Thursday 22/3 - 9h30-12

# IETF 101 ROLL

## Routing over Low-Power And Lossy Networks


**Chairs:**
Peter van der Stok
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

- Thursday 22/3 9h30-12: 150 minutes
- Remote Participation
  - Jabber Room: roll@jabber.ietf.org
  - Meetecho: http://www.meetecho.com/ietf101/roll
- Etherpad:
  - http://tools.ietf.org/wg/roll/minutes
- Audio Streaming:
- Minutes taker:
- Jabber Scribe:
- **Please sign blue sheets :-)**

# AGENDA - THURSDAY

| ROLL AGENDA IETF 101 | | |
|---|---|---|
| Thursday 22/3 9h30-12: 150 minutes | | |
| TIME | TOPIC | PRESENTER |
| 9:30 - 9:40 (10 min) | Introduction | Ines+Peter |
| 9:40 - 9:55 (15 min.) | Traffic-aware Objective Function<br>draft-ji-roll-traffic-aware-objective-function-00 | Georgio |
| 9:55 - 10:10 (15 min.) | RPL DAG Metric Container (MC) Node State and Attribute (NSA) object type,extension<br>draft-koutsiamanis-roll-nsa-extension-01 | Georgio |
| 10:10 - 10:40 (30 min.) | RPL Observations<br>draft-rahul-roll-rpl-observations-00 | Rahul |
| 10:40 - 11:25 (45 min.) | draft-ietf-roll-ccast-01<br>draft-thubert-roll-bier-01 | Carsten<br>Pascal |
| 11:25 - 11:55 (30 min.) | Enabling secure network enrollment in RPL networks<br>draft-richardson-6tisch-roll-enrollment-priority-00 | Michael |
| 11:55 - 12:00 (5 min.) | Open Mic | Everyone |

# Milestones

| Date | Milestone |
|---|---|
| Sep 2019 | Recharter WG or close |
| Jul 2019 | Initial submission of a Forwarder Selection Protocol for MPL to the IESG |
| Mar 2019 | Initial submission of a YANG model for MPL to the IESG |
| Jan 2019 | Initial submission of a proposal for Source-Route Multicast for RPL to the IESG |
| Dec 2018 | Initial submission of a proposal to augment DIS flags and options to the IESG |
| Aug 2018 | Initial submission of a root initiated routing state in RPL to the IESG |
| Jul 2018 | Initial submission of a solution to the problems due to the use of No-Path DAO Messages to the IESG |
| Jul 2018 | Initial submission of a reactive P2P route discovery mechanism based on AODV-RPL protocol to the IESG |
| Apr 2018 | Initial Submission of a proposal with uses cases for RPI, RH3 and IPv6-in-IPv6 encapsulation to the IESG |

# State of Active Internet-Drafts

| Draft | Status |
|---|---|
| **draft-ietf-roll-aodv-rpl-03** | IETF 101 |
| **draft-ietf-roll-dao-projection-03** | IETF 101 |
| **draft-ietf-roll-forw-select-00** | In pause |
| **draft-ietf-roll-useofrplinfo-22** | AD Evaluation, waiting for IESG Telechat |
| **draft-ietf-roll-dis-modifications-00** | To be continued |
| **draft-ietf-roll-mpl-yang-00** | Model to YANG doctors |
| **draft-ietf-roll-bier-ccast-01** | IETF 101 |
| **draft-ietf-roll-efficient-npdao-01** | IETF 101 |

# Related Internet-Drafts

- draft-richardson-6tisch-roll-enrollment-priority
  related to carrying new metrics to be fed into
  draft-richardson-6tisch-enrollment-enhanced-beacon

# IPRs

**Draft-ietf-roll-efficient-npdao-01: 2 IPRs**

**Draft-ietf-roll-dao-projection-03: 1 IPR**

**Comments?**

# Traffic-Aware Objective Function

**draft-ji-roll-traffic-aware-objective-function-01**

**Chenyang Ji**
**Remous-Aris Koutsiamanis: aris@ariskou.com**
**Georgios Z. Papadopoulos**
**Diego Dujovne**
**Nicolas Montavont**

**ROLL@IETF101**

# Standardisation Efforts

- **Objective Function → Preferred Parent**
  - **OF0**
  - **MRHOF**
  - ***Load balanced OF (LB-OF)***

# Problem statement

- **Using standard OFs (OF0, MRHOF) leads to unbalanced network:**
  - **Some nodes overloaded (forwarding)**

# Problem statement

- Using standard OFs (OF0, MRHOF) leads to unbalanced network:
  - Some nodes overloaded (forwarding)
  - Lower network and node lifetime

# Problem statement

- Using standard OFs (OF0, MRHOF) leads to unbalanced network:
  - Some nodes overloaded (forwarding)
  - Lower network and node lifetime
  - Higher packet losses (queueing)

# Problem statement

- **Using standard OFs (OF0, MRHOF) leads to unbalanced network:**
  - **Some nodes overloaded (forwarding)**
  - **Lower network and node lifetime**
  - **Higher packet losses (queueing)**
  - **Higher packet delay (queueing)**

# Examples (1)

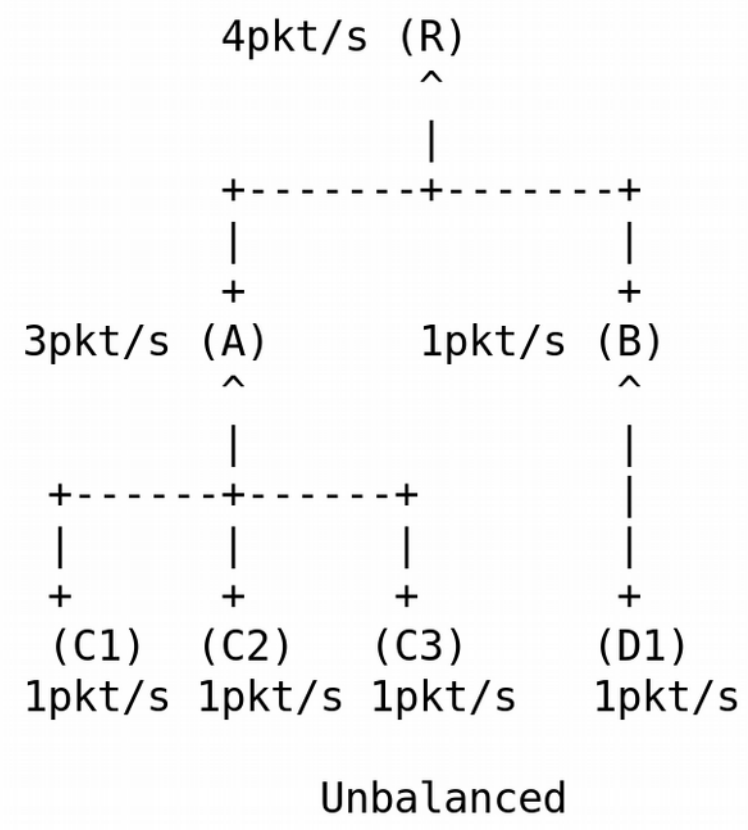


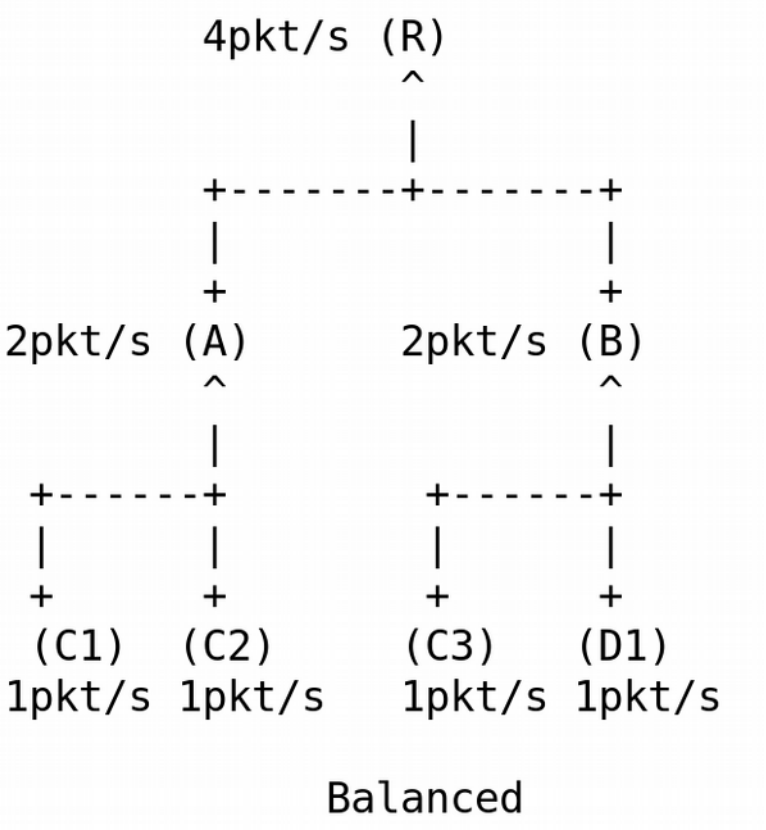


**Nodes with same TX requirements**

# Examples (2)

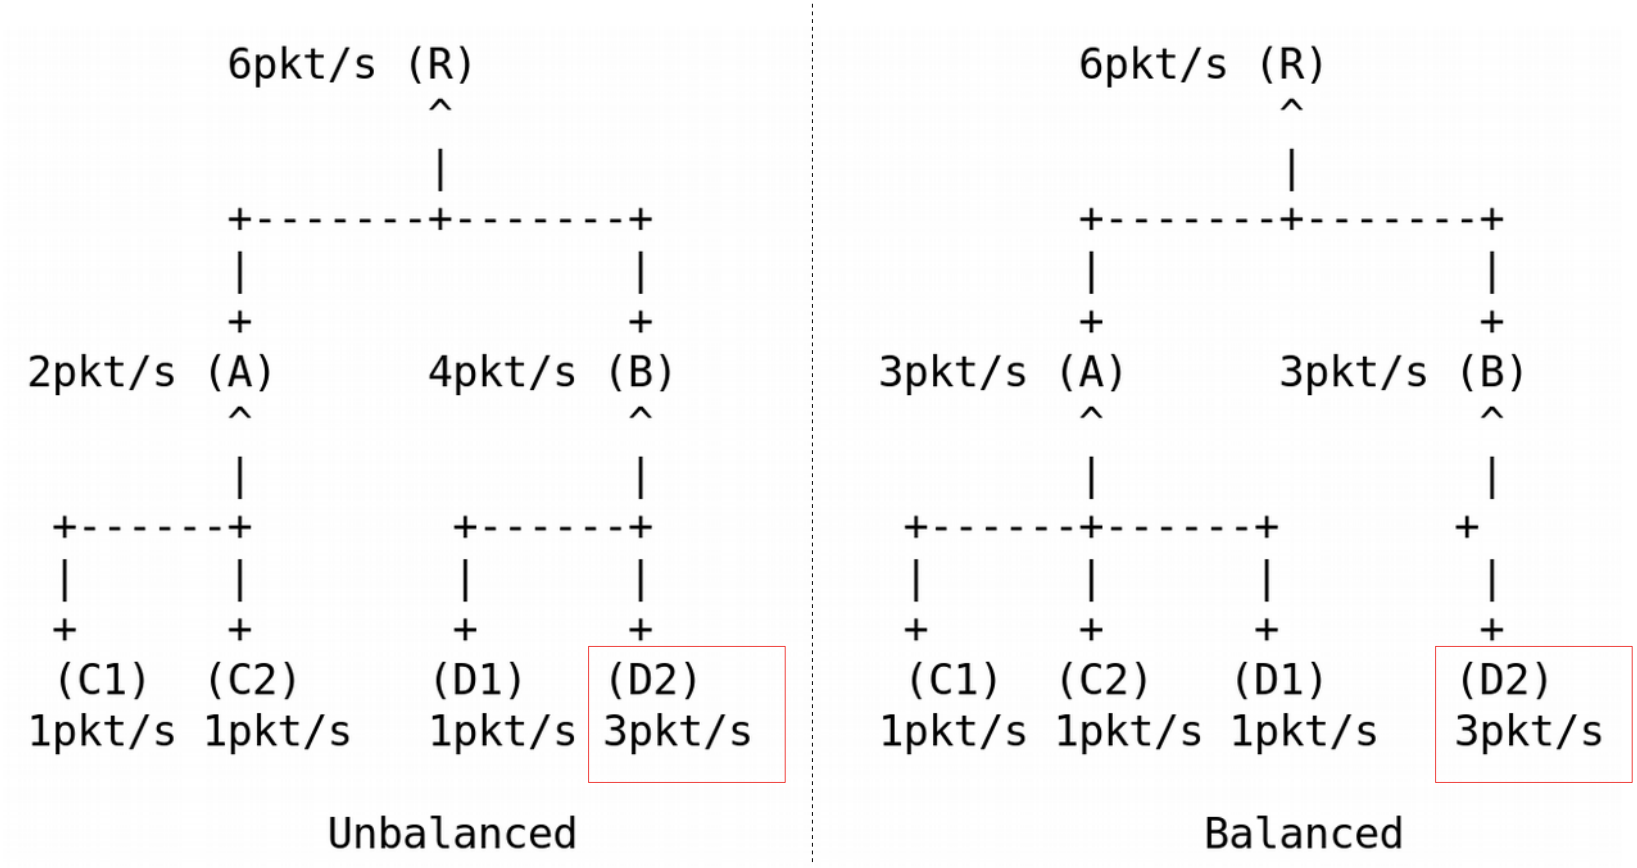


**Nodes with different TX requirements**

# Traffic-Aware OF

- New metric
  - Packet Transmission Rate (PTR) per node
    - Alternatively, cumulative
  - Data packets sent per time unit
    - Alternatively, octets per time unit?
- Traffic-Aware OF
  - Least PTR → Preferred parent

# DIO Format Example

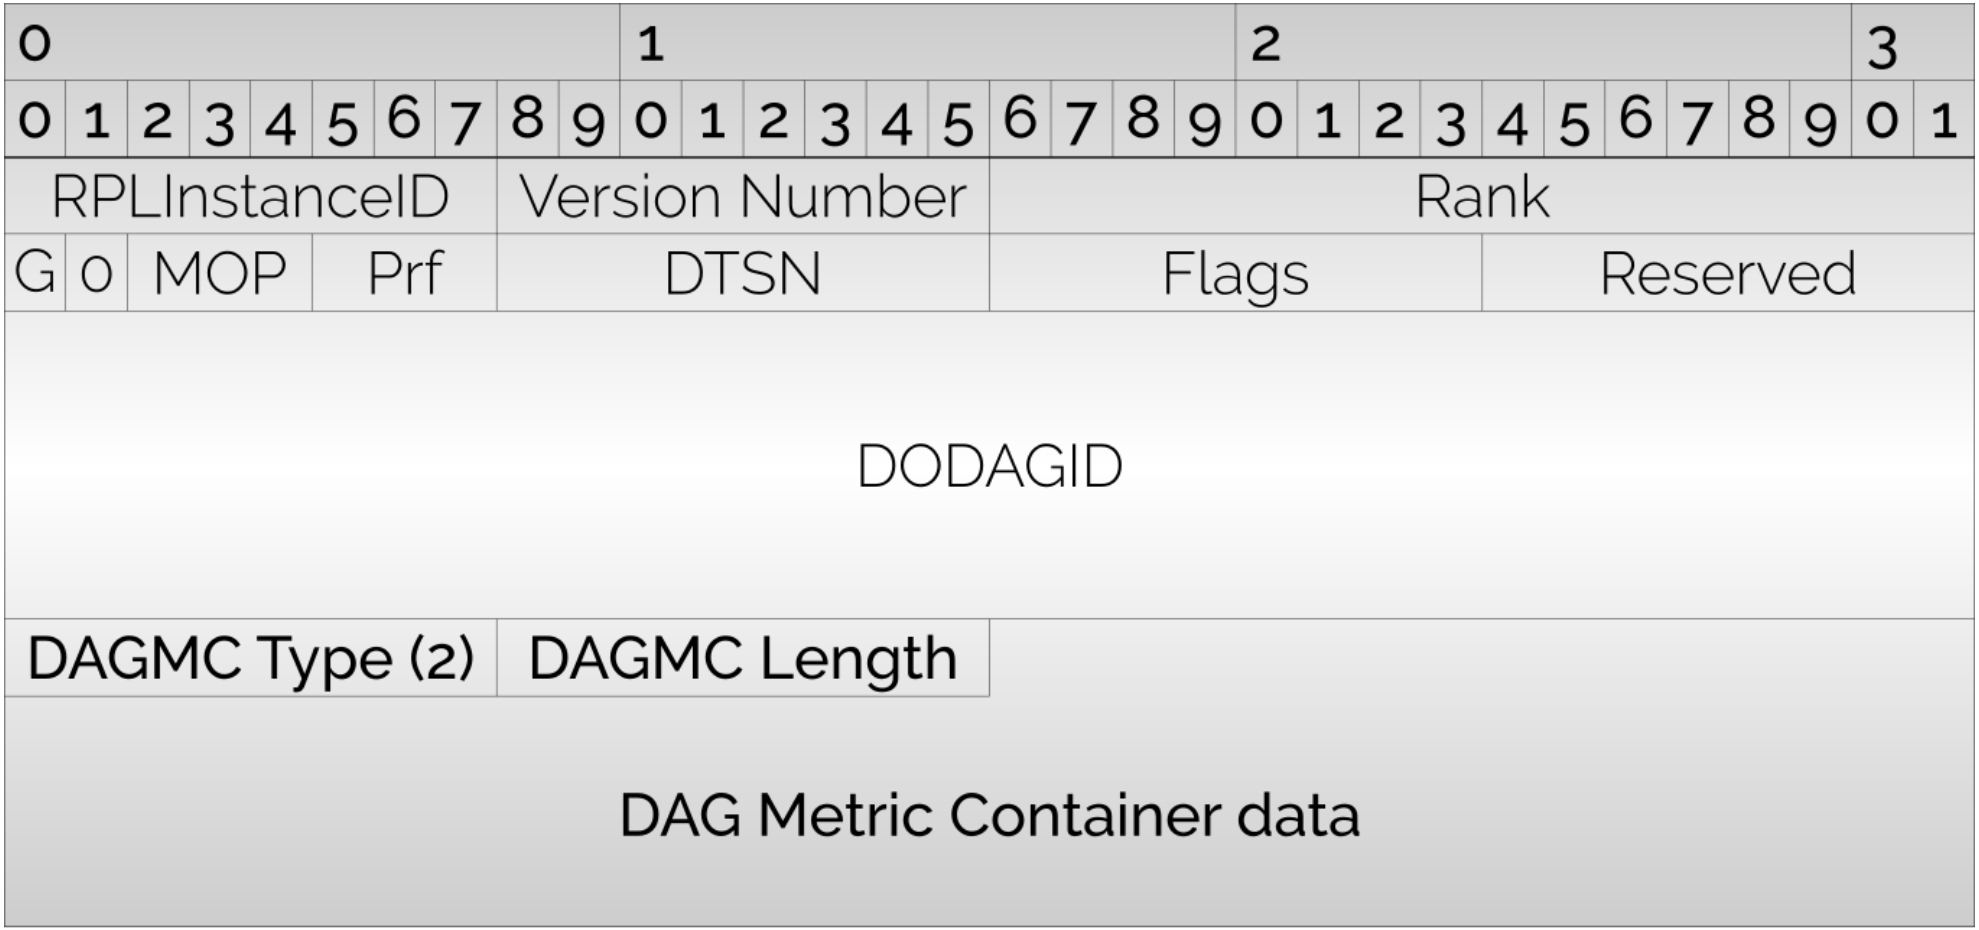

# Packet Transmission Rate

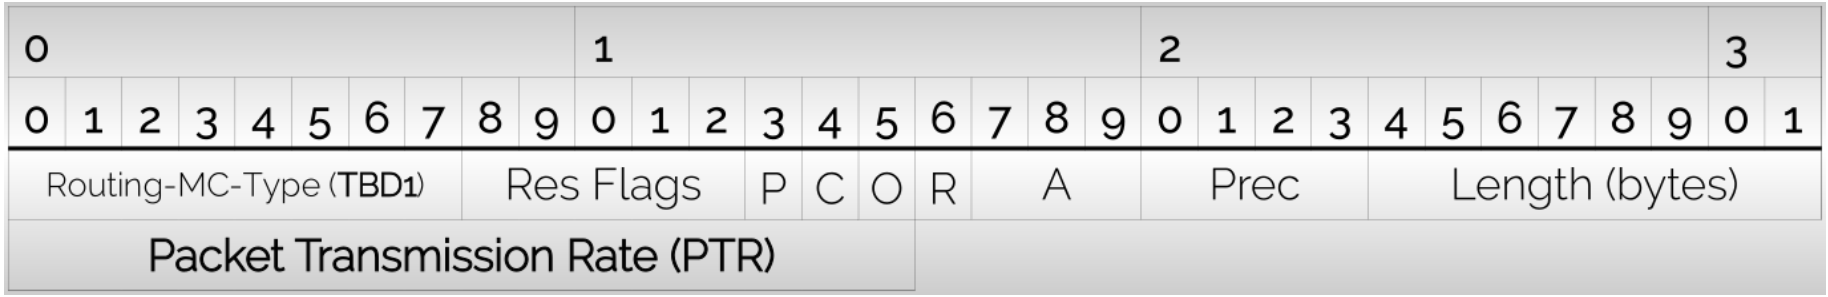


- **Node Metric Object (PTR)**
  - **2 octets – unsigned integer**

# Preliminary Results - Setup

- Contiki OS / COOJA
- 802.15.4 TSCH
- Sink node : node 1 (red)
- Senders : black nodes
  - Node 2, Node 8: 1 pkt / 1 sec
  - Rest node i: 1 pkt / i sec
- No MAC retransmissions
- Simulation time : 600s
- Scheduler : Orchestra
- Compare with
  - MRHOF
  - Load balancing OF

