order to keep their relative sizes balanced. For this, the approximate knowledge of size of the DODAG is an essential metric. Depending on the network policy, the size of the DODAG may or may not affect the minimum enrollment priority. It would be limiting its value to enforce that one is proportional to the other. This is why the current size of the DODAG is advertised separately in the new option.

As explained in [I-D.ietf-6tisch-enrollment-enhanced-beacon], higher values decrease the likelyhood of an unenrolled node sending enrollment traffic via this path.

RNFD: Fast border router crash detection in RPL

draft-iwanicki-roll-rnfd-00

Konrad Iwanicki (iwanicki@mimuw.edu.pl)

ROLL Interim, August 31st, 2021

Why consider LBR crashes?

An LBR:

- plays a central role in an LLN (DODAG root),
- is typically more involved than a constrained node,
- usually requires a tethered power supply (hard to back up in many deployments).

Failures in RPL

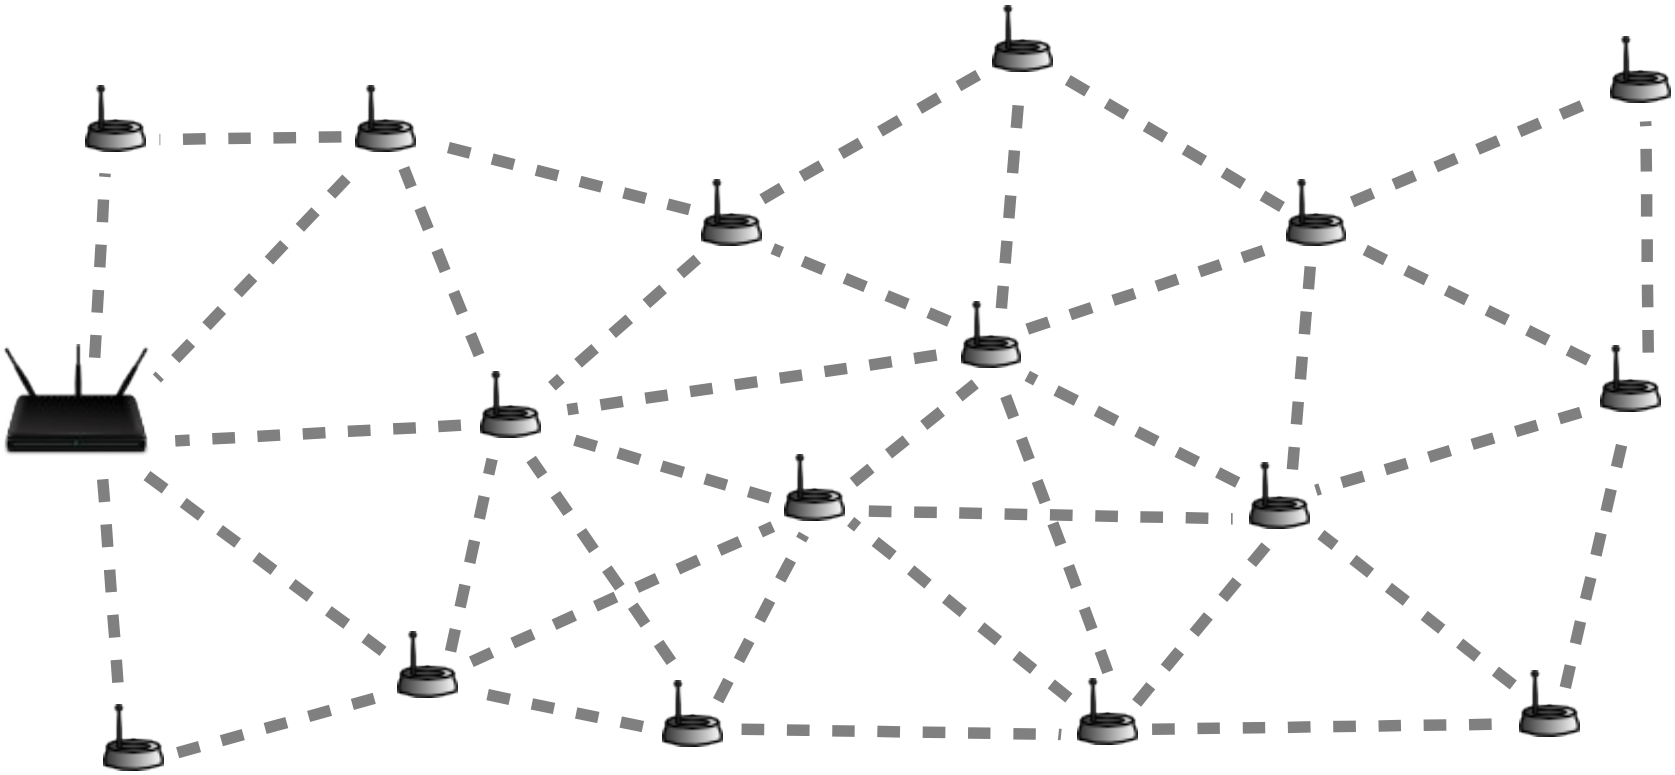

Failures in RPL

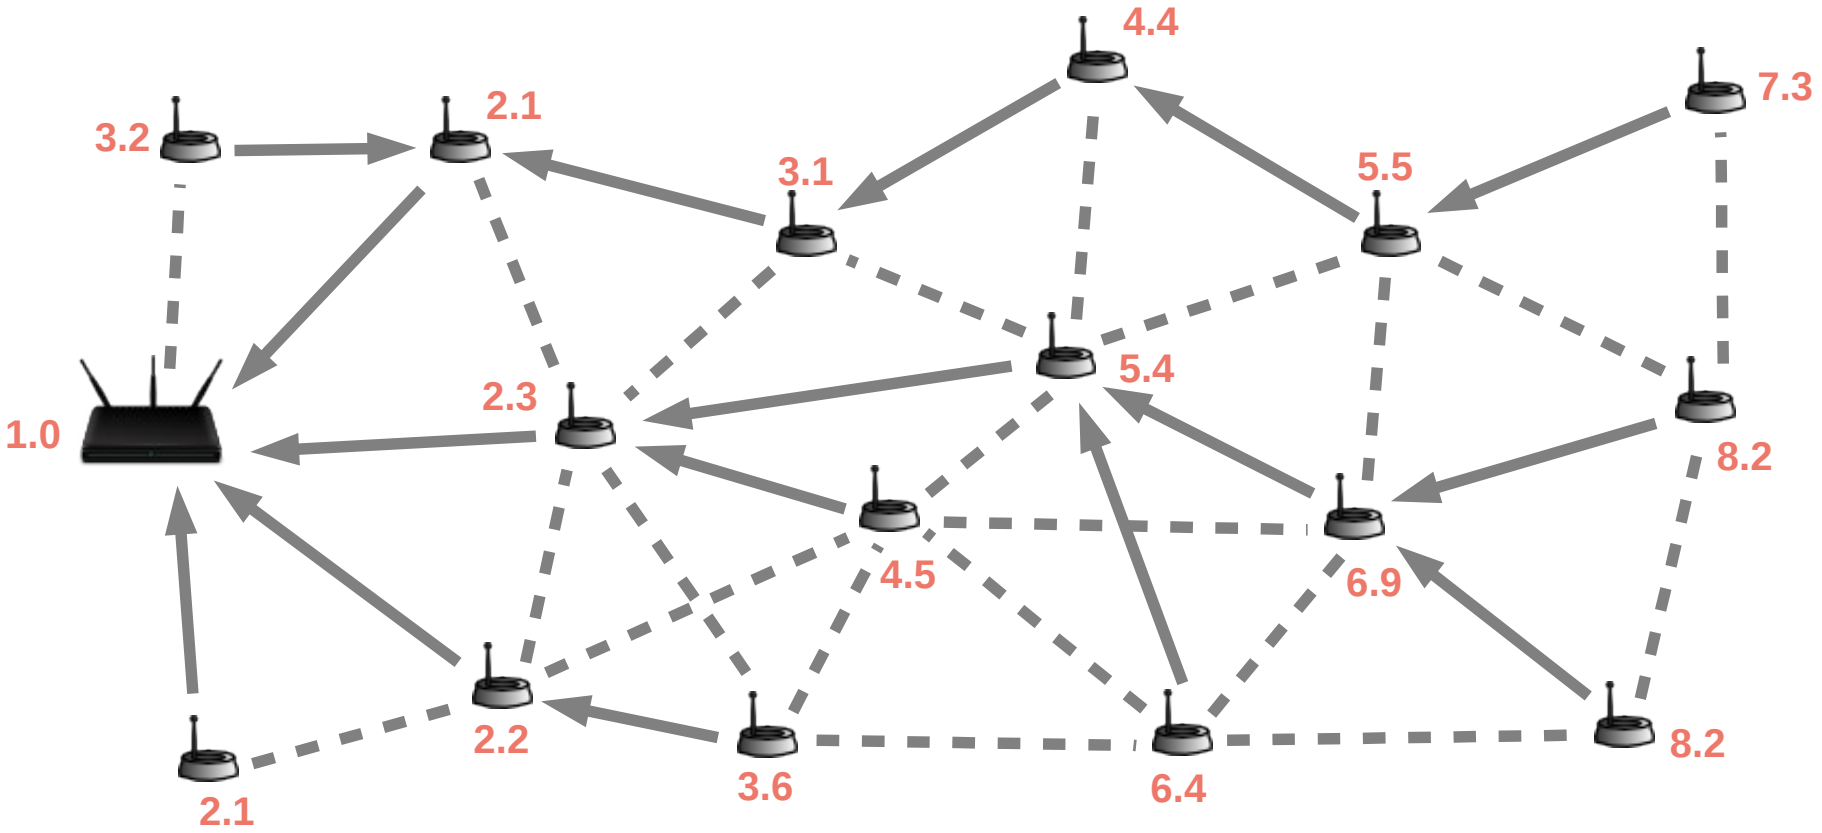

Failures in RPL

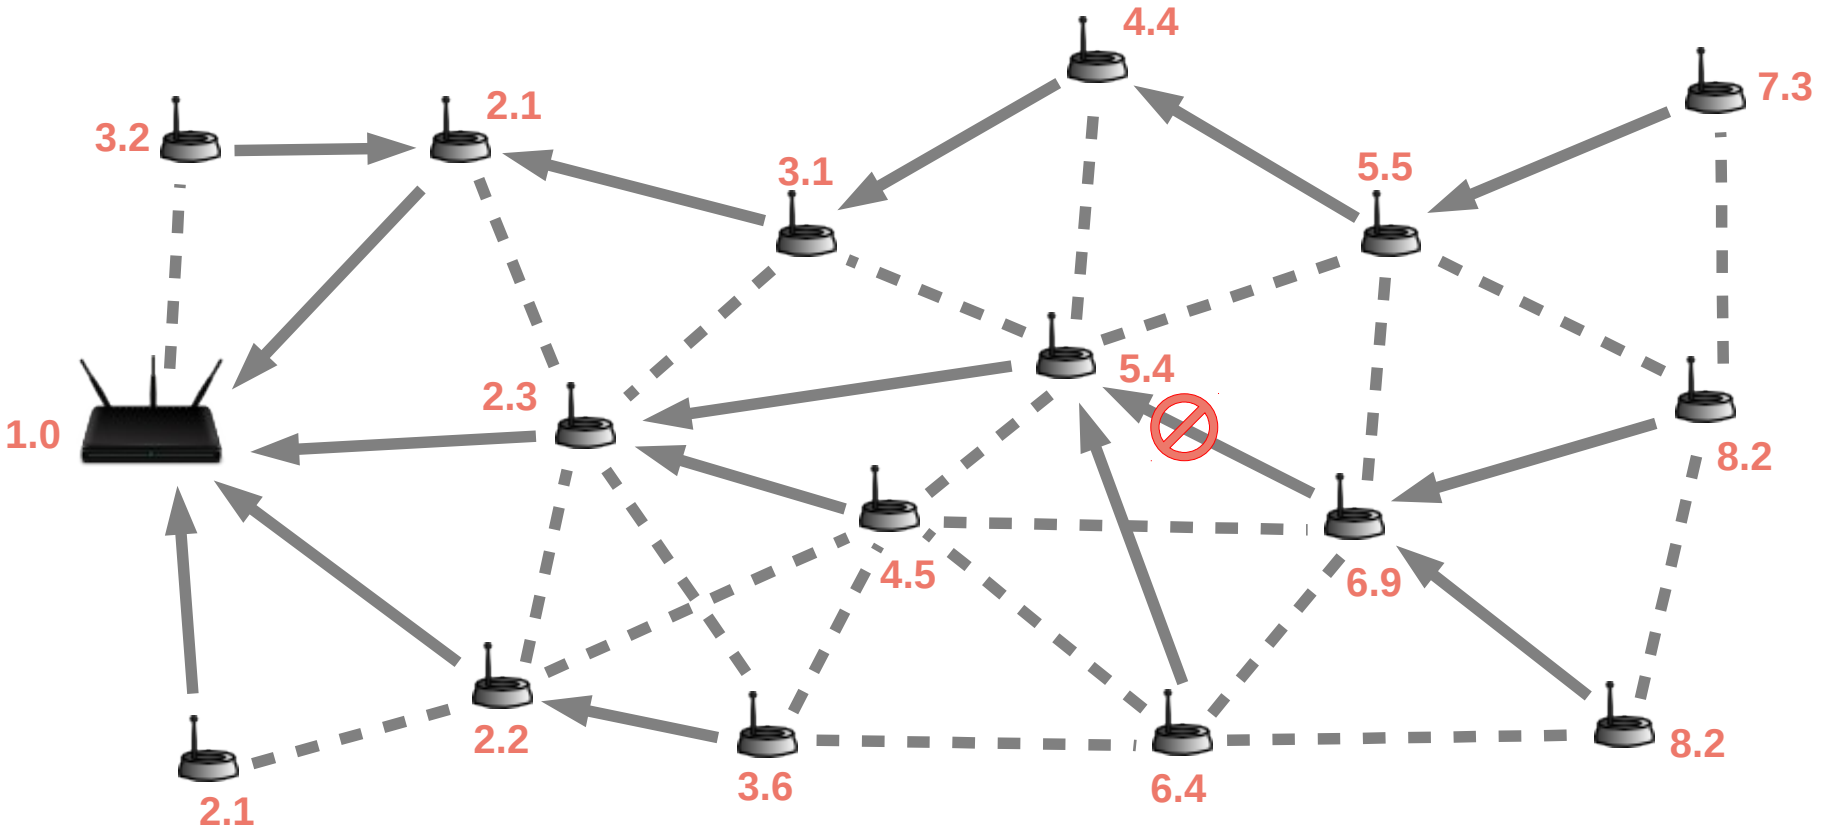

Failures in RPL

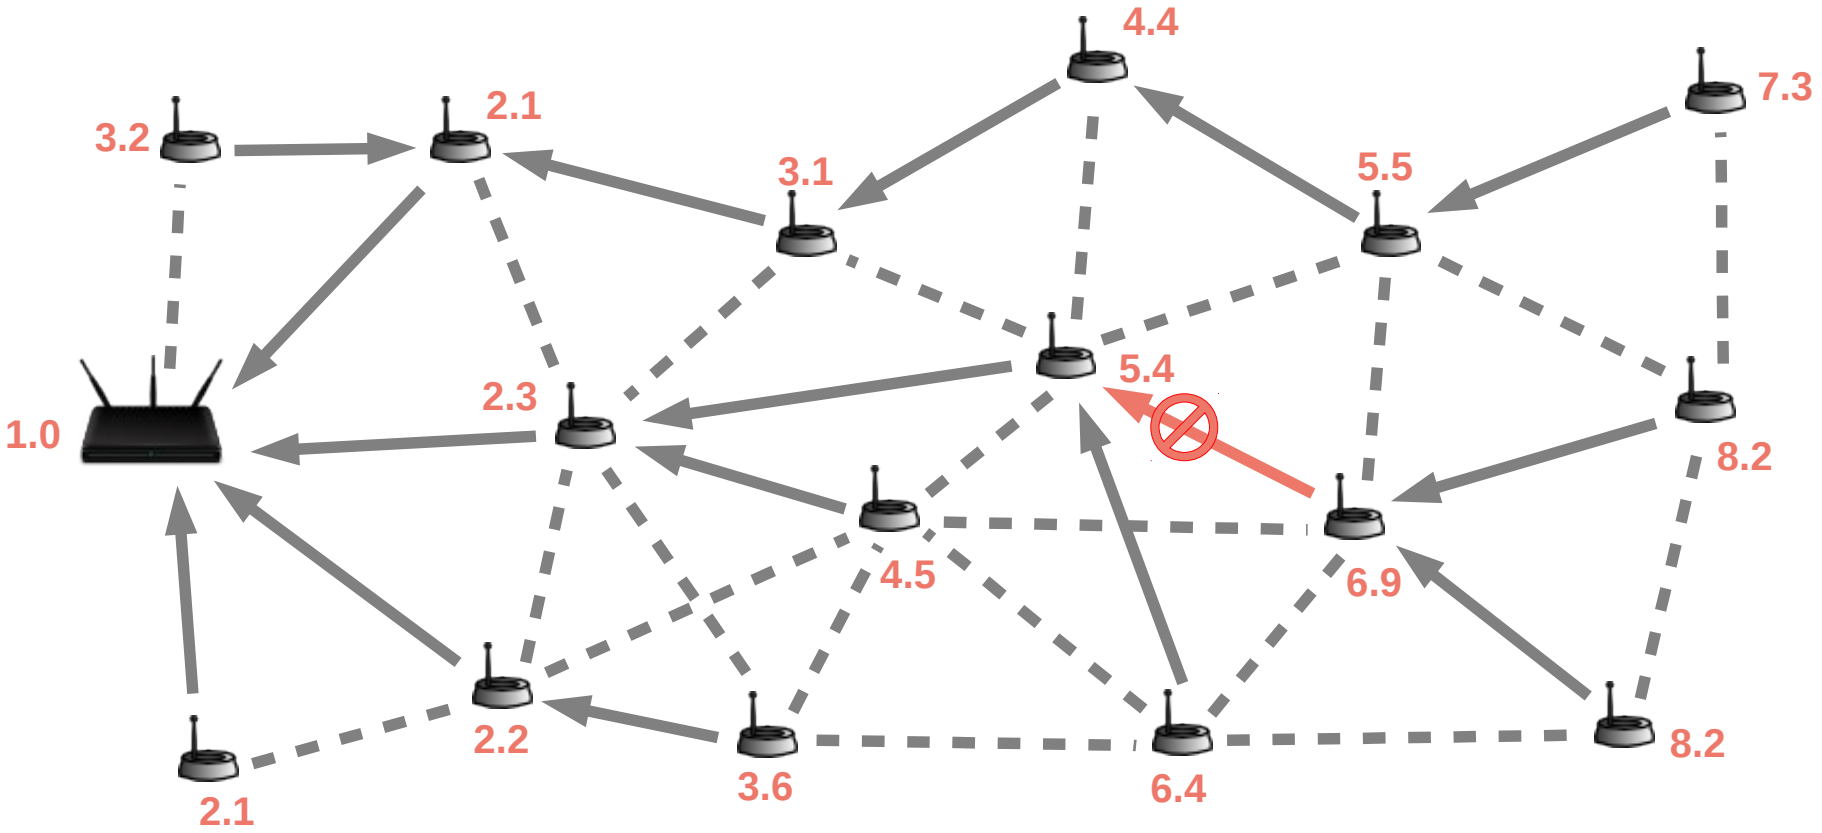

Failures in RPL

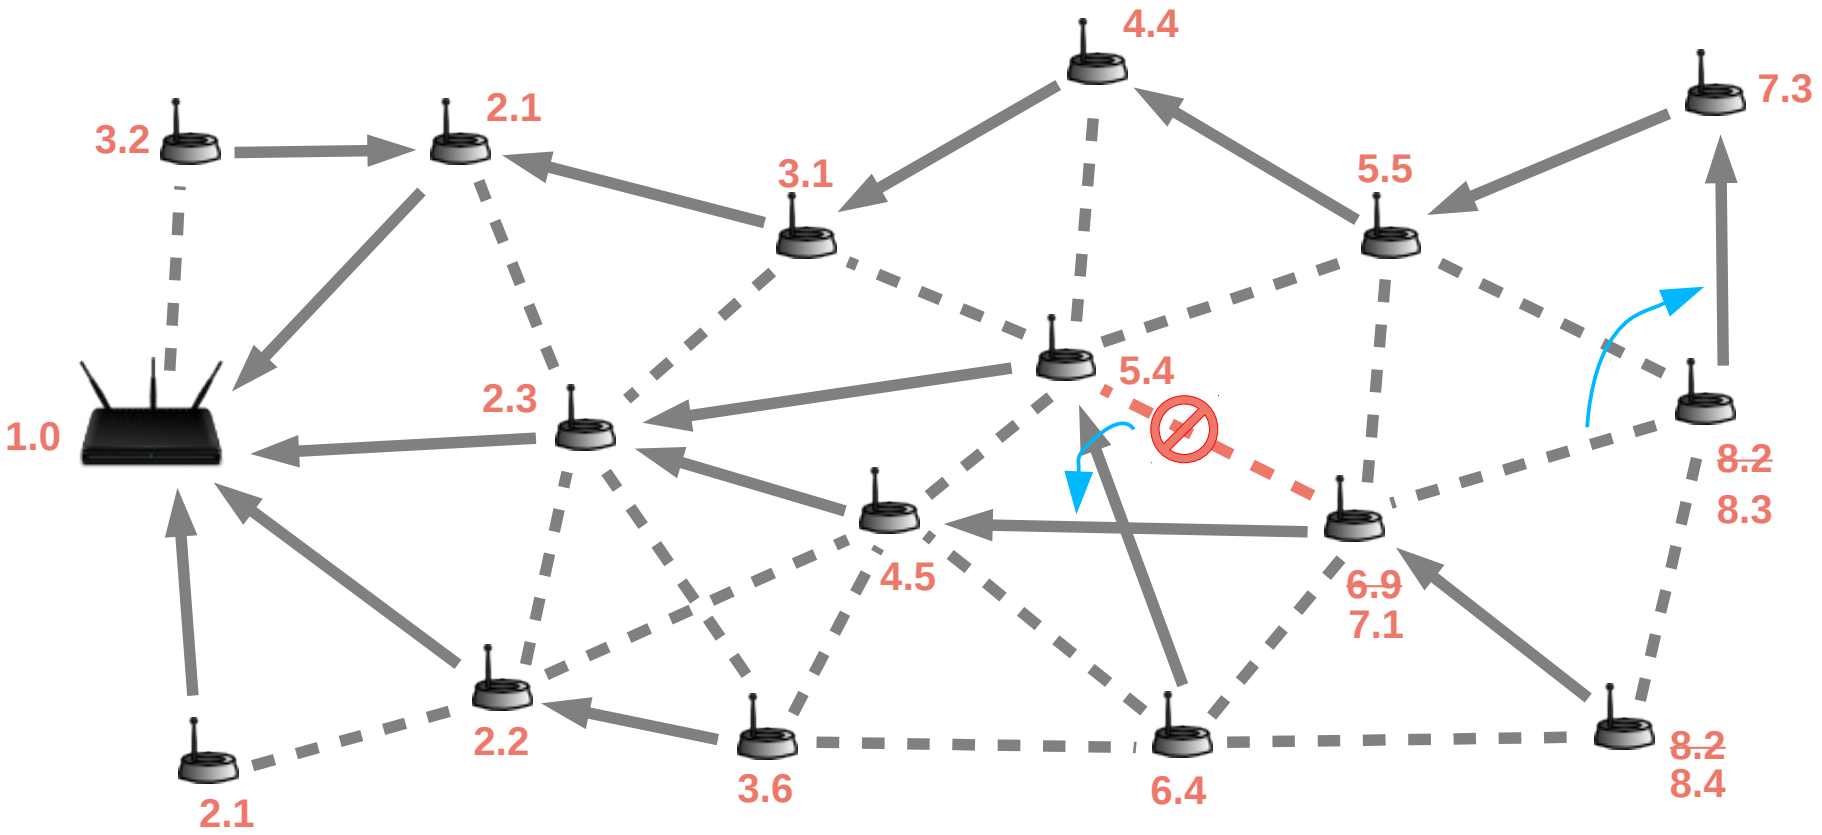

Failures in RPL

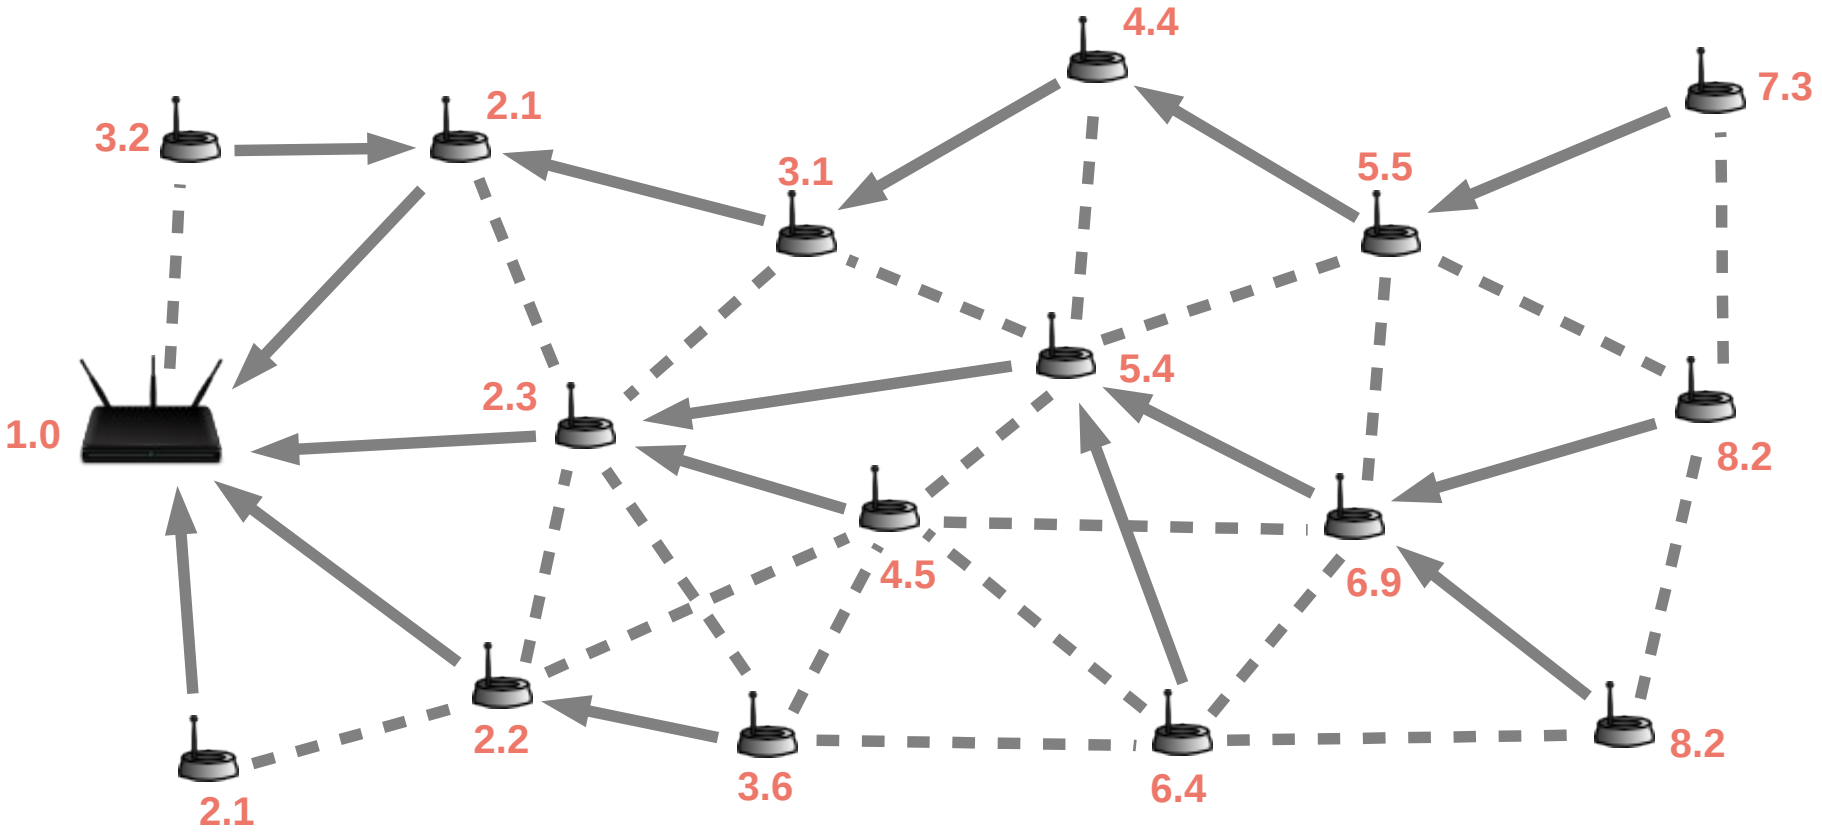

Failures in RPL

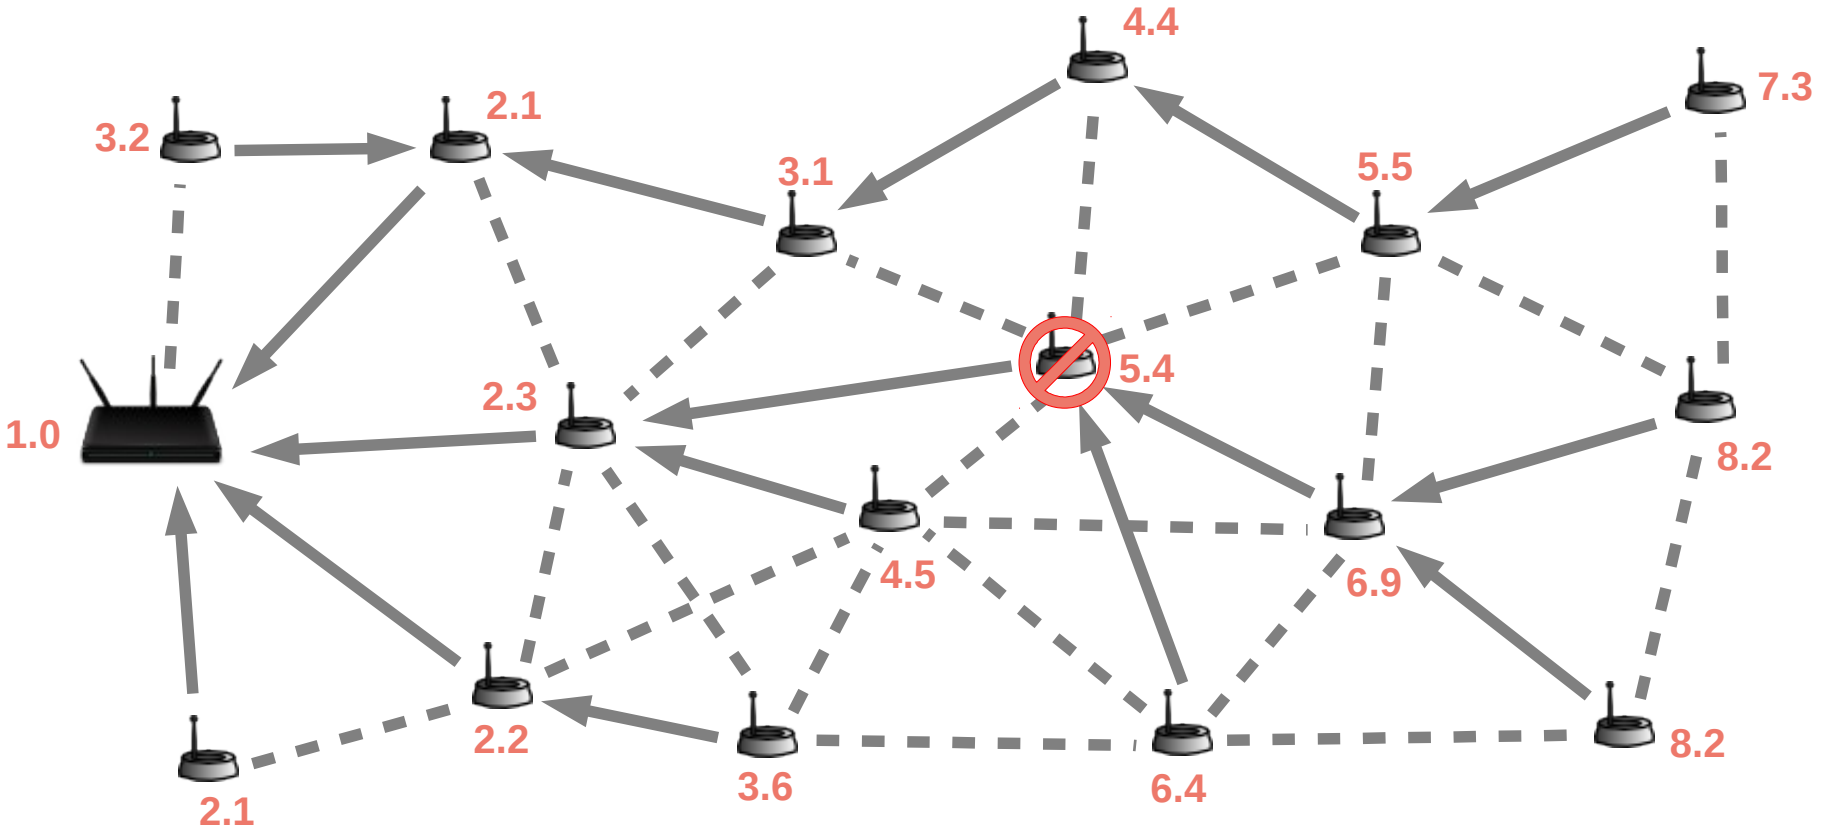

Failures in RPL

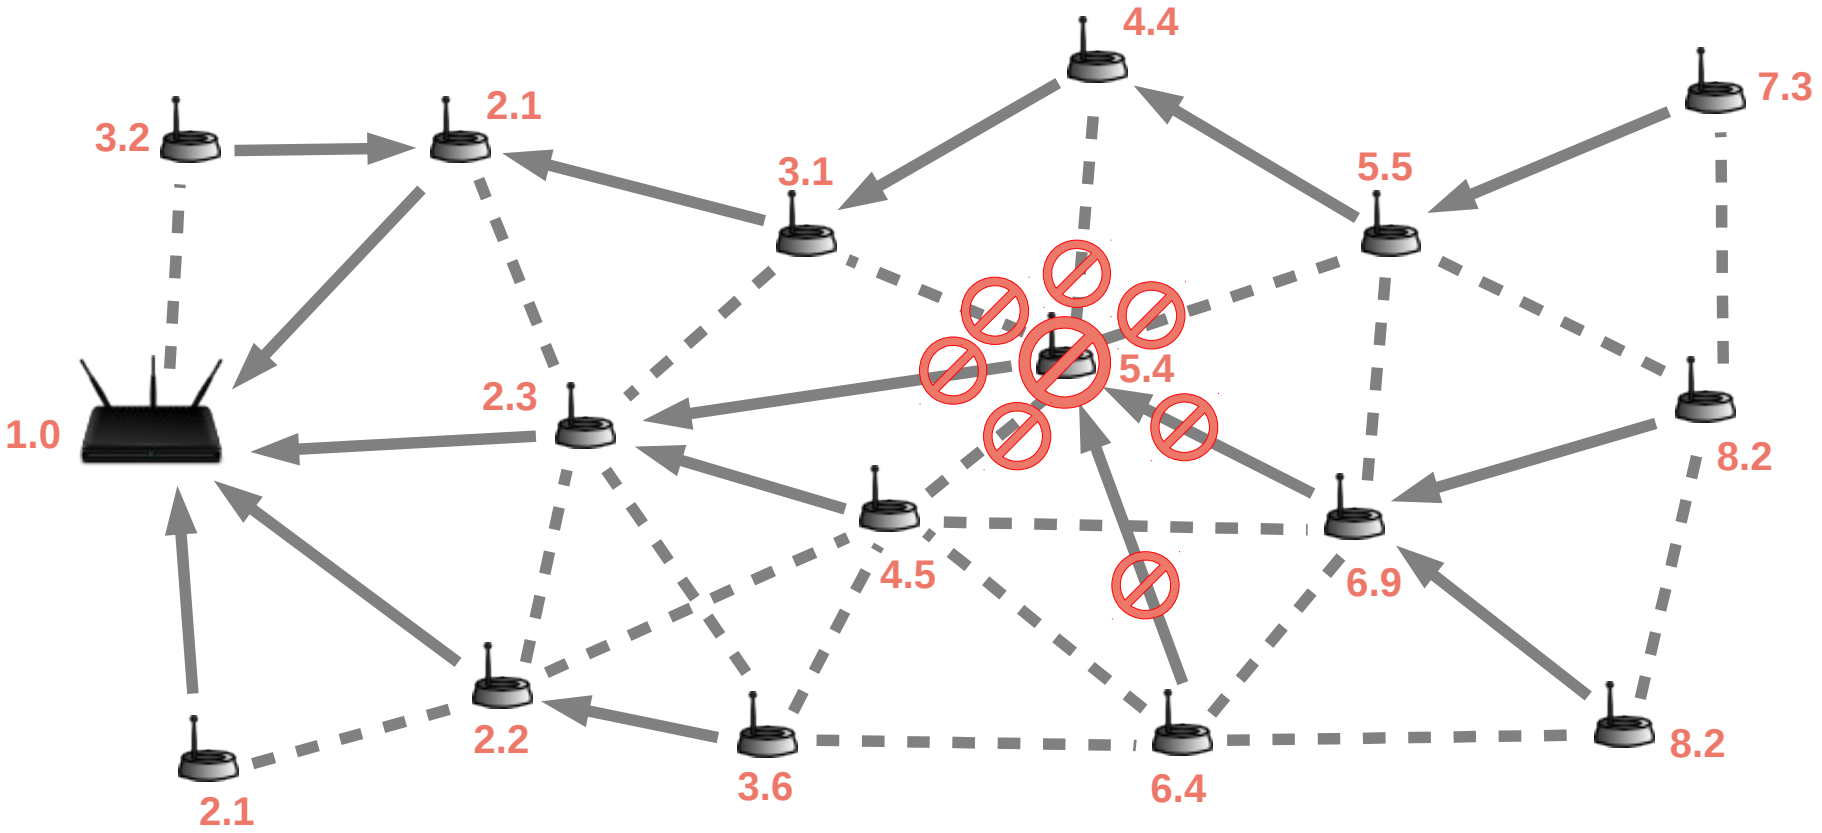

Failures in RPL

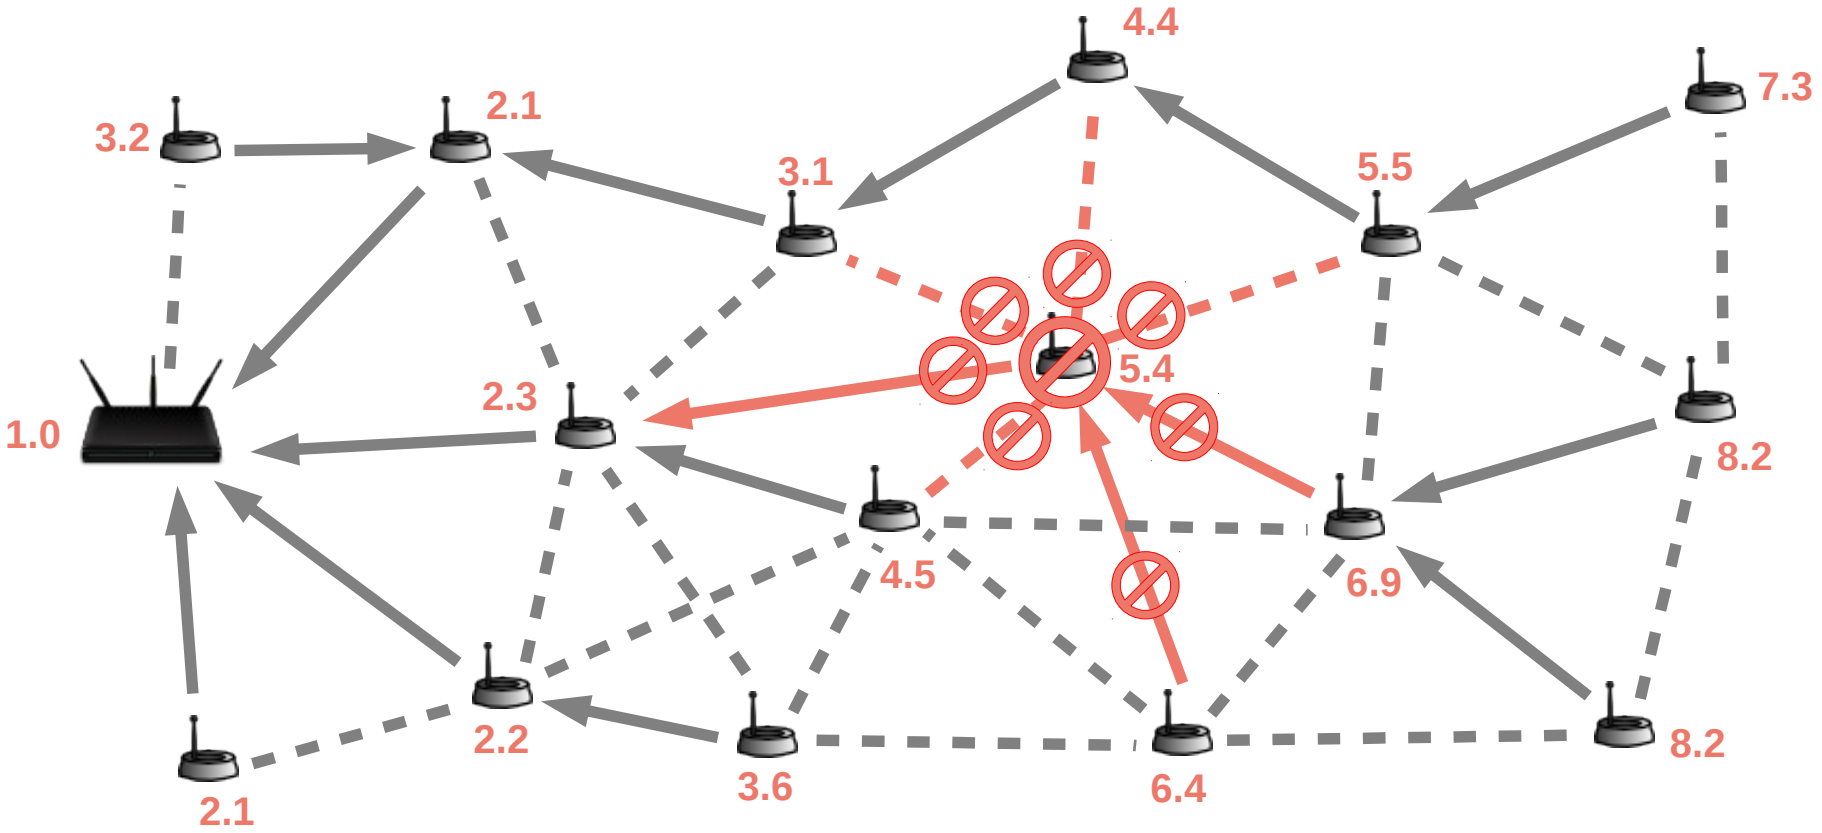