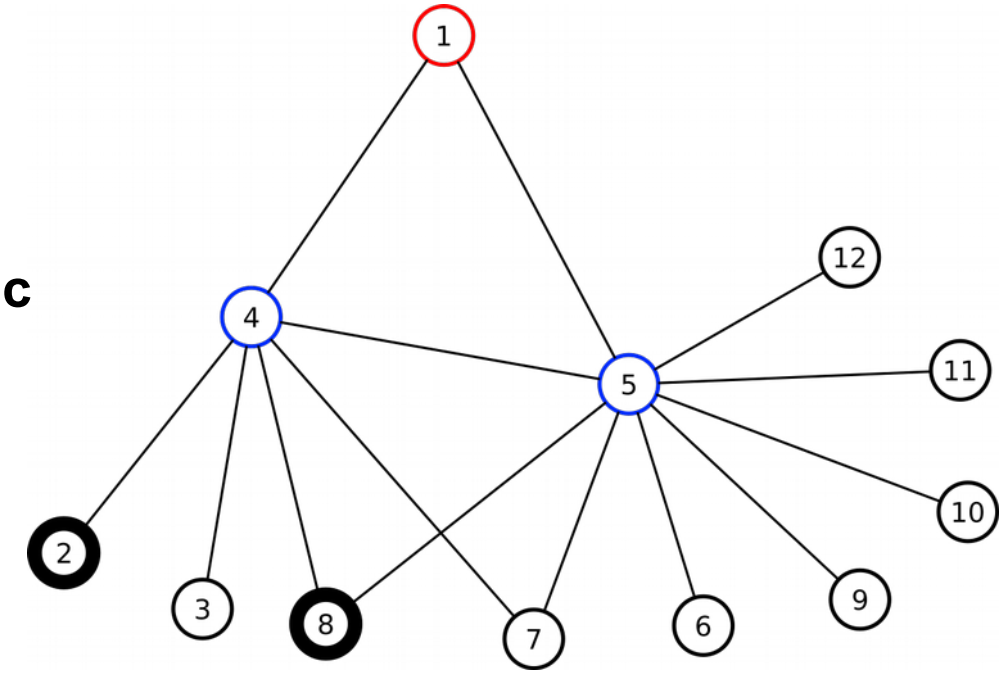

# Preliminary Results - PDR

- **Some improvements, especially with lower link qualities**

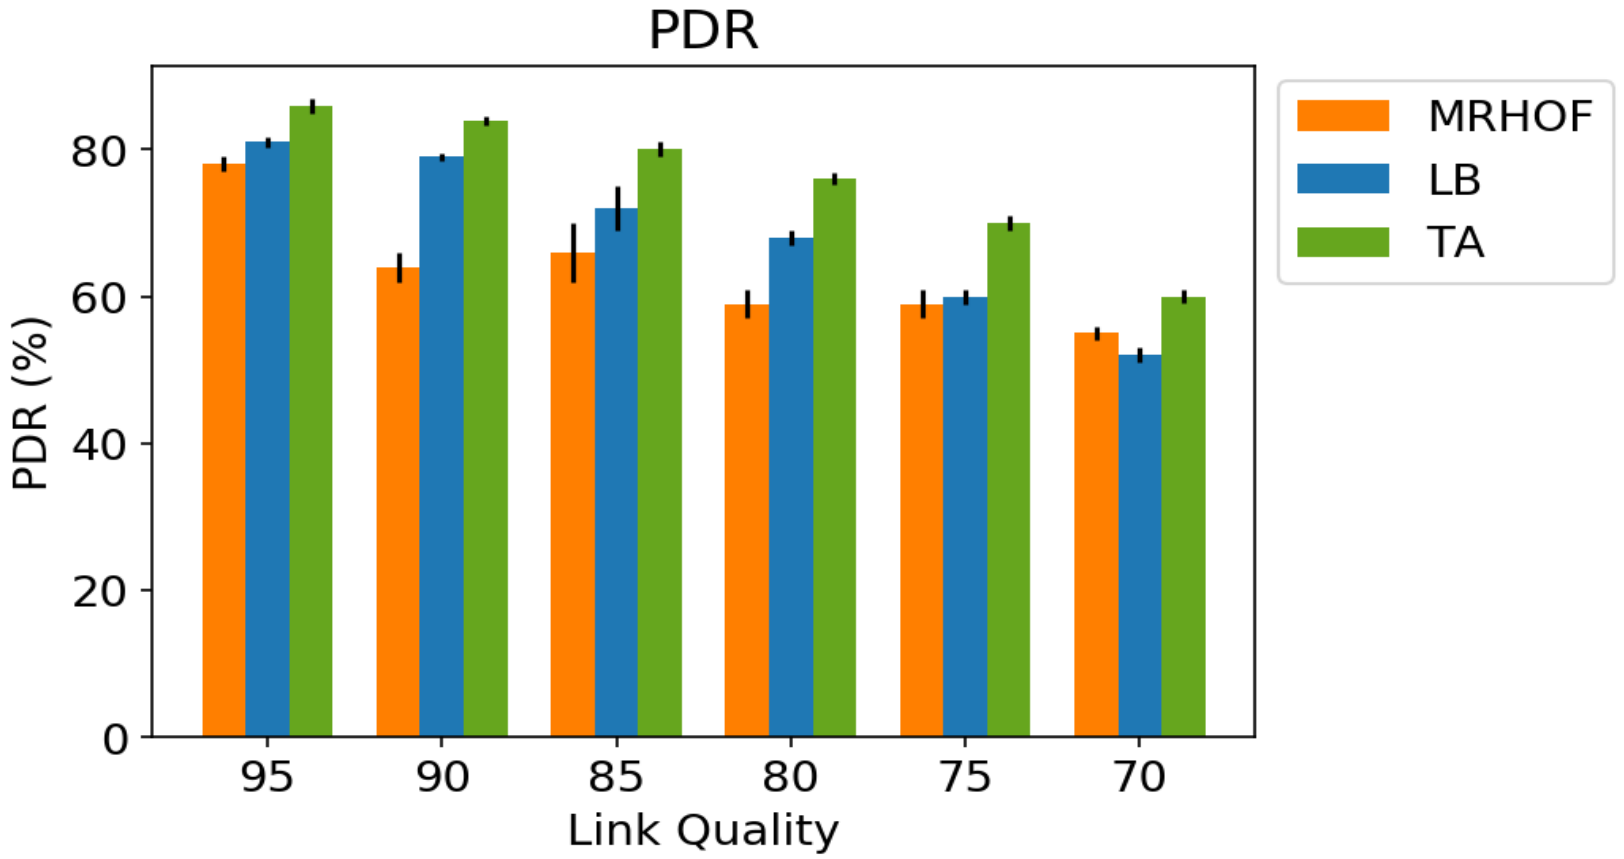

# Preliminary Results – Parent Changes

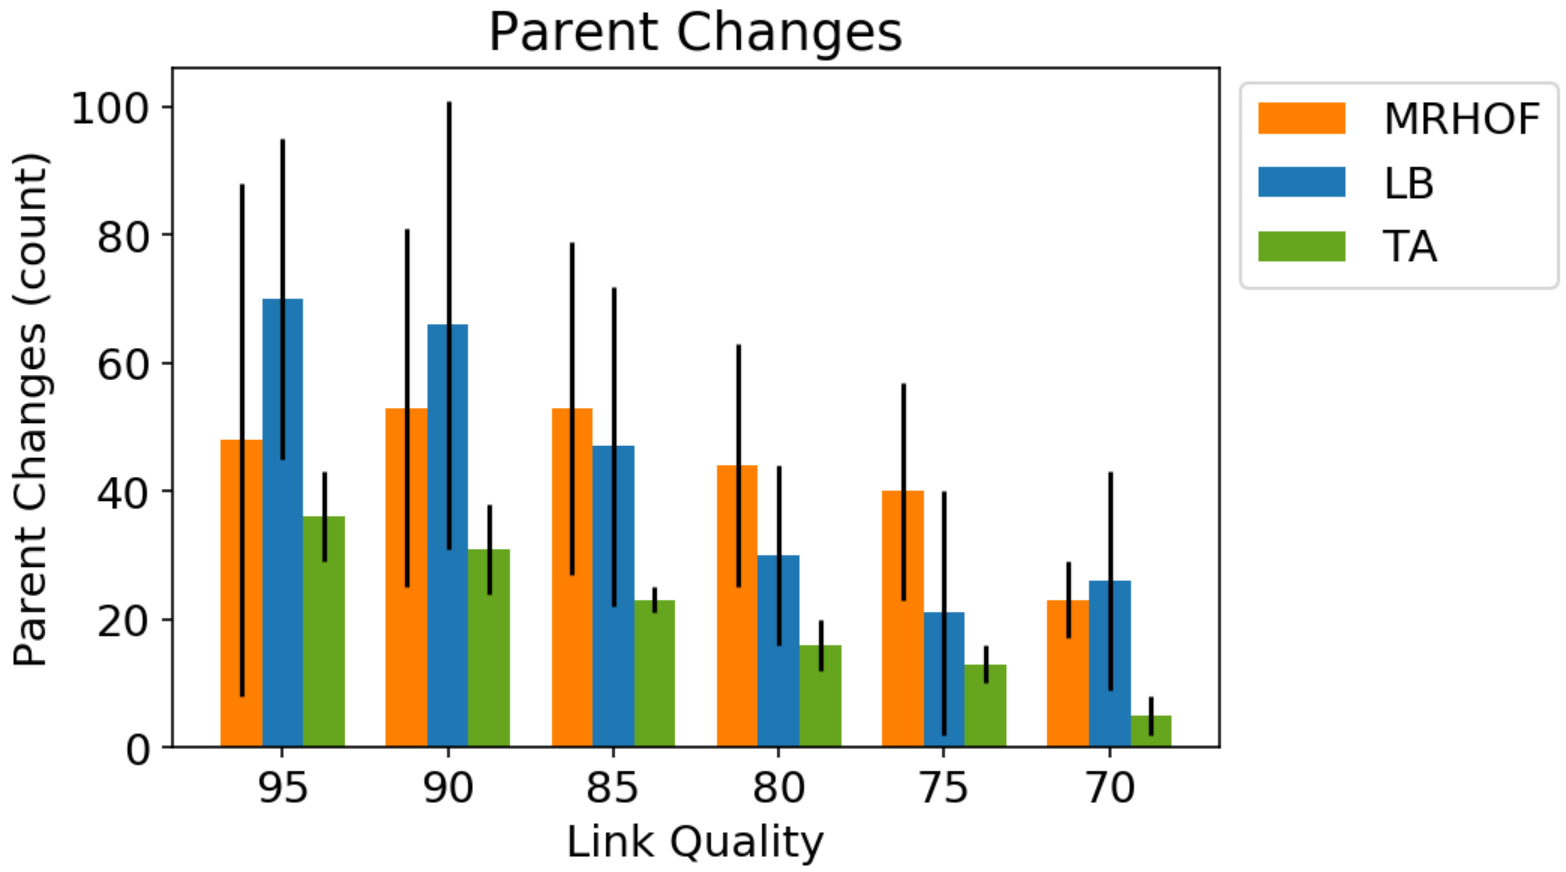

# Preliminary Results – DIOs

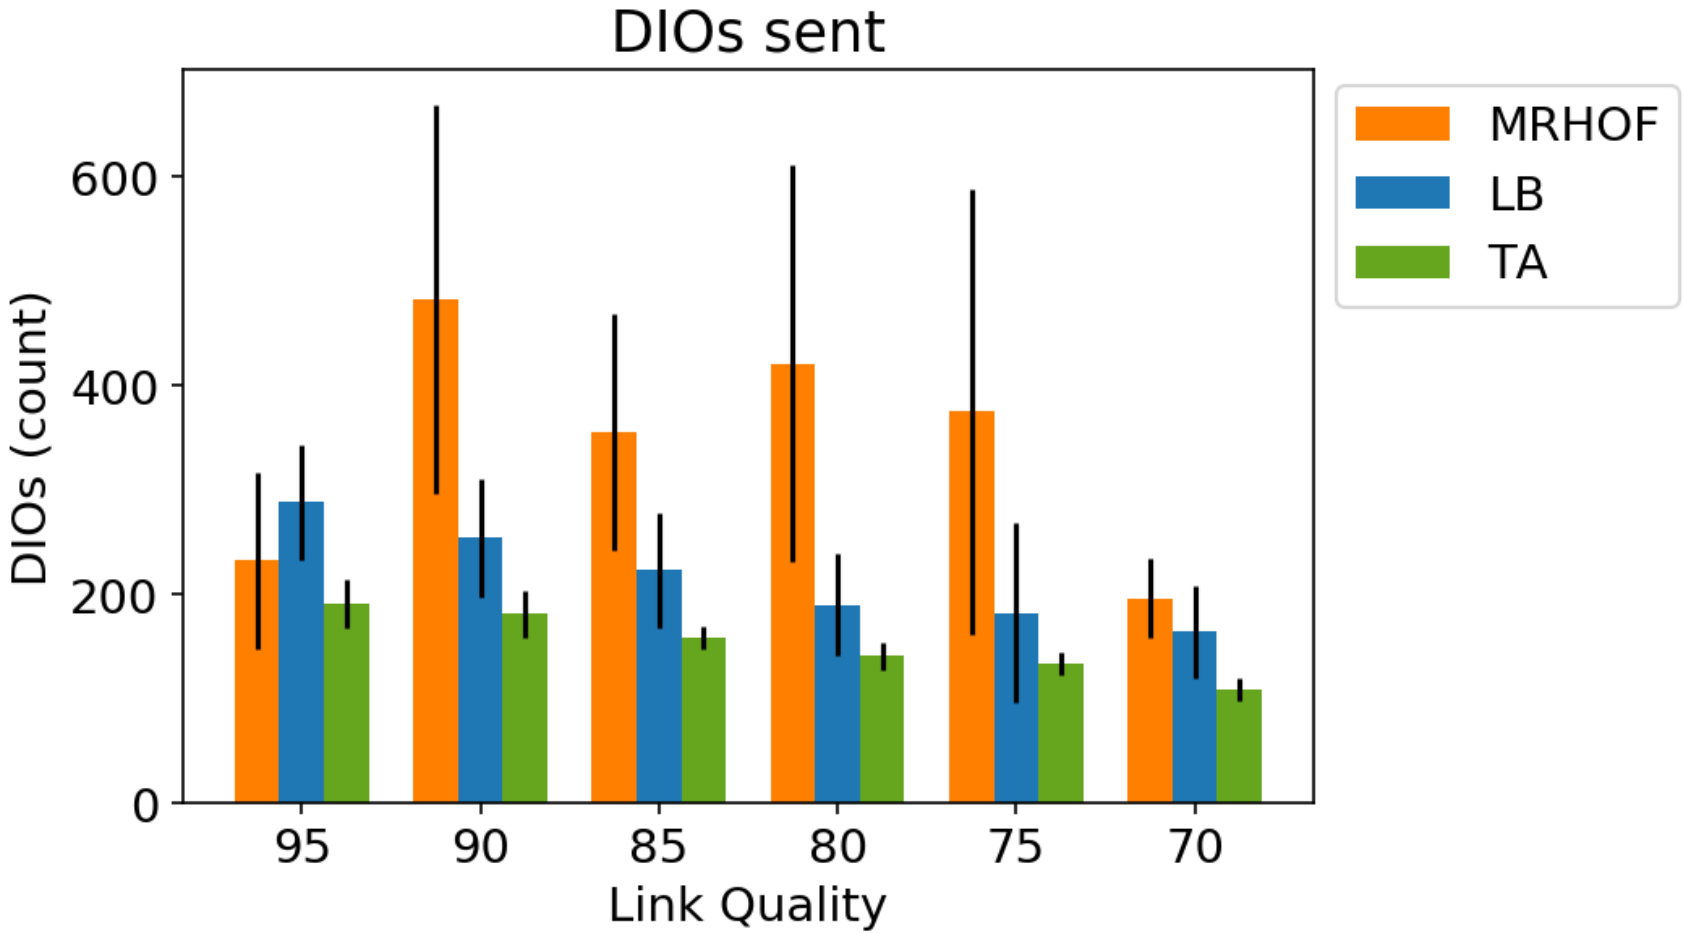

# Issues

- Homogeneity assumptions
  - Same data unit implied
    - Part of data sent?

# Issues

- Homogeneity assumptions
  - Same data unit implied
    - Part of data sent?
  - Same time unit implied
    - Part of data sent?

# Issues

- Homogeneity assumptions
  - Same data unit implied
    - Part of data sent?
  - Same time unit implied
    - Part of data sent?
  - Same forwarding max capacity
    - Normalize PTR to capacity, or
    - Send max capacity in data

# Issues

- Homogeneity assumptions
  - Same data unit implied
    - Part of data sent?
  - Same time unit implied
    - Part of data sent?
  - Same forwarding max capacity
    - Normalize PTR to capacity, or
    - Send max capacity in data
  - Same energy consumption

# Issues

- Homogeneity assumptions
  - Same data unit implied
    - Part of data sent?
  - Same time unit implied
    - Part of data sent?
  - Same forwarding max capacity
    - Normalize PTR to capacity, or
    - Send max capacity in data
  - Same energy consumption
- Locality of information
  - Full path PTR more useful?

# Feedback

- cf. 6TiSCH → Addition to EB/IE?
- Other / overlapping concerns?

**Many thanks to Jianqiang Hou (Derek) for comments**

# RPL DAG Metric Container (MC) Node State and Attribute (NSA) object type extension

**draft-koutsiamanis-roll-nsa-extension-01**


**Remous-Aris Koutsiamanis: aris@ariskou.com**
**Georgios Z. Papadopoulos**
**Nicolas Montavont**
**Pascal Thubert**


**ROLL@IETF101**

# New since -00

- Some editorial changes (diagrams, wording)
- Fully specify DAG MC field semantics
- Wireshark dissectors

# Toward Determinism

- **Reliable communication;**

# Toward Determinism

- **Reliable communication;**
- **Low jitter performance;**

# Toward Determinism

- Reliable communication;
- Low jitter performance;
- Packet Replication Elimination
  - Replication
  - Elimination
  - Promiscuous overhearing

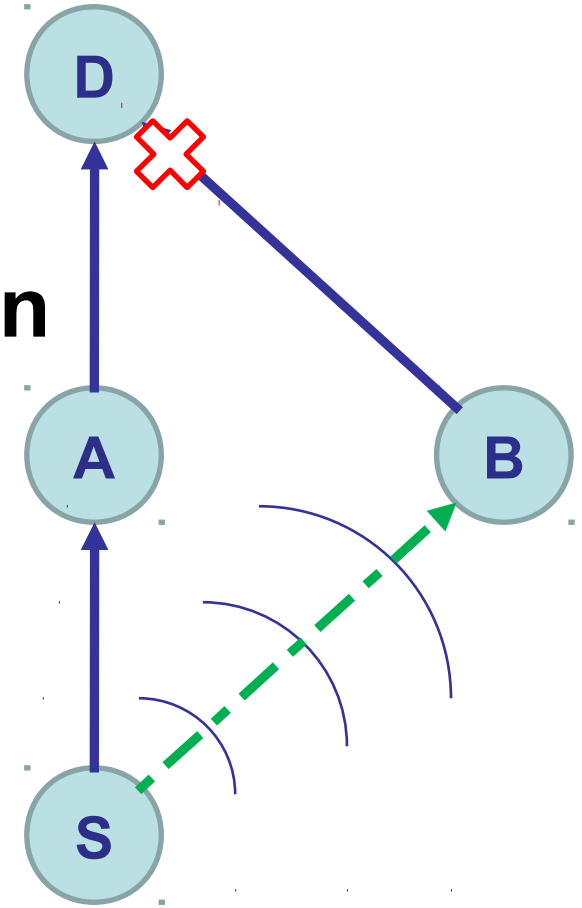

# Requirements [1]

- **Alternative Parent Selection;**
  - RPL DODAG Information Object (DIO) message format SHOULD be extended
  - routing protocol should be extended to allow for 6TiSCH nodes to select AP(s)
- **Promiscuous Overhearing;**
  - 6top Protocol should be extended to allow a cell reservation with two receivers
  - 6P ADD Request Format should be transmitted either twice or once in multicast
- **Cells without ACKs;**
  - only one parent MUST acknowledge the data packet
  - Or an efficient way for double ACKS
- **Packet Elimination.**
  - Tagging Packets for Flow Identification

[1] G. Z. Papadopoulos, N. Montavont, and P. Thubert, "Exploiting Packet Replication and Elimination in Complex Tracks in 6tisch LLNs," Working Draft, IETF Secretariat, Internet-Draft draft-papadopoulos-6tisch-pre-reqs-00, July 2017.

# Requirements [1]

- **Alternative Parent Selection;**
  - **RPL DODAG Information Object (DIO) message format SHOULD be extended**
  - **routing protocol should be extended to allow for 6TiSCH nodes to select AP(s)**
- **Promiscuous Overhearing;**
  - **6top Protocol should be extended to allow a cell reservation with two receivers**
  - **6P ADD Request Format should be transmitted either twice or once in multicast**
- **Cells without ACKs;**
  - **only one parent MUST acknowledge the data packet**
  - **Or an efficient way for double ACKS**
- **Packet Elimination.**
  - **Tagging Packets for Flow Identification**

**[1] G. Z. Papadopoulos, N. Montavont, and P. Thubert, "Exploiting Packet Replication and Elimination in Complex Tracks in 6tisch LLNs," Working Draft, IETF Secretariat, Internet-Draft draft-papadopoulos-6tisch-pre-reqs-00, July 2017.**

# Alternative Parent Selection

**Draft enables Alternative Parent Selection mechanism**

**Allows selecting alternative parent with common *ancestor***

# DIO Format Example

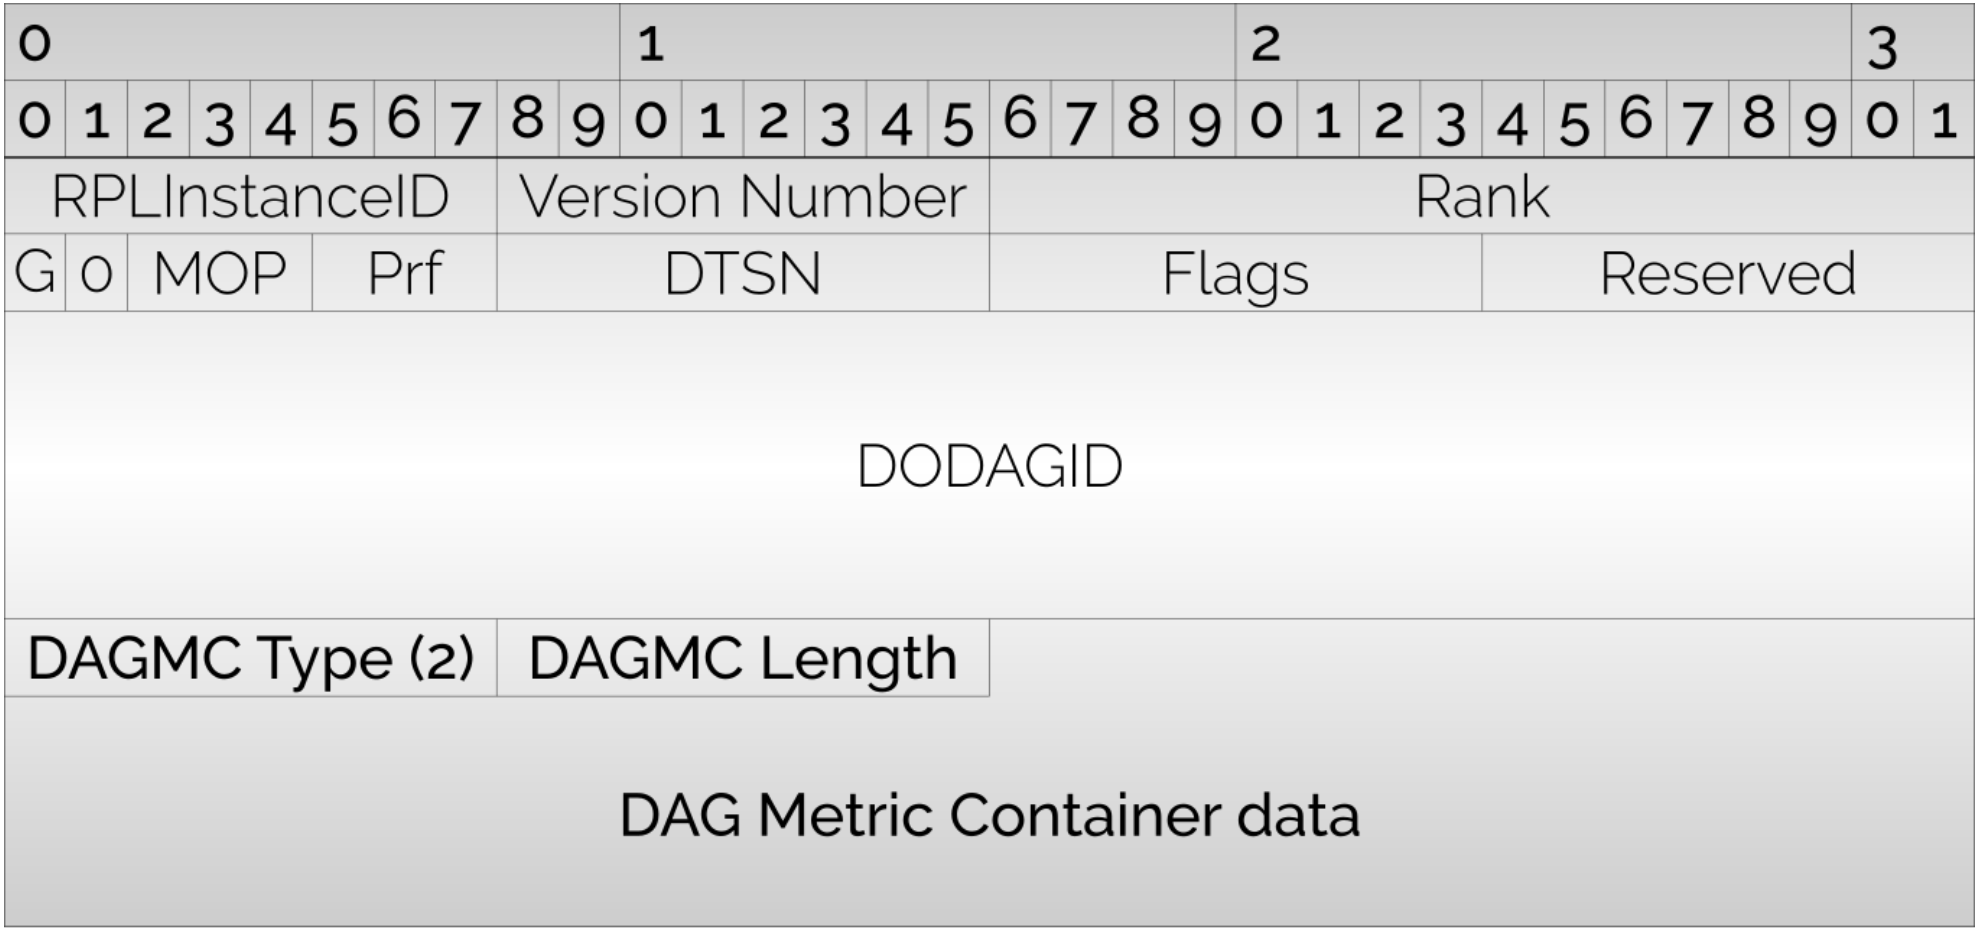

# MC/NSA Format Example (1)

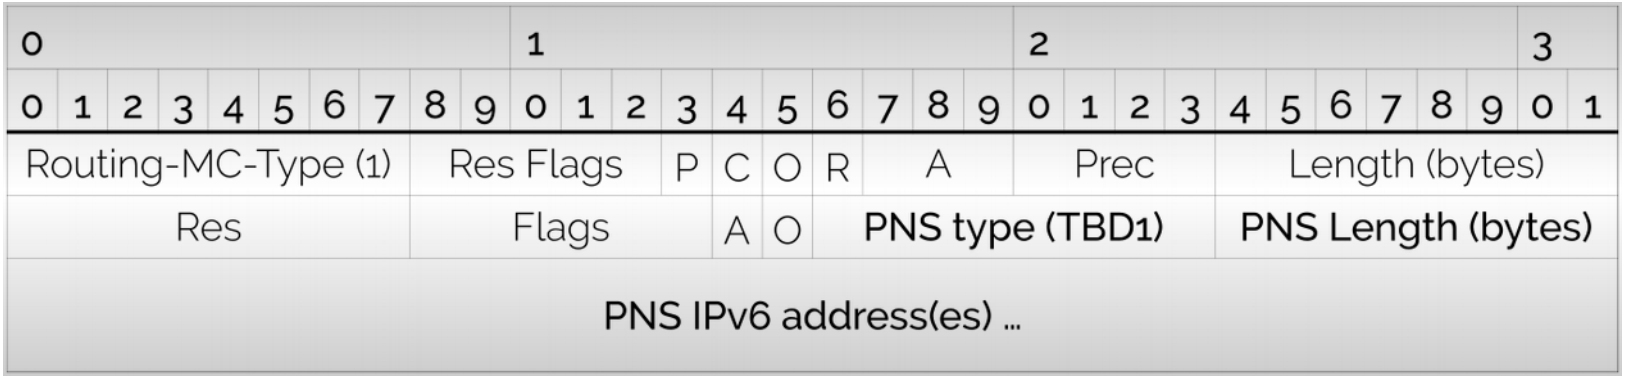


- **Parent Node Set (PNS)**
  - **NSA Option**
  - **PNS type = 1 (8 bits)**
  - **PNS Length = # of PNS addresses x IPv6 address size (8 bits)**
  - **PNS IPv6 addresses = 1 or more IPv6 addresses**

# MC/NSA Format Example (2)

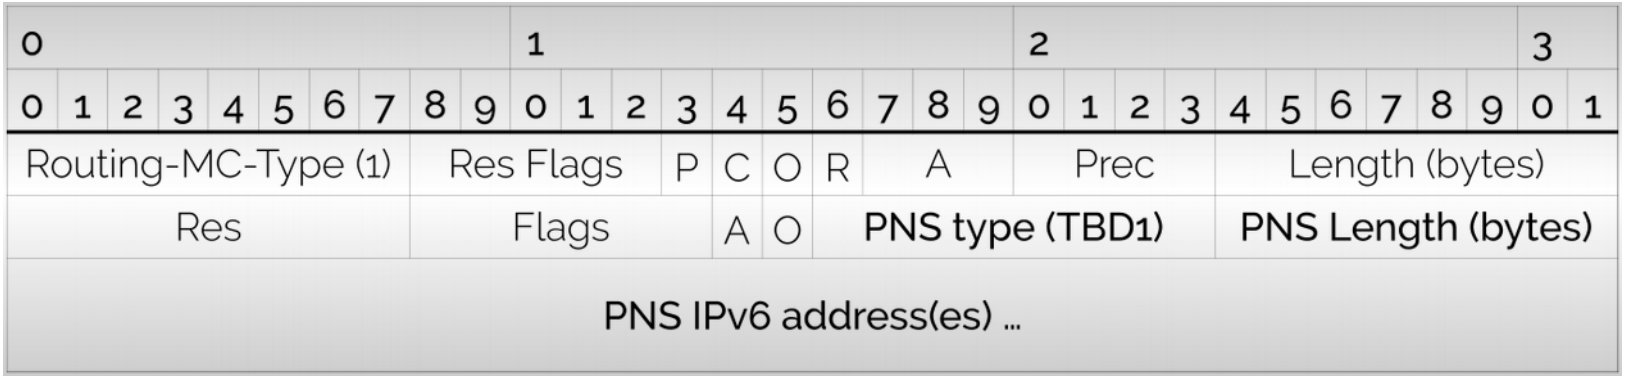


- **MC fields**
  - **P =0 (constraint)**
  - **C = 1 (constraint)**
  - **O = as normally, for optionality, as per RFC6550**
  - **R = 0 (constraint)**
  - **A = 0 (contraint)**
  - **Prec = as normally, for precedence, as per RFC6550**
- **NSA fields**
  - **A = as normally**
  - **O = as normally**

# MC/NSA Capture

- **Implemented in Contiki OS**
- **Modified Wireshark**

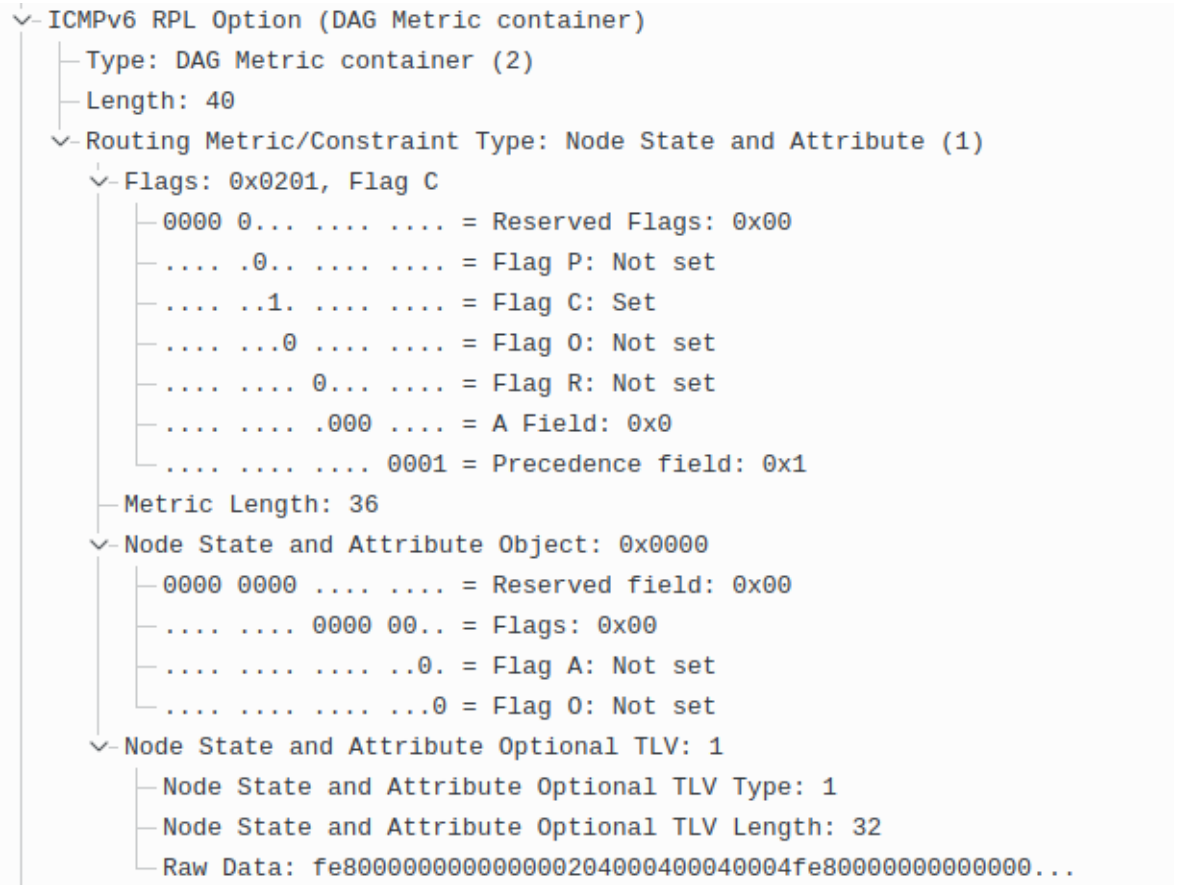

# Example

- RPL DAG
  - S→A→D
  - B→E

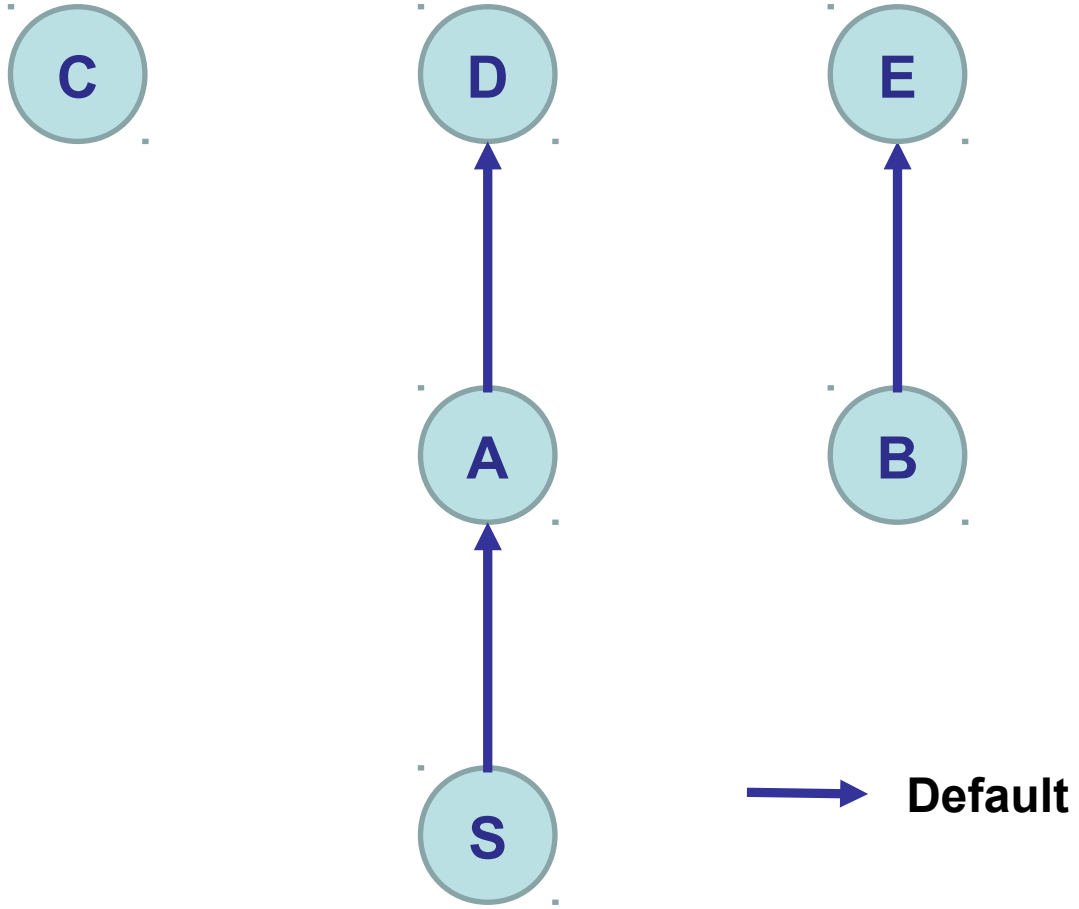

# Example

- **Parent set S:**
  - **{A, B}**

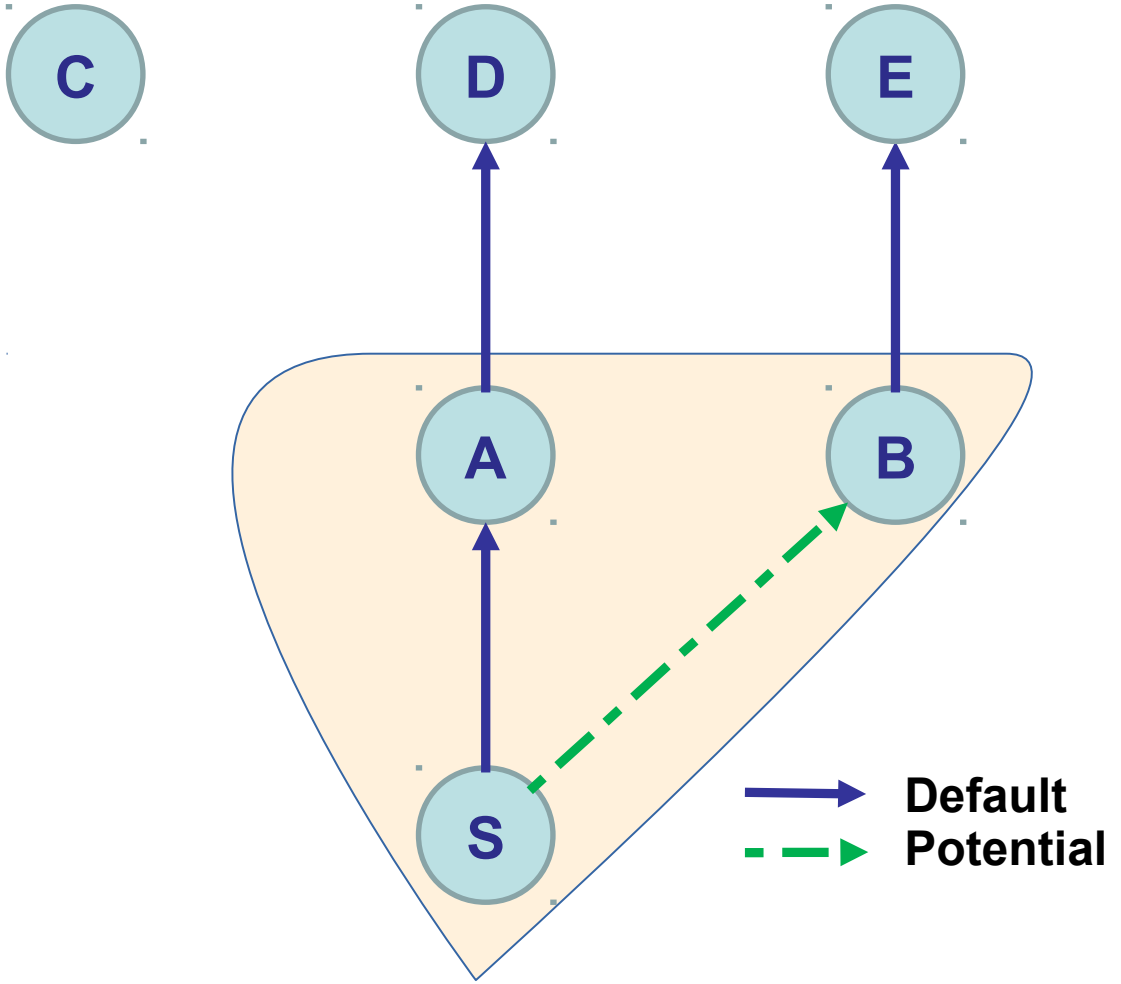

# Example

- **Parent set S:**
  - **{A, B}**
- **Parent Set A:**
  - **{D, C, E}**

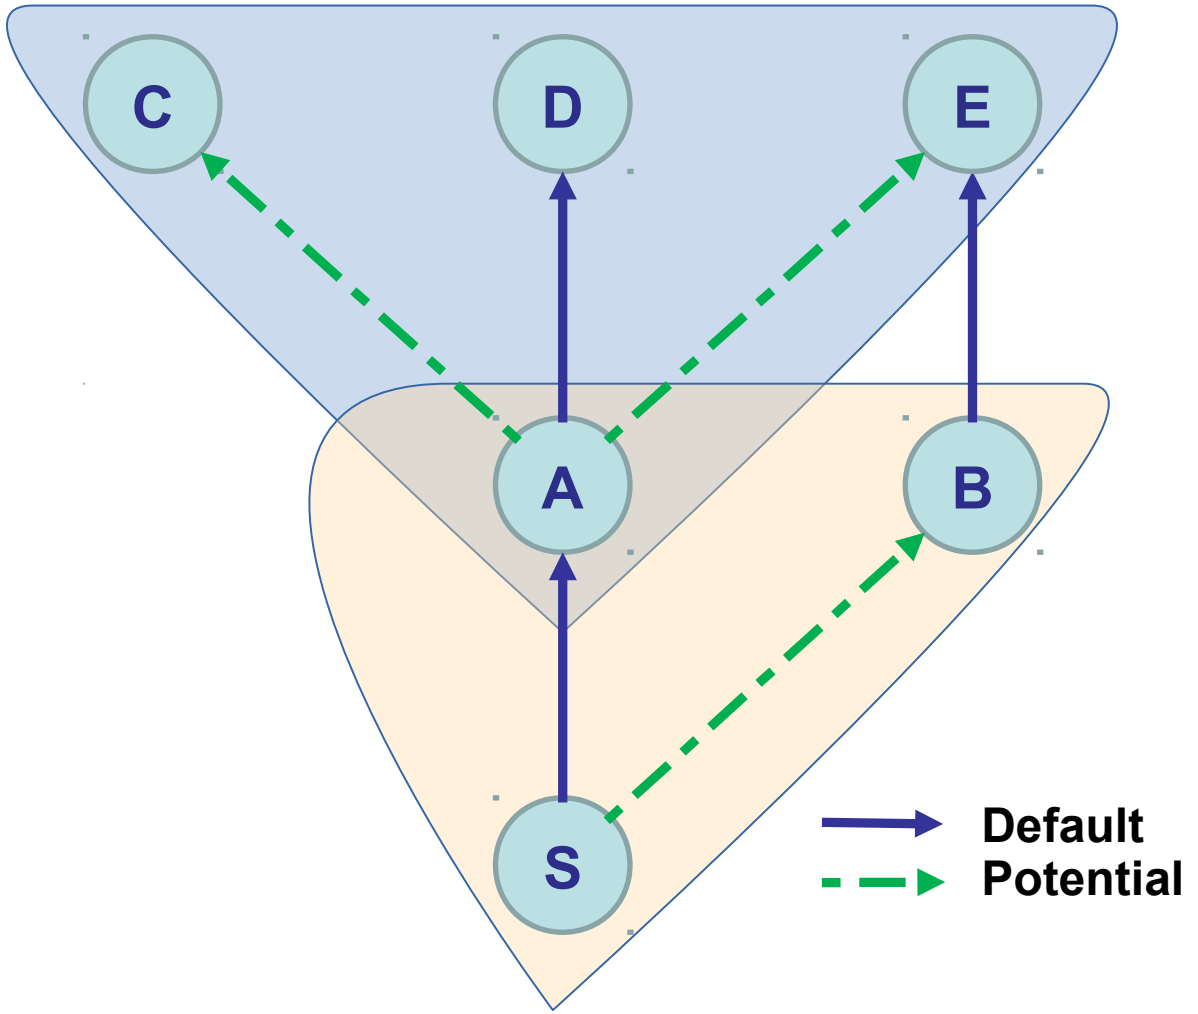

# Example

- **Parent set S:**
  - **{A, B}**
- **Parent Set A:**
  - **{D, C, E}**
- **Parent set B:**
  - **{E, D}**

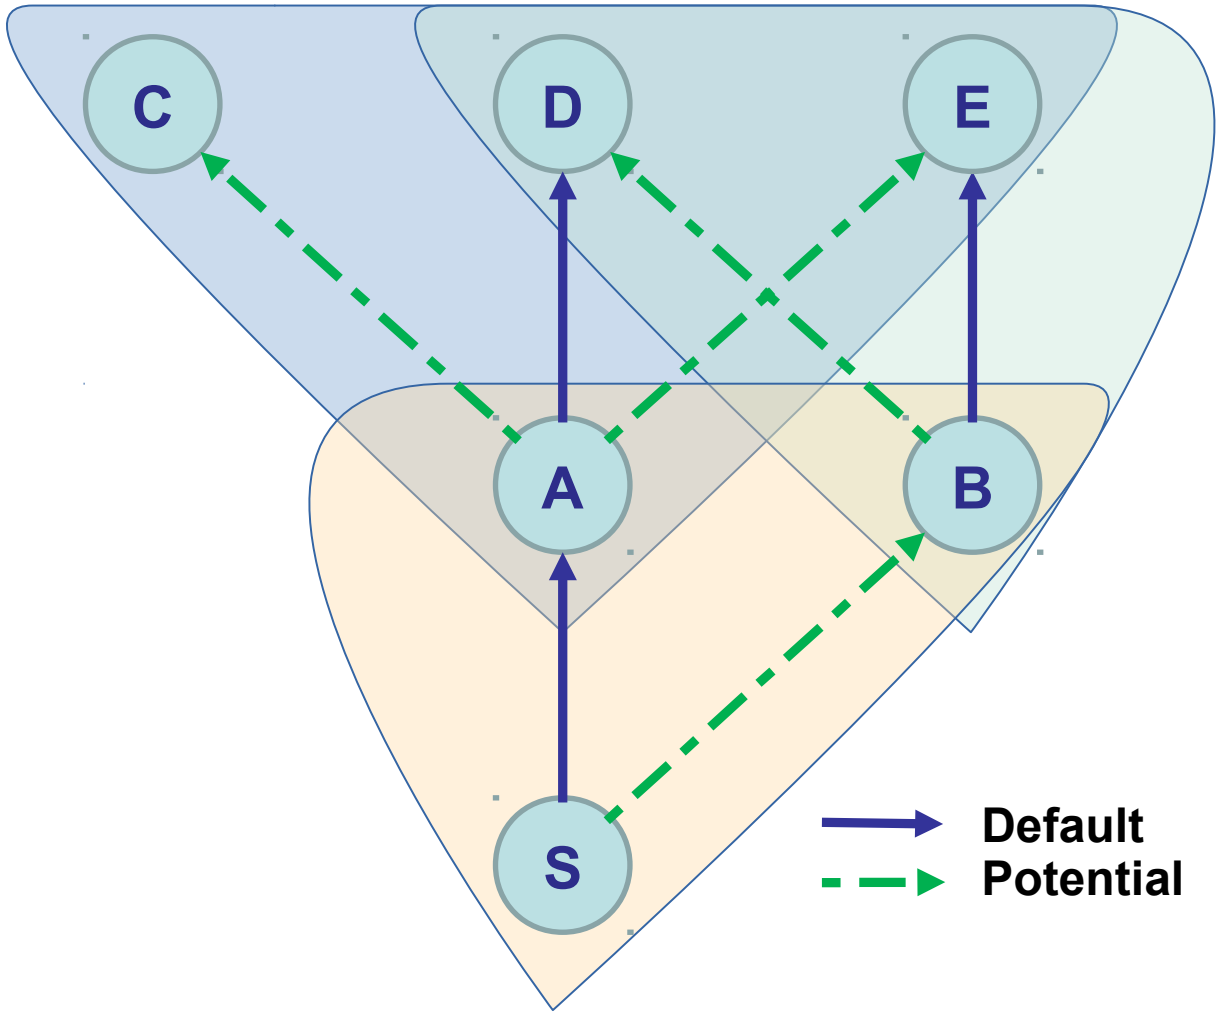

# Example

- A's DIO
  - Parent Set A: {D, C, E}

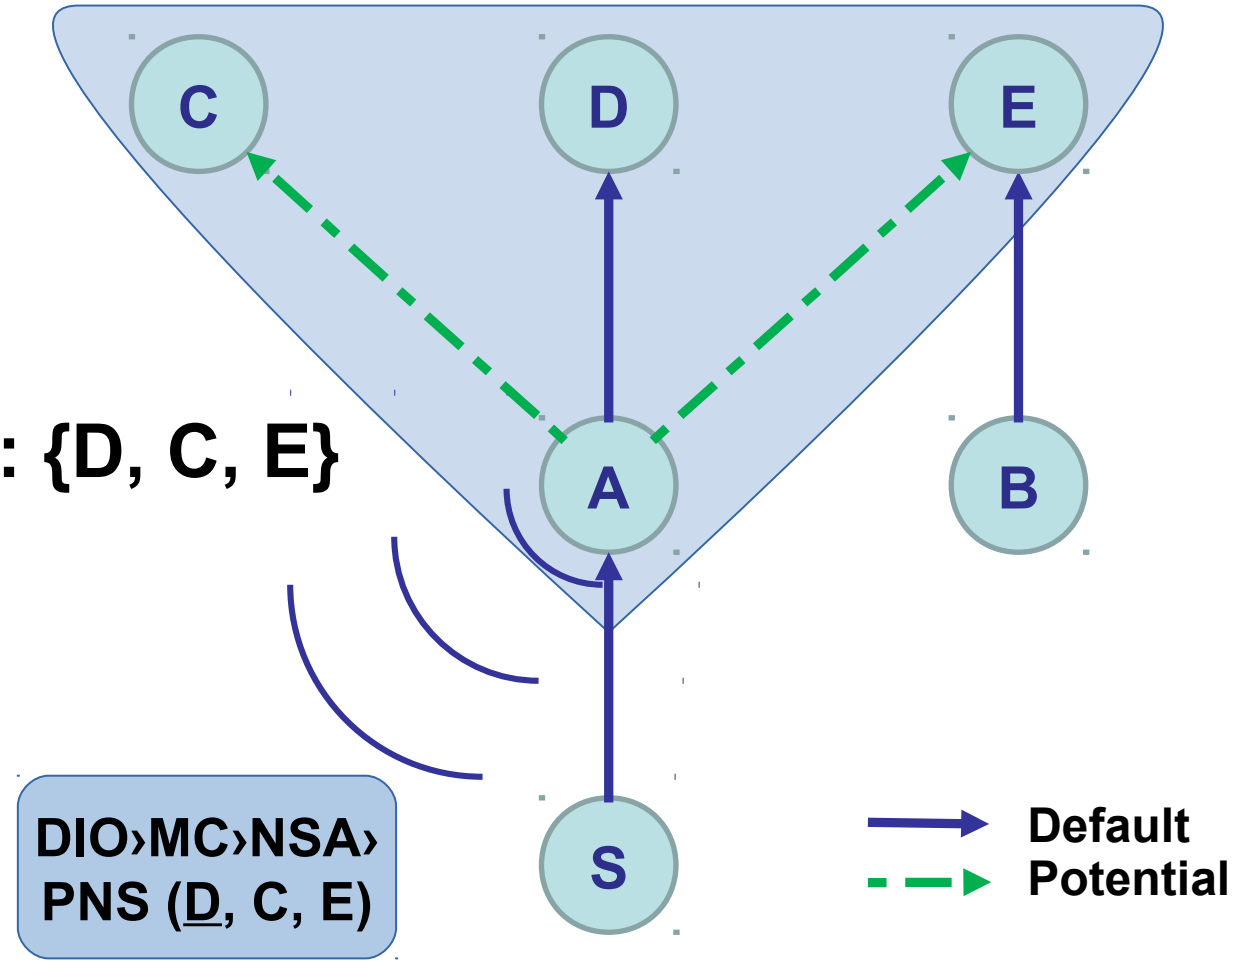

# Example

- B's DIO
  - Parent set B: {E, D}

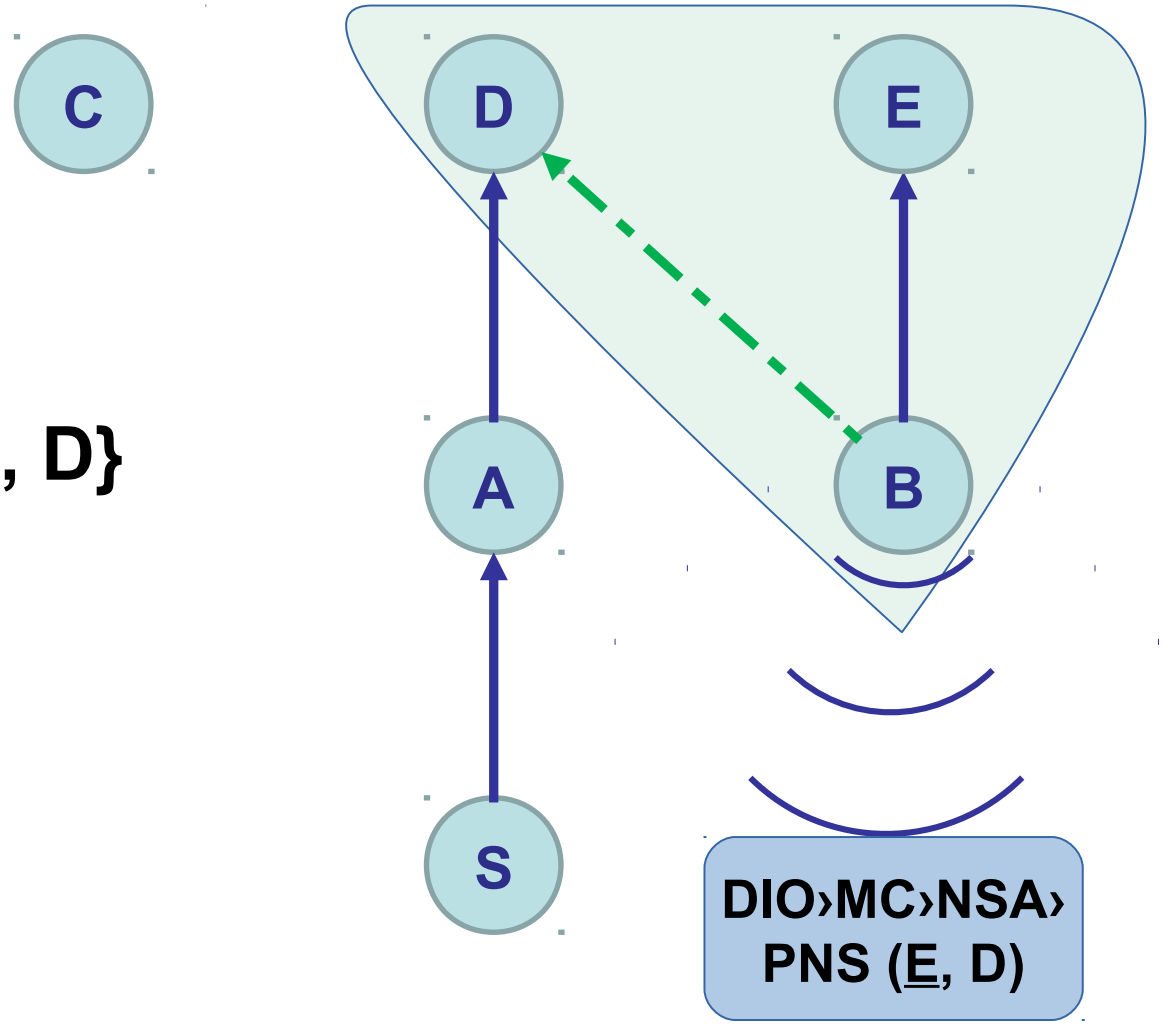

# Example

- S via A:
  - Default Grand Parent:
    - D
- S via B:
  - Grand Parent Set:
    - {E, D}
- D is in {E, D}