Failures in RPL

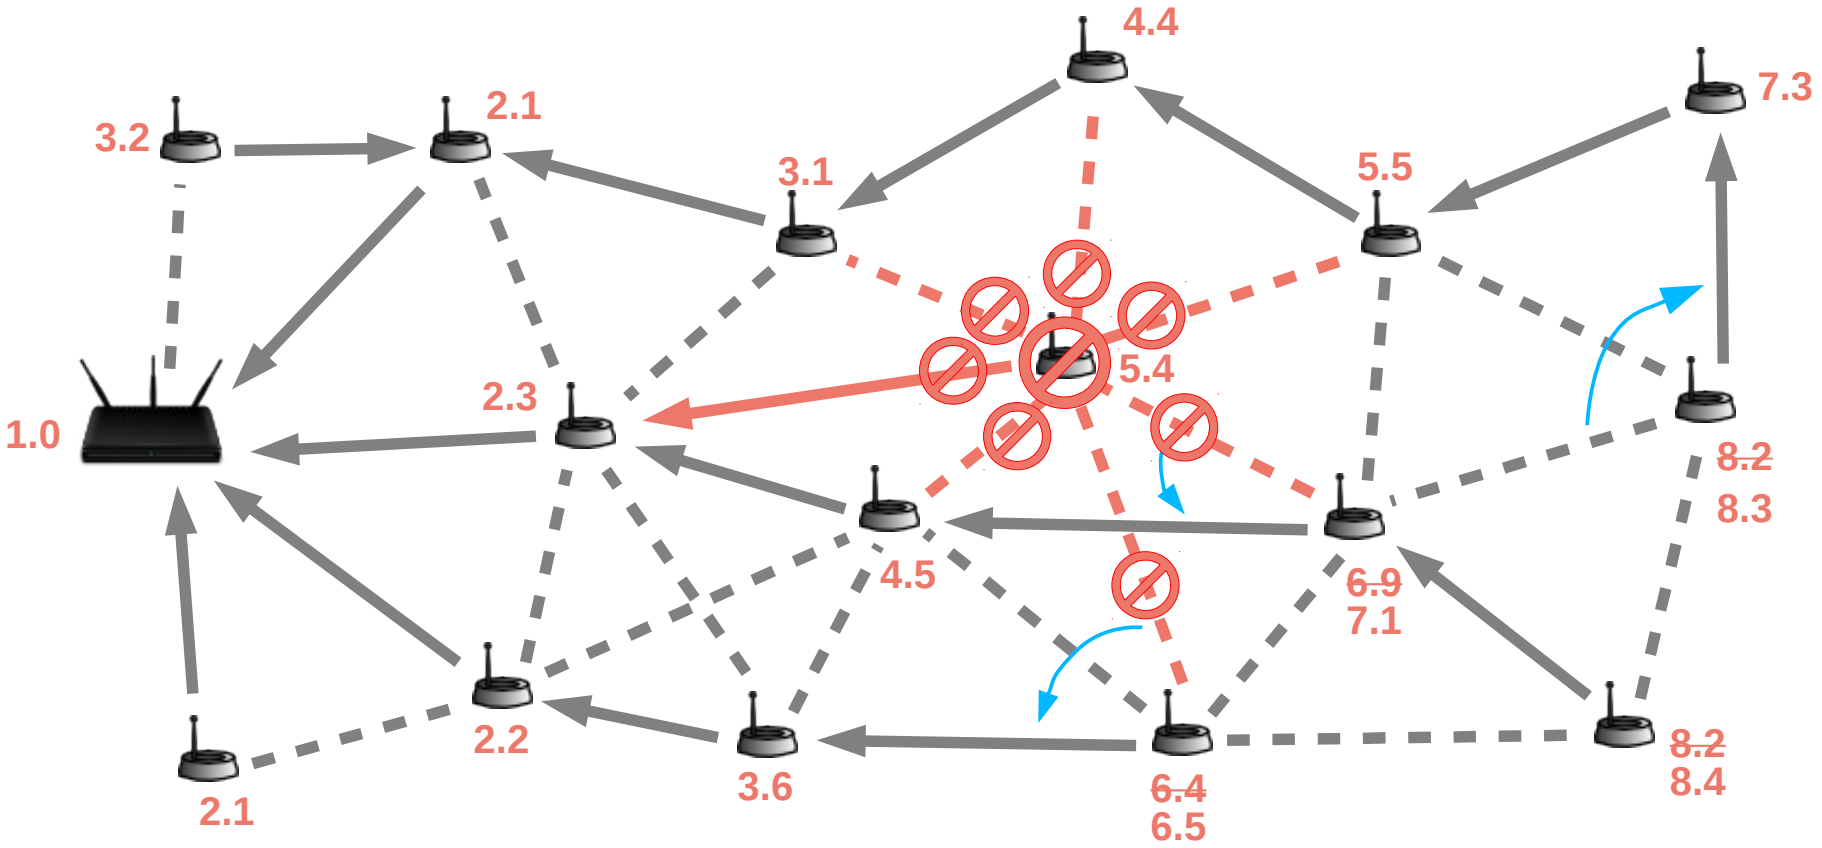

Failures in RPL

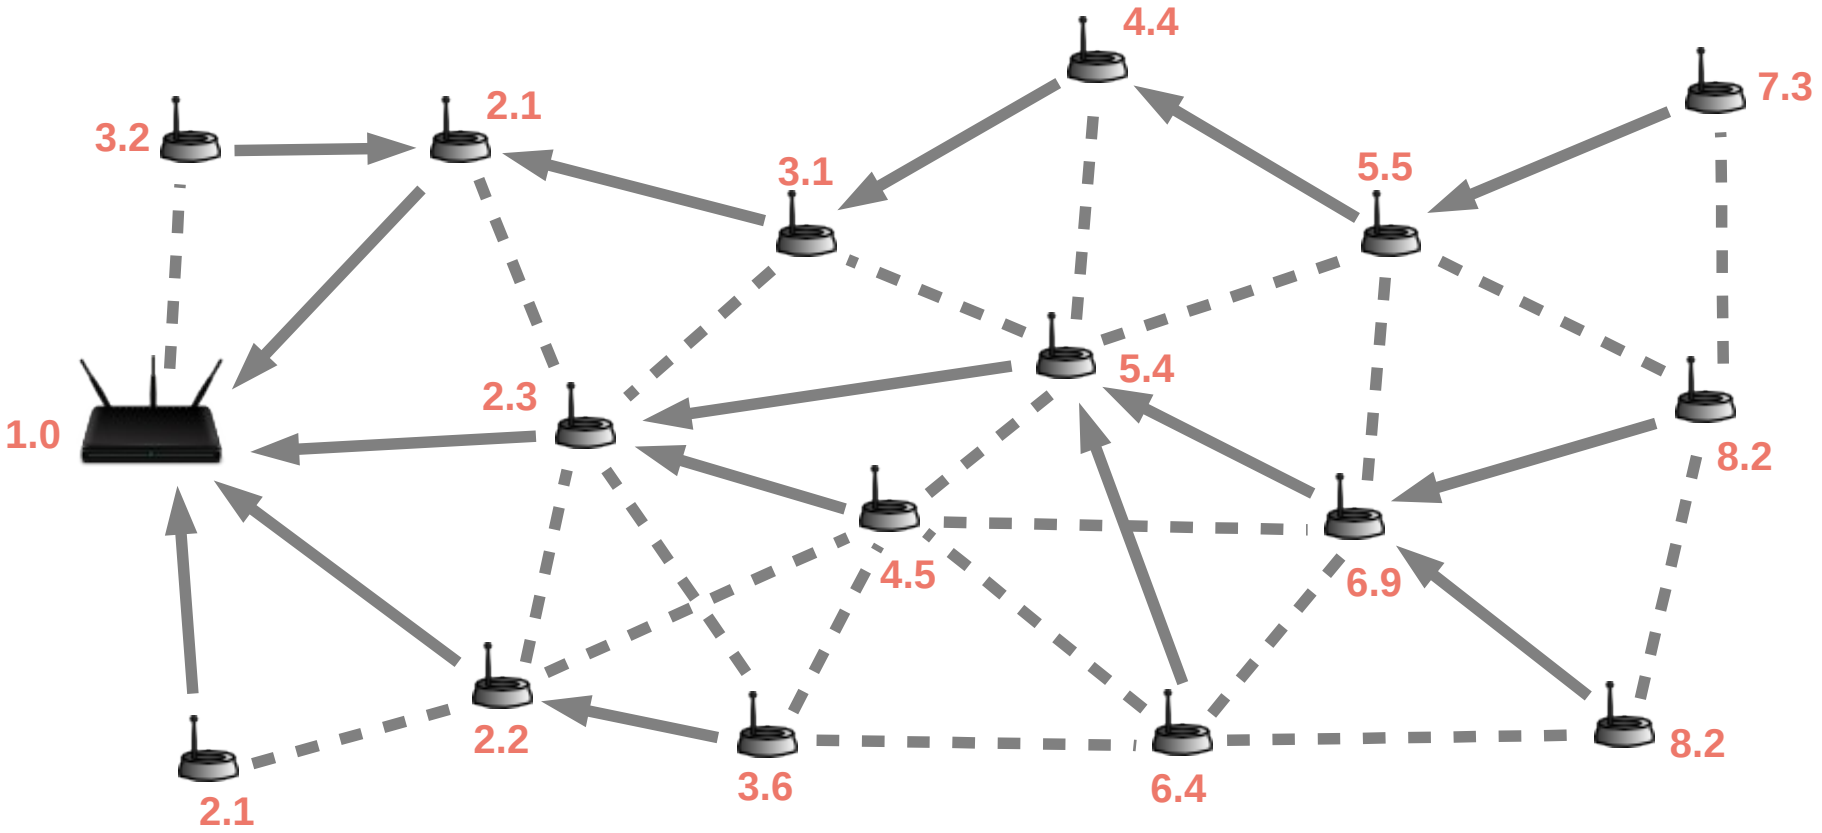

Failures in RPL

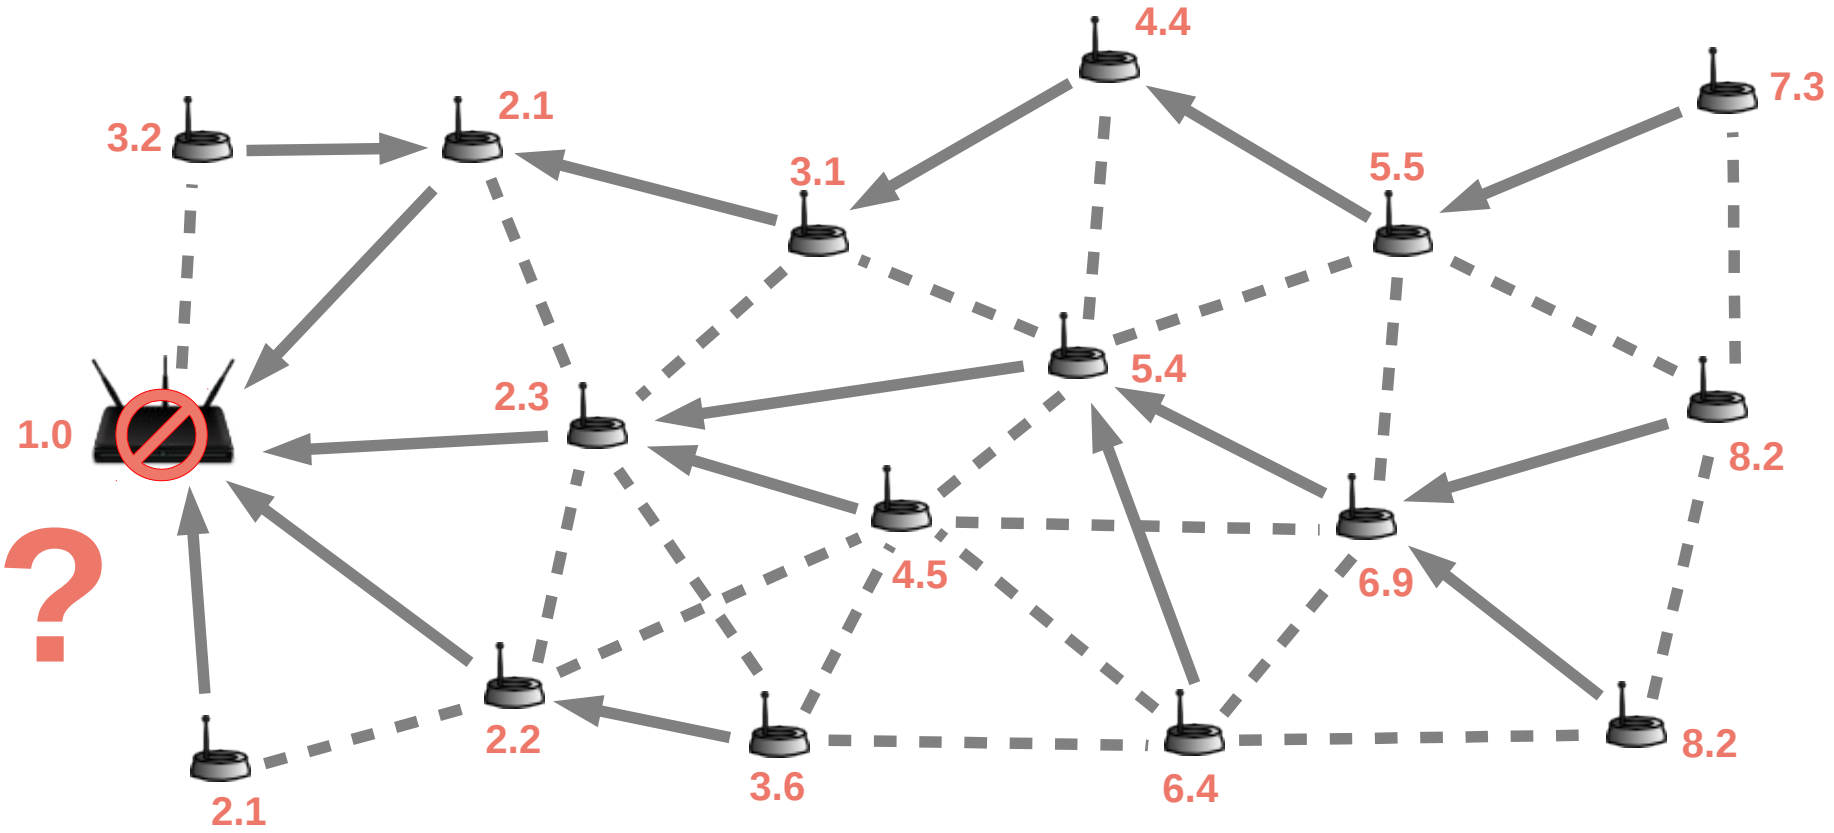

Failures in RPL

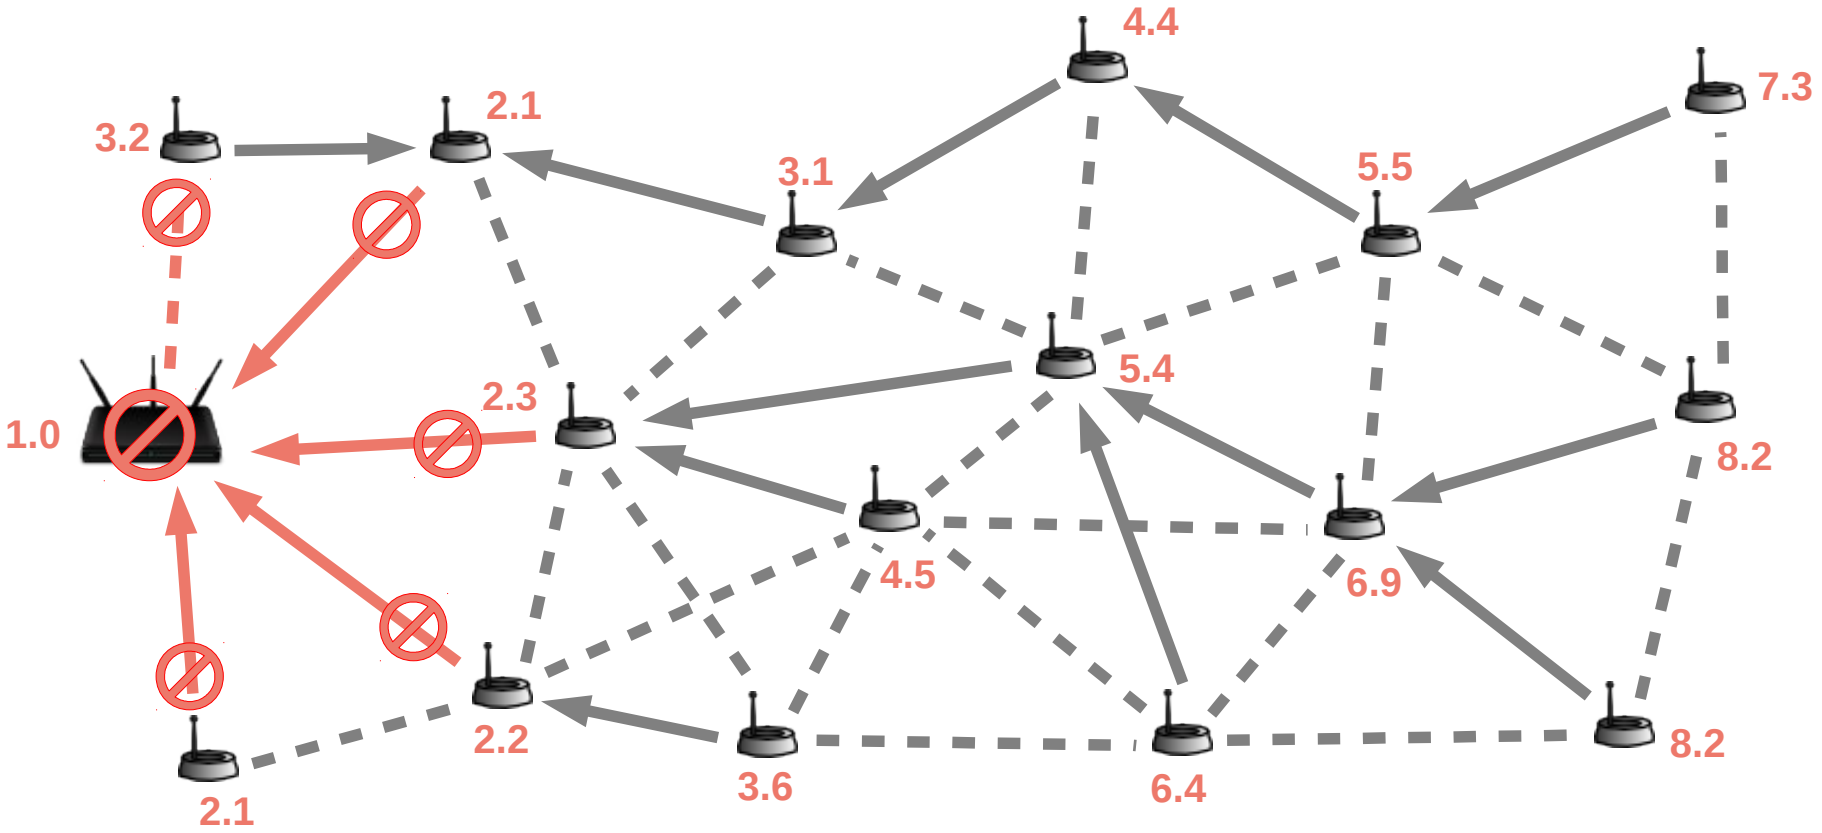

Failures in RPL

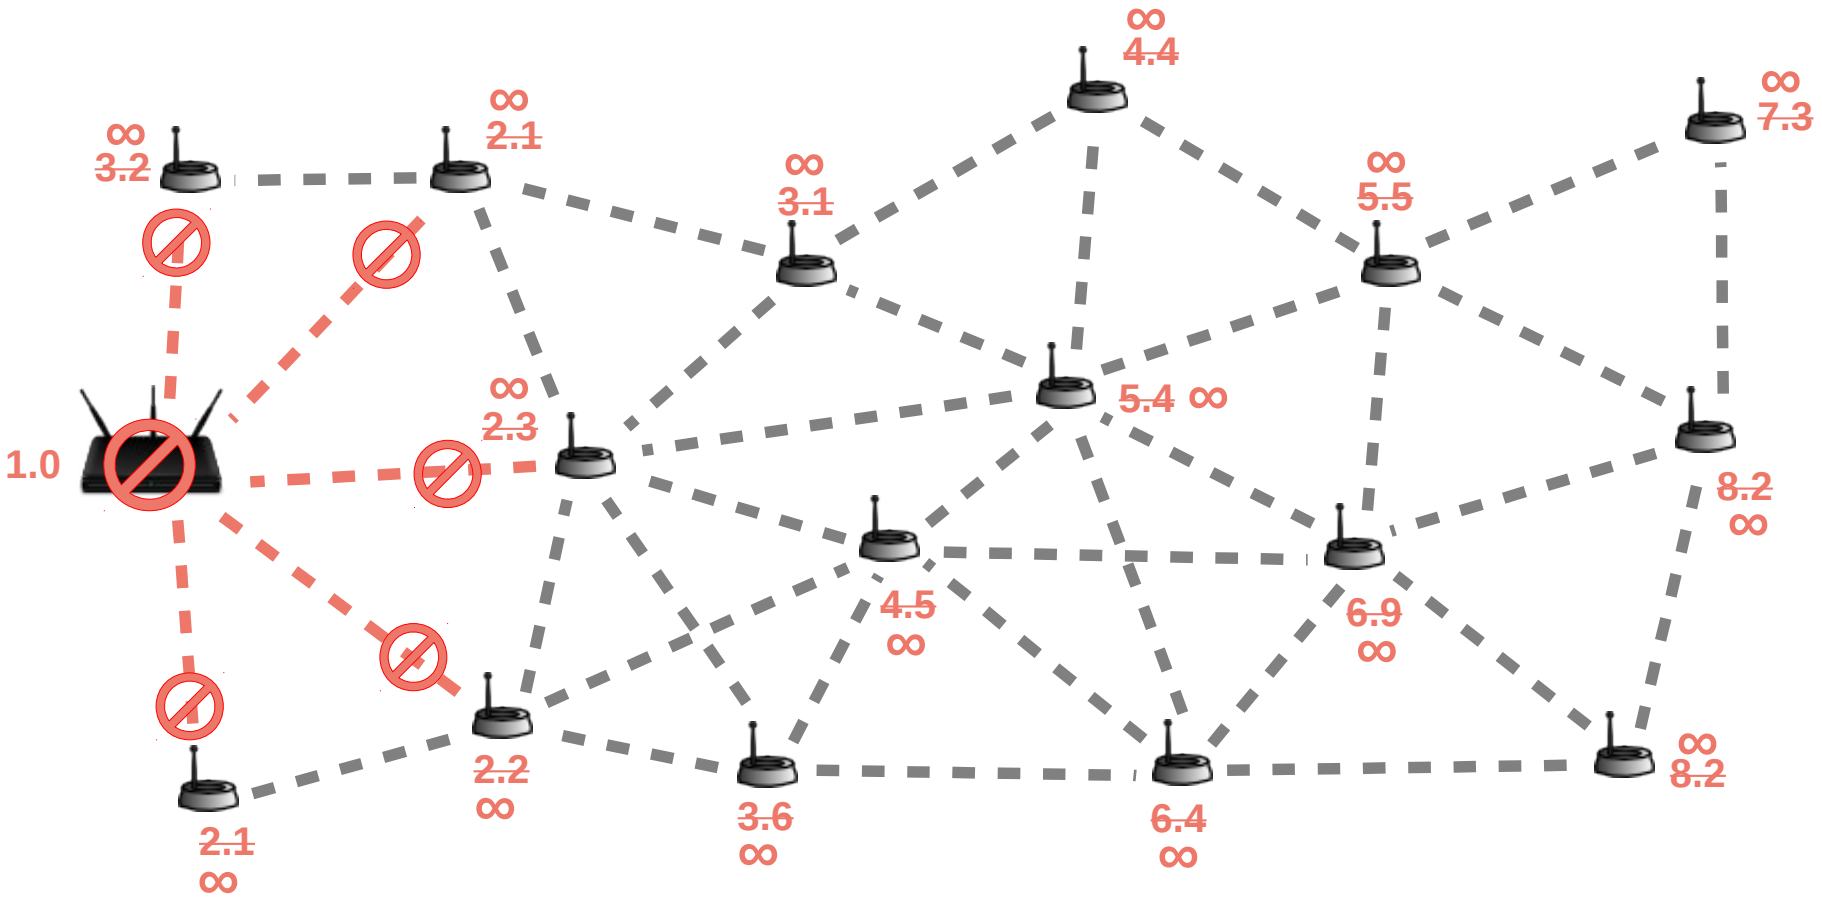

What happens in practice under an LBR crash?

- Some RPL stacks (with major bugs) enter a chaotic state in which an LLN simply collapses: explosion in control traffic.
- Some others (with minor bugs) do not detect the failure (in reasonable time): node ranks grow unbounded; control traffic is heavier than normally.
- Some are correct but still they require considerable:
  - time and
  - traffic.

  to handle an LBR crash.

Why is detecting an LBR crash in RPL suboptimal?

- All links to the dead LBR have to be detected as down by the LBR's neighbors.
  - Otherwise, the LBR's neighbor with such a link may incorrectly advertise a valid path.
- Link crash detection is typically reactive:
  - In low-data-rate applications, it may take a while.
- Learning by all nodes that none of their links may contribute to a path to the LBR is slow and requires traffic:
  - repeated parent changes due to local repair attempts,
  - routing loops due to inconsistencies between nodes,
  - Trickle timer resets upon parent changes and loop detection.

RNFD Goals

- RNFD = Root Node Failure Detector
- Goal: to minimize
  - time and
  - traffic

  required to detect a crash of an LBR (a DODAG root).