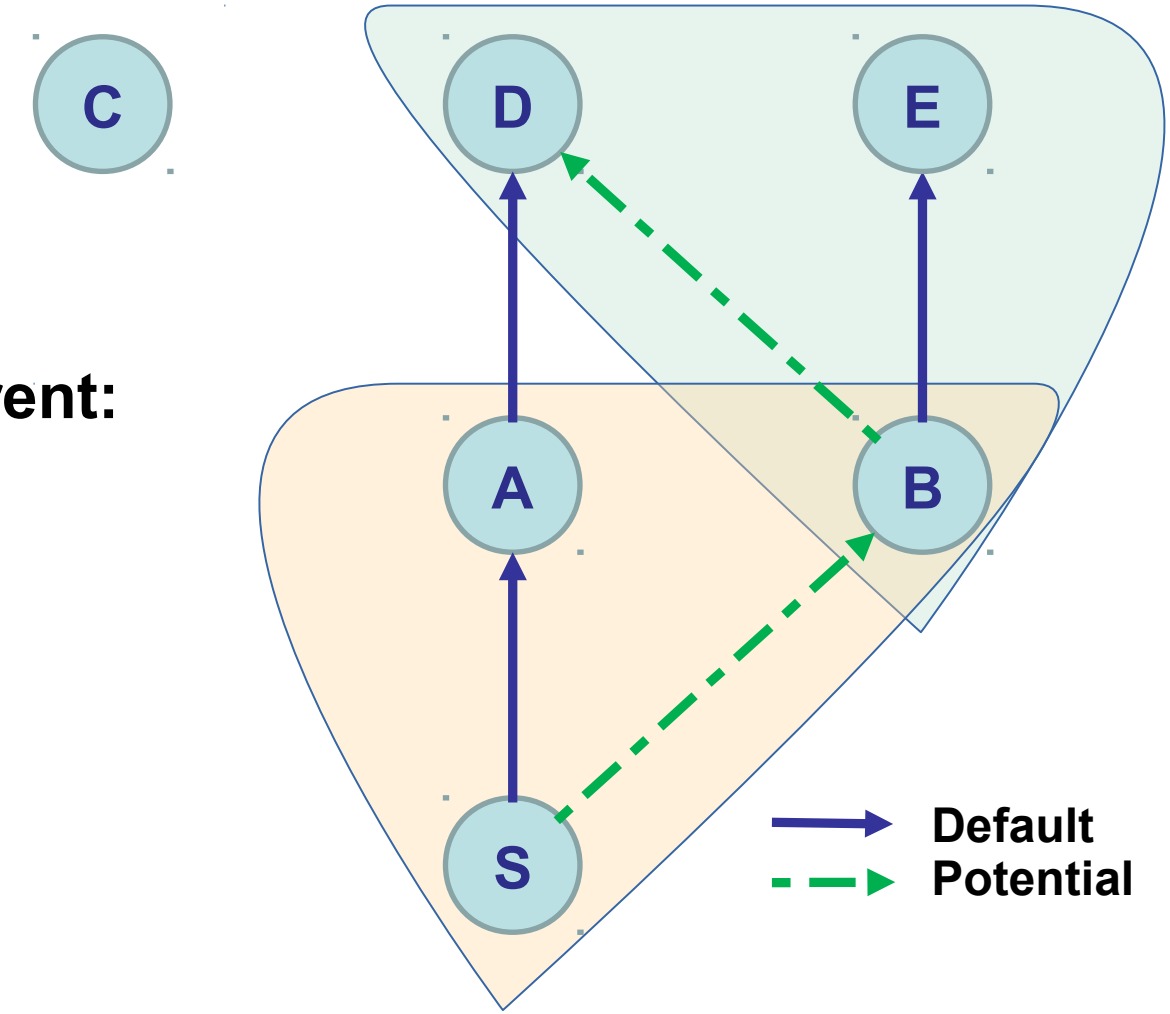

# Example

- S → B
  - Alternative Parent

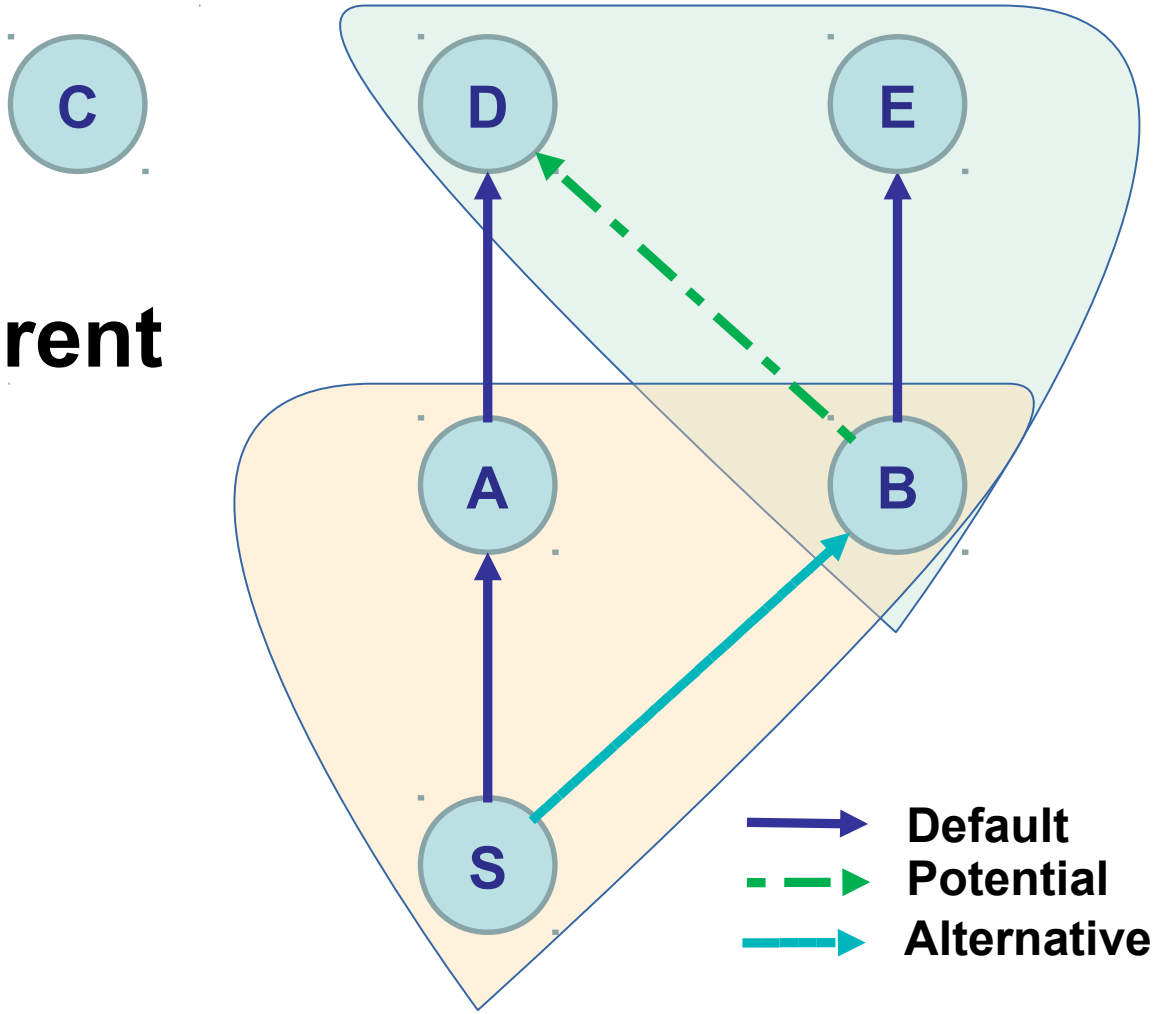

# Example

- Similarly:
- Alternative Parents:
  - A → C
  - B → D

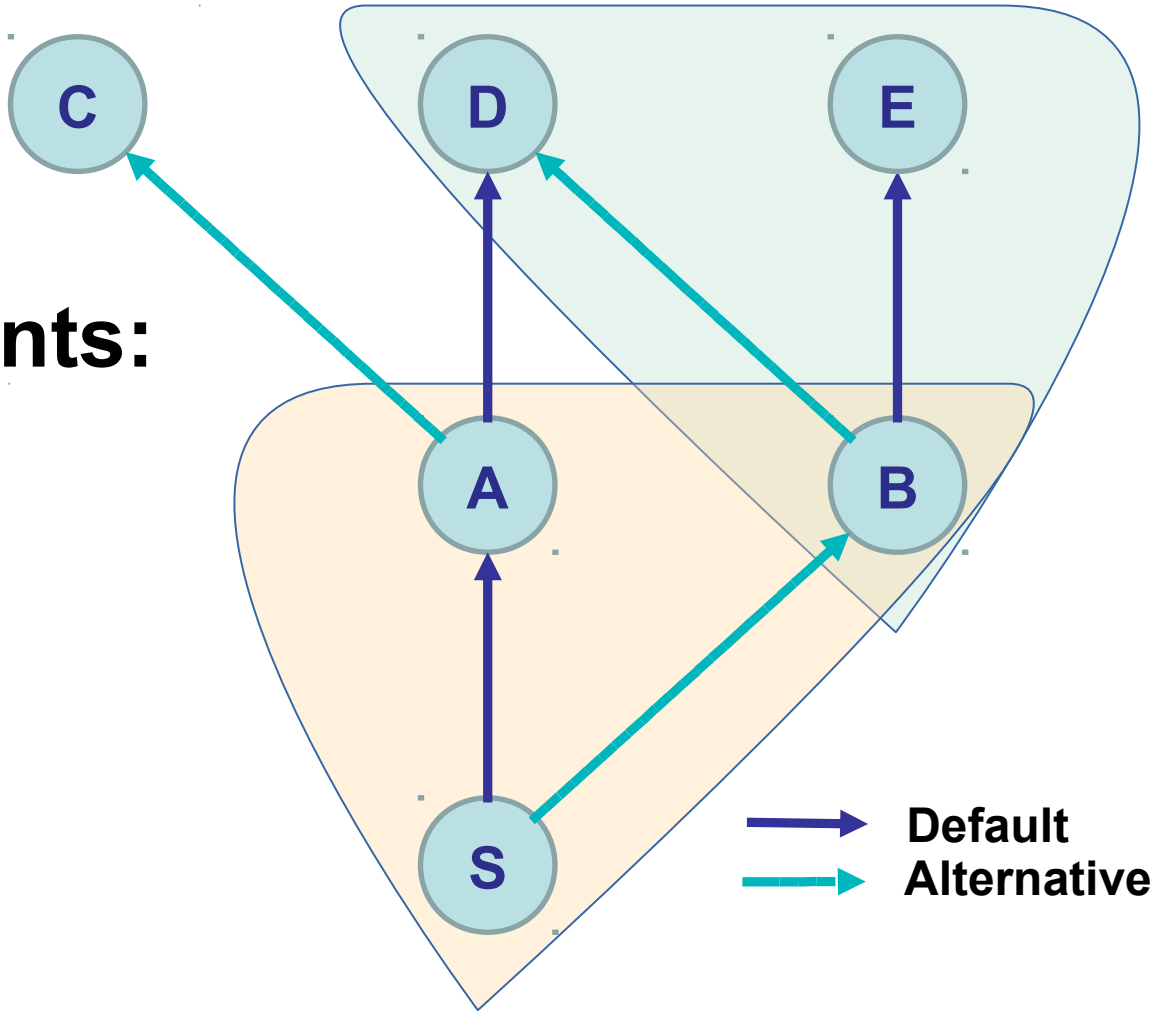

# Parent Selection - DIO Messages

- **Parent Set A:**
  - **{D, C, E}**
- **Parent set B:**
  - **{E, D}**

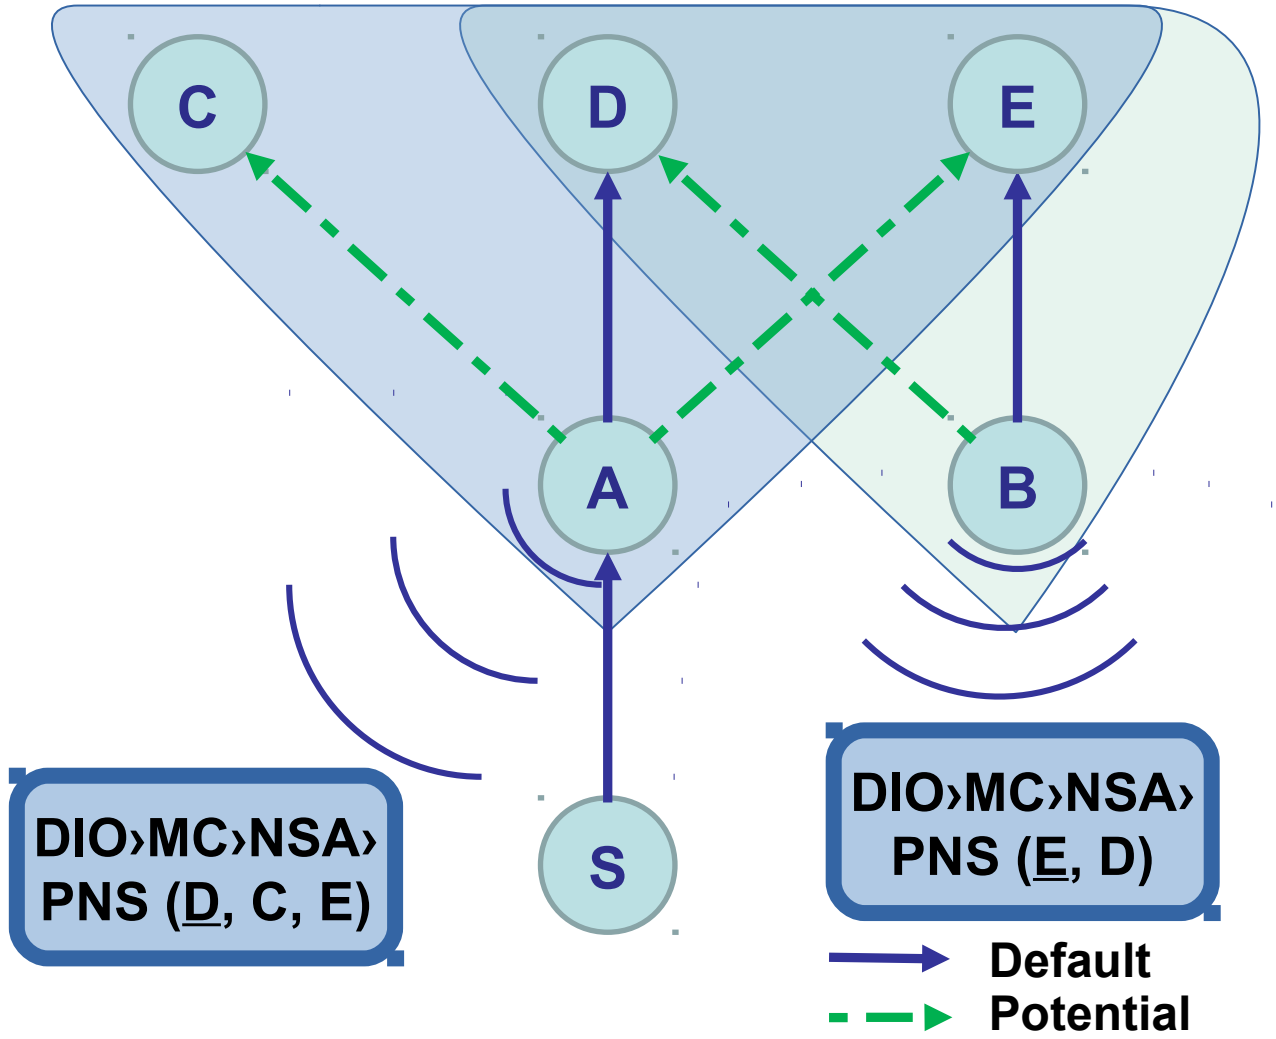

# Issues

- **Compression: IPv6 addresses**

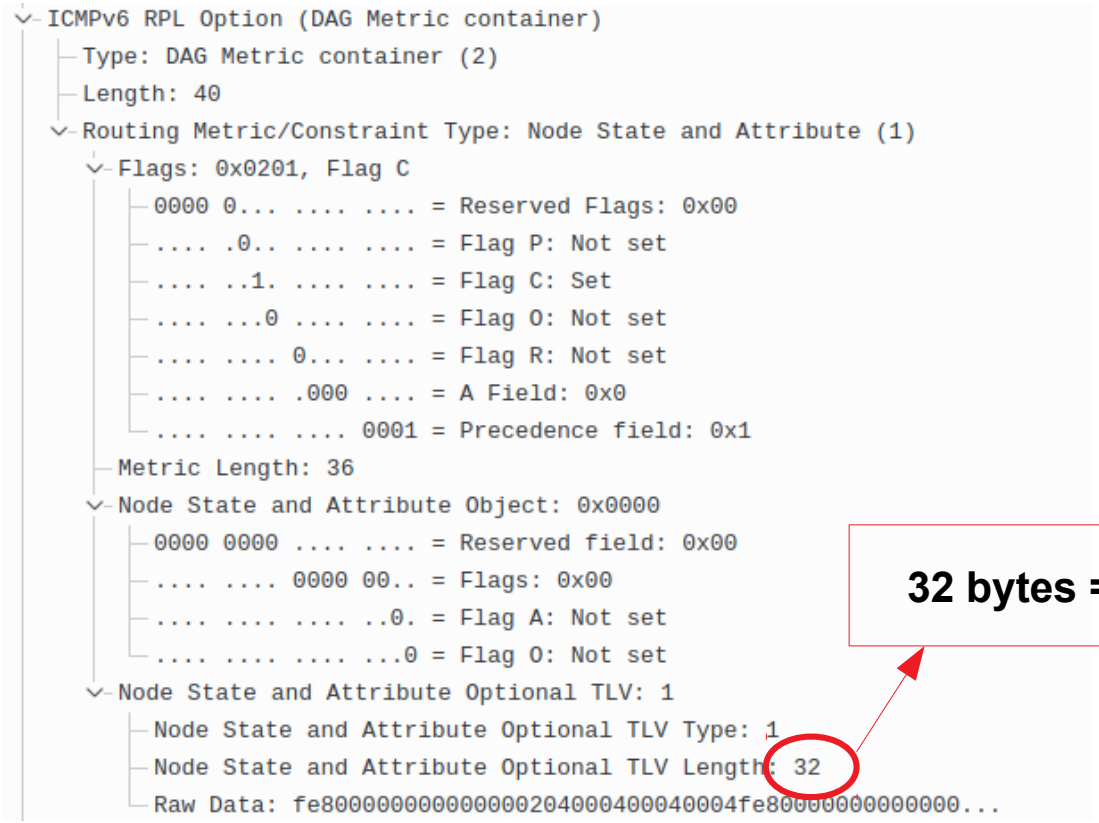

# Ideas

- Alternative use: Diversification
  - Preferred parent selection
    - Select PP for disjoint path
  - Multipath
    - Select alternative parent(s) for disjoint path

# Feedback

- Volunteers to REVIEW the draft;
- Is it relevant in ROLL WG?

# RPL DAG Metric Container (MC) Node State and Attribute (NSA) object type extension

**draft-pkm-roll-nsa-extension-00**

**Remous-Aris Koutsiamanis**
**Georgios Z. Papadopoulos: georgios.papadopoulos@imt-atlantique.fr**
**Nicolas Montavont**
**Pascal Thubert**

# Backup

**Compression is the potential problem but Solution with MAC addresses**

**and Alternative Parent Selection are out of the scope of this draft**

# comments

1. Something about DNA-based Alternative Parent Selection: why? To allow Overhearing procedure from AP and DP
2. non-congruent paths?
3. multiple parallel or disjoint paths ?

# RPL Observations

draft-rahul-roll-rpl-observations-00

- Rahul, Rabi, YueFeng@ Huawei

IETF101, London

# Background for the observations …

- Background for the observations
  - Observations were made during solution implementation based on 802.15.4
  - Mostly for storing mode of operation
- We had some sort of implementation in place for the problems
  - But we don't believe our solutions are necessarily optimal/best

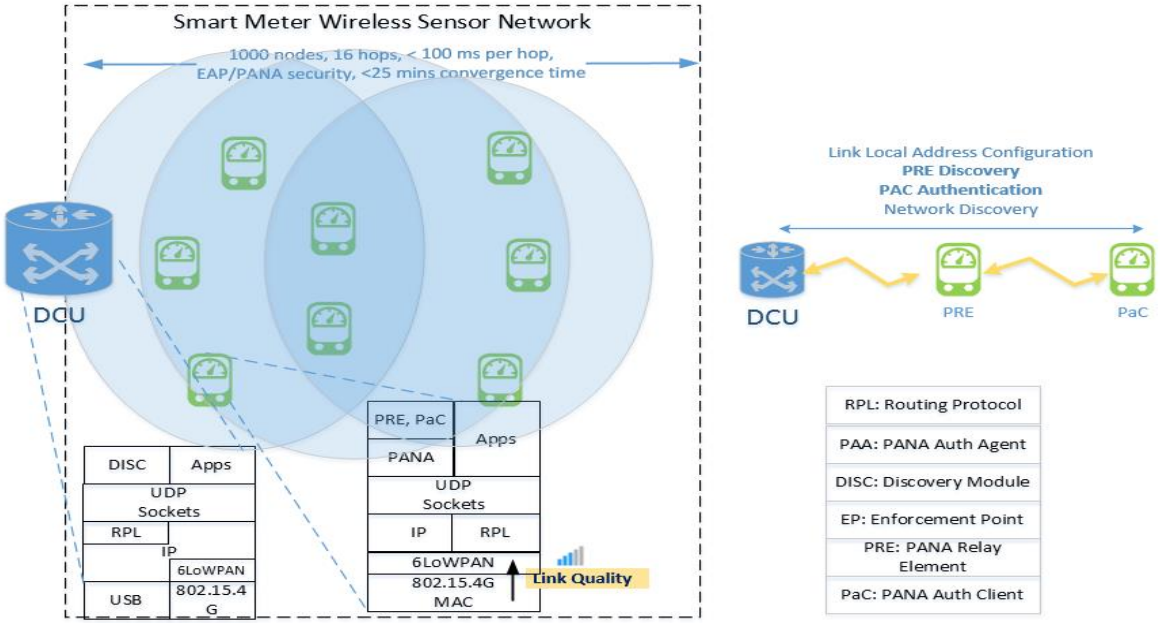

# The Problem: DTSN in storing mode

- Problems to handle
  - Dependent nodes route update
  - Impacts downstream route availability
- DTSN (DAO Trigger Sequence Number)
  - Decides if DAO should be sent
  - Decider element for RPL Control Traffic
- Problem in storing MOP only
- Tradeoff downstream route-availability vs control overhead

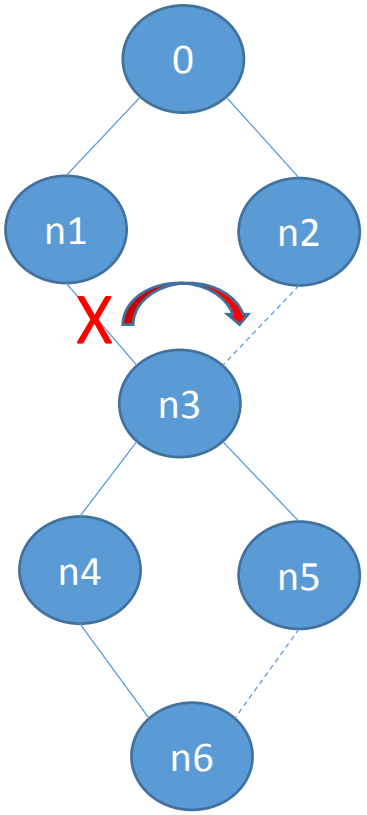

# Implementer's Dilemma1

- Should DTSN be incremented with every DIO trickle timer interval?
  - What happens if you do?
    - DAO traffic is too high
  - What happens if you do don’t?
    - DAO redundancy is too low. High probability of DAO not reaching BR.
    - With increase in hops, the probability of DAO success drops sharply.

# Dilemma2

- On parent switch, should node increment its DTSN ?
  - Yes, of-course. Otherwise how would child nodes update their paths.
  - Should child nodes in turn even reset DIO trickle timer and increment DTSN?
    - How would sub-child updates their paths?

# DAO-ACK: Multiple interpretations

## Hop-by-hop ACK

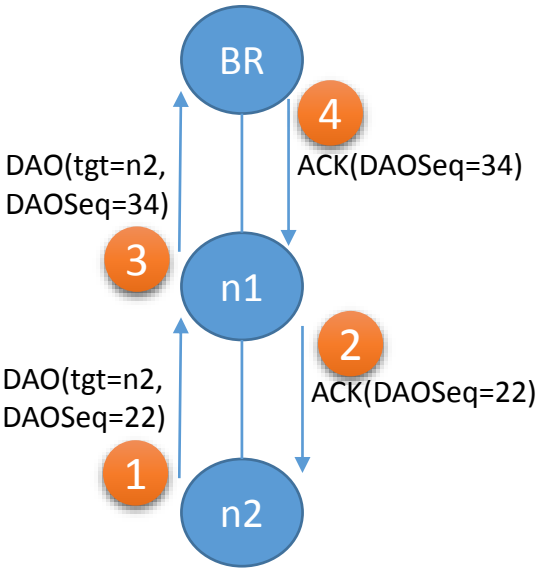


Pros:

1. No additional RAM
2. Very easy handling. No state.

Cons:

1. Does not help target determine if the DAO has reached BR.

RIOT implements this.

## End-to-End ACK

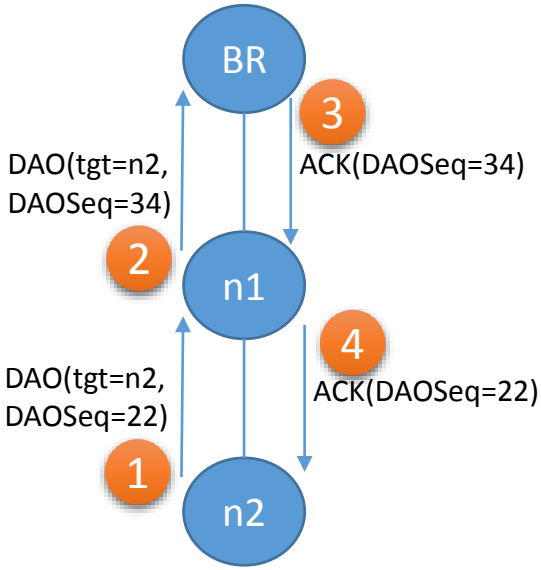


Pros:

1. Helps target determine if the DAO has reached BR.

Cons:

1. 6LRs need to maintain DAOSeq in routing entry. Thus 1B per routing entry.
2. Managing DAO-ACK timeout is non-trivial.

Contiki(new version) implements this.

# DAO-ACK and aggregated targets: How to ACK?

- DAO-ACK is for DAO message, not individual targets within DAO
  - Also, ACK cannot carry any options as per existing RPL spec
- If multiple targets in a DAO and if subset of targets fail, then how to ACK?

- RPL is not clear on how to handle aggregated targets
  - It certainly allows, but does not do failure handling…
  - RIOT implementation currently sends aggregated targets.
  - Contiki does not work with aggregated targets.
  - Thus interop between them is not possible today (at multi-hops)…

# DAO Retransmission and DAO-Ack

- DAO-Ack is important because
  - Only way for node to know that the E2E path is established…
- In hop-by-hop case
  - What happens if DAO/DAO-Ack fails on ancestor links?
- Can we ACK end-to-end using global IP address?
  - No RAM requirement
  - Reduced handling on the intermediate 6LRs
  - ACKs can't be aggregated in this case

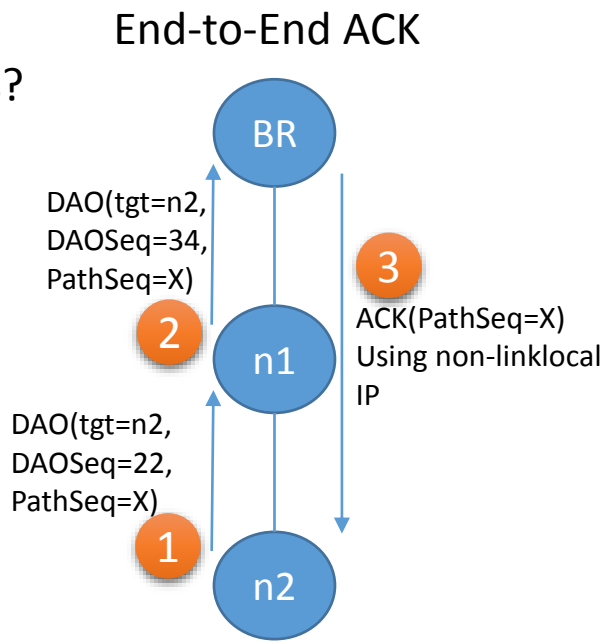


Lot already discussed on ML (Oct 2015):
https://www.ietf.org/mail-archive/web/roll/current/msg09469.html

# Handling node reboots

- RPL State information needs to be maintained across node reboots
  - For e.g.
  - DODAGVersionNumber
  - DTSN
  - PathSequence
- Losing this state across reboot could result in serious loss of connectivity
- Implications of using persistent storage
  - Implications of flash endurance on network lifetime
- 6LRs are particularly impacted

# Handling resource unavailability

- Neighbor cache table and Routing Table
- Handling neighbor cache entry full scenario
  - Basic handling there currently
    - DAO-NACK and NA status!=0…
  - It's not enough though…
  - How to avoid connecting to same neighbor in the future?
- Handling routing table full scenario
  - No multi-level proactive feedback, i.e. what happens if the ancestor node does not have space?
- DIO does not signal resource availability currently…

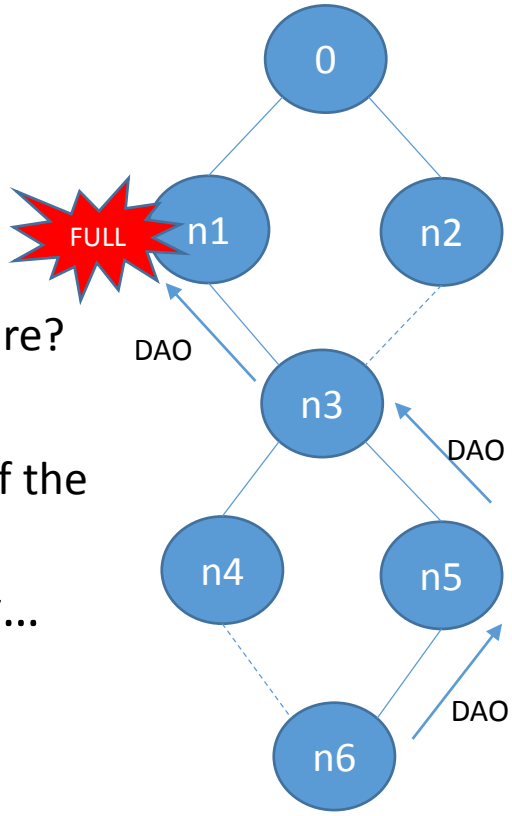

Thank you

# Constrained-Cast: Source-Routed Multicast for RPL

draft-ietf-roll-ccast-01

Carsten Bormann – IETF 101

# Send Bloom Filter with packet, match OIF

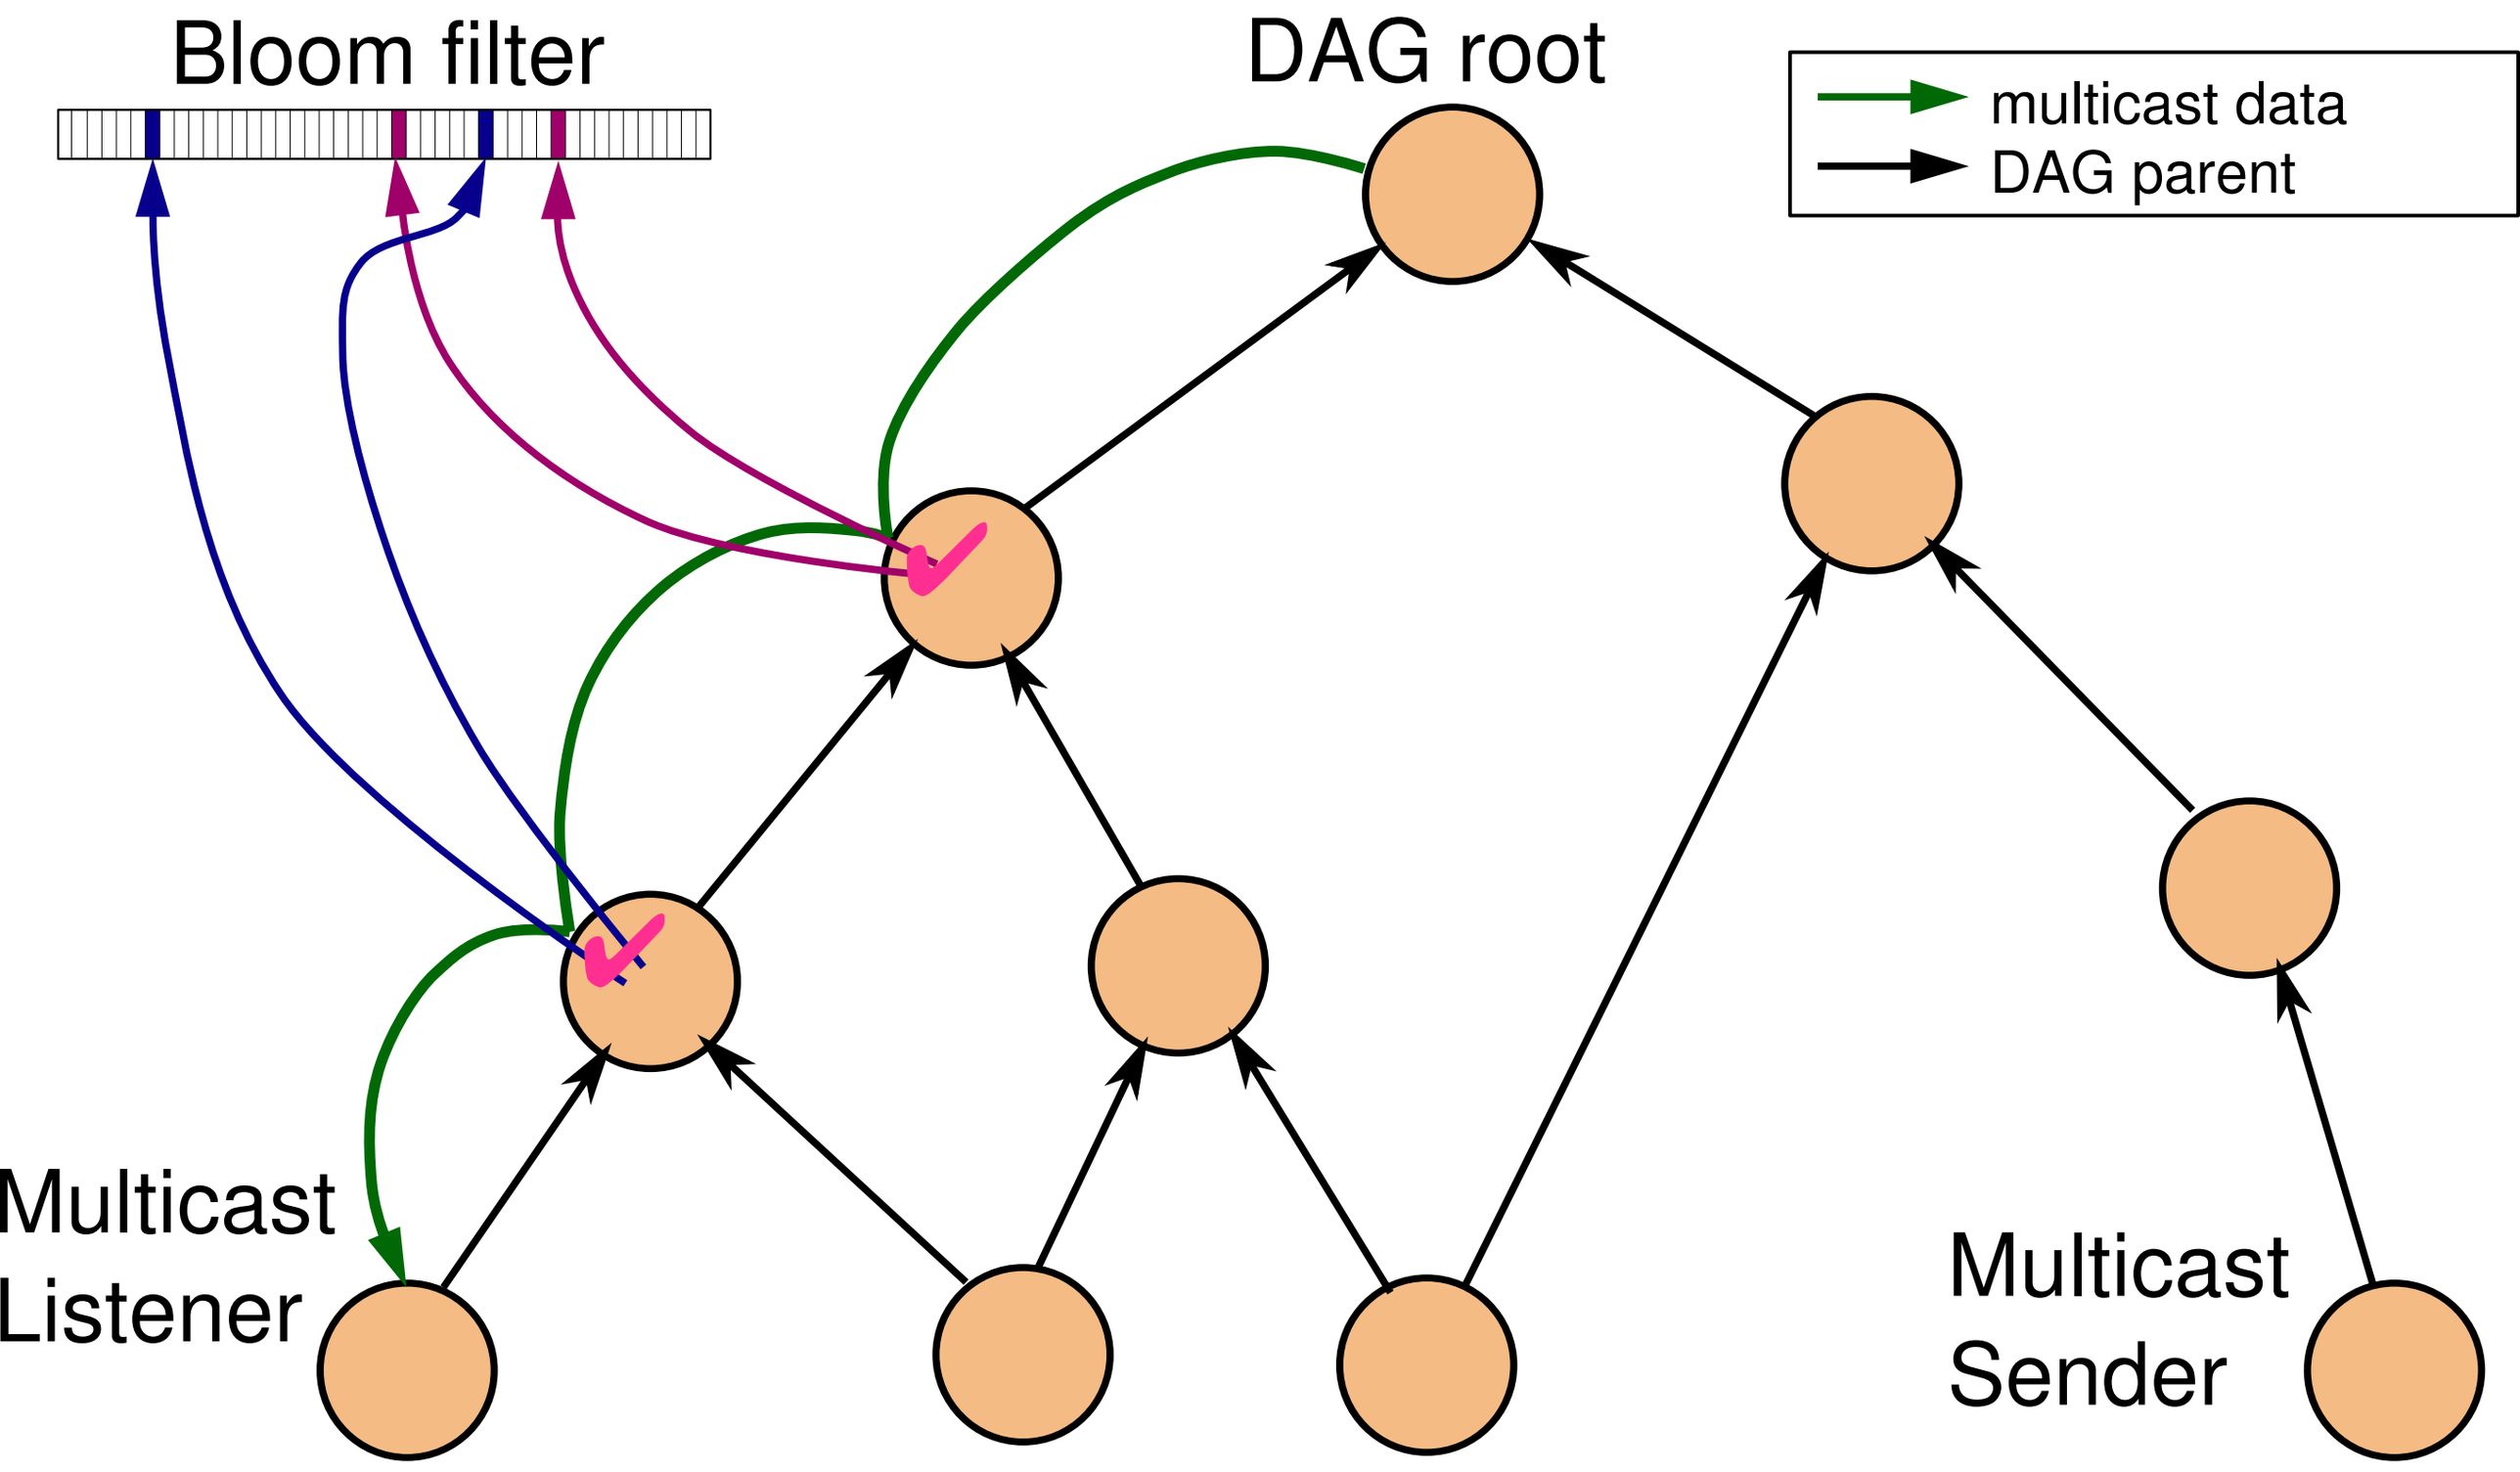

# False positives?

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

# False positive causes spurious transmission

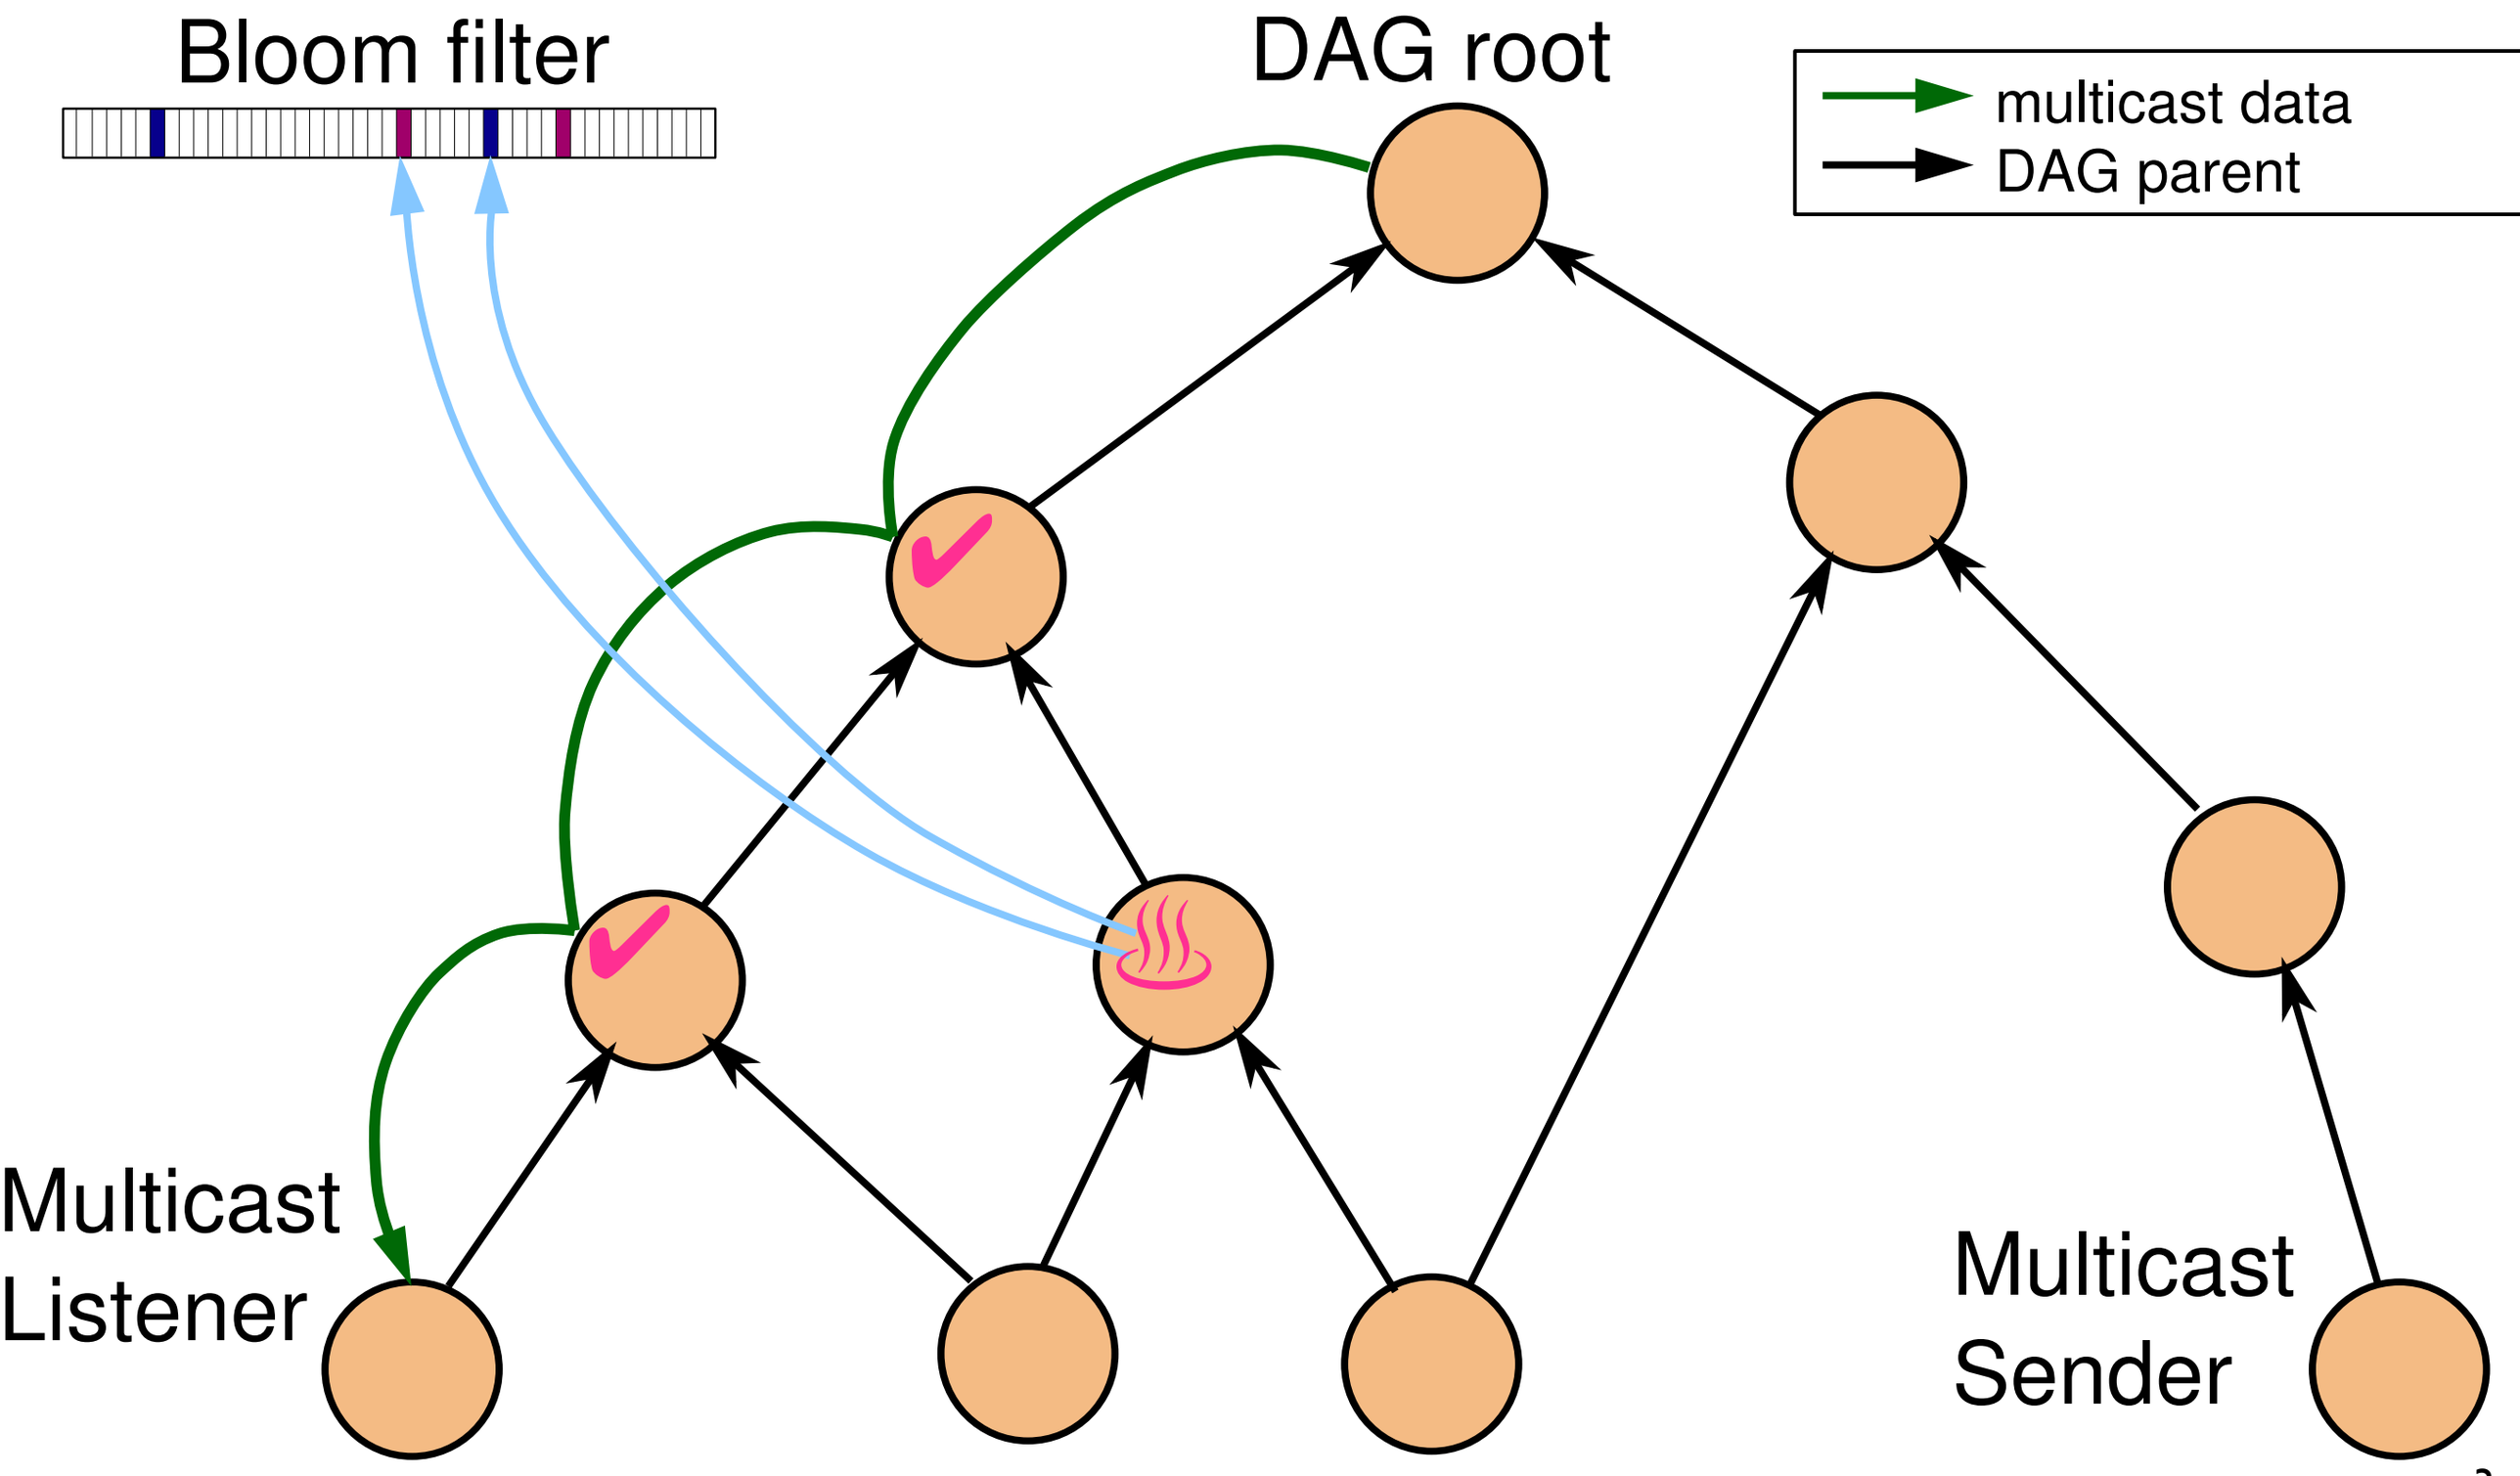

# How bad are false positives?

- 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.:

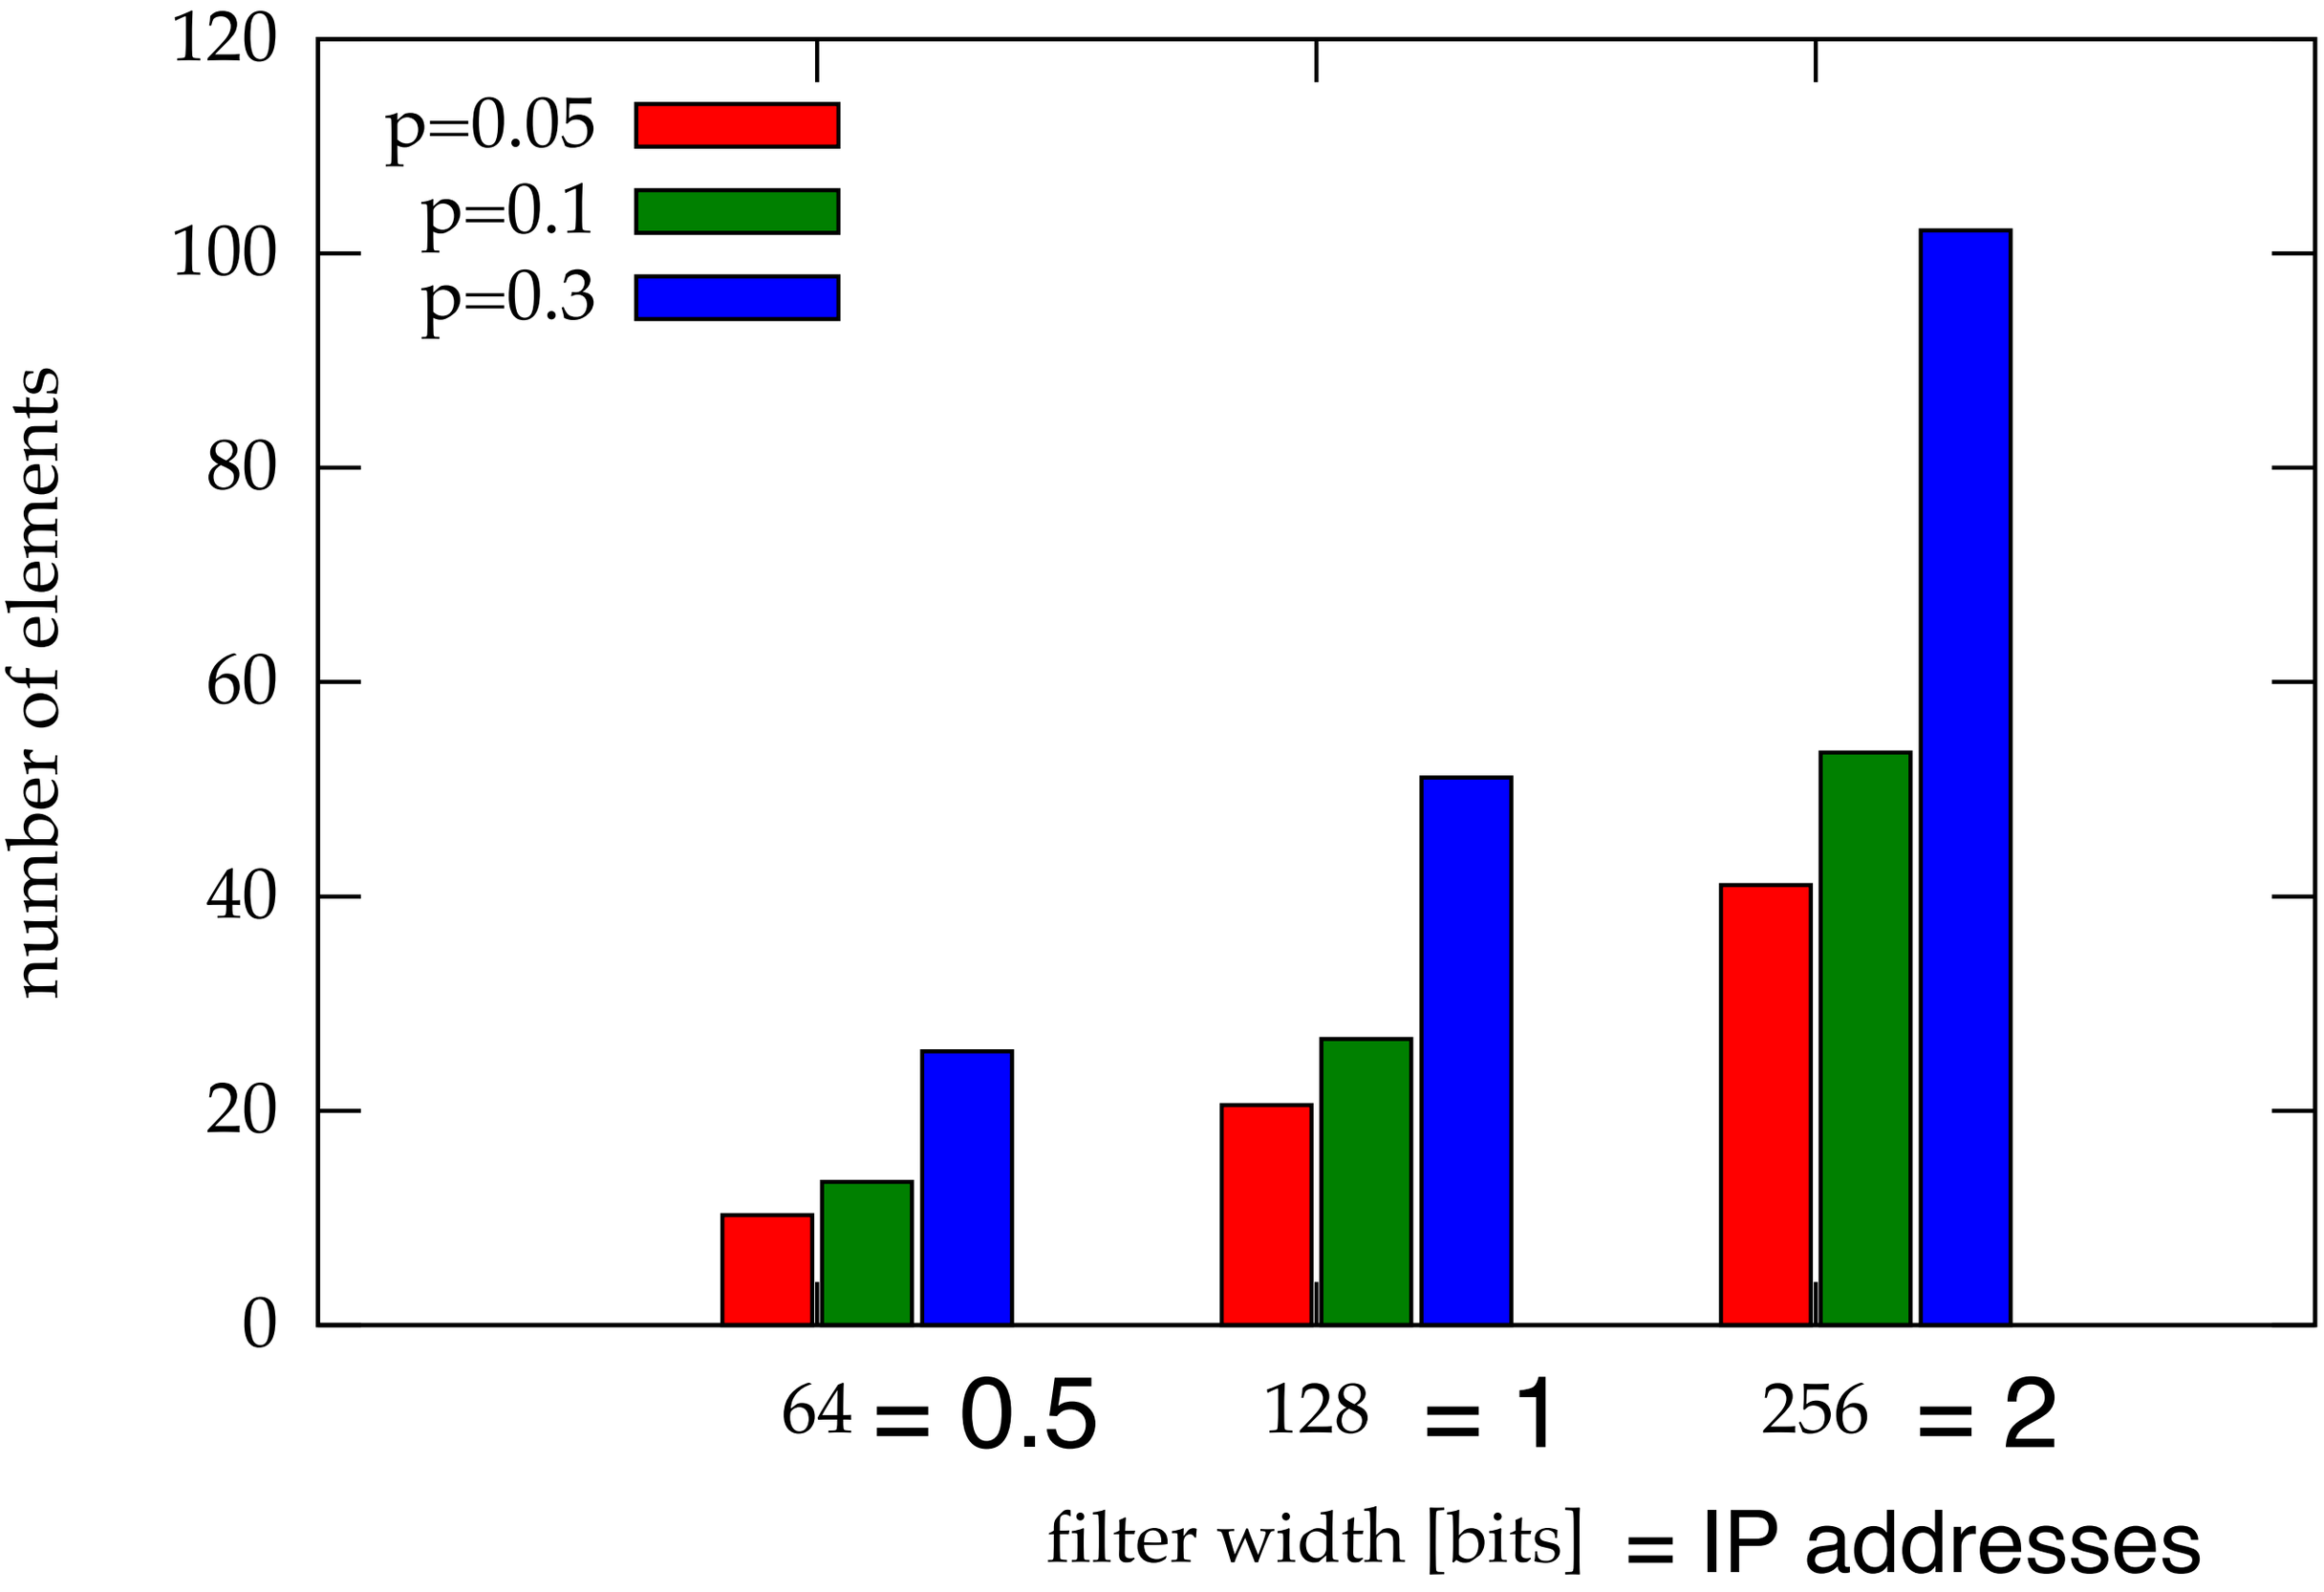

# BOCK and ROLL

draft-thubert-6lo-fragment-recovery-00

P.Thubert

IETF 101

London

# Context on BIER

- BIER Architecture now RFC 8279
  - Bits in a bitmap indicate selected end-points
  - Encapsulation defined in RFC 8296 for MPLS
- BIER-TE Architecture WG Doc draft-ietf-bier-te-arch
  - Bits in a bitmap indicate segments (hops) on the way
  - draft-huang-bier-te-encapsulation-extension-00
- Fast Reroute with “Protection Methods for BIER-TE”
- PREF with “BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM”

# BIER @ ROLL

| | BIER (Storing mode) | BIER-TE (Non-Storing) |
|---|---|---|
| Bit-by-Bit | draft-thubert-roll-bier-01 | |
| Bloom Filter | | draft-ietf-roll-ccast-01 |

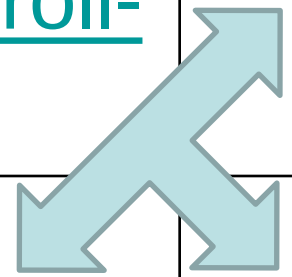

https://tools.ietf.org/html/draft-thubert-6lo-bier-dispatch

# Discussion: which / how many drafts

- Four possibilities => Four drafts?
  - Bloom filters do not depend on Mode of Operation
  - In storing mode, a DAO carries one bitstring, and the state in nodes is one bitstring per child => huge benefit
  - Non-storing Mode consumes more bits to express all the possible links
  - draft-thubert-roll-bier-01 is written as to collapse all 4 cases but most of the Bloom work is left to be added

# Discussion: BIER Work

- Allocation of a bit to a new Address
  - update 6lo ND?
- Packet compression
  - update 6LoRH? e.g.,draft-thubert-6lo-bier-dispatch
- Bits in DAO messages
  - update RPL?
- Bitmap Overflow management

# Discussion: BIER-TE Work

- Selection of  unique LinkID
- Allocation of a bit to a Link?
- Packet compression
  - update 6LoRH?
- Bits in DAO messages
  - update RPL?
- Bitmap Overflow management

# Discussion: Bloom Filter

- Specification of hash functions
- What gets hashed: IPv6 @? LinkID?
- Distribution/update of the hash functions
- Packet compression
  - update 6LoRH?