- Possible empirical improvements:
  - time = a few times, an order of magnitude less,
  - traffic = a few times less.

RNFD Design Principles

- Explicitly coordinating LBR monitoring between nodes.
- Avoiding probing all links to the dead LBR.
- Proactive checking for a possible LBR crash when some nodes suspect such a failure may have taken place.
- Maximizing independence of RPL.

Node Roles in RNFD

- Sentinel – DODAG root's neighbor that monitors the DODAG root's status.
  - There are typically multiple of them.
  - Not every neighbor of the root has to be Sentinel.
- Acceptor – any node that is not Sentinel and only accepts their observations.
  - The DODAG root itself is also Acceptor.

Node Roles in RNFD

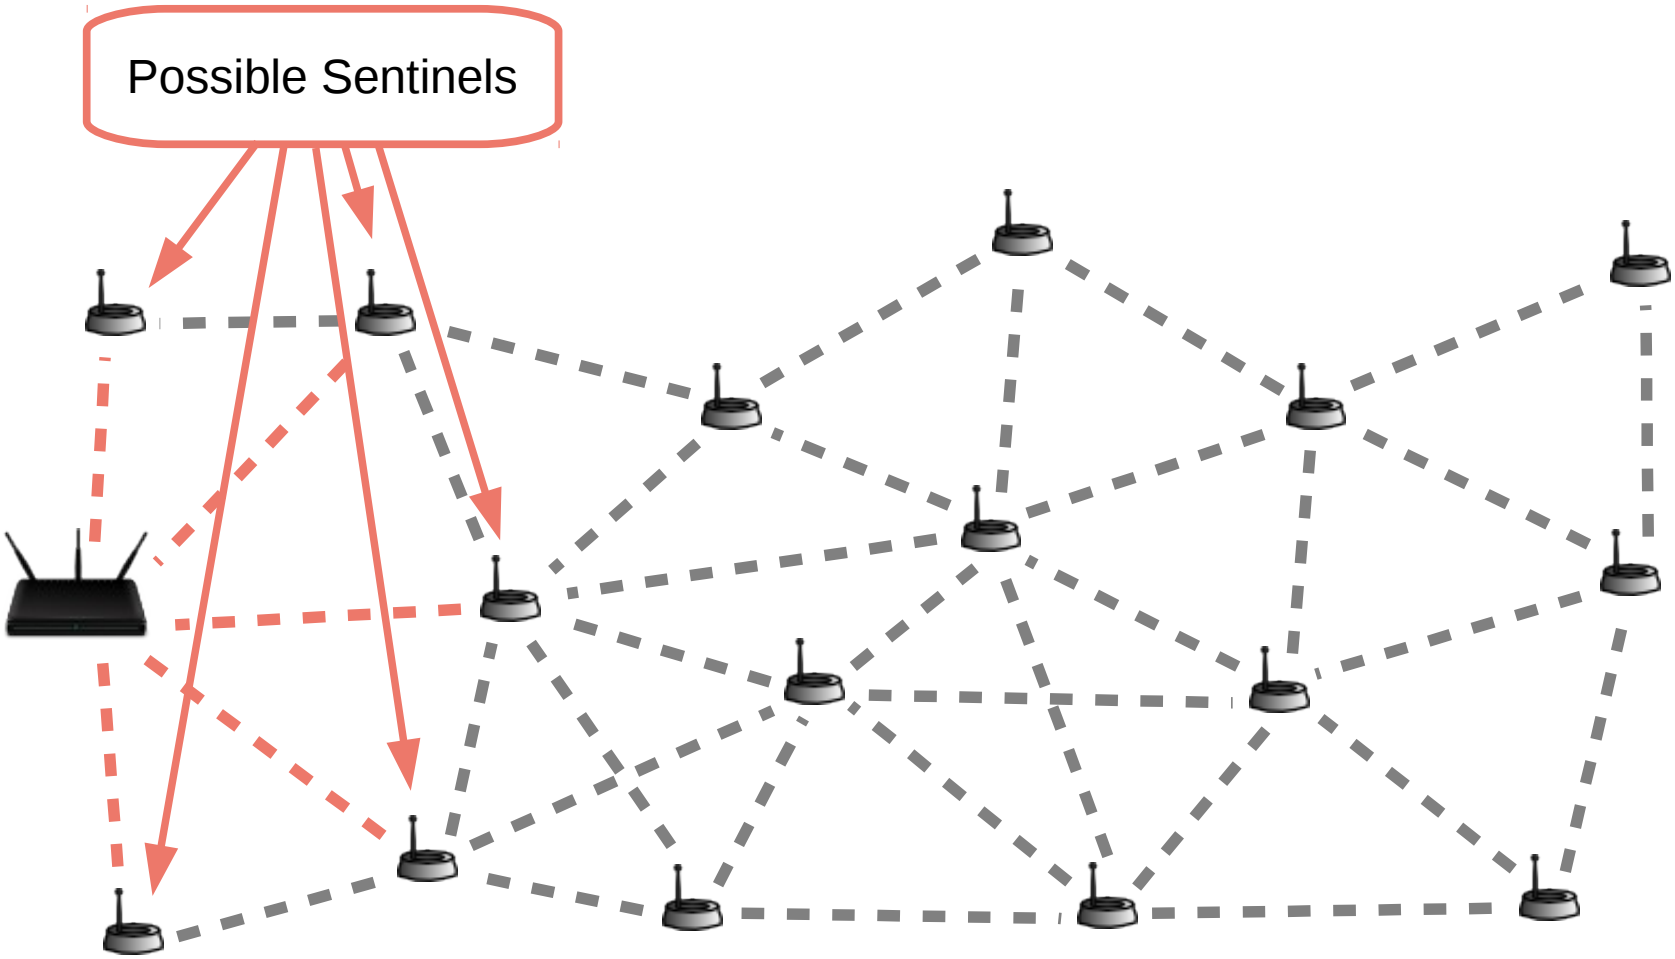

Local Node Data in RNFD

- LORS = Locally Observed DODAG Root's State
  - Describes the node's local knowledge of the DODAG root's status.
  - Attains well-defined values with well-defined transitions.
- CFRC = Conflict-Free Replicated Counter
  - Counts sentinels exhibiting some property.
  - CFRCs can be compared (partial order).
  - CFRCs can be merged (order- and duplicate-insensitive).
  - Each node keeps two CFRCs:
    - Positive CFRC – counts Sentinels that have considered or still consider the root node as alive in the current DODAG Version
    - Negative CFRC - counts Sentinels that have considered or still consider the root node as dead in the current DODAG Version.