# Past IETF presentation

P.Thubert

IETF

Prague

# BIER / RPL

Pascal Thubert

IETF 99

Prague, July 2017

# Unreliable BIERPL

# Example topology

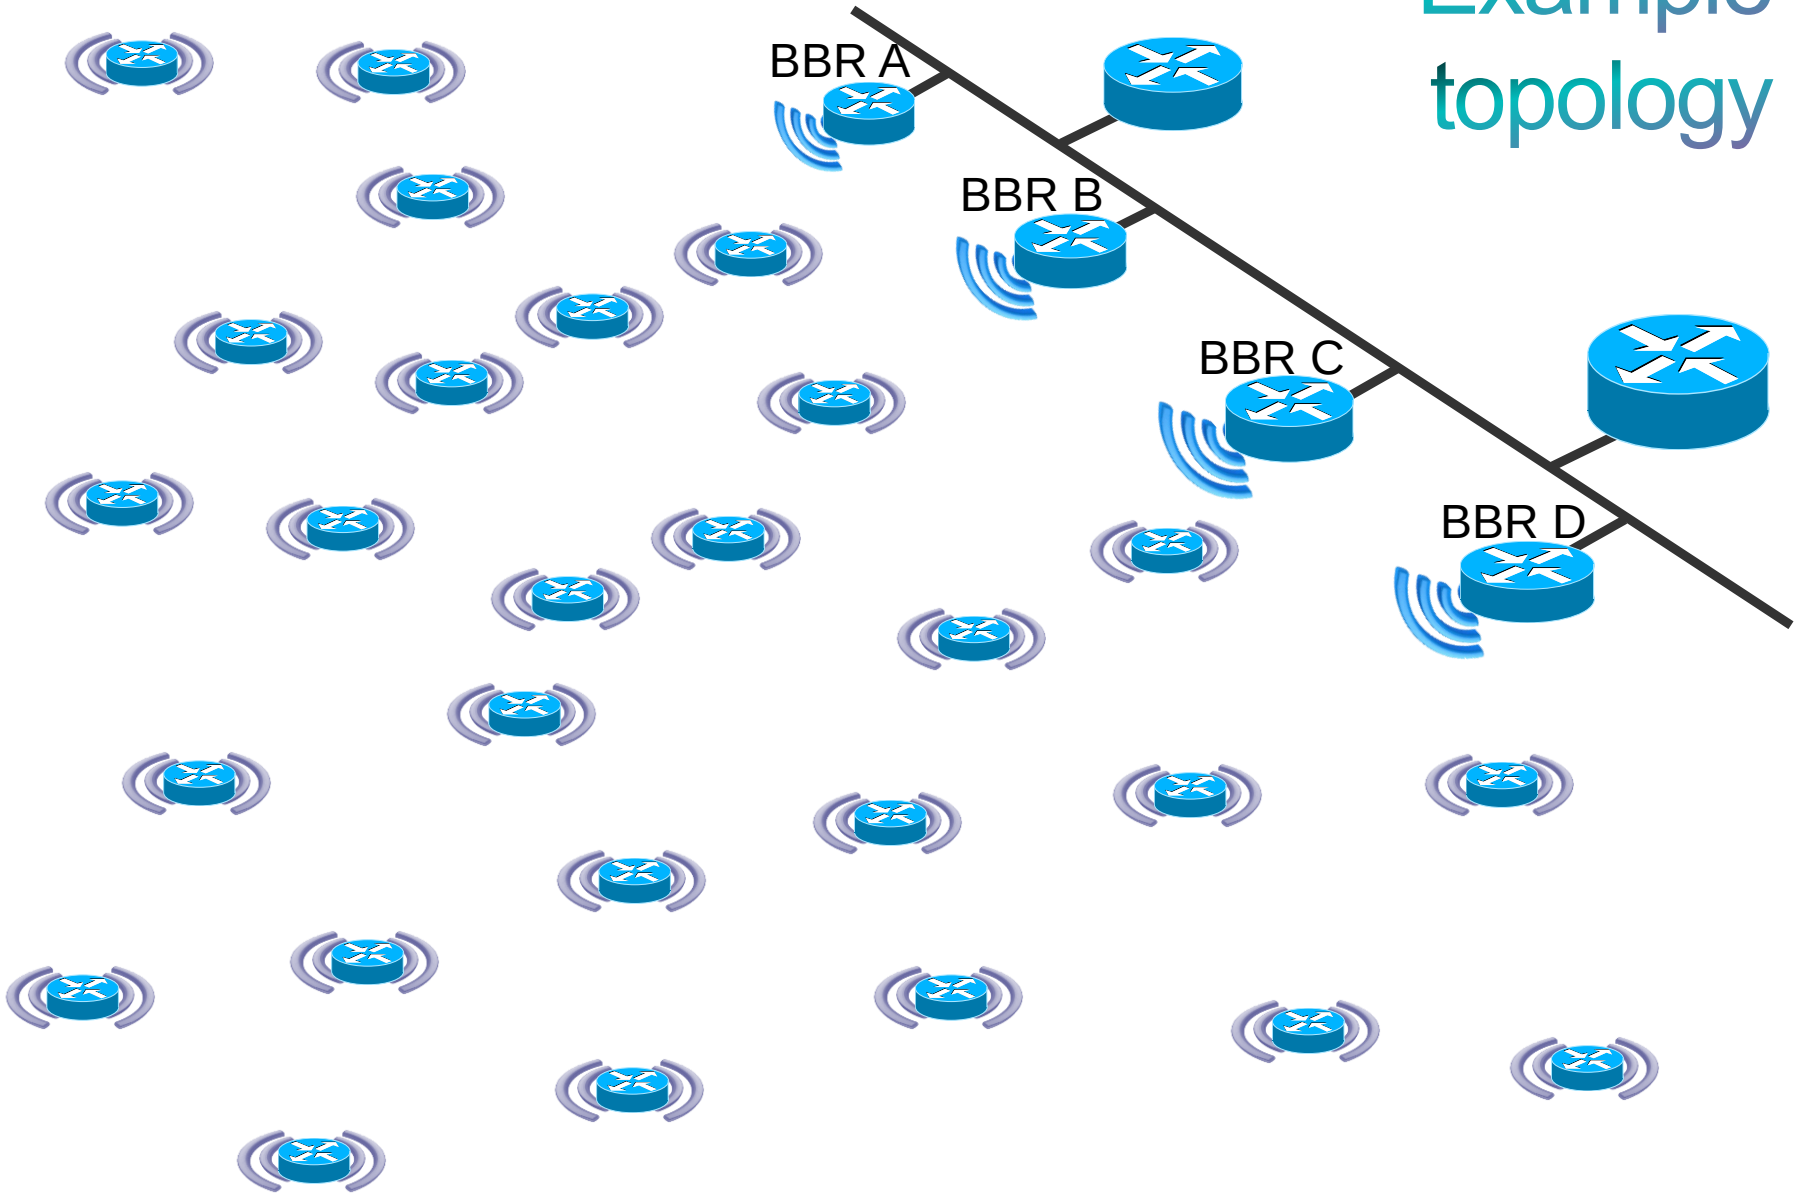

# An example of RPL DODAGs

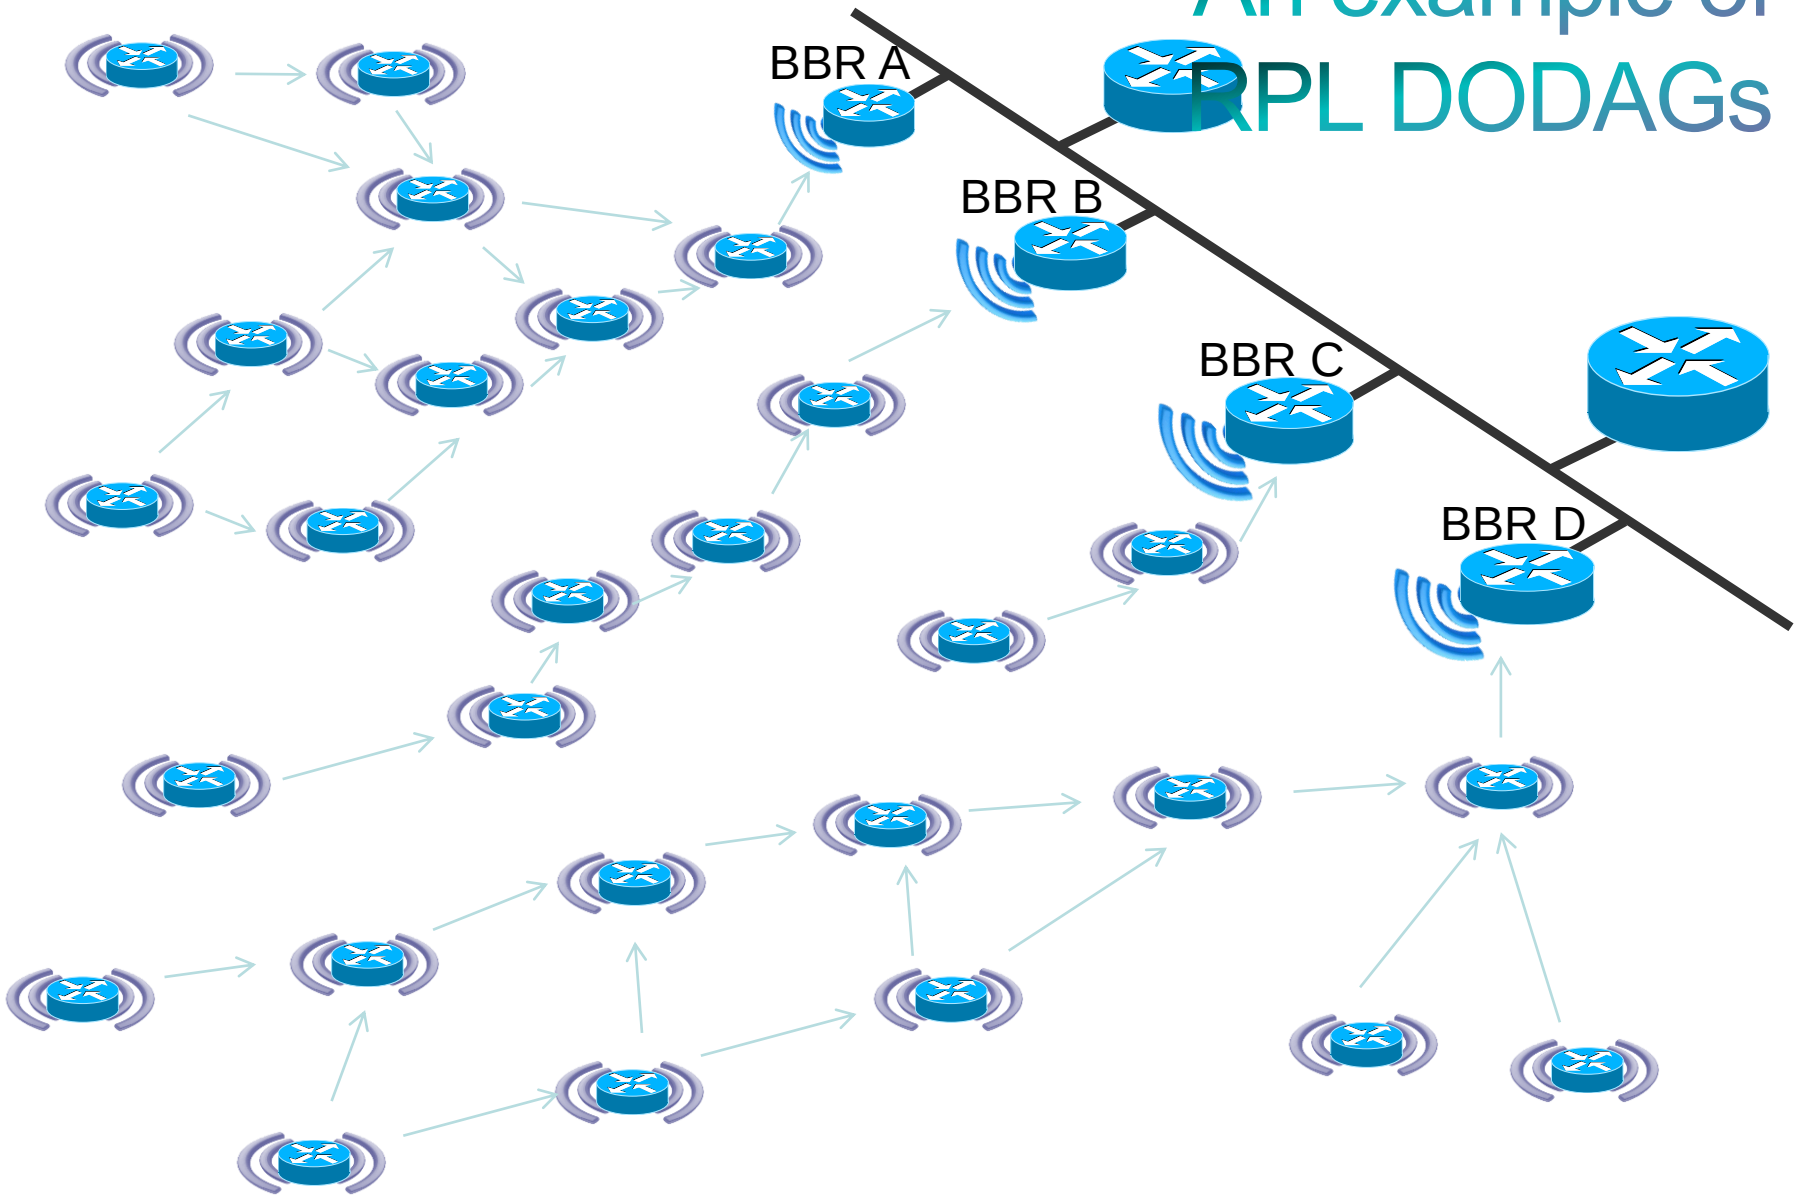

# Preferred tree Within DODAG

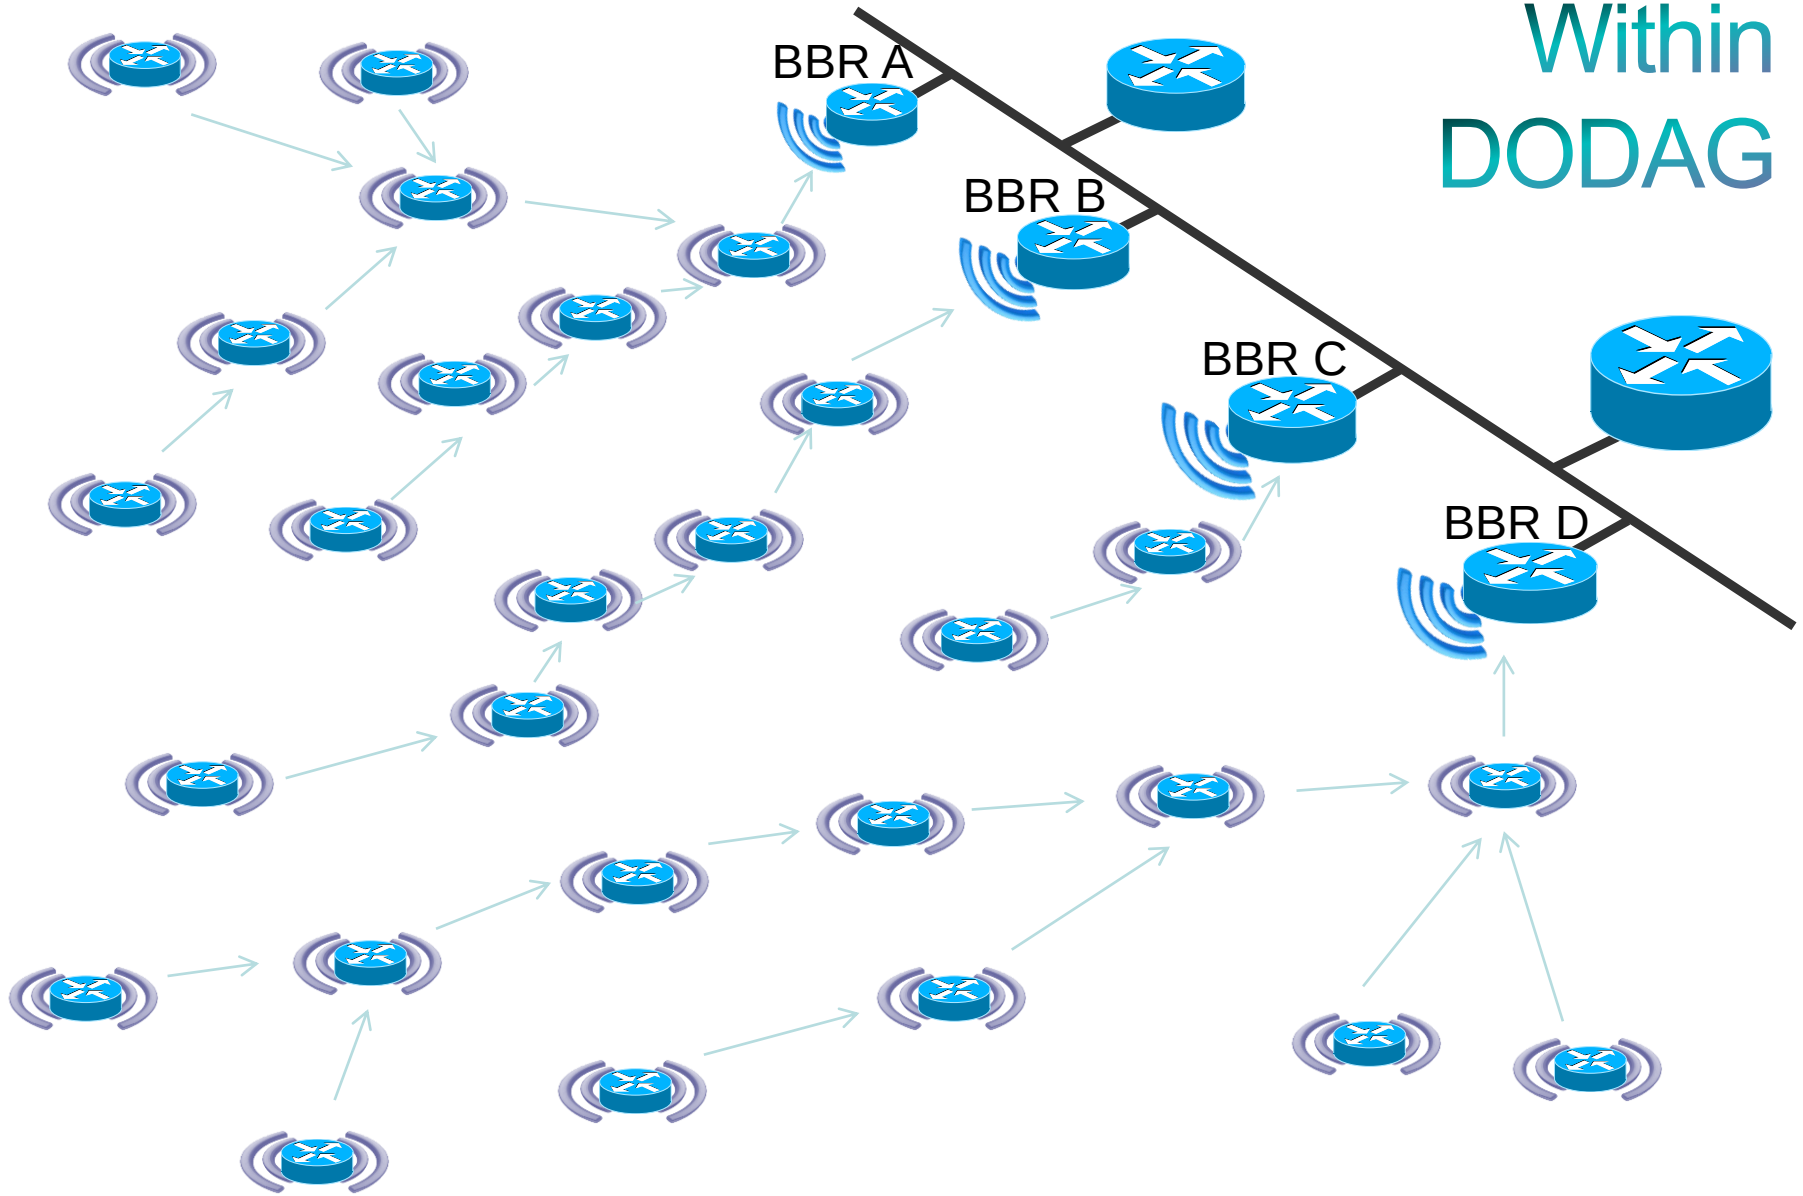

# Classical addressing

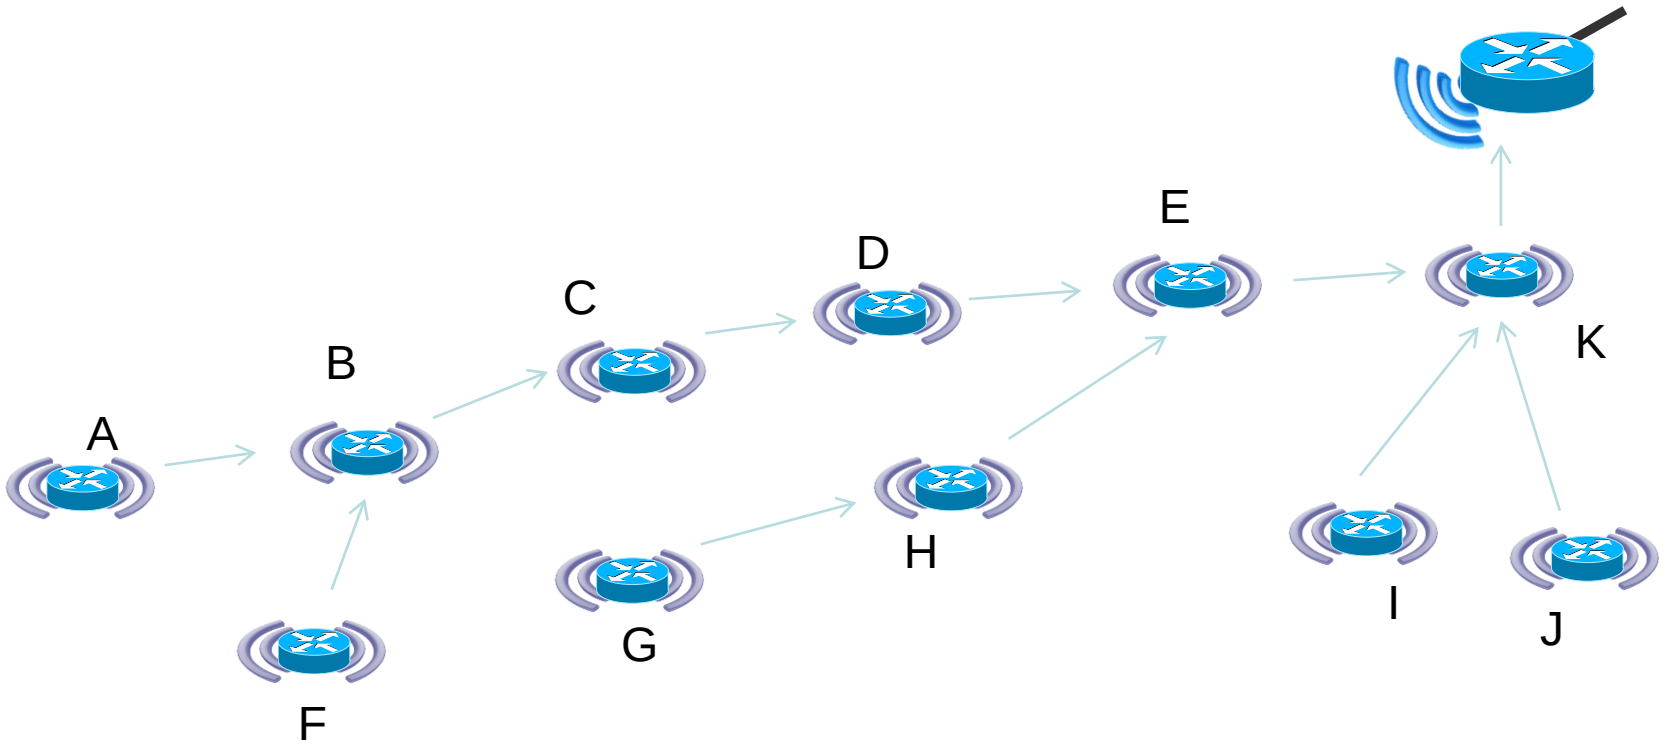

# Classical non storing routing

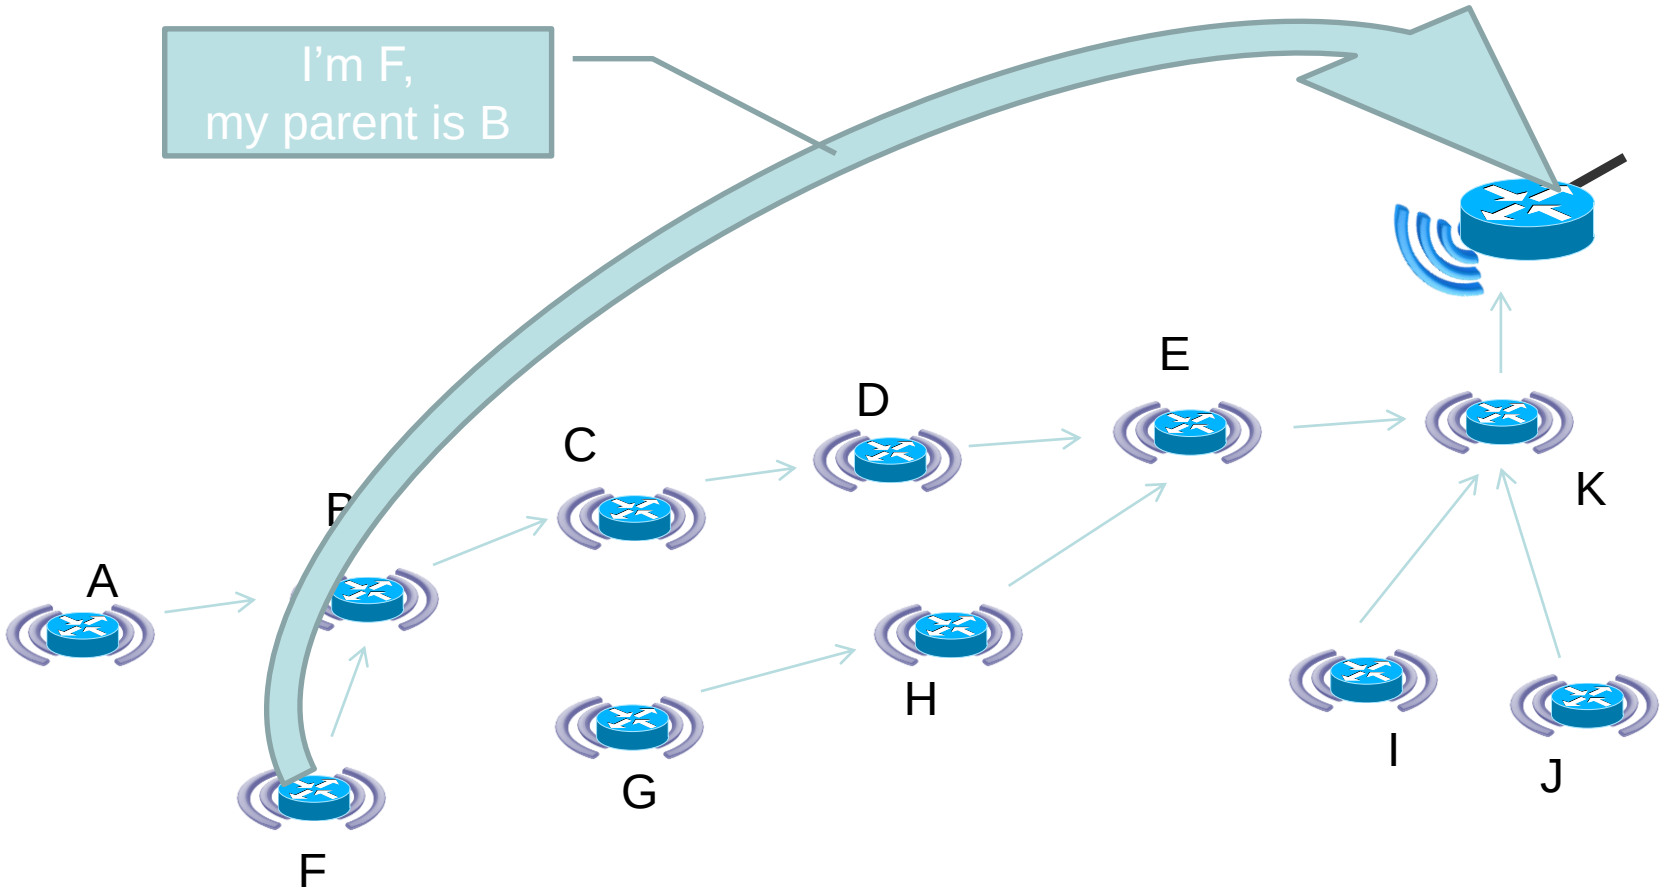

# Classical non storing routing

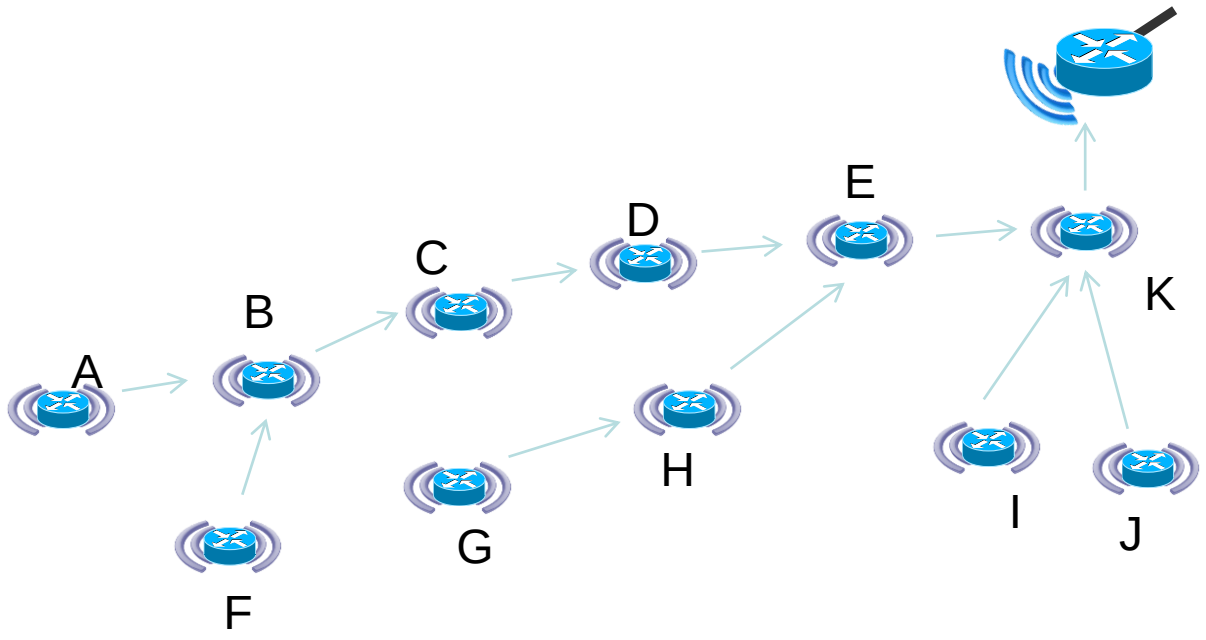


| Target | Transit |
|---|---|
| J | K |
| I | K |
| E | K |
| D | E |
| C | D |
| B | C |
| F | B |
| A | B |
| … | |

# Allocating bit positions

Option 0 - > static

Option 1 -> autoconf. RAs carry the current bitmap of allocated bits like they carry 6lowpan context info, and nodes pick a free bit. Collisions are handled as part of 6lowpan ND, like DAD.