Data Communicated in RNFD

- RNFD Option added to link-local RPL Control Messages.

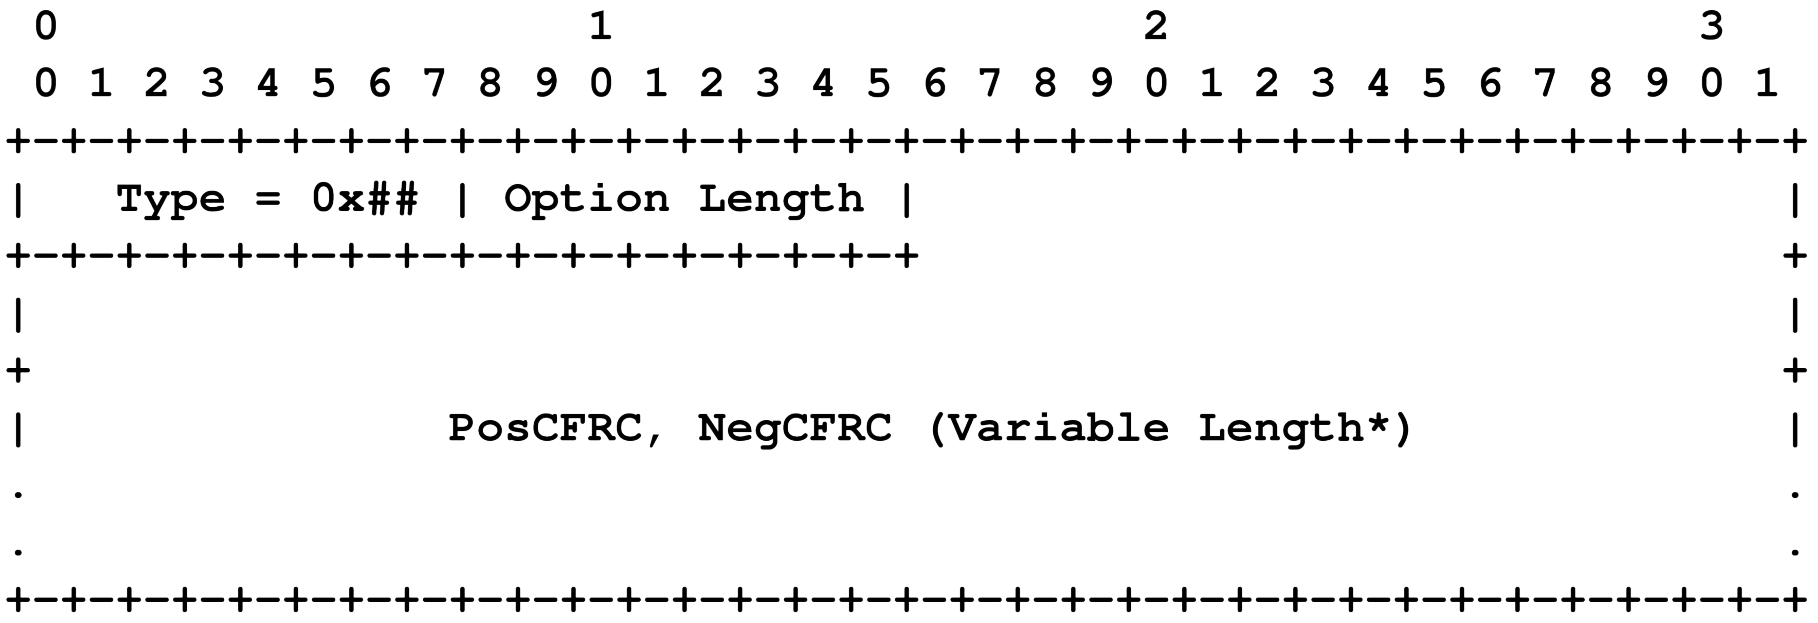

- Upon its reception, a node merges the carried CFRCs with its corresponding local ones.

LORS Transitions in RNFD

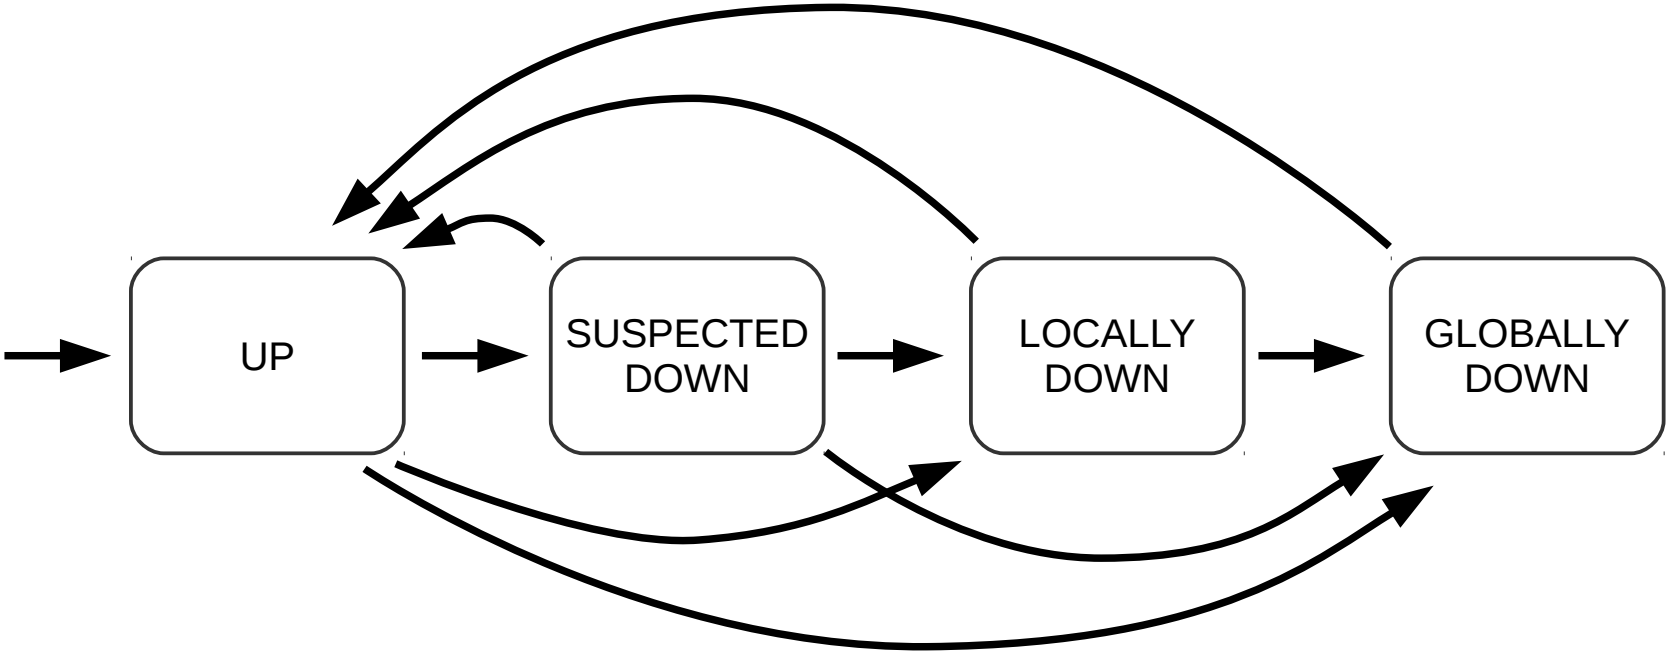

LORS Transitions in RNFD

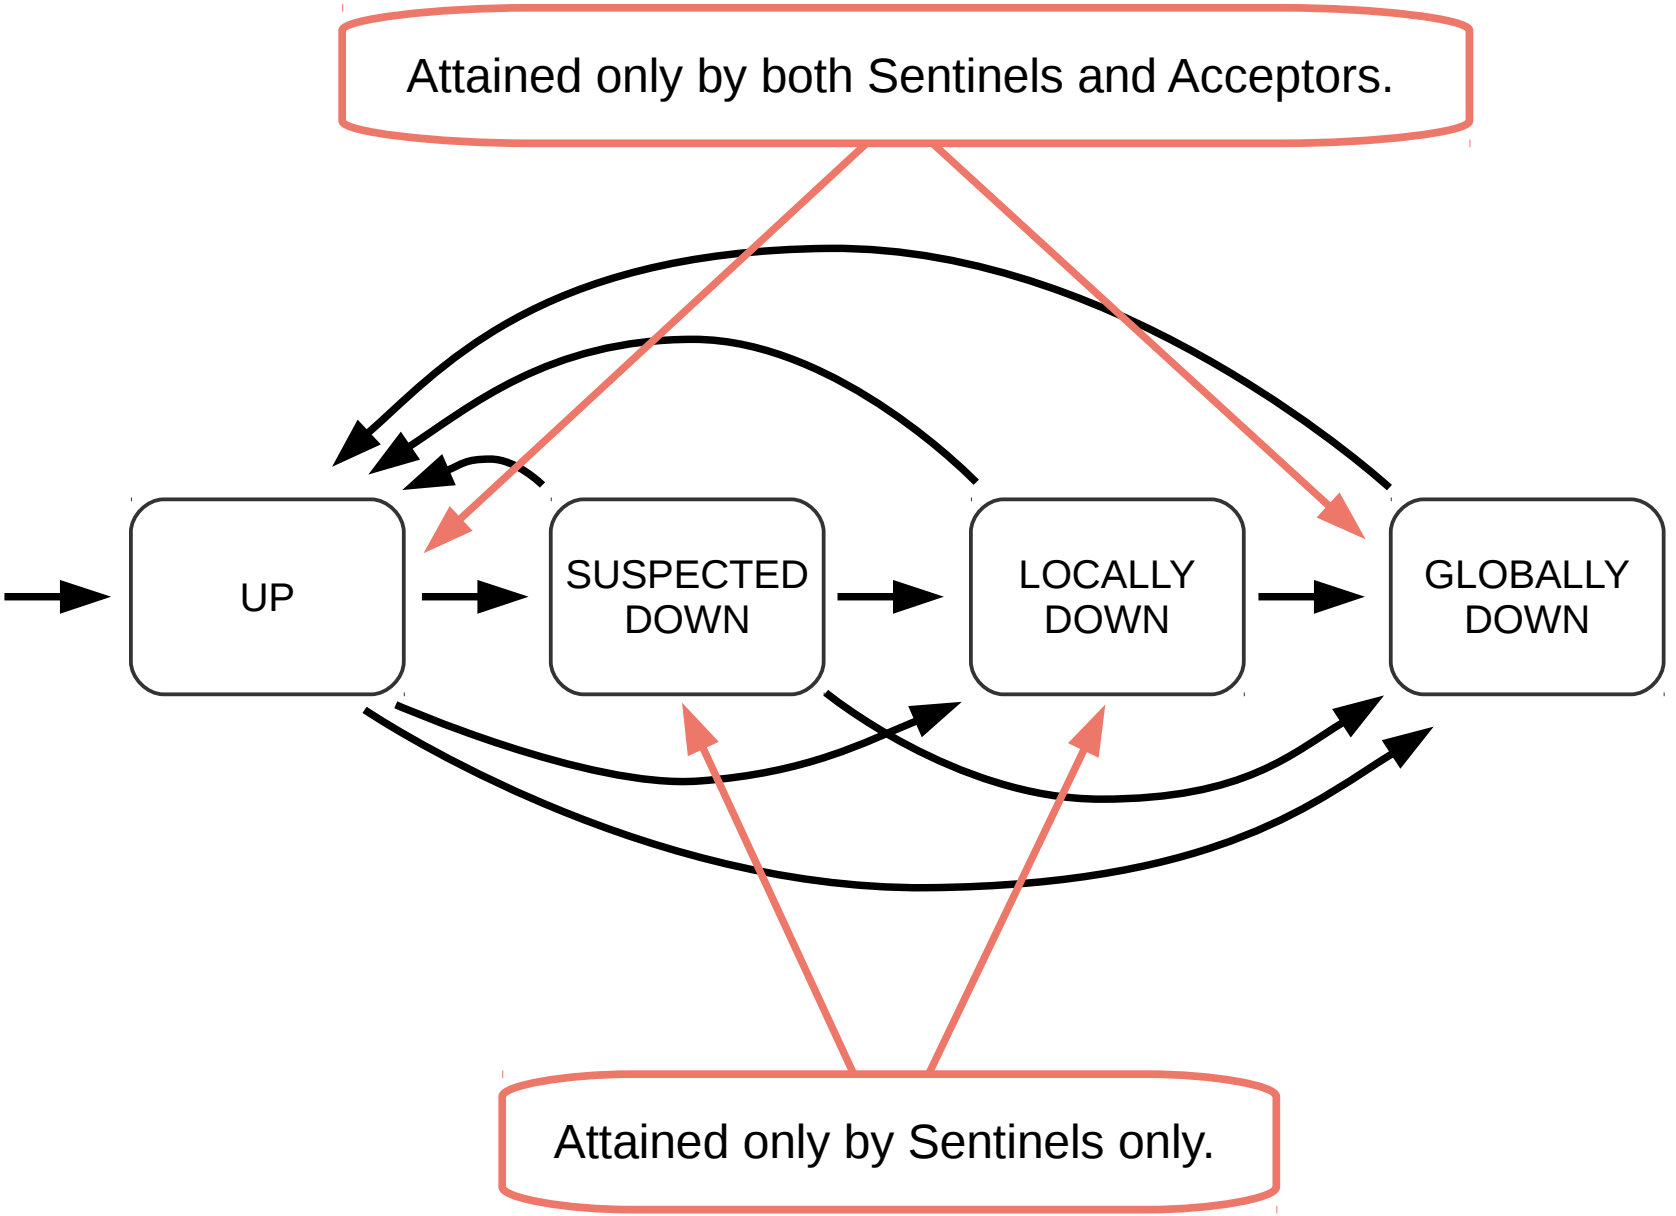

LORS Transitions in RNFD

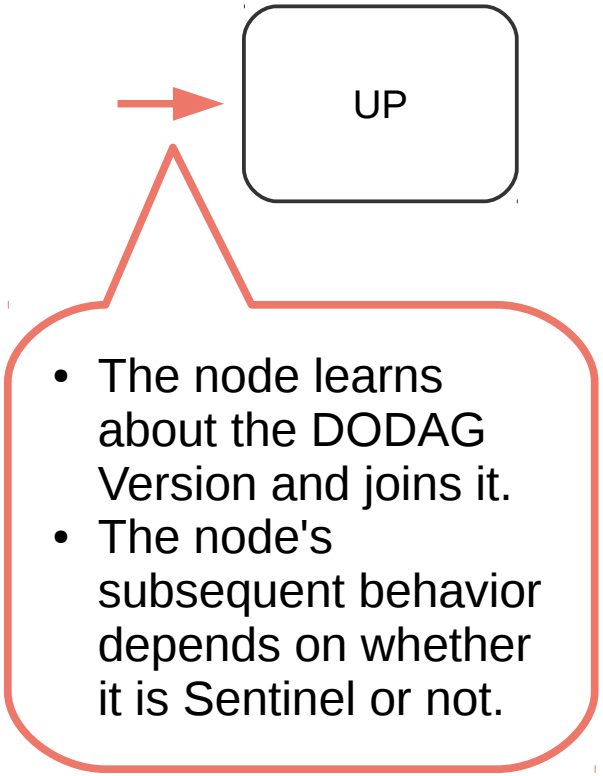

LORS Transitions in RNFD

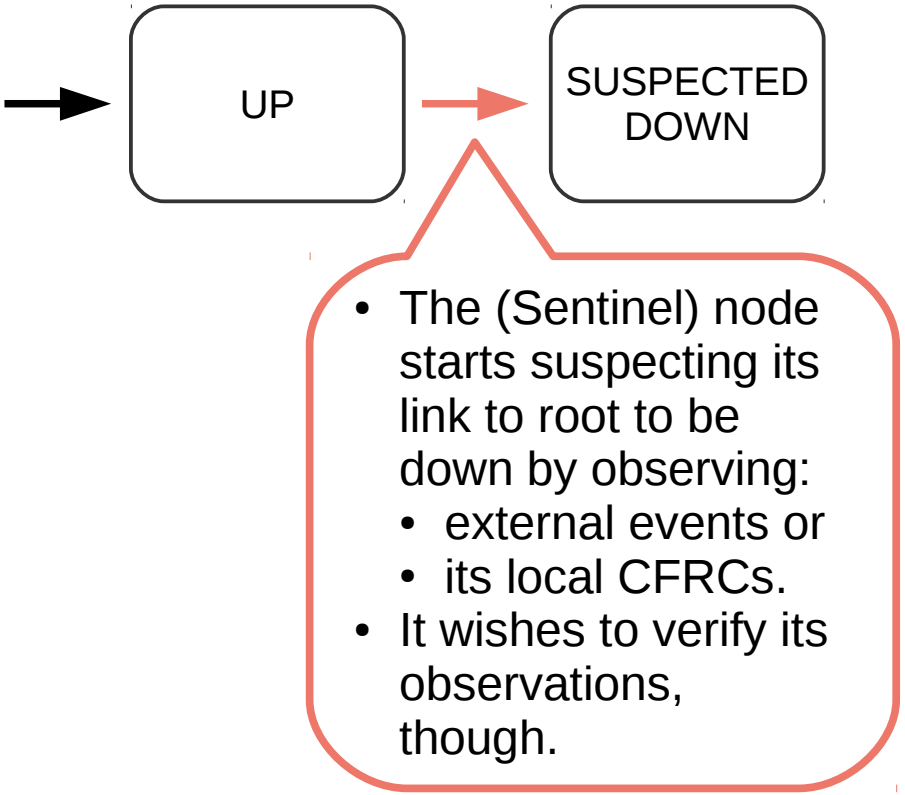

LORS Transitions in RNFD

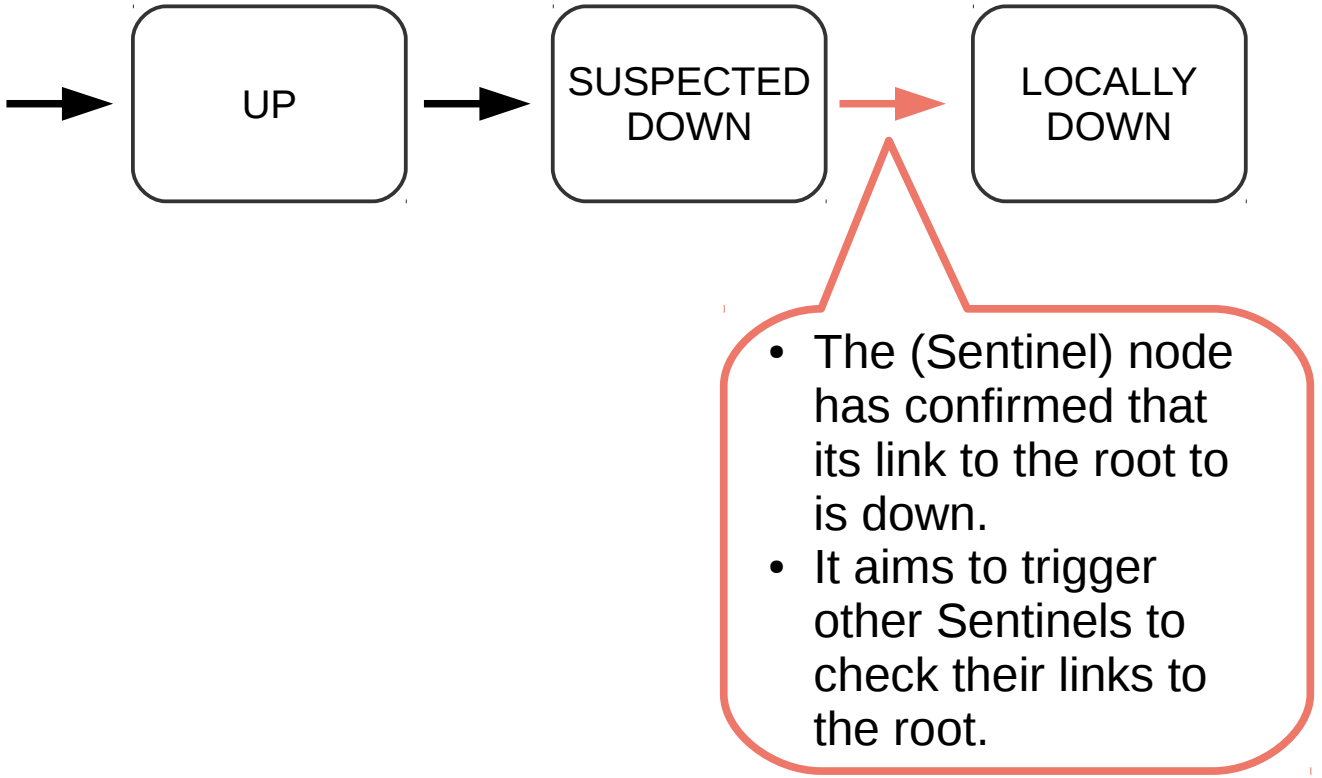

LORS Transitions in RNFD

- The (Sentinel) node may skip the verification if it has enough confidence that its link to the root is indeed dead.

LORS Transitions in RNFD

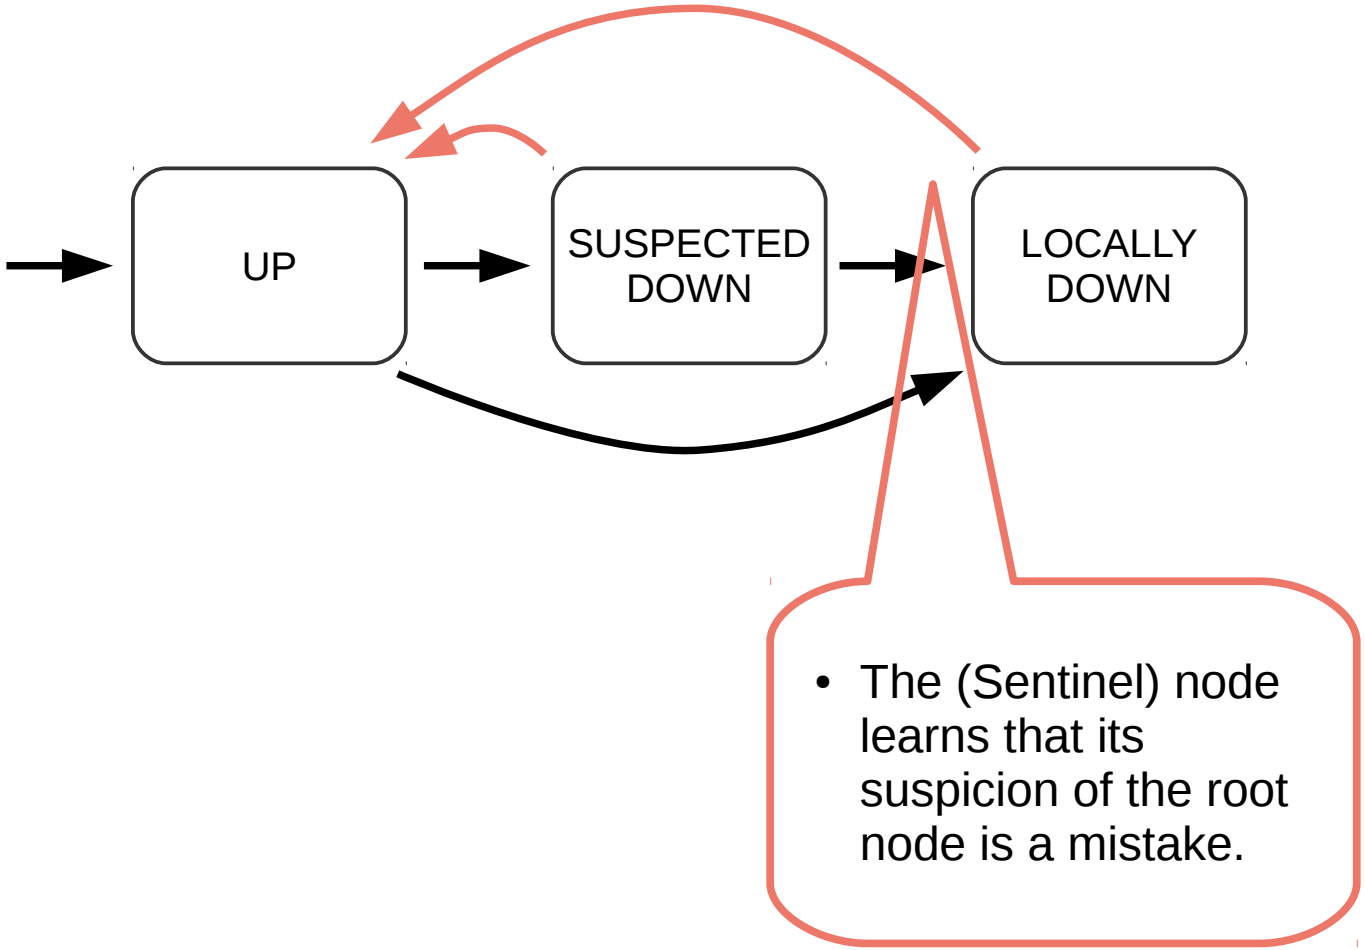

LORS Transitions in RNFD

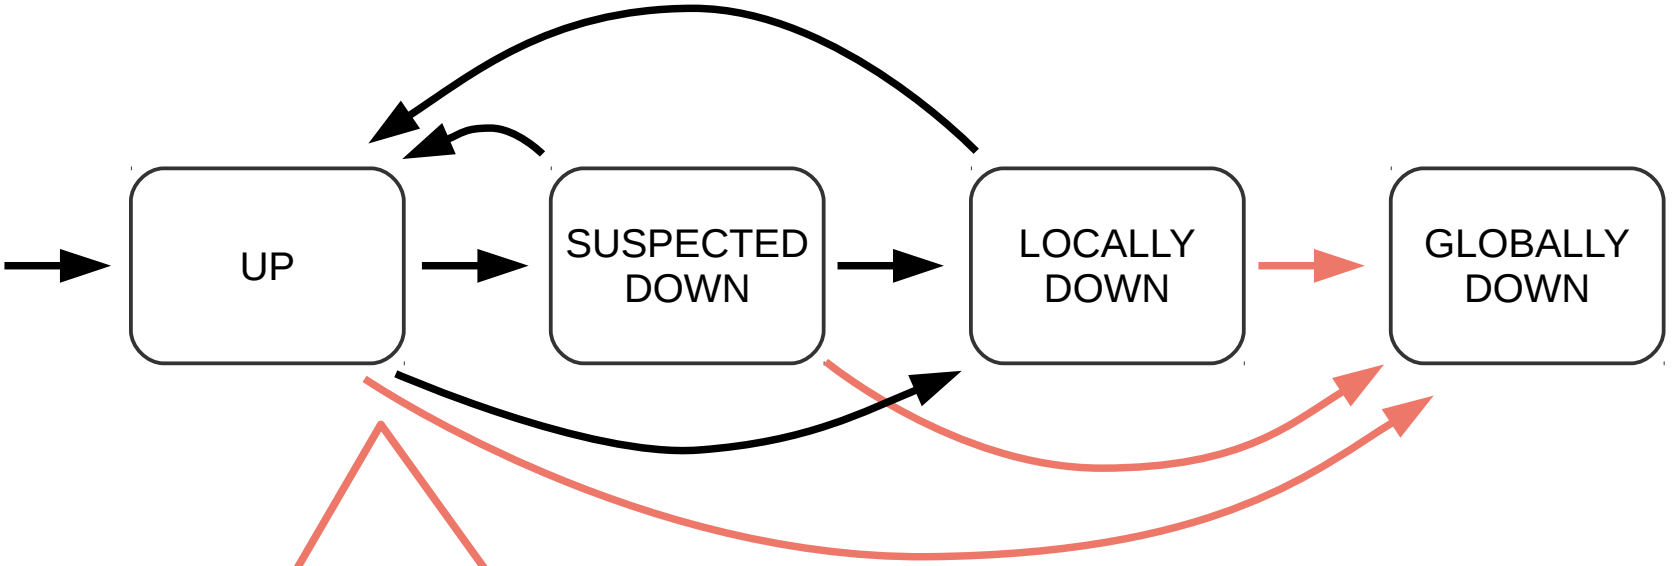

- The node learns that sufficiently many Sentinels consider their links to the root to be down.
- A global decision is thus made that the root itself is down.

LORS Transitions in RNFD

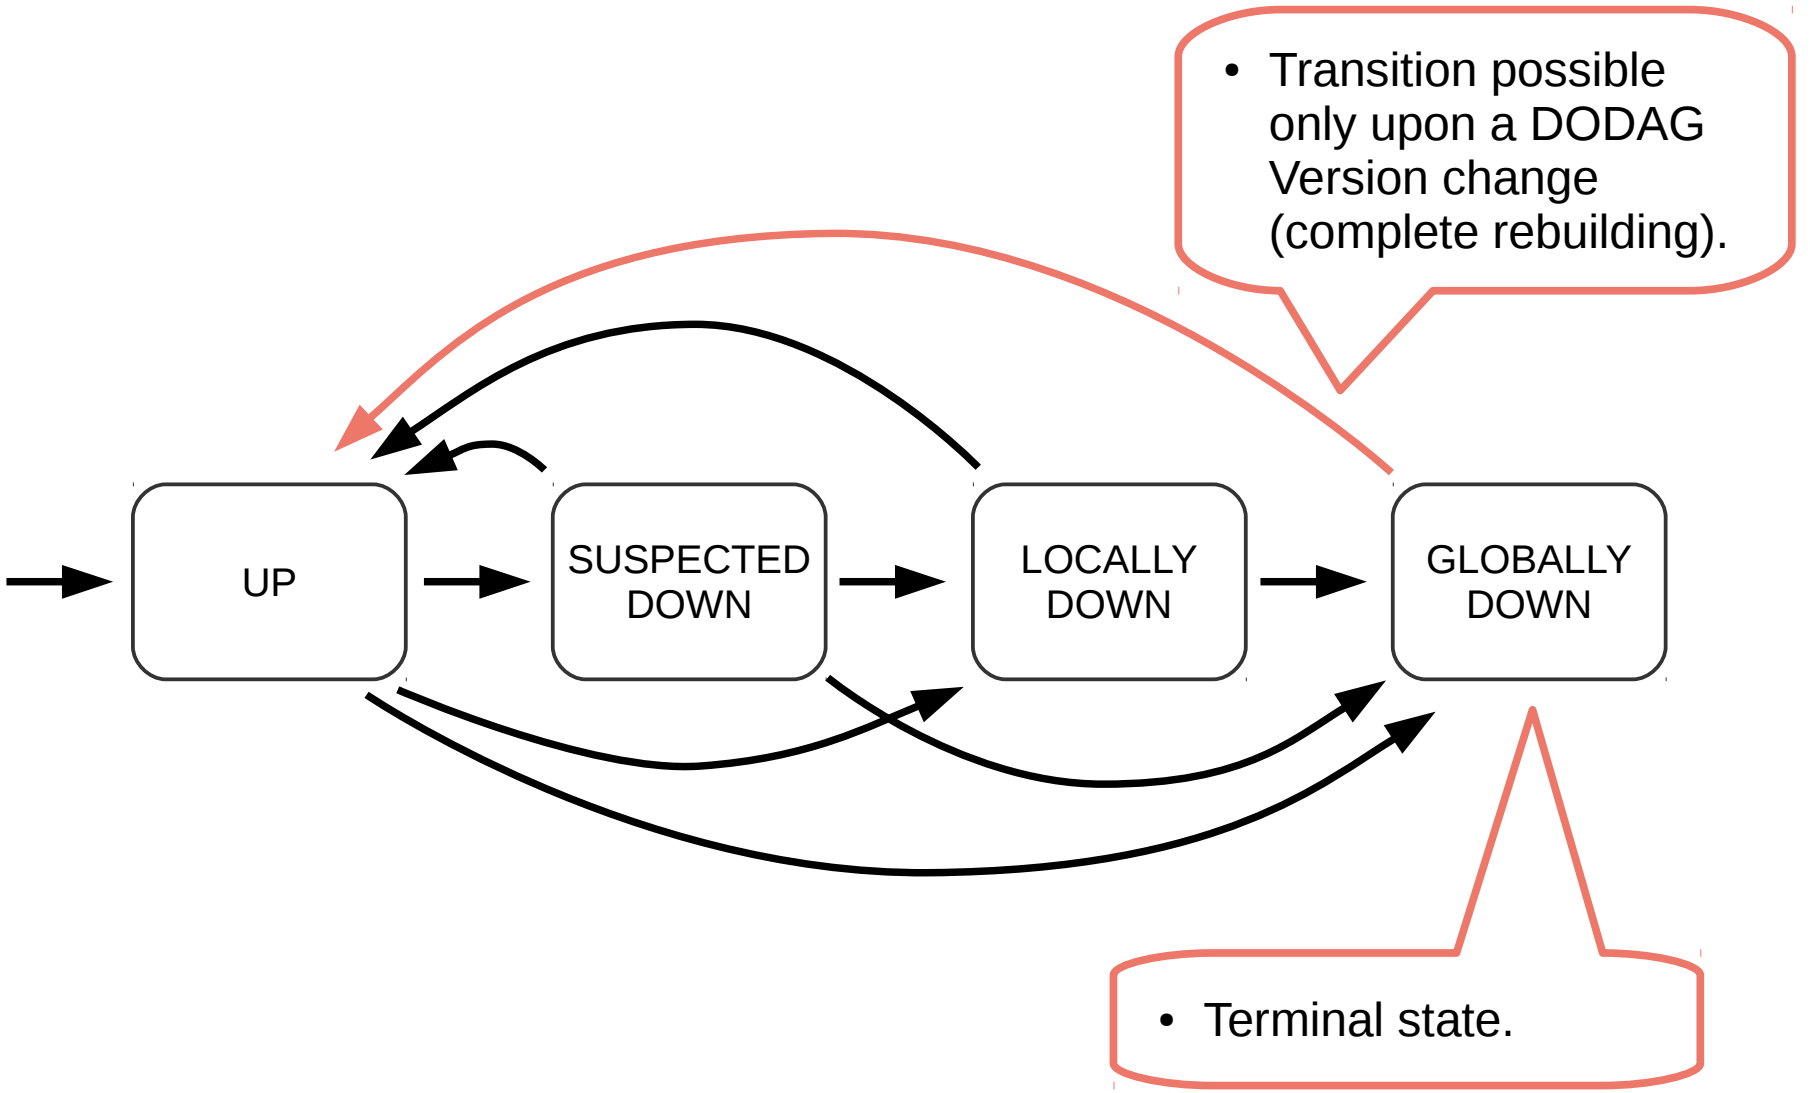

Status of the draft

- Submitted as draft-iwanicki-roll-rnfd-00 on April 6th, 2020.
- Since then several clarifications in the GitHub repo thanks to WG comments.
- Further contributions welcome.

Questions?

draft-iwanicki-roll-rnfd-00

Konrad Iwanicki
iwanicki@mimuw.edu.pl

Root initiated routing state in RPL

draft-ietf-roll-dao-projection

P. Thubert, Ed.; R.A. Jadhav, M. Gillmore

Pascal Thubert

Interim Aug. 2021

ROLL Virtual Meeting

Status to the draft

- Published -17 to -19 since last interim
- Sibling address in SIO based on RFC 8505 discovery
- Simplification: PDR always from Track Ingress, which provides TrackID
- Clarification on maintenance: how to repair Track Segments (storing mode) and Legs (non-Storing e2e)
- Clarification: Policies to inject traffic on Tracks is OOS
- Reviews by S.V.R. Anand and Toerless (in progress)

On Maintenance (Anand's review)

- Anand's point on making the connectivity to the root more reliable.
- No-Path P-DAO indicated by lifetime of 0
- Can indicate a section of a Segment
- Updating can be lossless but possible misordering; e.g.:

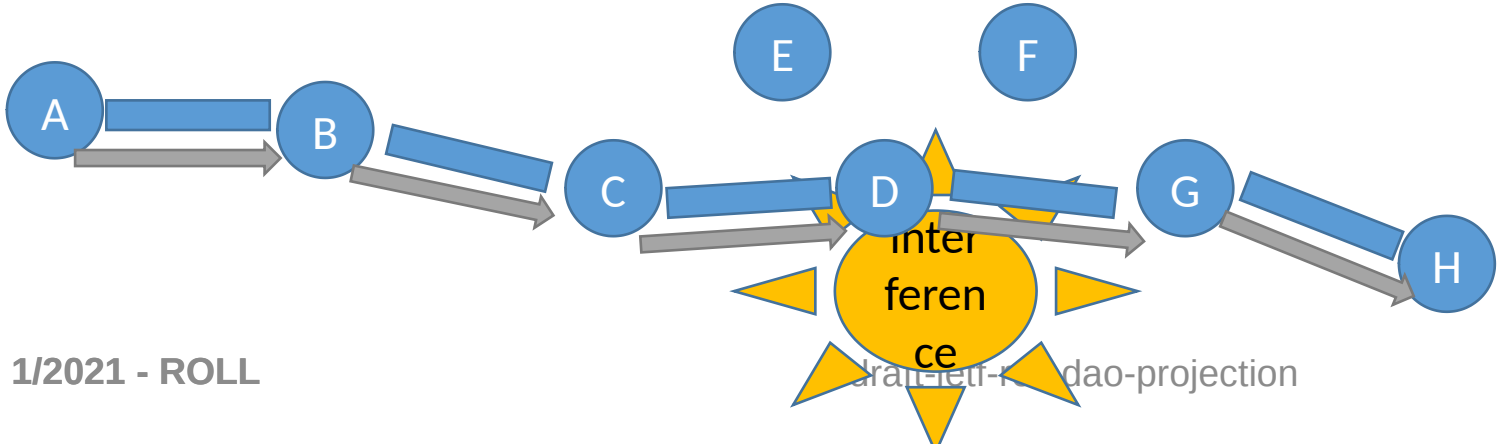

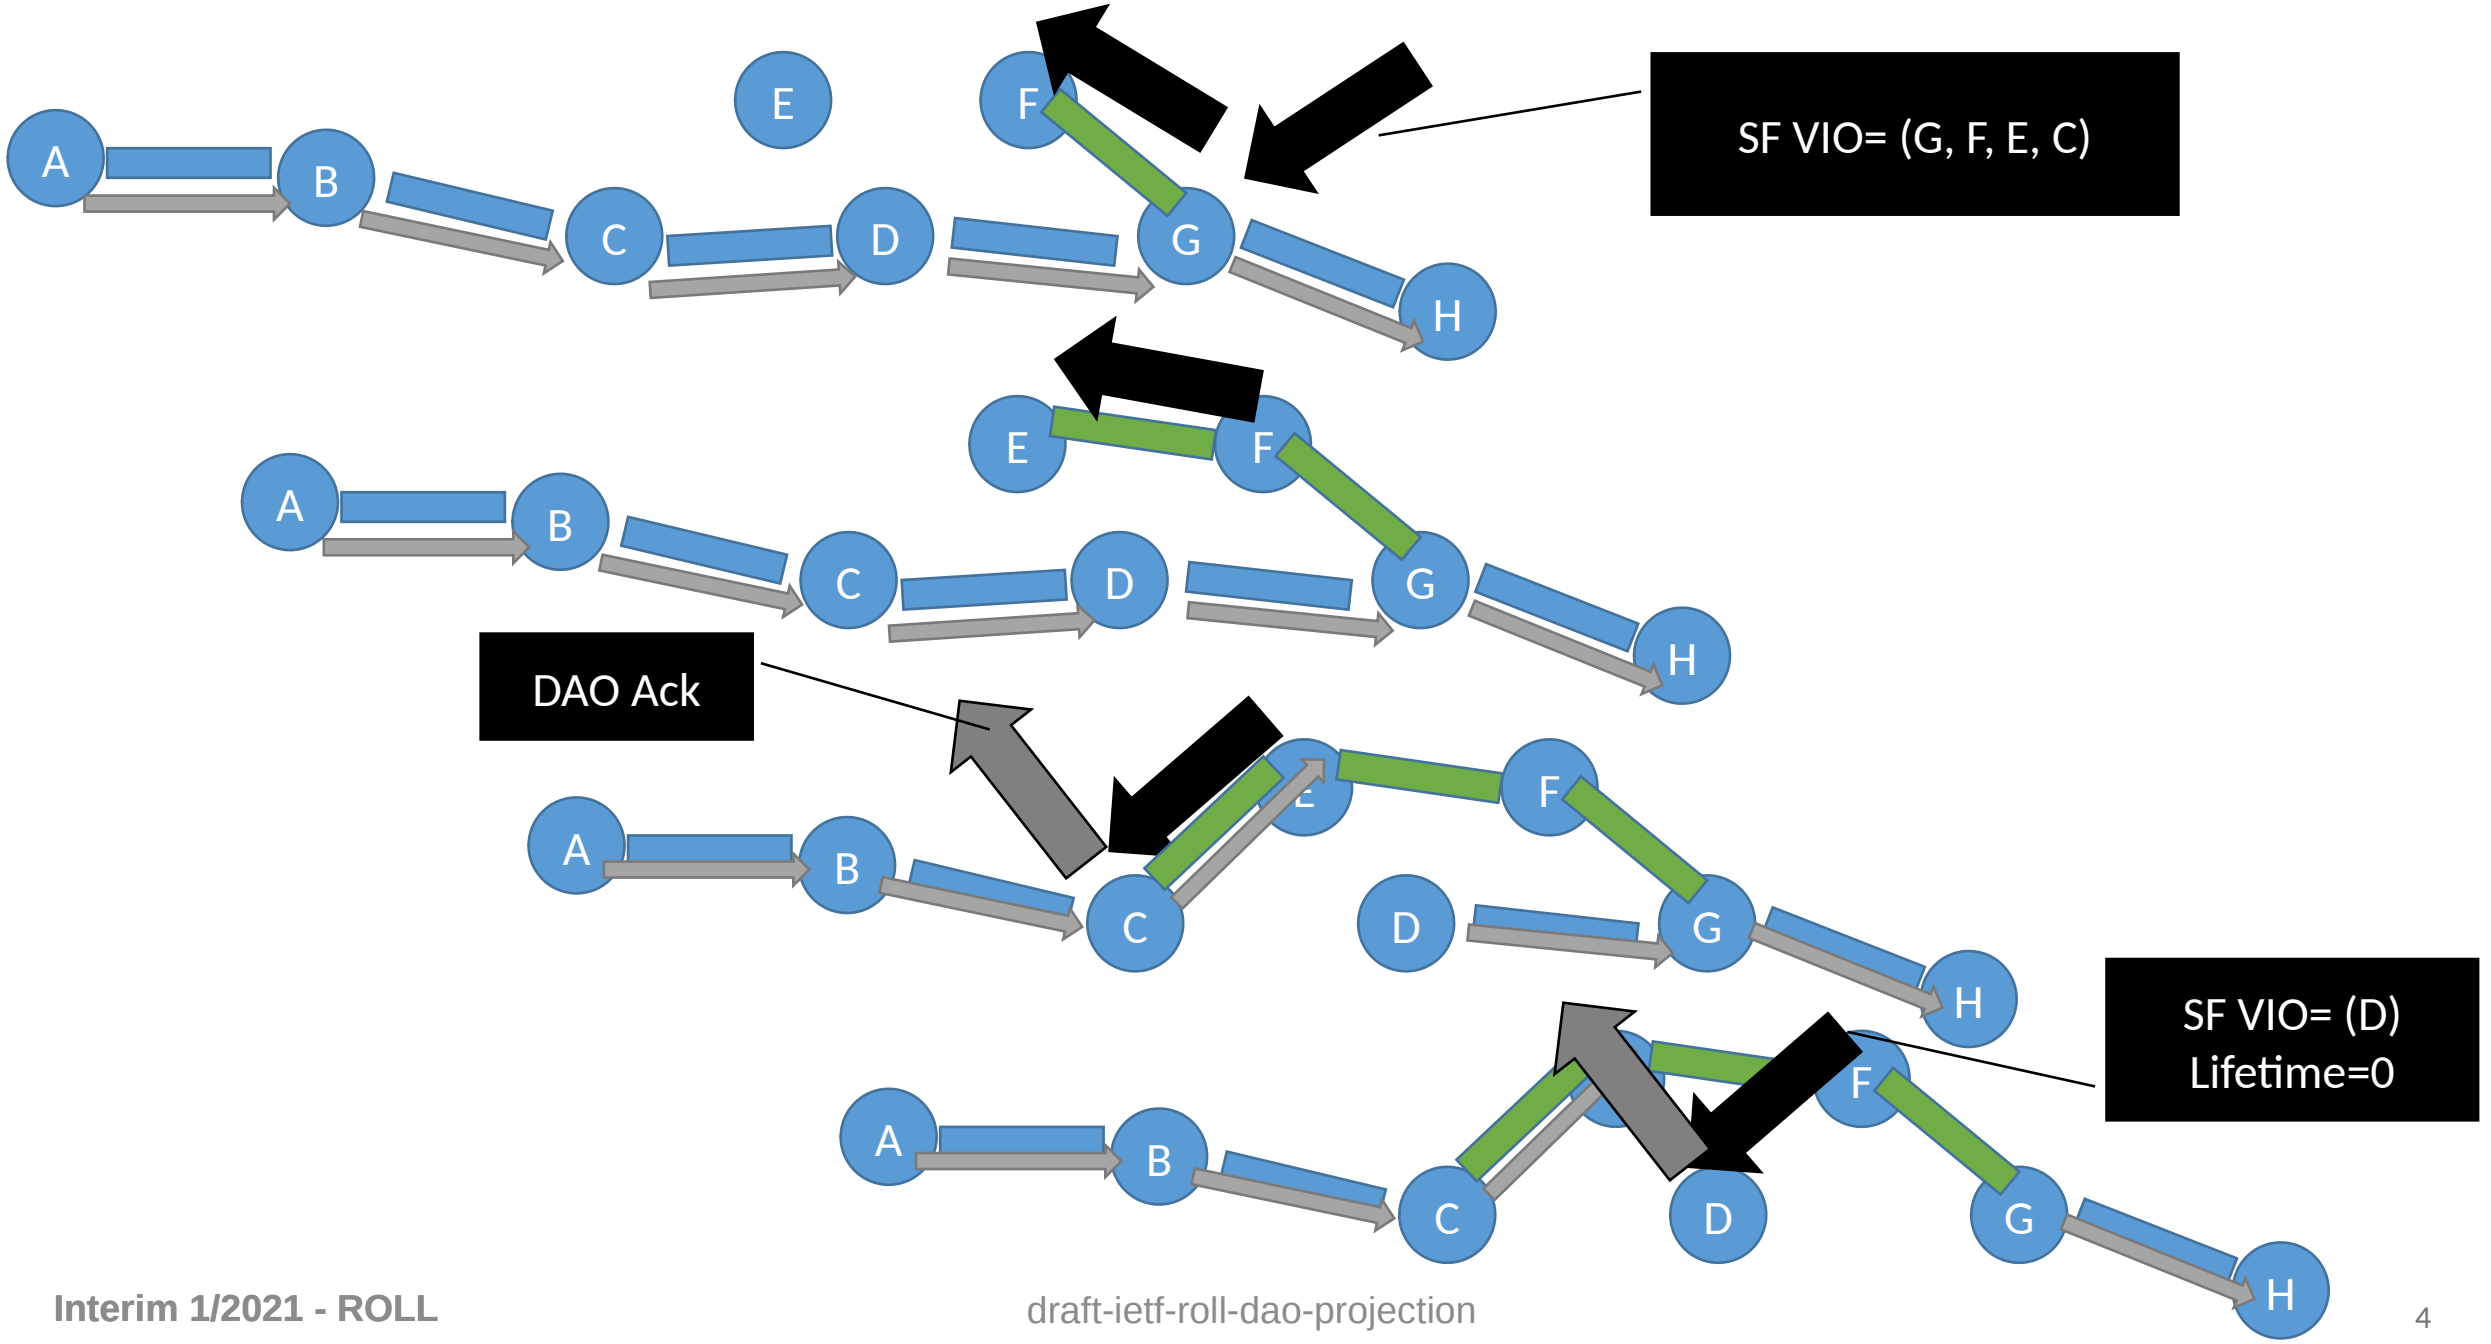
A
B
C
D
E
F
G
H
SF VIO= (G, F, E, C)
DAO Ack
SF VIO= (D)
Lifetime=0

On Track

- A Track is a RPL local instance installed by P DAO.
- Goal: to be equivalent to Non-Storing Mode main DODAG
- The spec only builds multi-legged Tracks (parallel or crossing)
  - Root is Track Ingress, Root address + DODAGID identify the Track
  - A Leg is signaled by a Non-Storing-Mode P-DAO message
  - Track Ingress encapsulates external packets (as in useofrplinfo)
  - Track Ingress places the SRH in the packet in source routed tracks
  - There cannot be non-storing segments (only Tracks within Tracks)
- Storing Mode P-DAO signals Segment of a Track or of main DODAG
  - Does not need re-encapsulation
  - Unless implicit Track => Do we support that ?