Option 2 -> the 6LBR assigns a bit and returns it on the DAR/DAC exchange

Note: Upon mobility to new 6LBR, a new bit has to be assigned.

# Possible Optimizations

Derive from the short address in 802.15.4

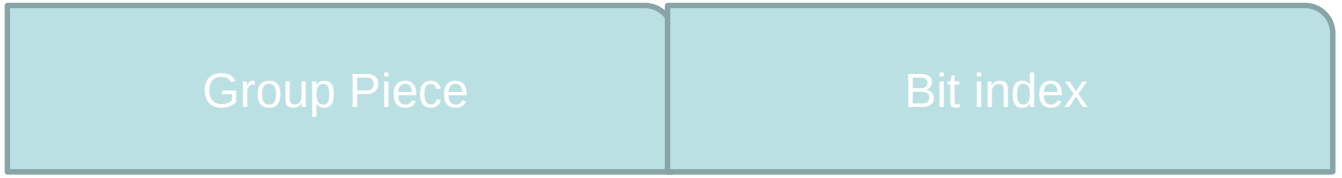


Use 2 bits per device to index 3 addresses

| Bits | Address |
|---|---|
| 00 | Not a target |
| 01 | Address 1 |
| 10 | Address 2 |
| 11 | Address 3 |

# Result of Allocation of bit positions

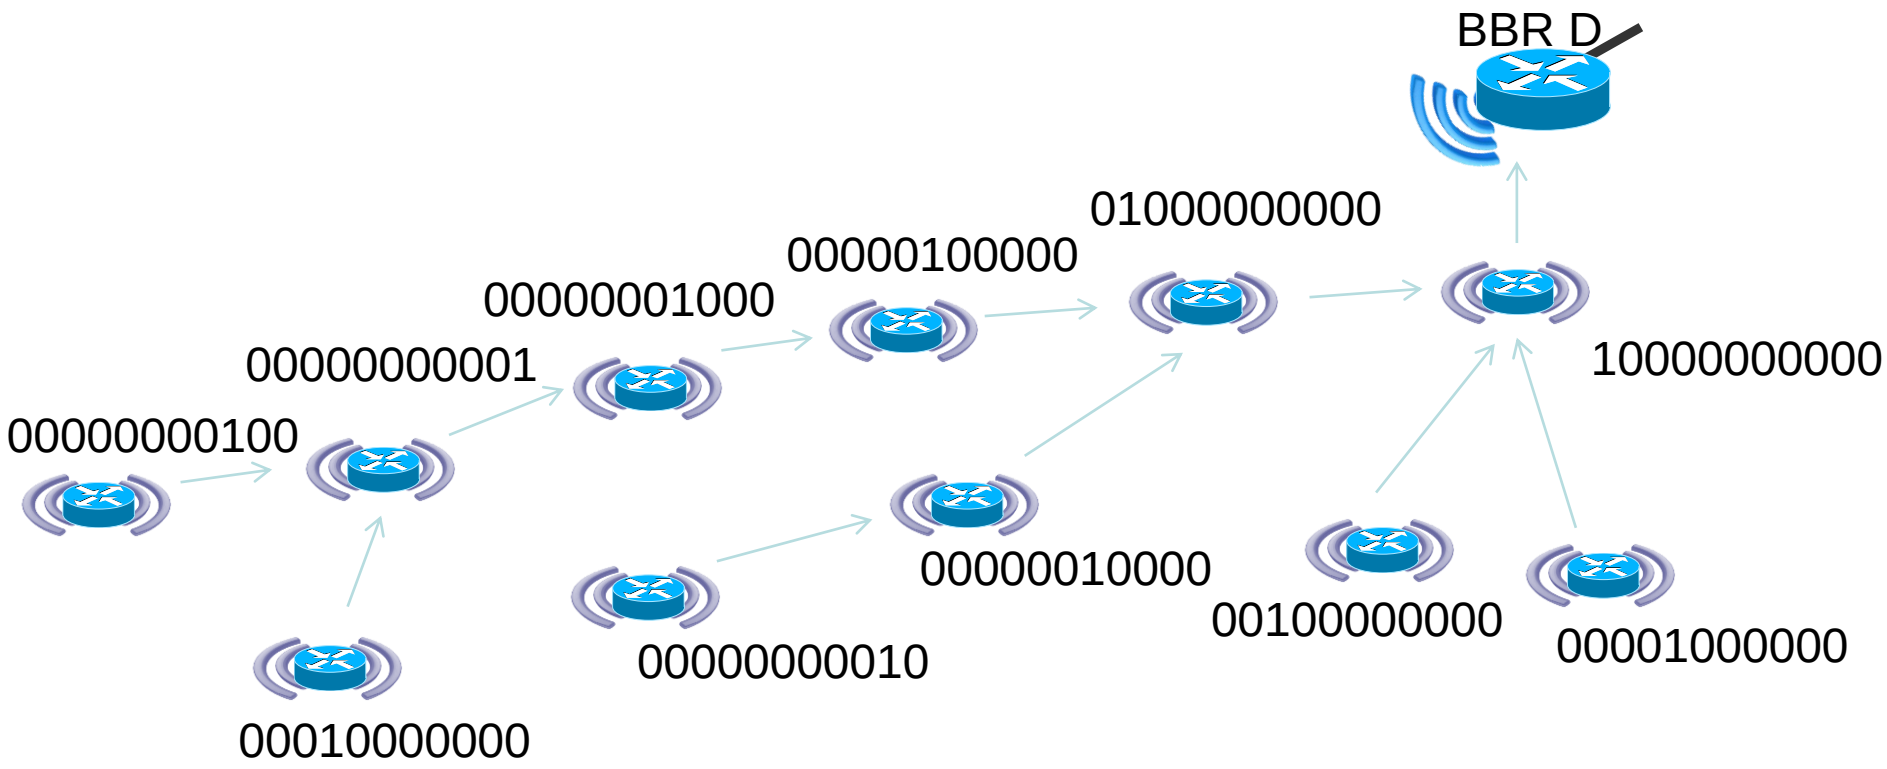

# Address to bit position table stored in root

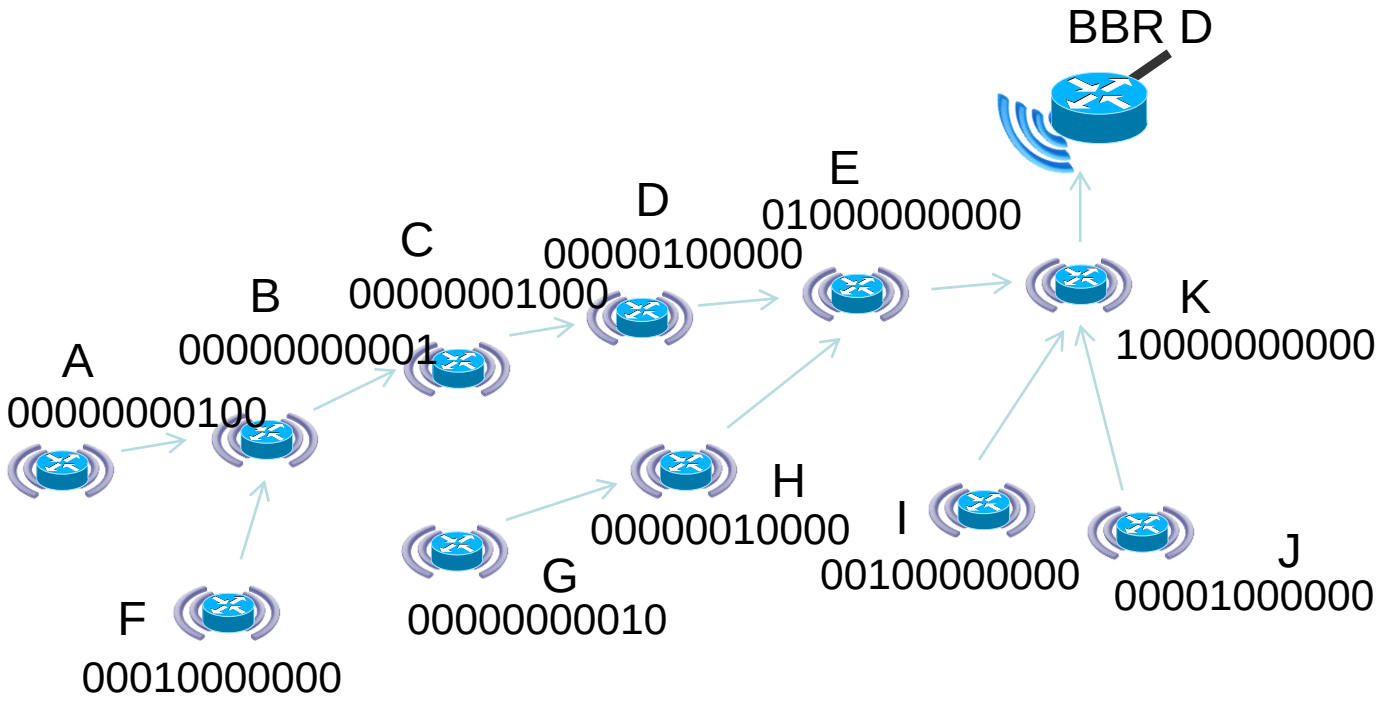


| Address | Bit |
|---|---|
| A | 9 |
| B | 11 |
| C | 8 |
| D | 6 |
| E | 2 |
| F | 4 |
| G | 10 |
| H | 7 |
| I | 3 |
| J | 5 |
| K | 1 |

# Aggregating bits in DAO operation

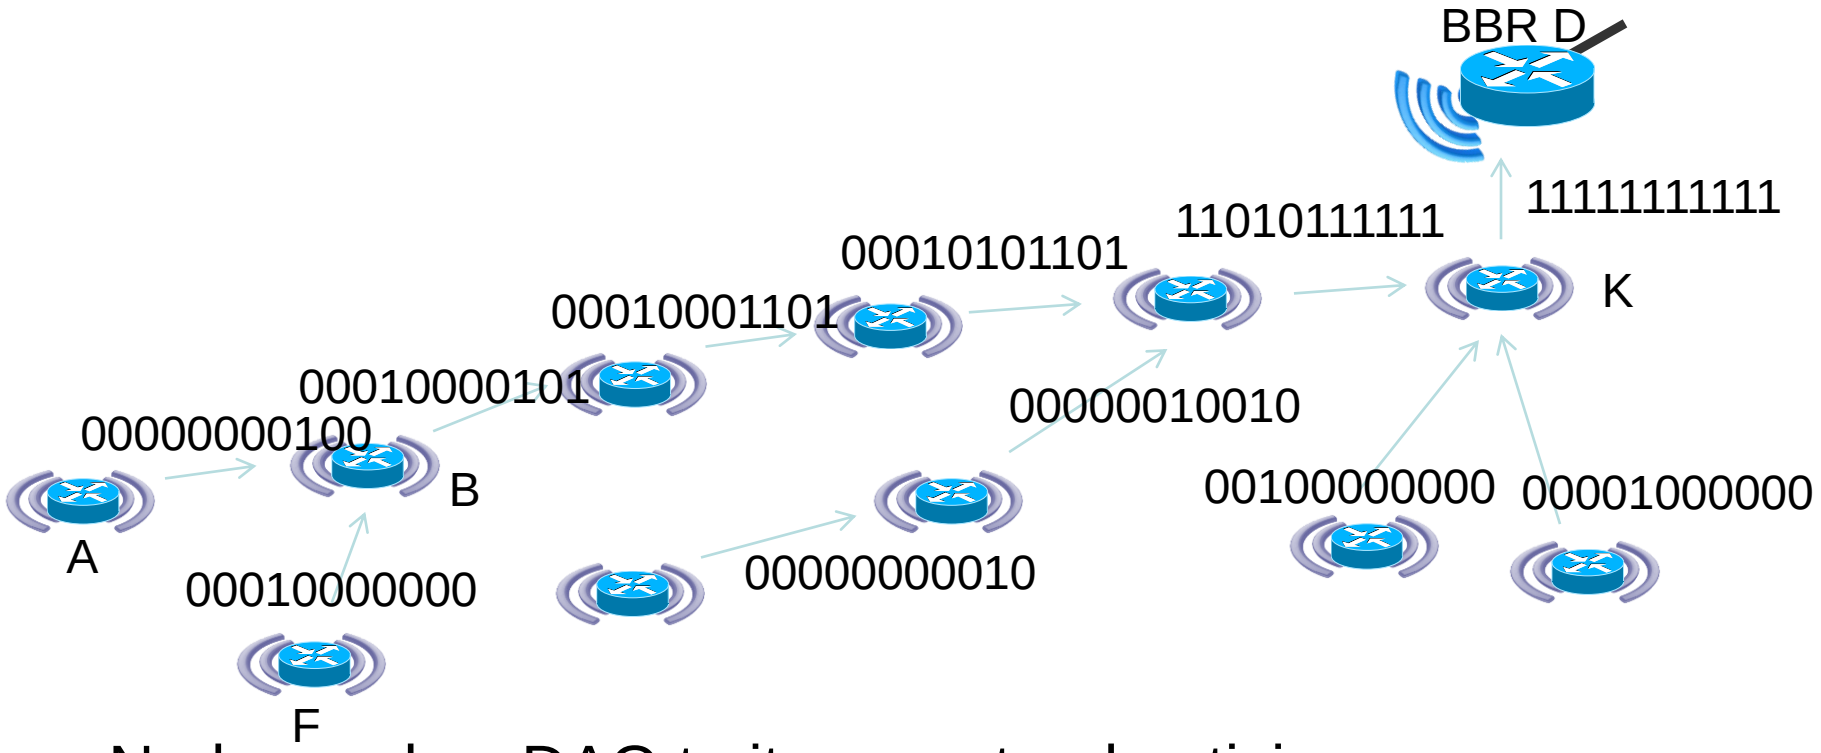


Node sends a DAO to its parent, advertising

n

Node’s bitmap = (OR child I’s bitmap) OR Node’s bit

i=1

# e.g. B's DAO operation

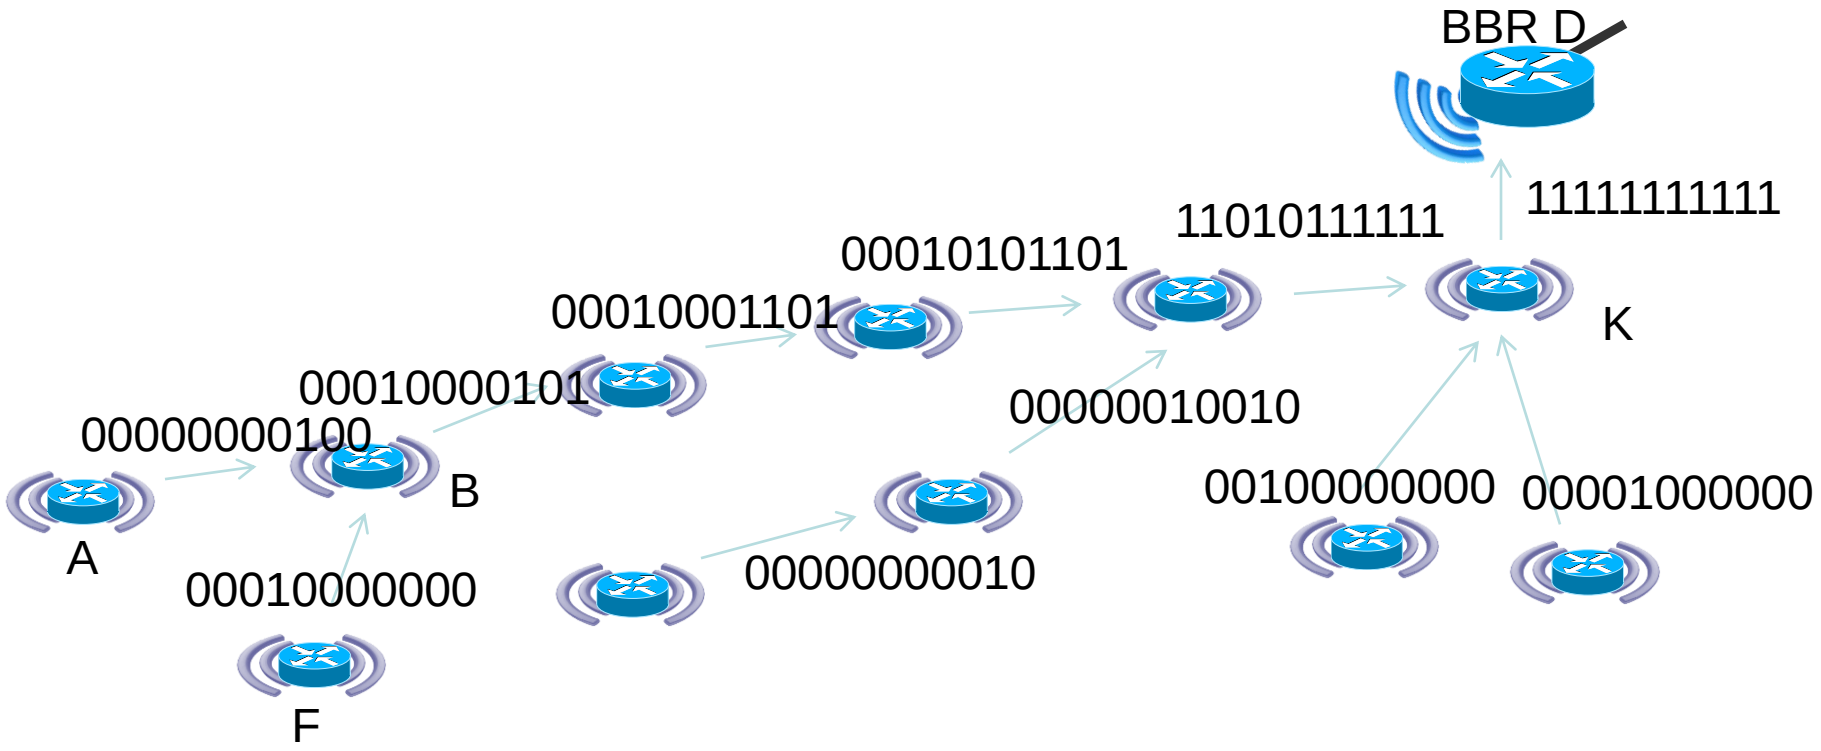


B sends a DAO to its parent, advertising B's bitmap = (A's bitmap OR F's bitmap OR B's bit)

# State in B

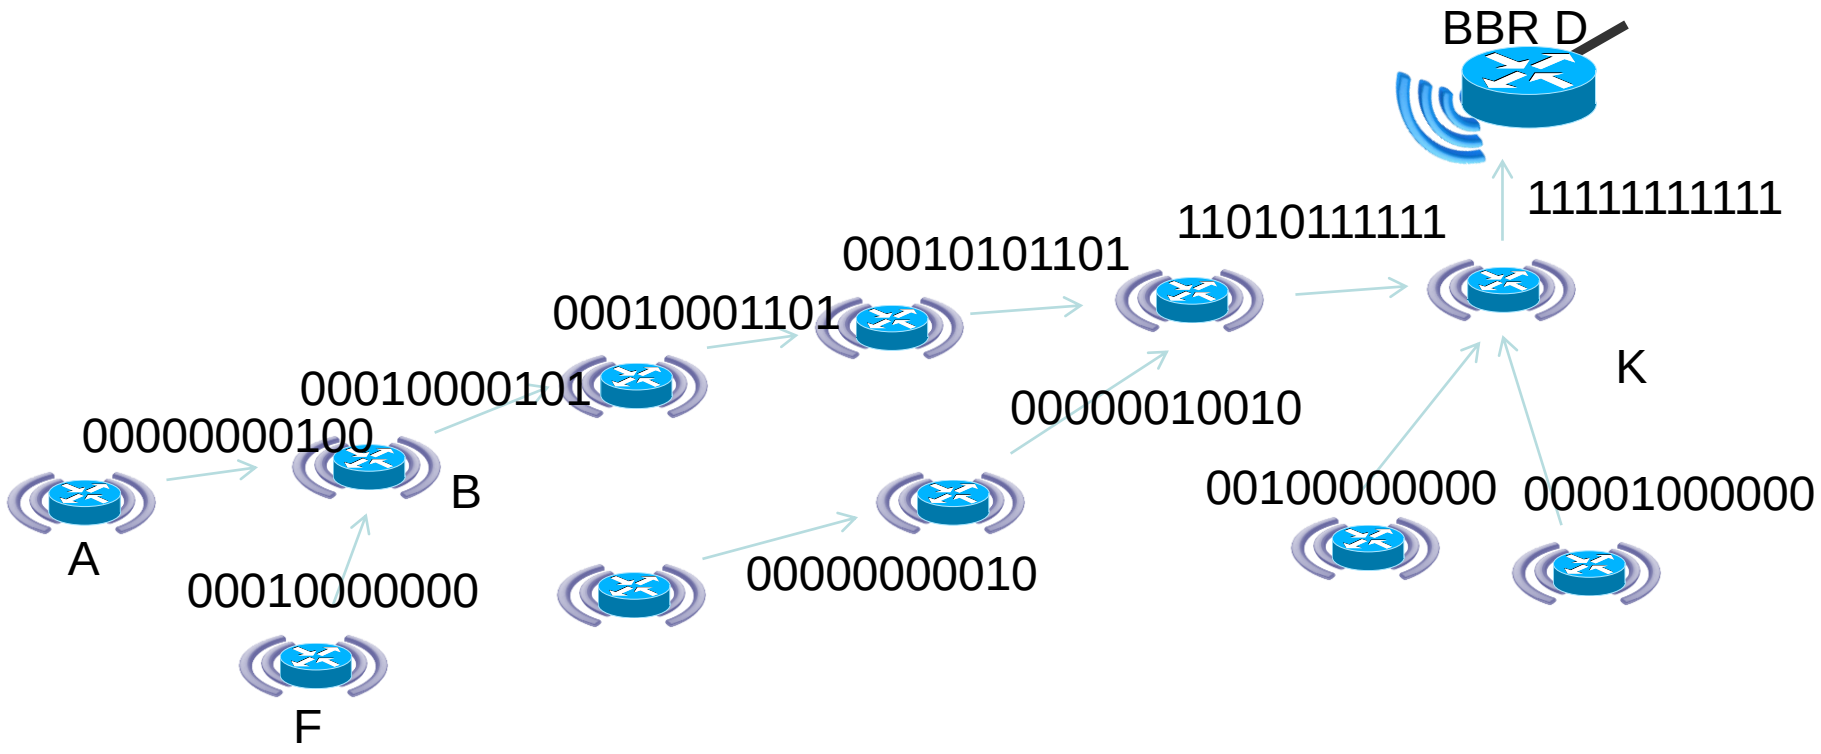


| Child | BitMap |
|---|---|
| A | 00000000100 |
| F | 00010000000 |

# State in K

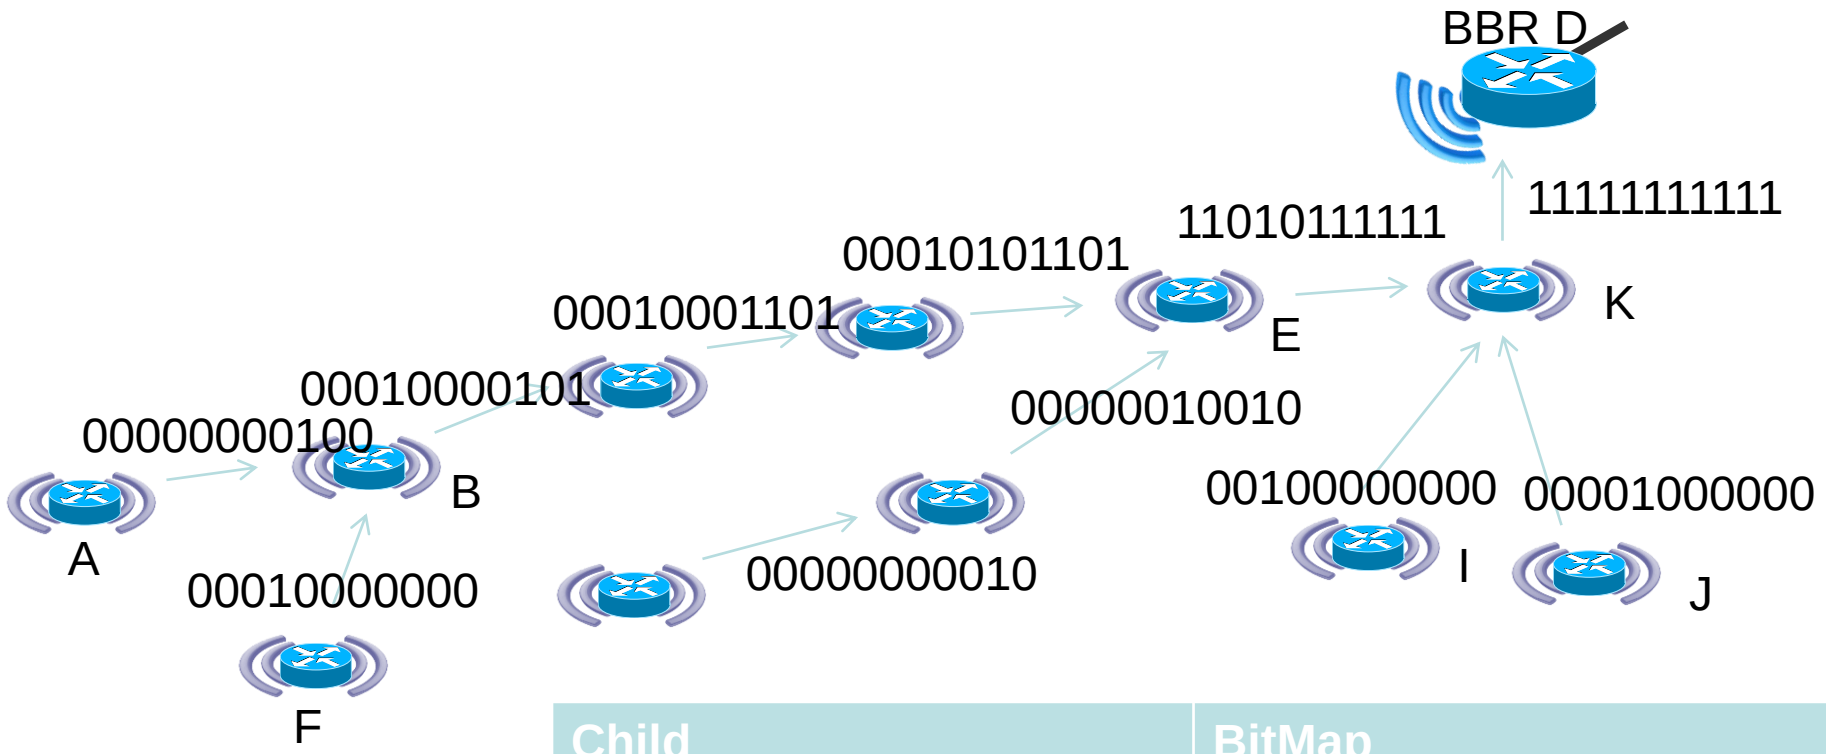


| Child | BitMap |
|---|---|
| E | 11010111111 |
| I | 00100000000 |
| J | 00001000000 |