Discussion: building real DODAGs

- With -19, a complex track is multi-legged, e.g., 2 Legs below
- Allows 1+n

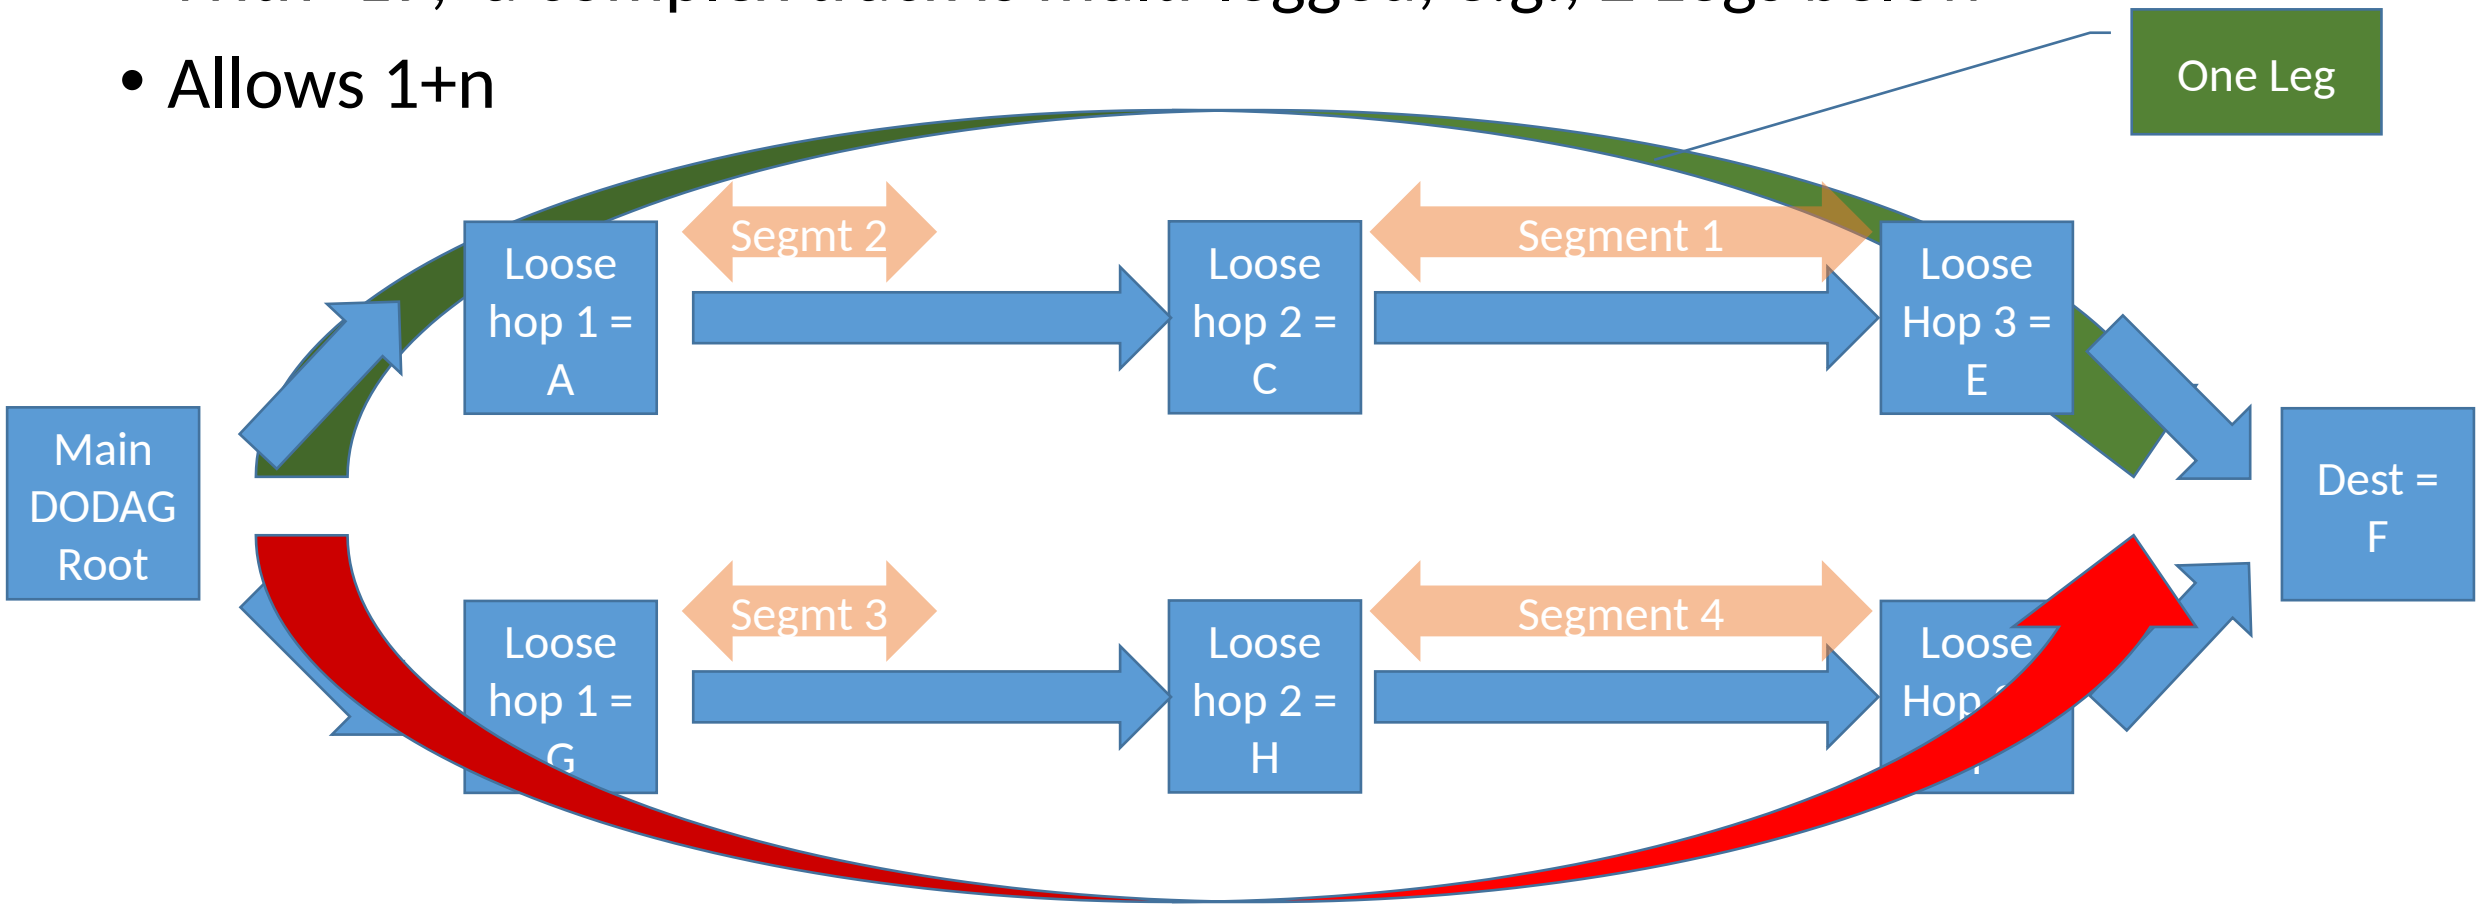

Discussion: building real DODAGs

- NO North-South Segment

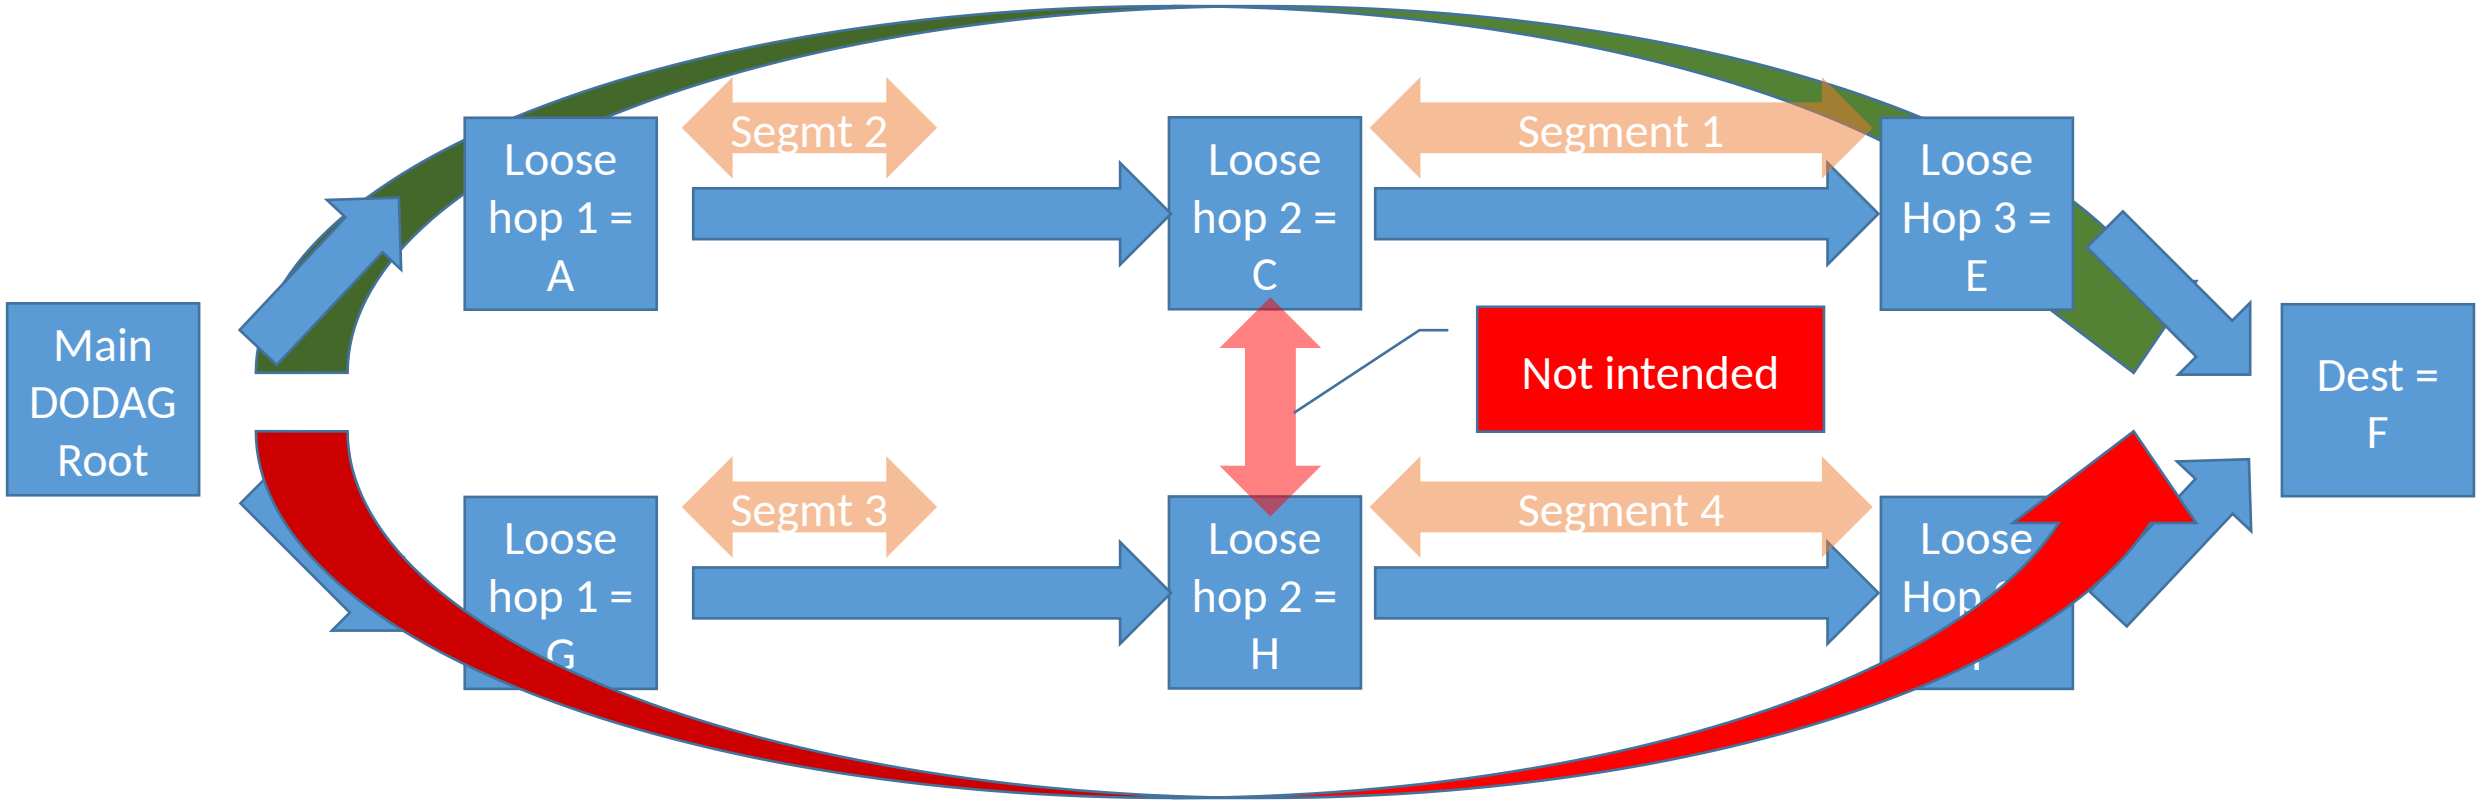

Discussion: How to do this?

- Proposal: use RFC 6550 non-storing Target and Transit to indicate loose parent child relationship, many of them in one P-DAO

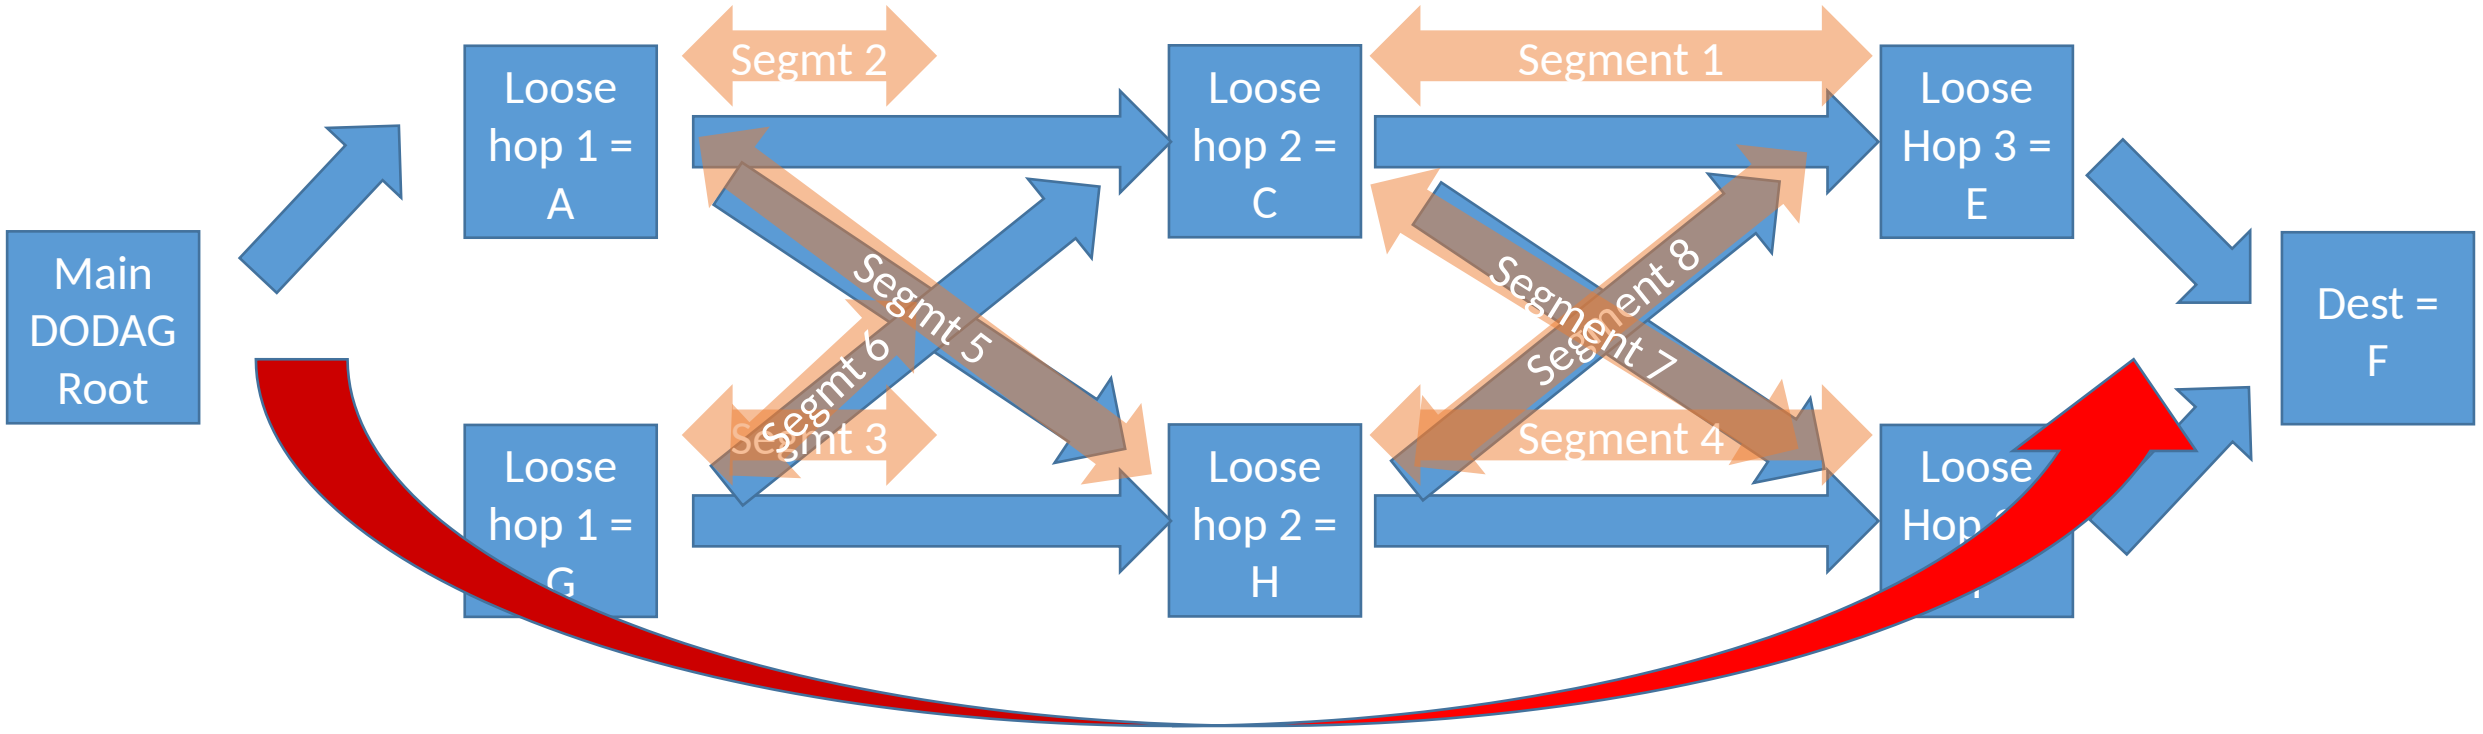

Proposal on the table: building real DODAGs

- Allow the root to pack 1*(1*target 1*transit) in P_DAO
- Provides the Ingress similar information to what a main Root has in non storing mode
- Ingress makes same computations as main Root and gets the instance topology as a DODAG.

Slides from previous meetings

Status to the draft (cont)

- RPI modified to indicate P-Route
- Extending RFC 6553 and RFC 8138

```
         0                   1                   2
         0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        |1|0|E| Length  | 6LoRH Type 7  | RPLInstanceID |
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- New P-RPI-6LoRH, both elective and non-elective forms

Encapsulation Rules

- Source of outer header MUST be Track Ingress- think DODAG Root
- RPL Instance ID in RPI MUST indicate TrackID (if not main DODAG)
- SR-VIO: Loose from Track Ingress, excluded, to Egress, included
  - Copied Verbatim in inserted SRH-6LoRH,
  - Requires encapsulation (can be recursive)
- SF-VIO: Strict from Segment Ingress to Egress, both included
  - No Encapsulation if Source and RPI both match Segment definition
  - A Segment is an Implicit Track if P-DAO Ingress == 1st SF-VIO entry
- TBD: matching rules, Flow Info option, when to tunnel?

P-DAO construction

- RPL Target Options can be factorized
- But there is one and only one VIO (SF-VIO or SR-VIO)
- So the Ack management is easier
- VIO sent to egress; SR-VIO sent to ingress
- Track ID is a RPL local instance ID
- Taken from the Track Egress Name Space

P-DAO Format

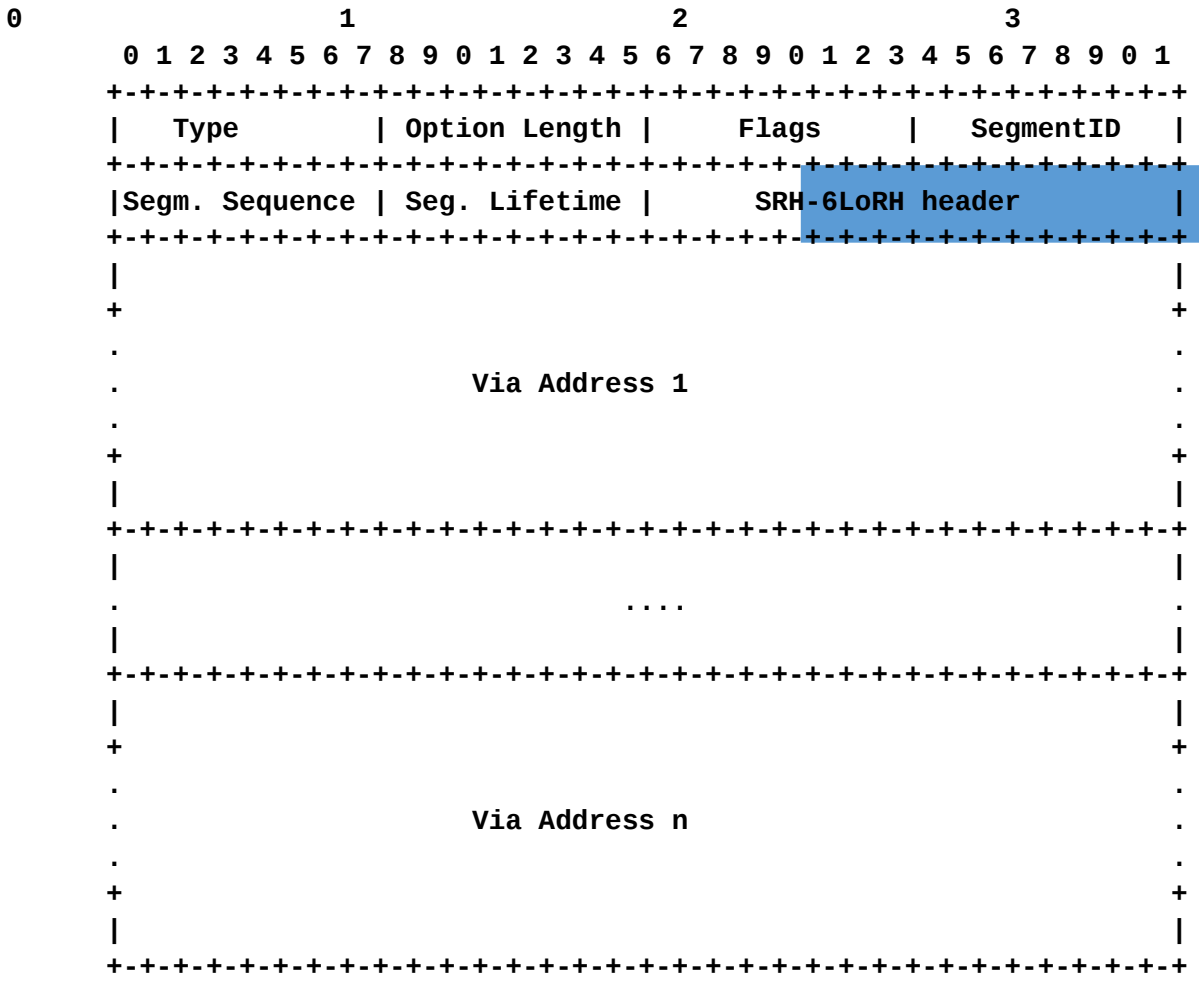

May be more than one in Non-storing Mode

Must be optimized in Non-storing Mode, to be used as is in packets

Topology awareness

- Initially out of scope
- Now we have non storing mode + Sibling info option
  - Acronym conflict with RPL’s Solicited Information Option
- Needed for profiles >= 3
- Which sibling to advertise is still out of scope
  - Separate draft?

Profile 1:
Compress SRH in main DODAG with strict SM Segments

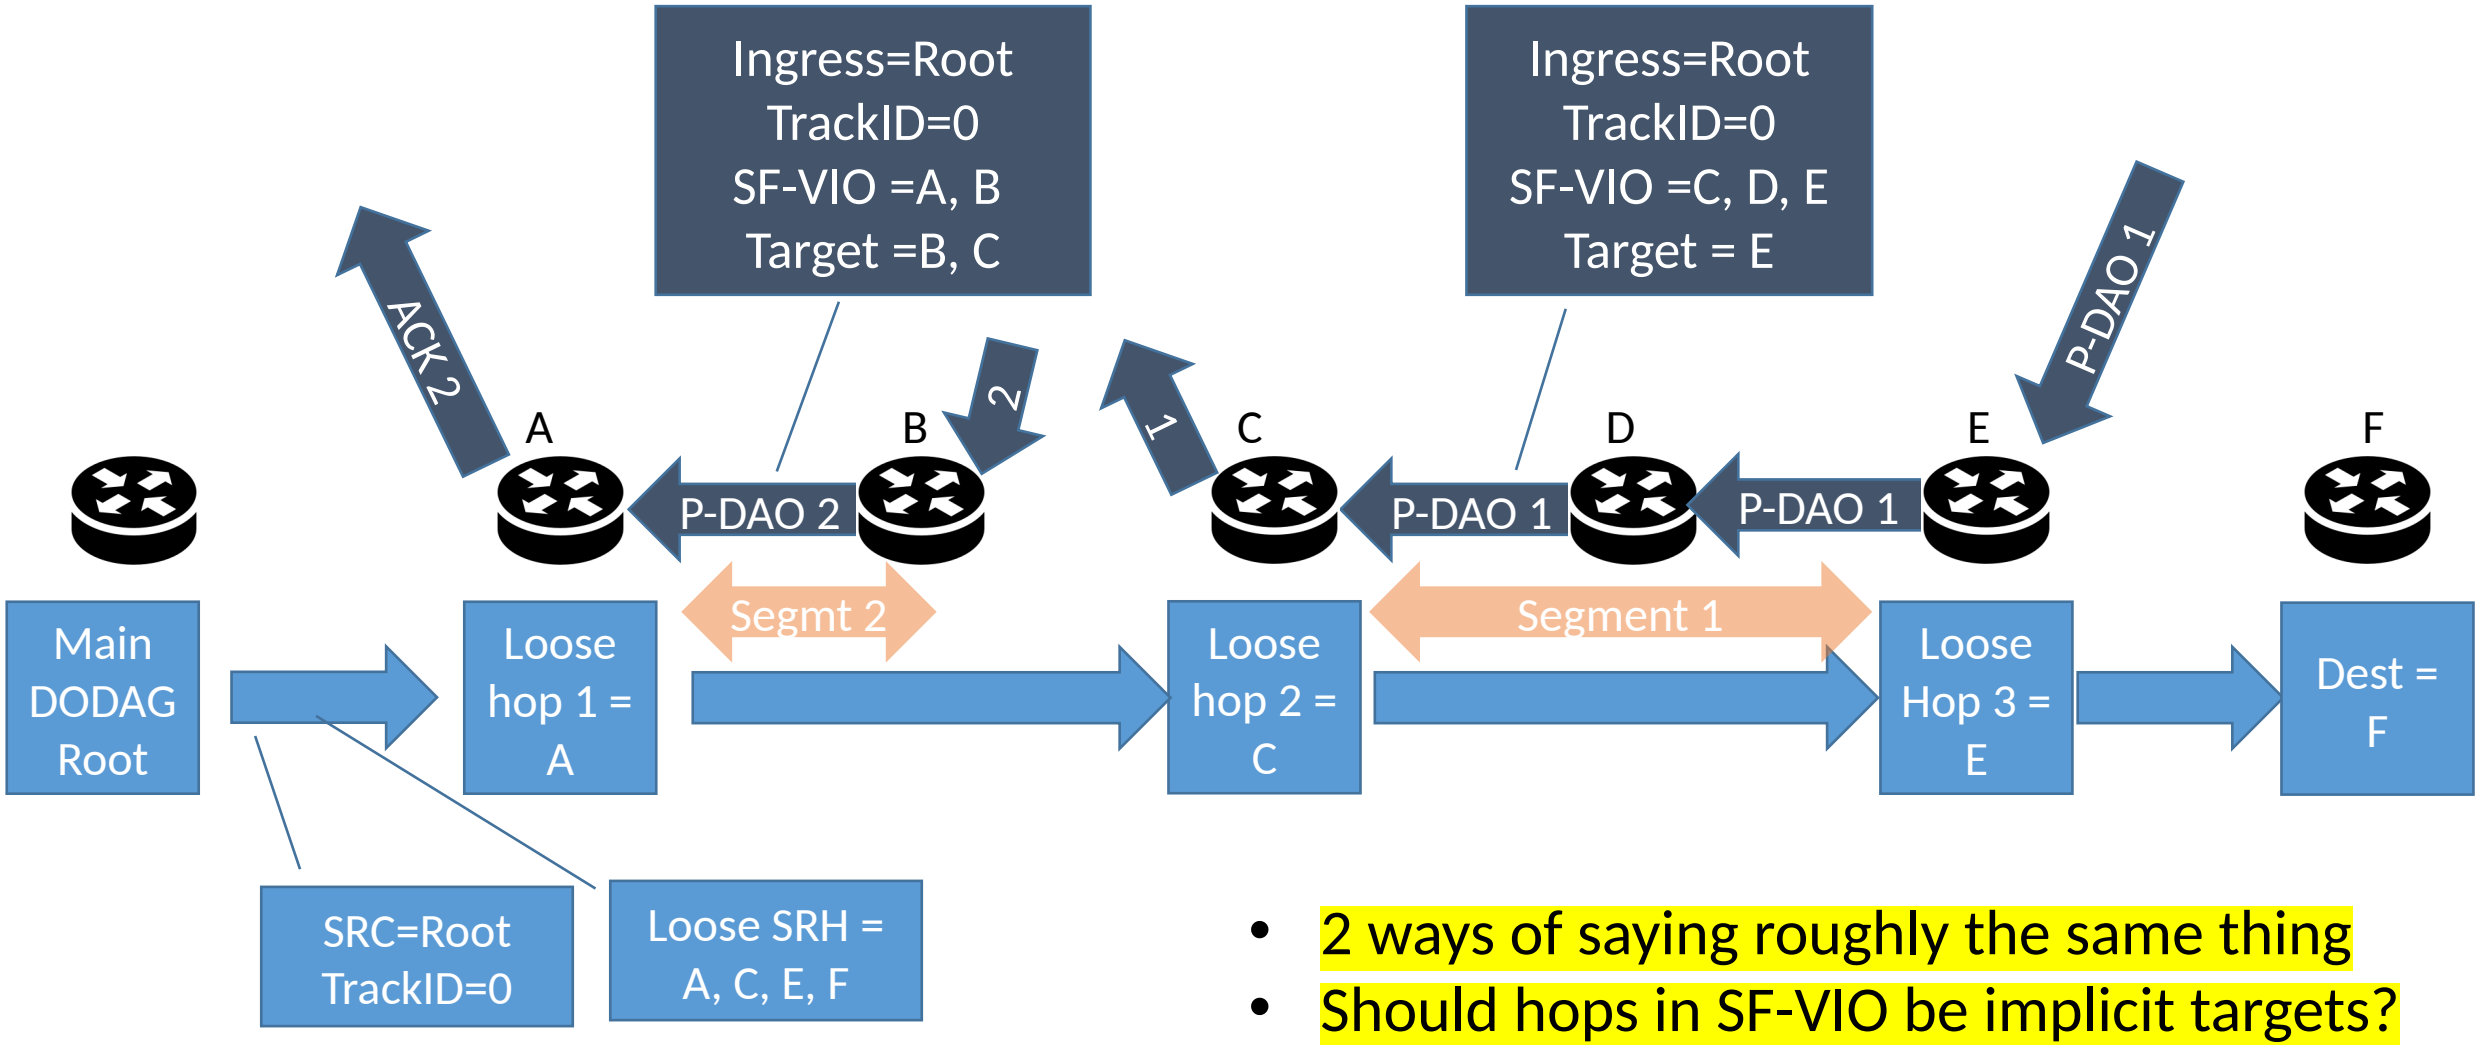

- 2 ways of saying roughly the same thing
- Should hops in SF-VIO be implicit targets?

Profile 2:
Compress SRH in main DODAG with Strict NSM Tracks

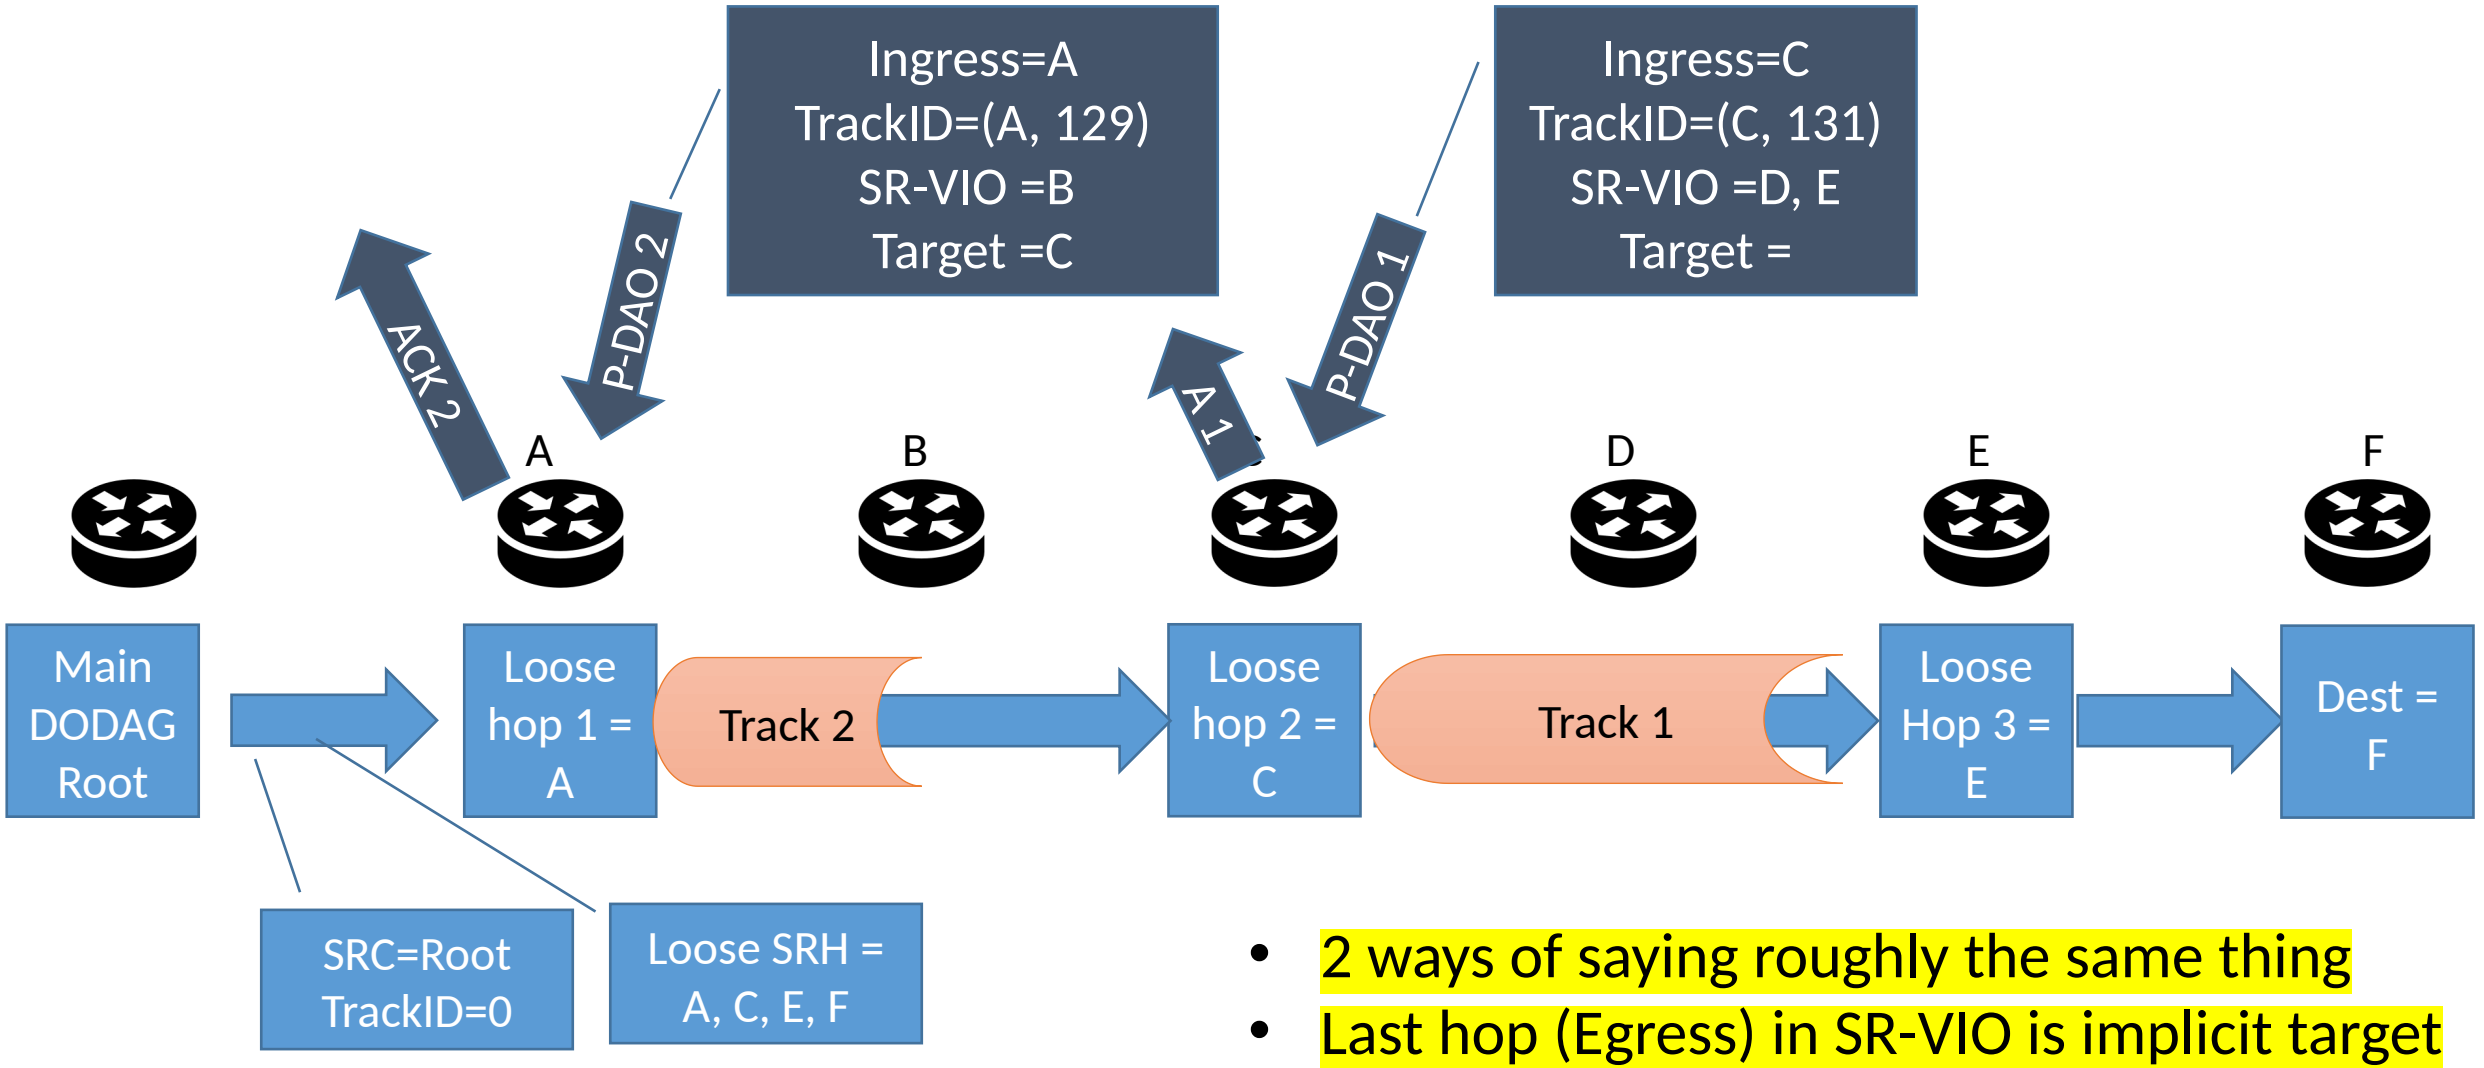

- 2 ways of saying roughly the same thing
- Last hop (Egress) in SR-VIO is implicit target

Profile 3: Implicit Track with Strict SM Segments,

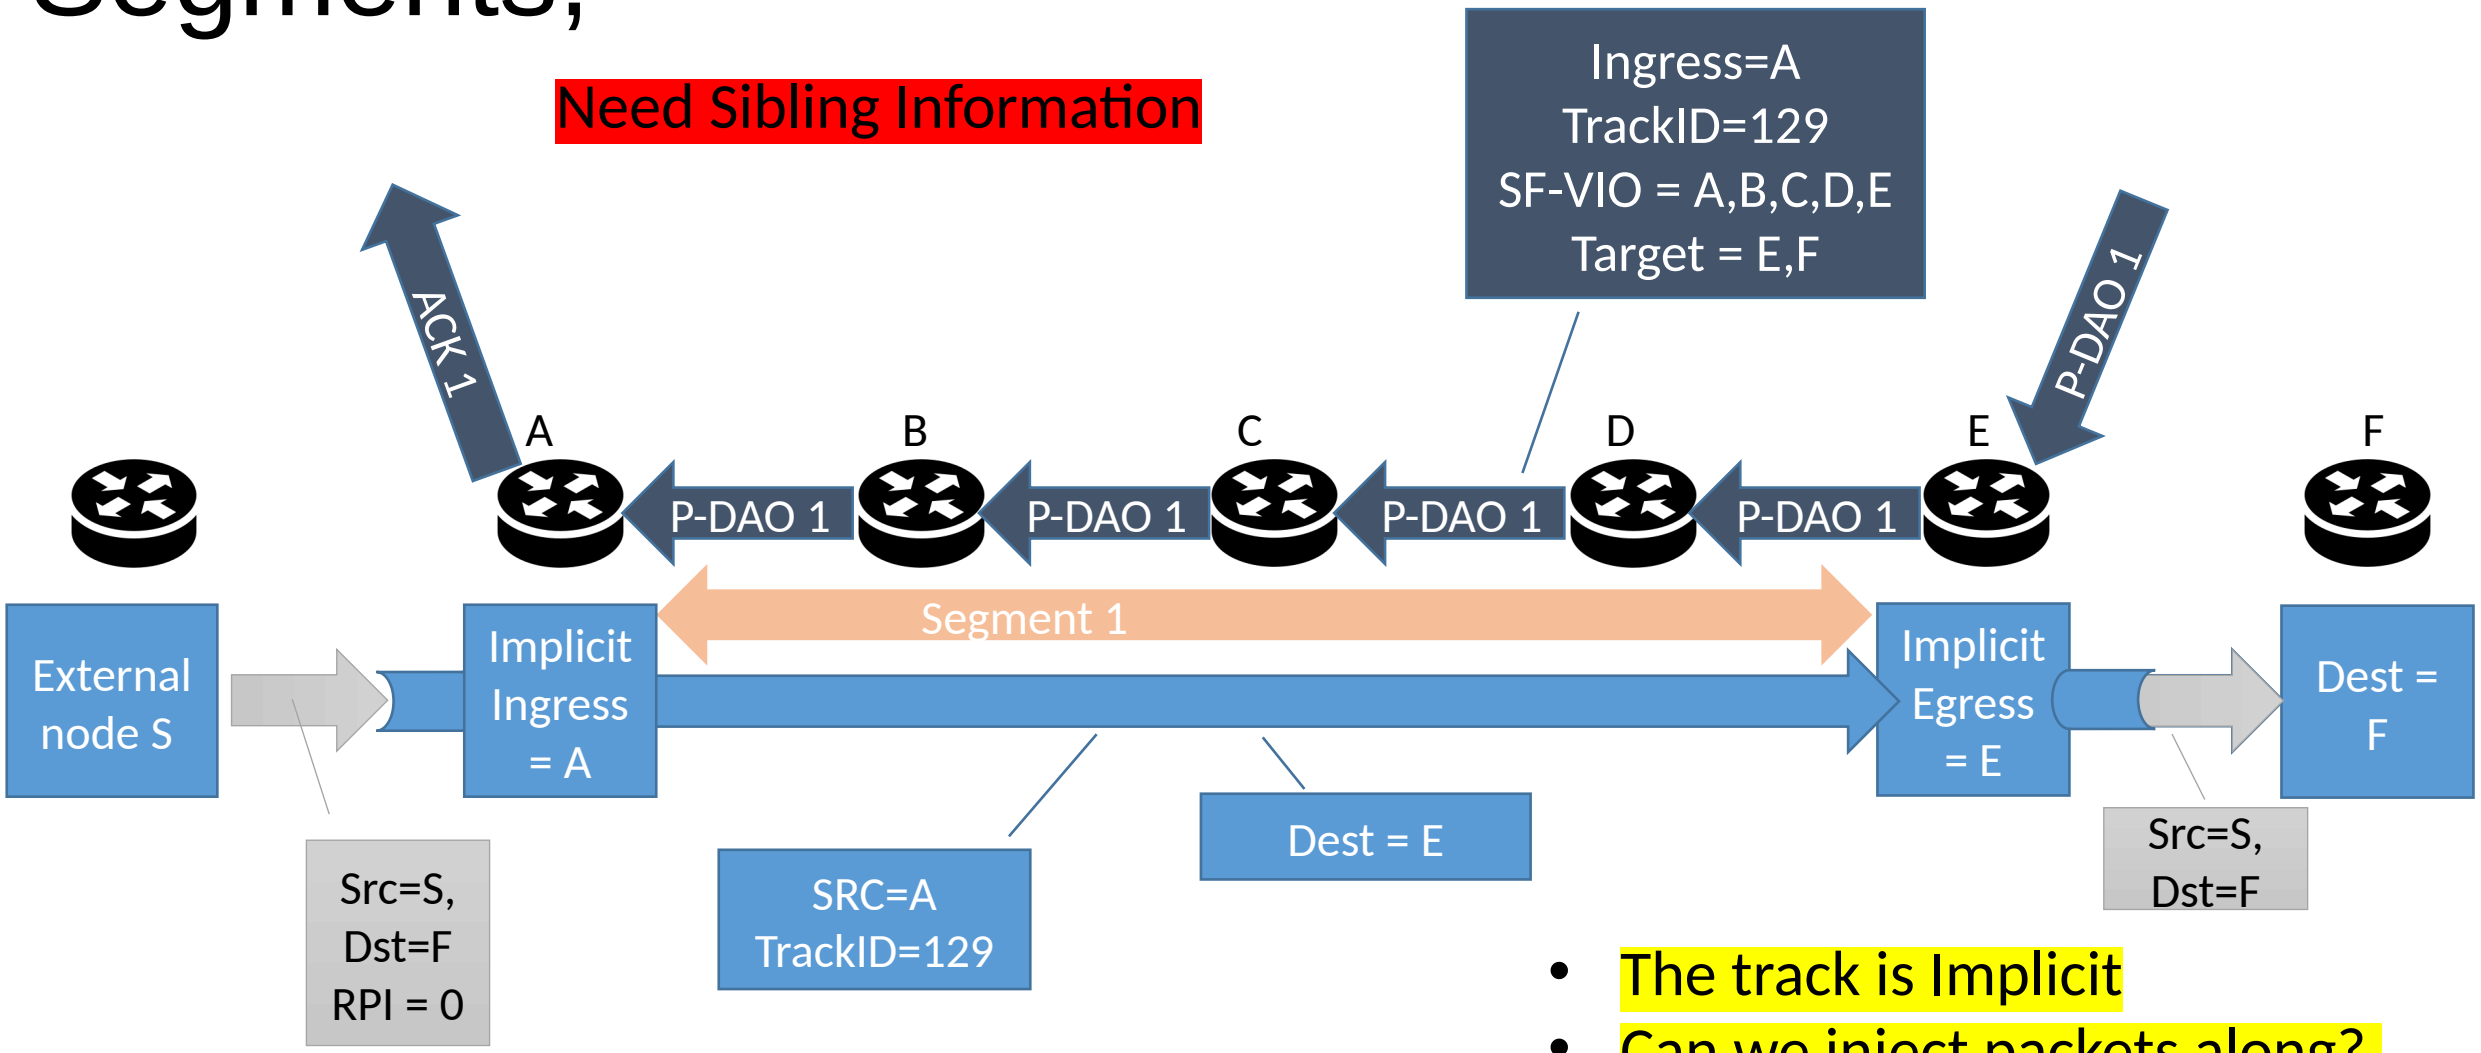

- The track is Implicit
- Can we inject packets along?

Profile 4: Strict NSM Explicit Track

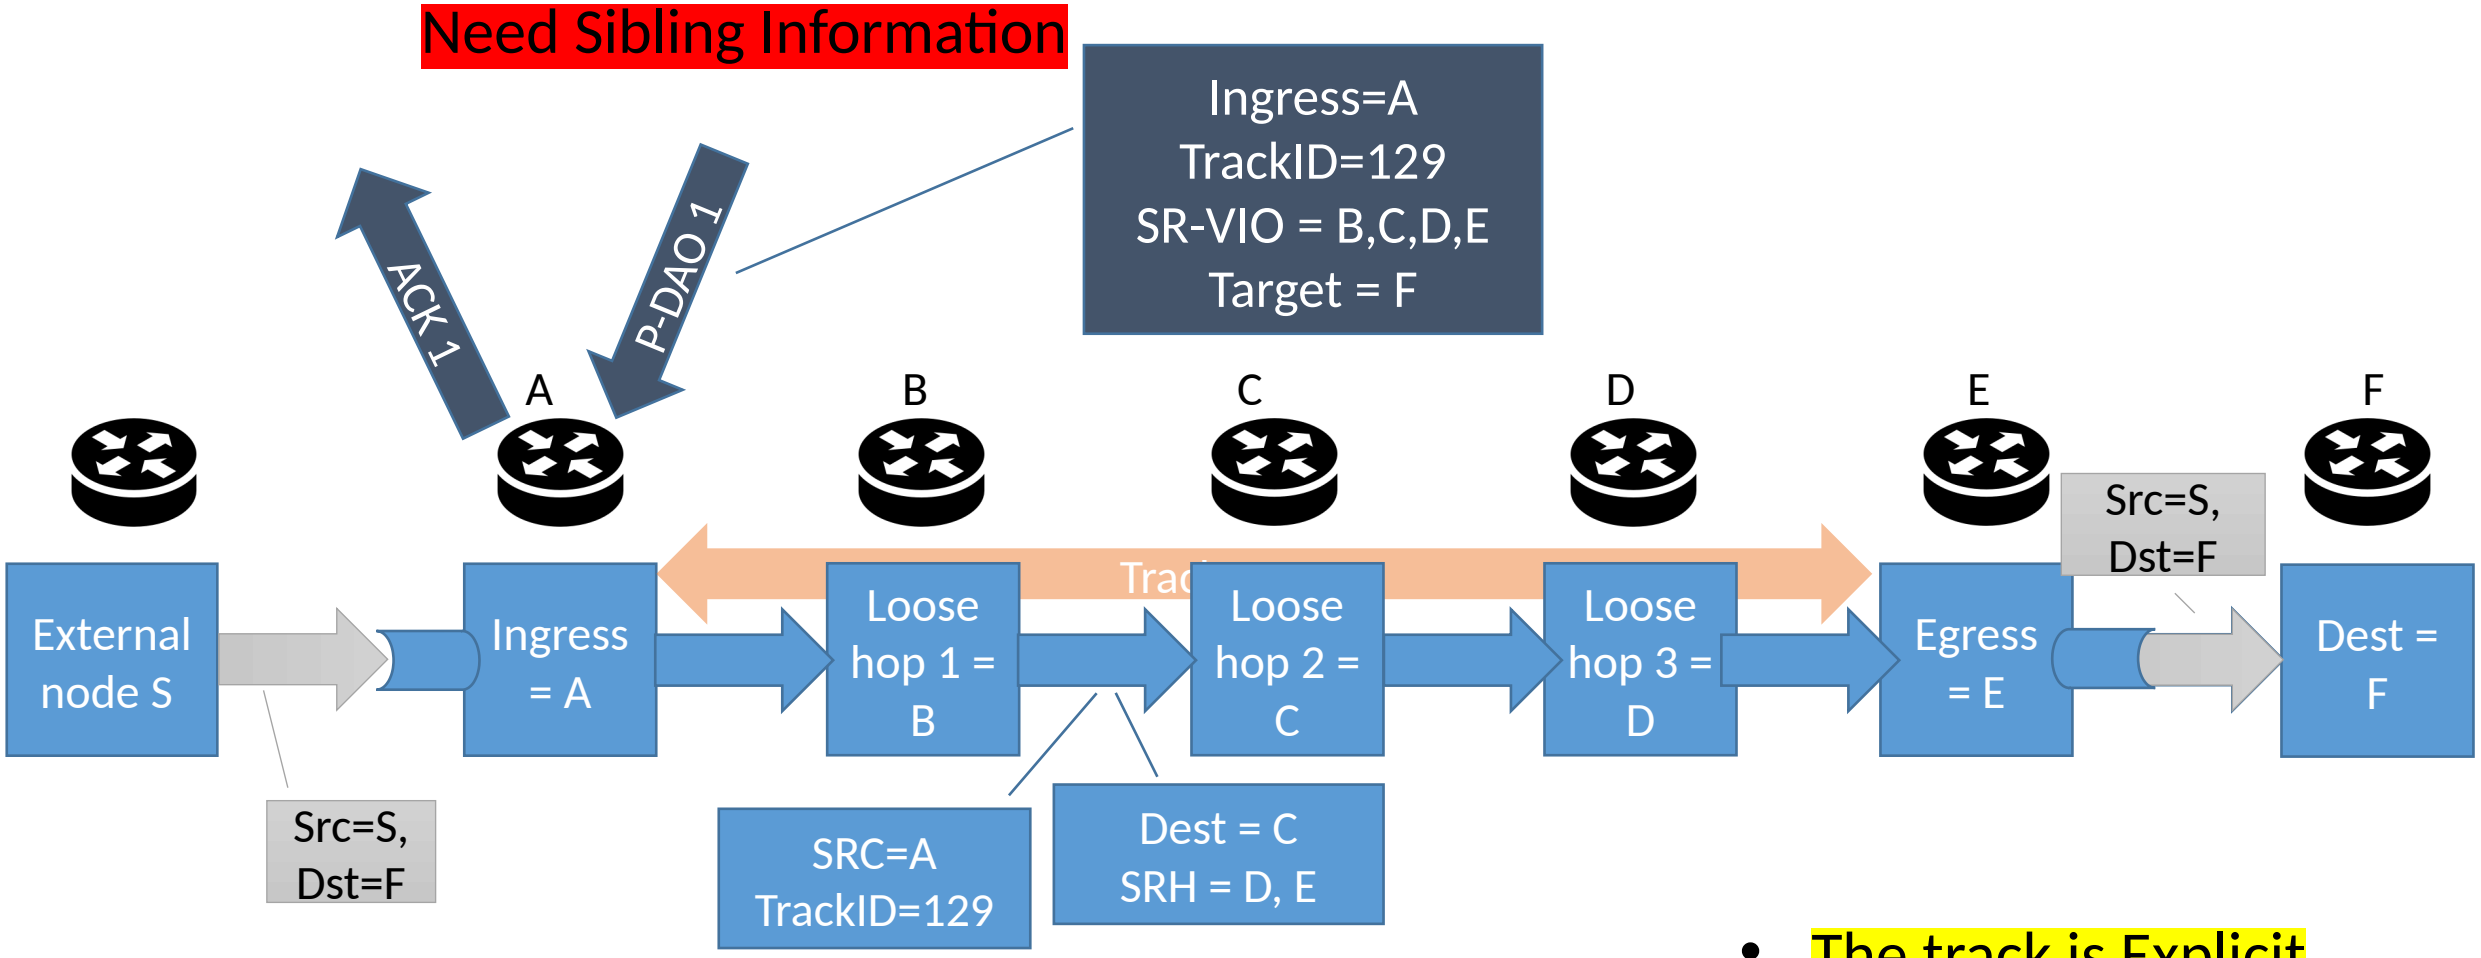

- The track is Explicit
- Same encap as profile 2

Profile 5: Compress SRH in Track with Strict SM Segments

Need Sibling Information

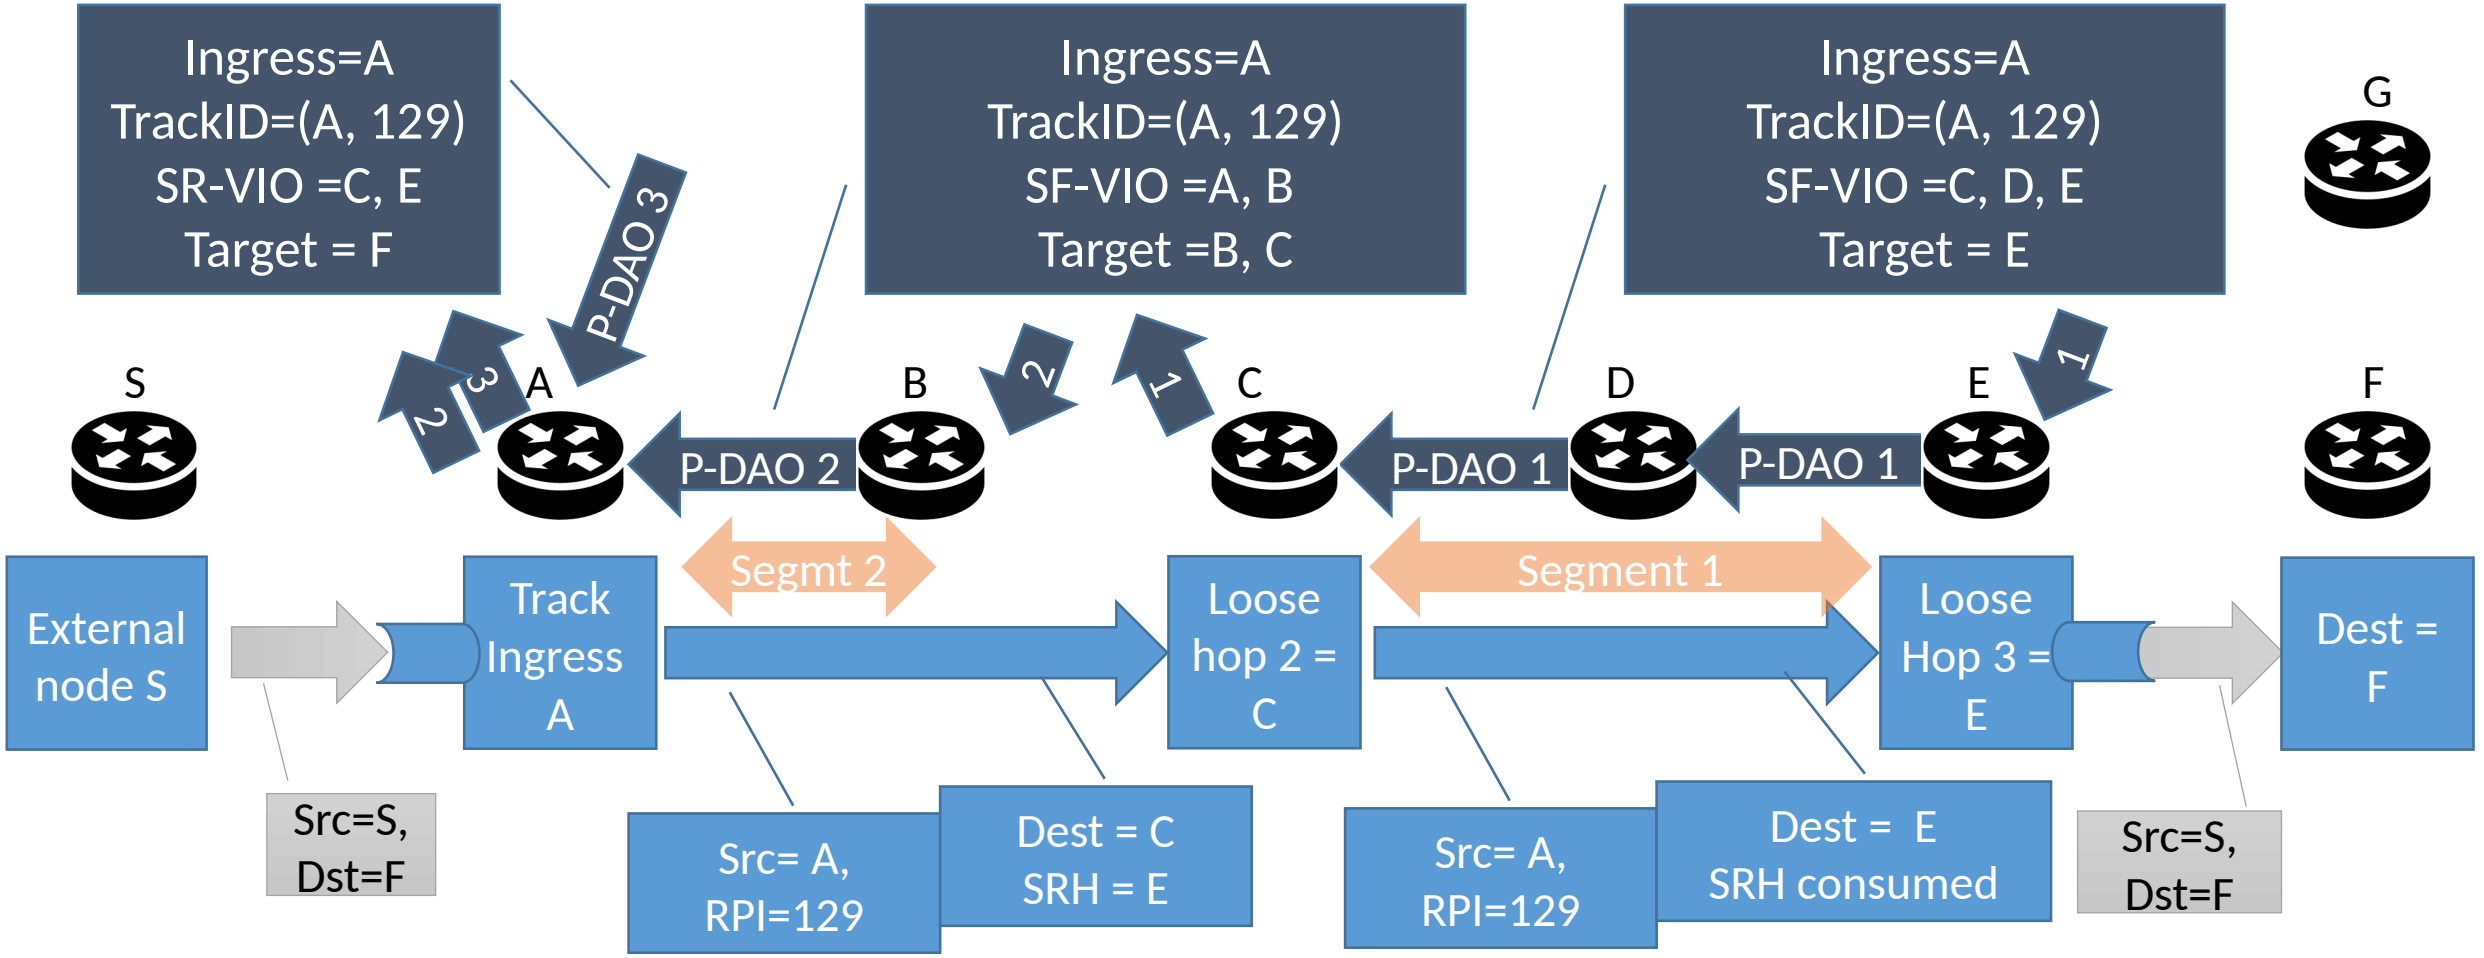

- Same as Profile 1, but for Track

Profile 6:
Compress SRH in Track with NSM Tracks (Recursive?)

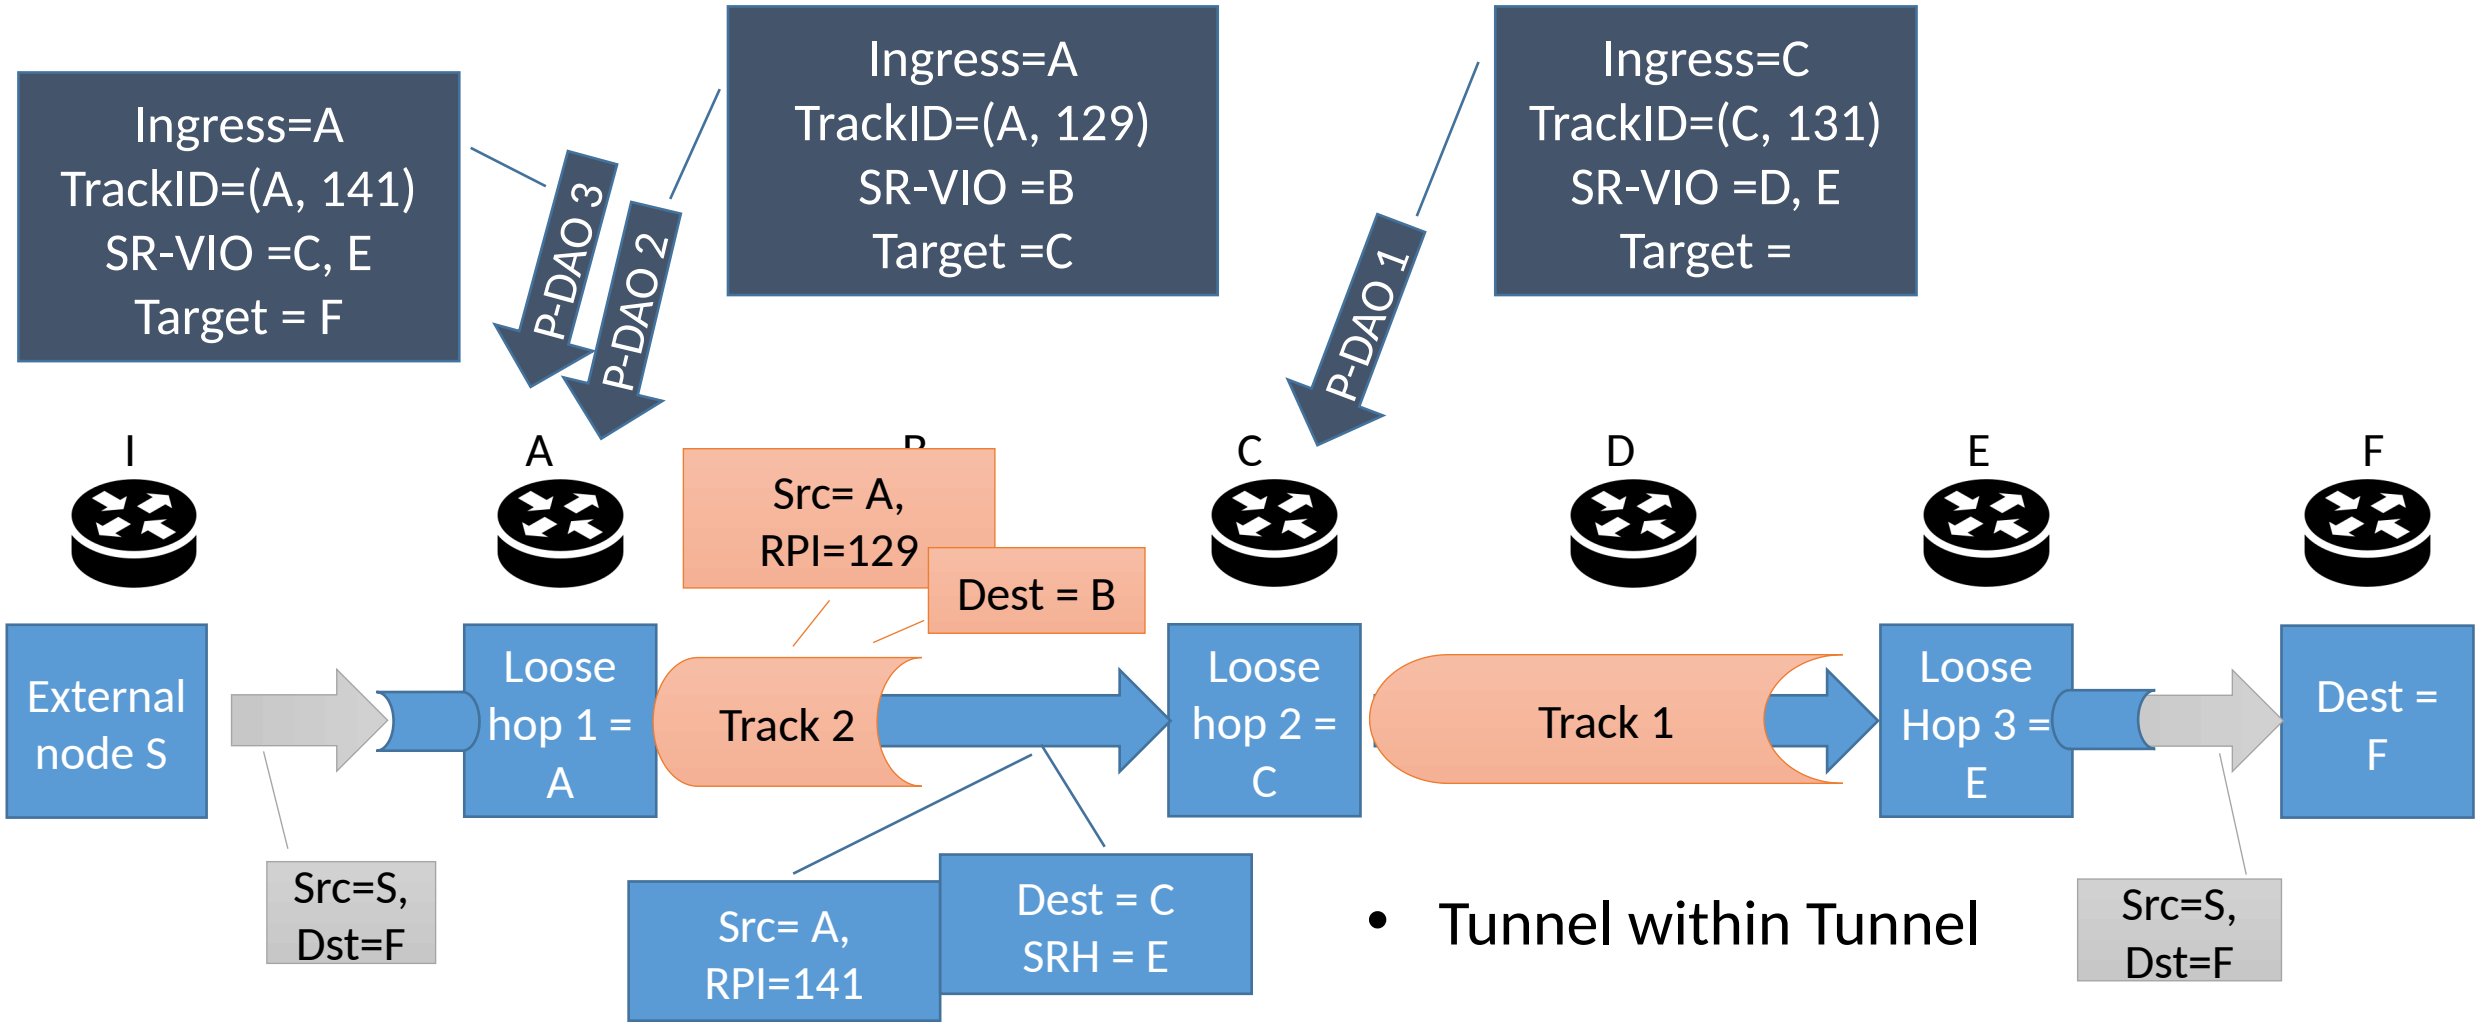

- Tunnel within Tunnel

Huimin’s comments / suggestions

- Lifetime unit: ReqLifetime, Track lifetime, and Segment Lifetime are defined as 8 bits. And their lifetime Unit is obtained from the DODAG configuration option. It will lead to inflexibility as all tracks in the PAN use the same lifetime unit. We propose to define lifetime unit separately for each track (for example adding a 2-bit flag to indicate second, minute, hour, day). Details can be discussed later.
- Now the TrackID has the same meaning as Local RplInstanceID. How does a node judge whether the received message is a P-DAO message or Local RPL instance DAO message? Is it possible to define a flag in the P-DAO message?
- The P-DAO track/segment is single-directional. I suggest to add the possibility for creating bi-directional segments/tracks. We can add a flag in the PDR message to indicate the requested track is single-directional or bi-directional.
- I suggest to add a flow of message exchanges for “PDR, PDR-ACK, P-DAO, P-DAO ACK” in the draft.

Other to be done

- Loop avoidance
- Who sends PDR? If it was destination, then it could select the trackID from its name space
- ND (RFC 8505) to maintain sibling neighbor state
- Be very specific if Ingress and Egress are listed in VIOs
  - Ingress to indicate which source address to use
  - Egress to build the full SRH 6LoRH

BAckup

Open Floor