# Message to n nodes

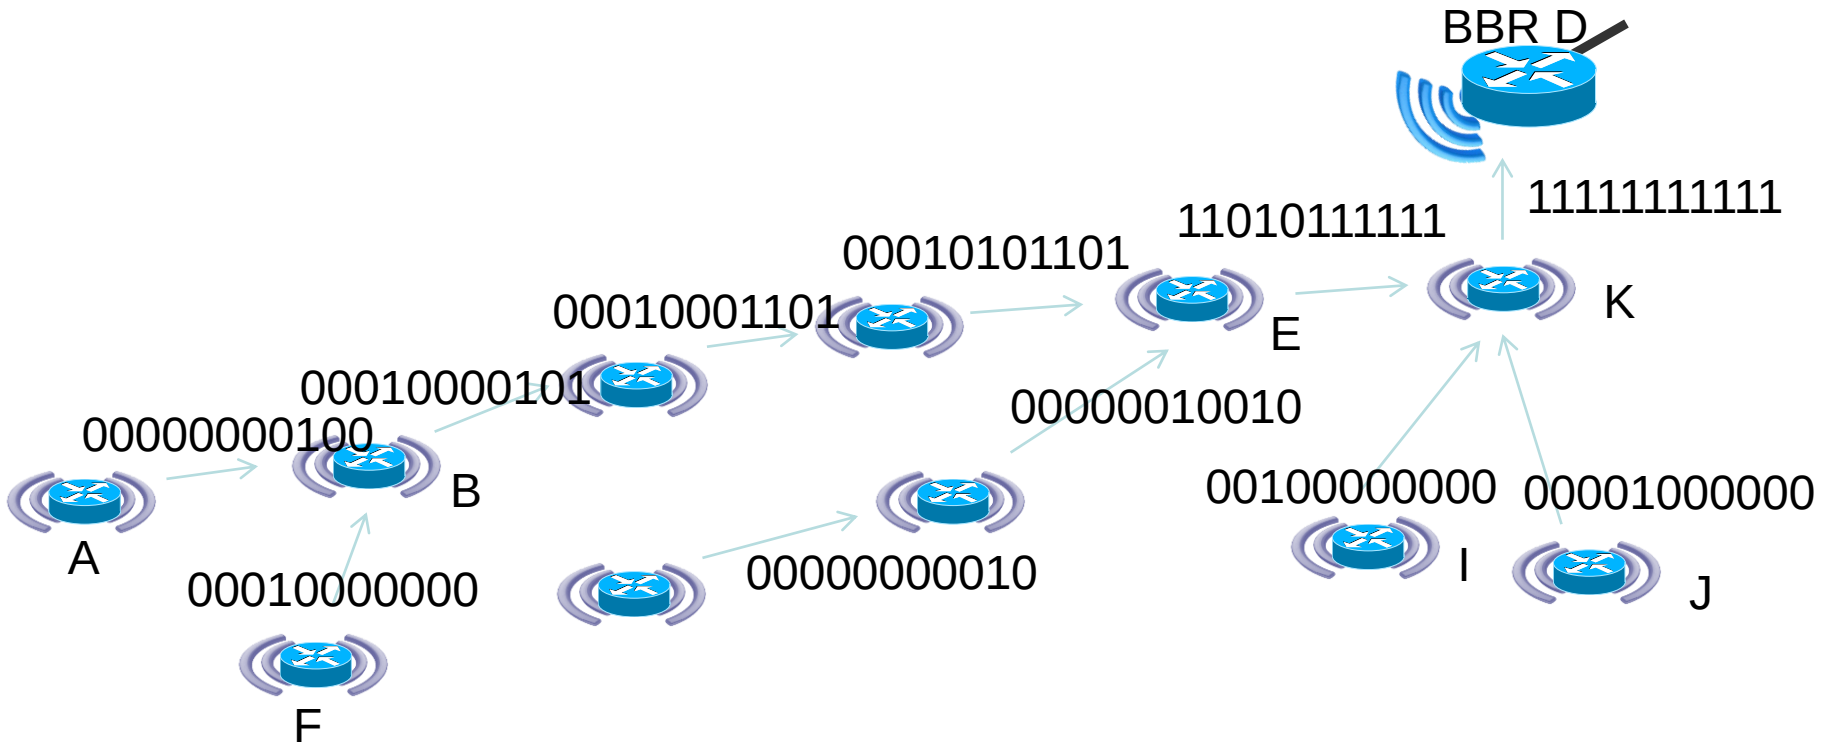


Root computes destination bitmap as

$$\text{Dest bitmap} = \bigvee_{i=1}^{n} (\text{OR node } i\text{'s bit})$$

# Forwarding operation

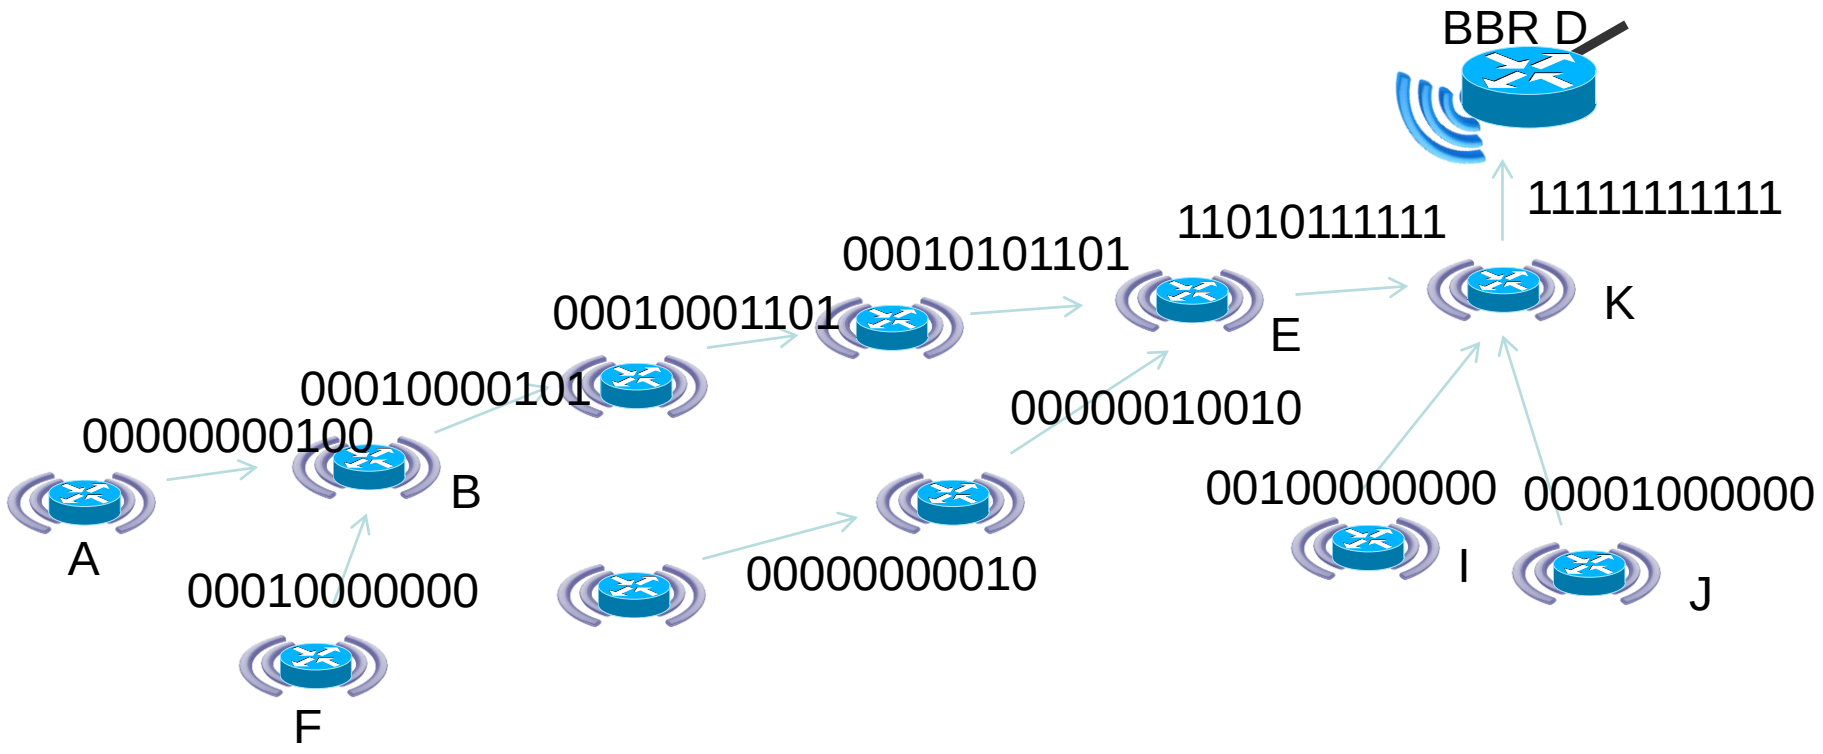


Node computes (Dest bitmap AND child's bitmap) for all children

When result is TRUE (non-zero), node copies the packet as a MC level unicast to child.

# Alt Forwarding operation

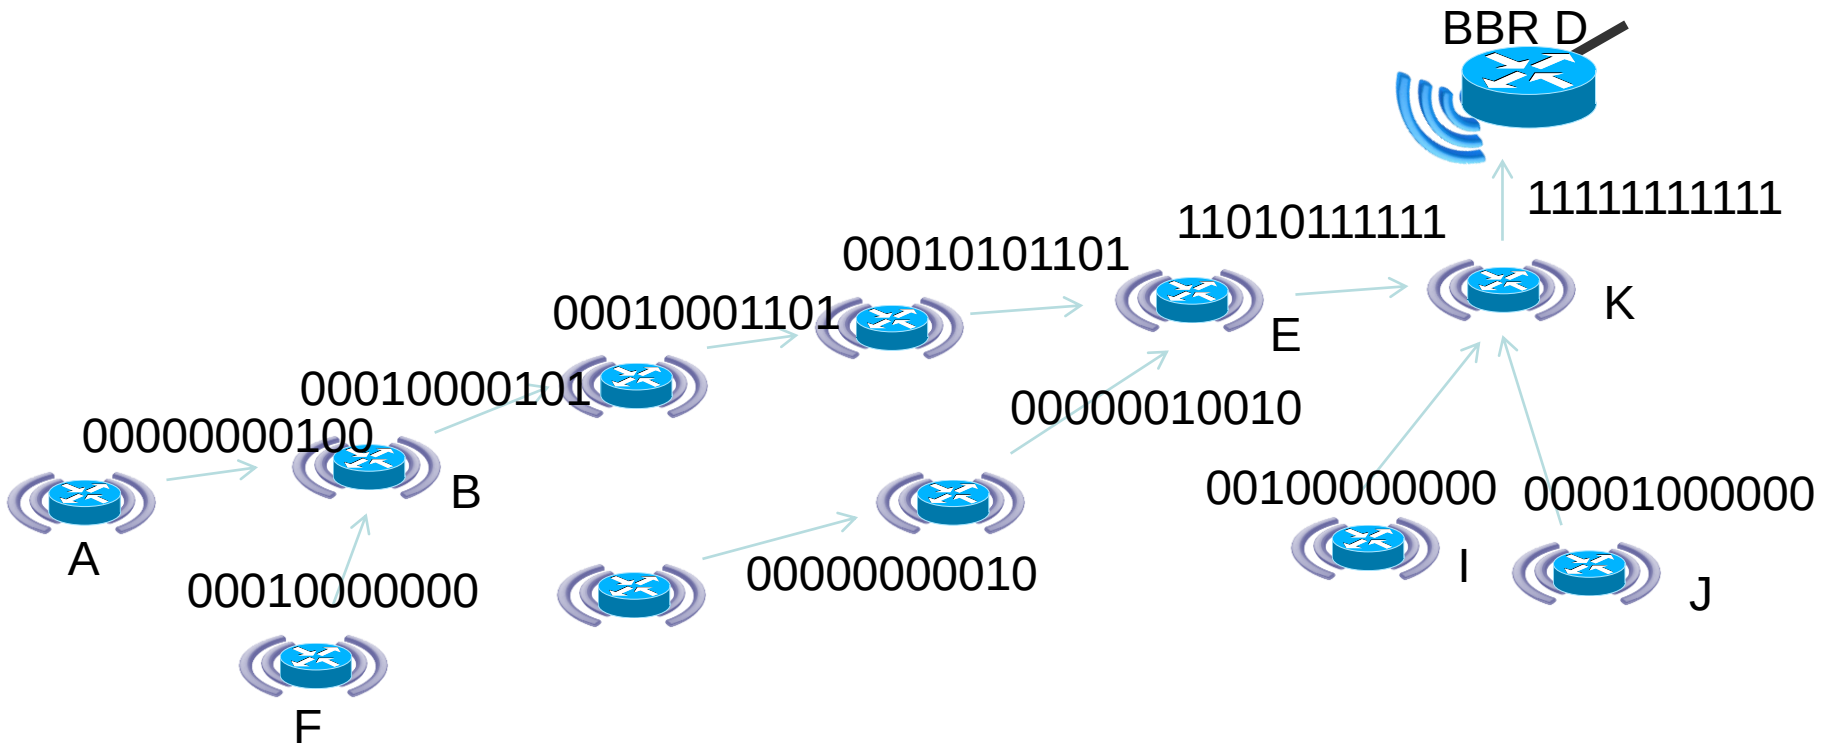


If most of the children are targetted, it makes sense to broadcast the message to all children. In that case, receiving children perform the OR operation with the bitmap they advertise in DAO and drop on receive is the result is not TRUE

# e.g. Message to A, F and J

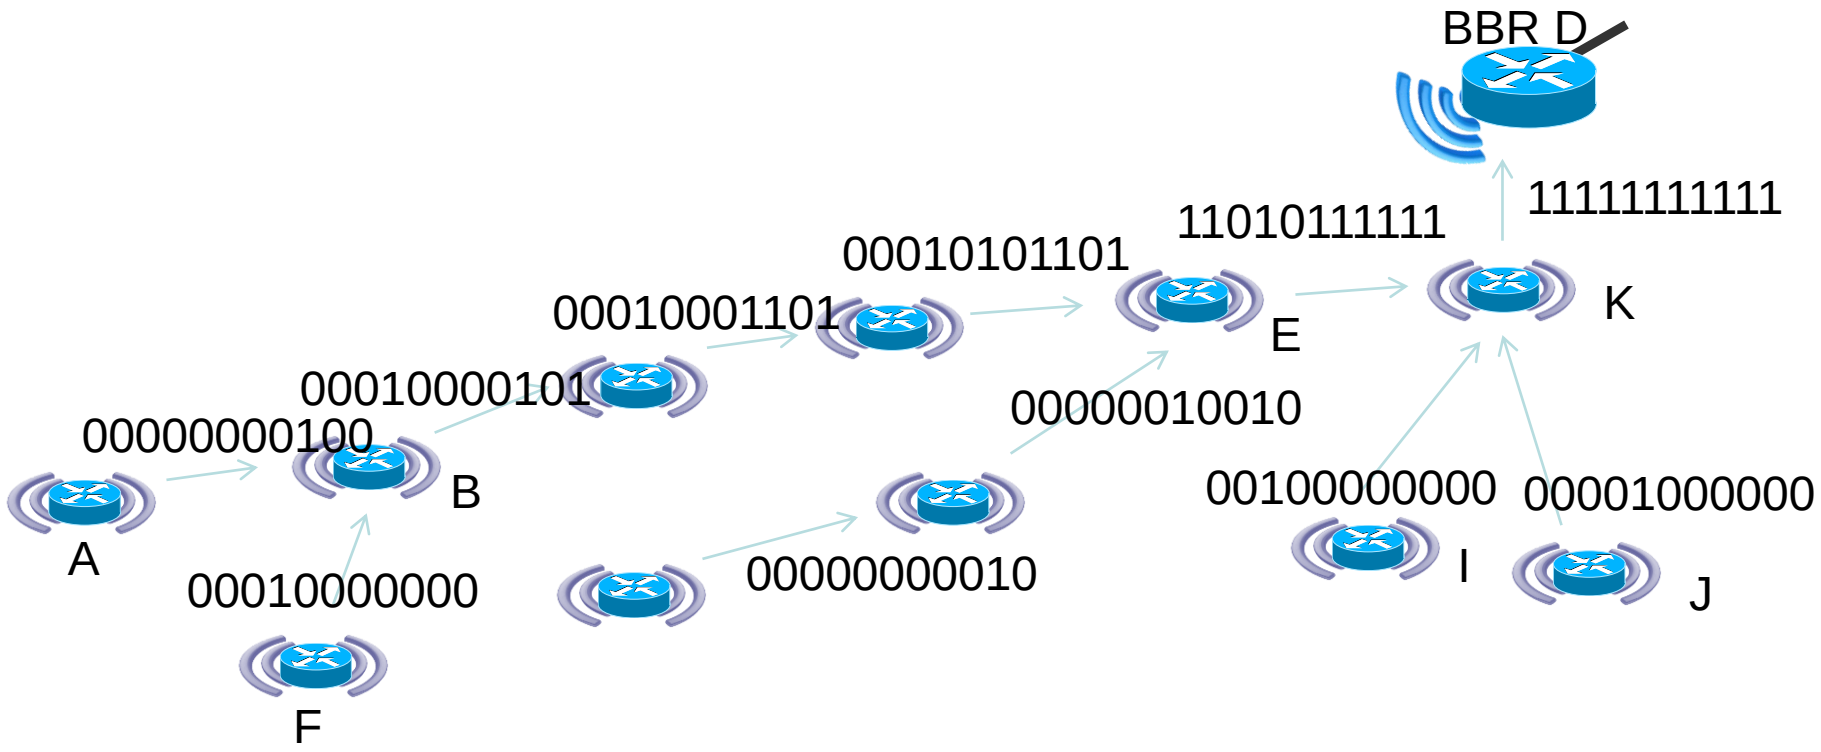


Root computes destination bitmap as
Dest bitmap = (A's bit OR F's bit OR J's bit)
= 00011000100

# K’s Forwarding operation

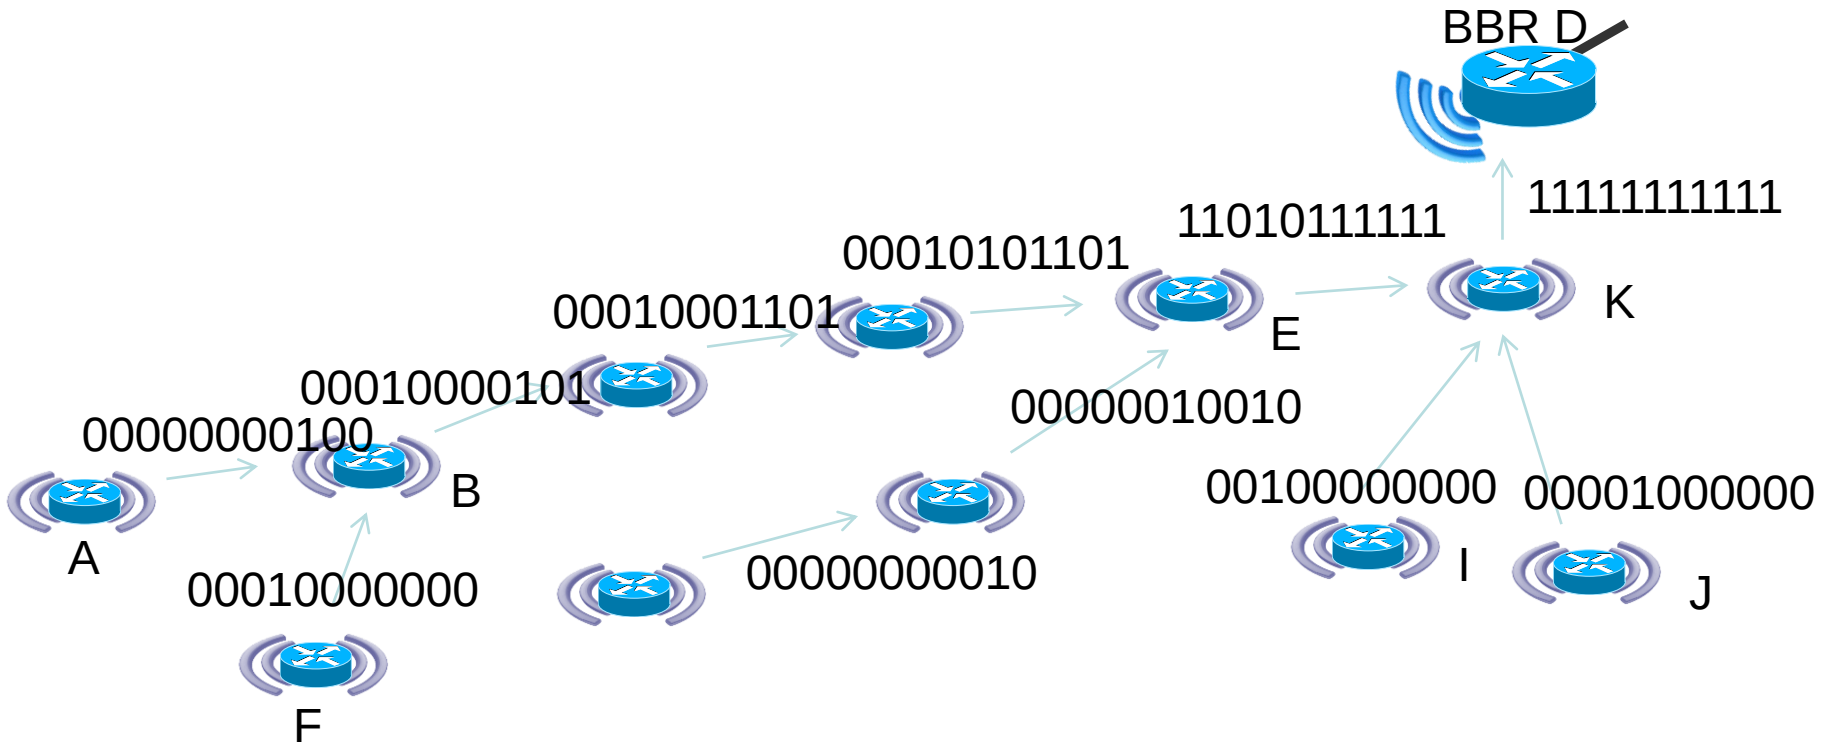


Dest bitmap = 00011000100
(Dest bitmap AND J’s bitmap) = 00001000000 -> MAC unicast to J
(Dest bitmap AND I’s bitmap) = 00000000000 -> NO copy to I
(Dest bitmap AND F’s bitmap) = 00010000100 -> MAC unicast

# B's Forwarding operation

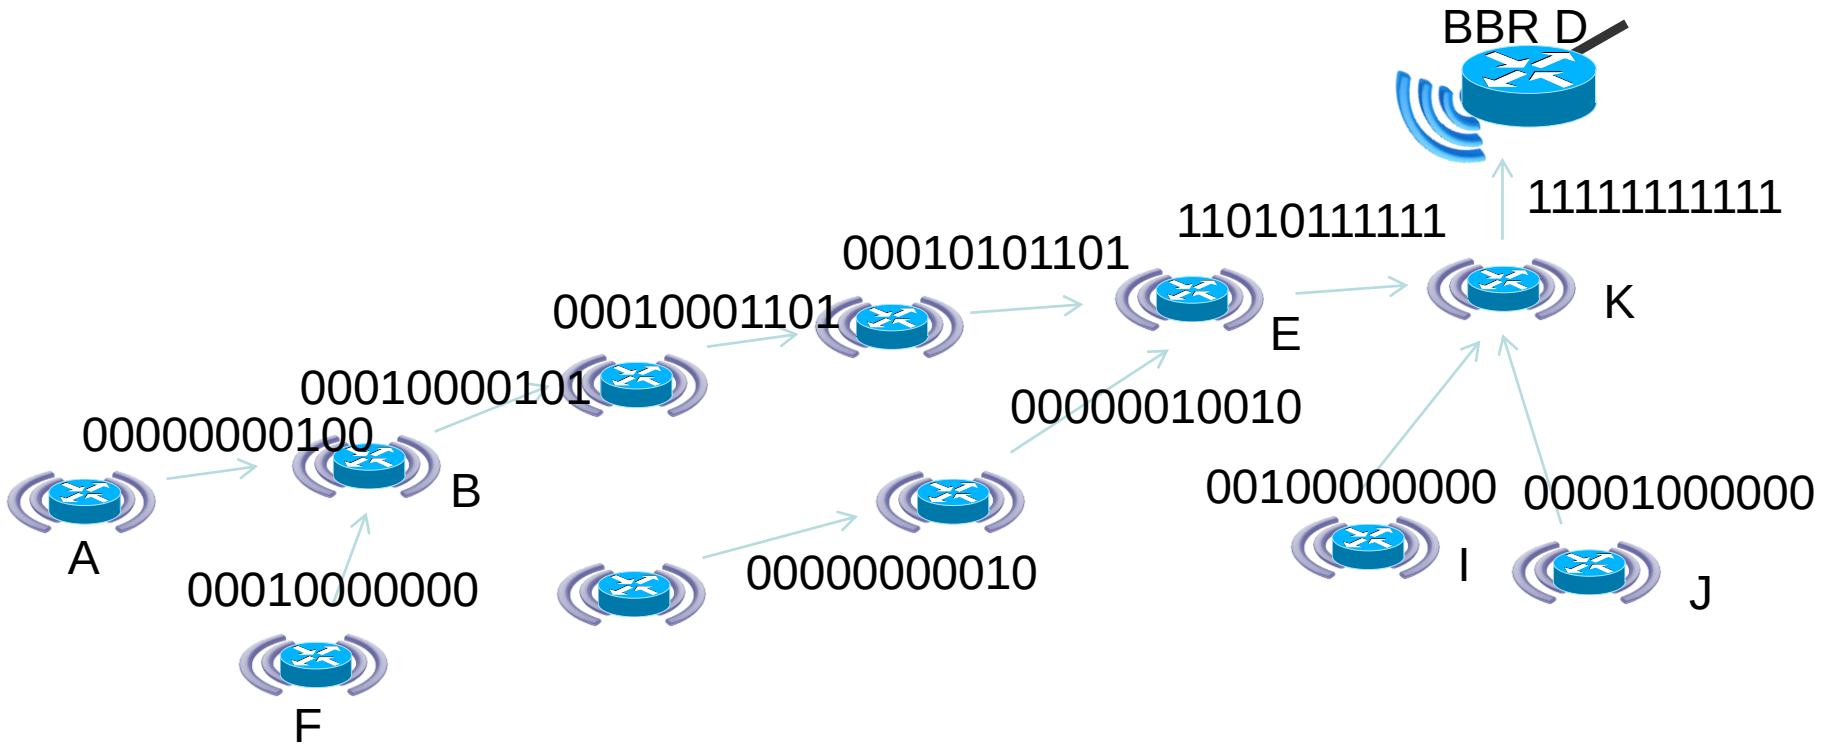


Dest bitmap = 00011000100
In B: (Dest bitmap AND B's bitmap) = most bits set -> B broadcasts
In A: (Dest bitmap AND A's bitmap) = 00000000100 -> A accepts

# Reliable BIERPL

# e.g. reliable mcast message to A, F and J

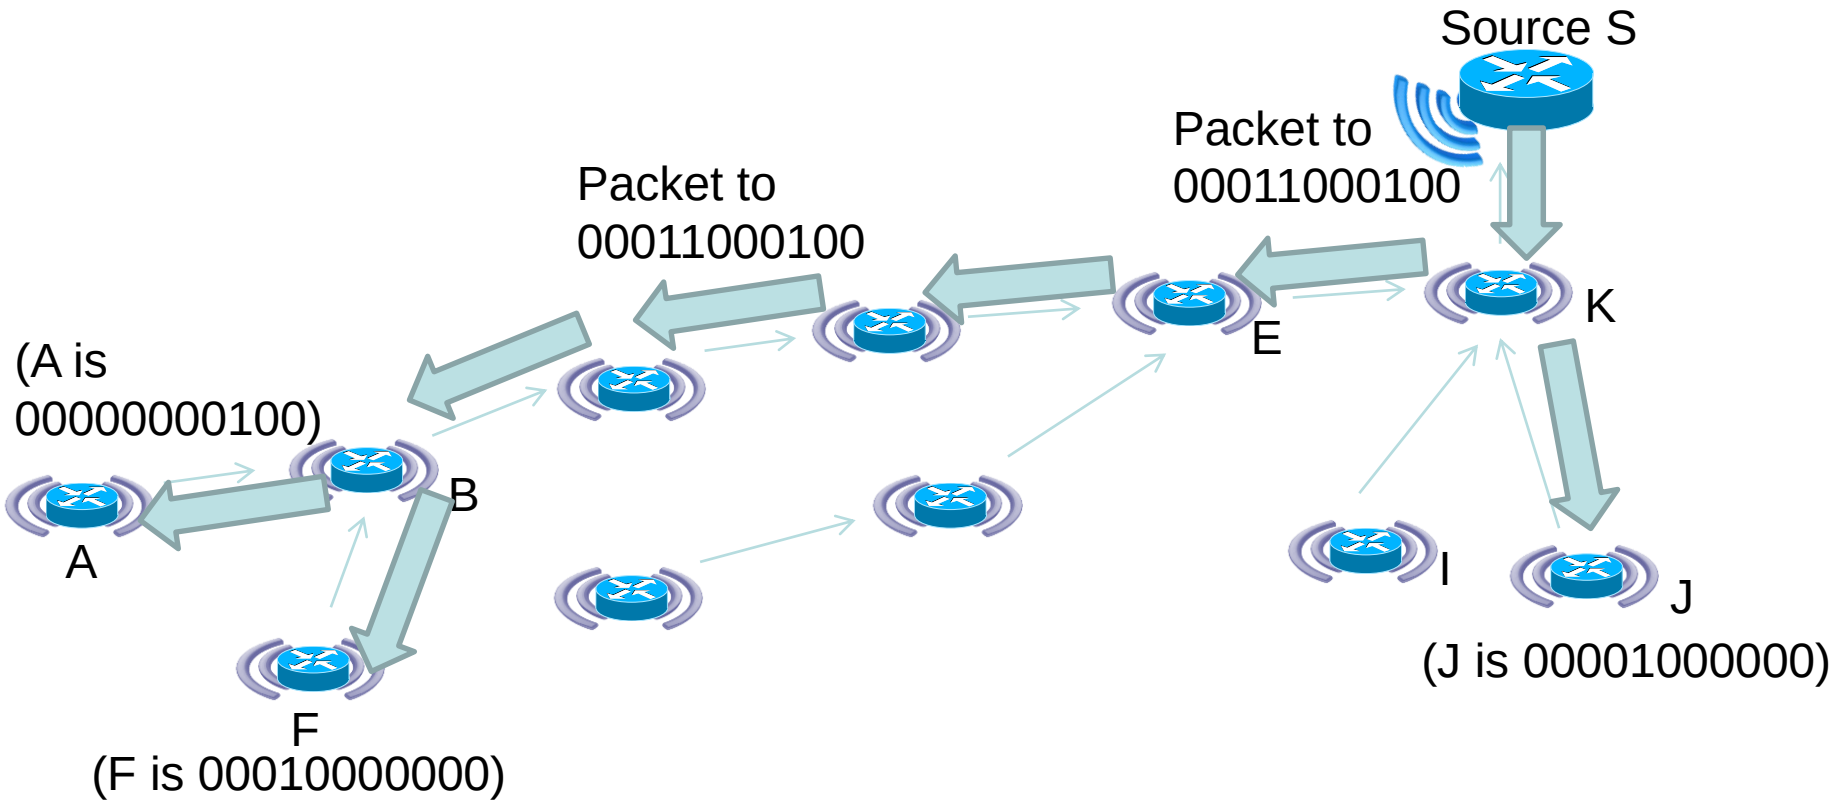


Root computes destination bitmap as
Dest bitmap = (A's bit OR F's bit OR J's bit) =
00011000100
Forwarding expected to follow

# e.g. reliable mcast message to A, F and J

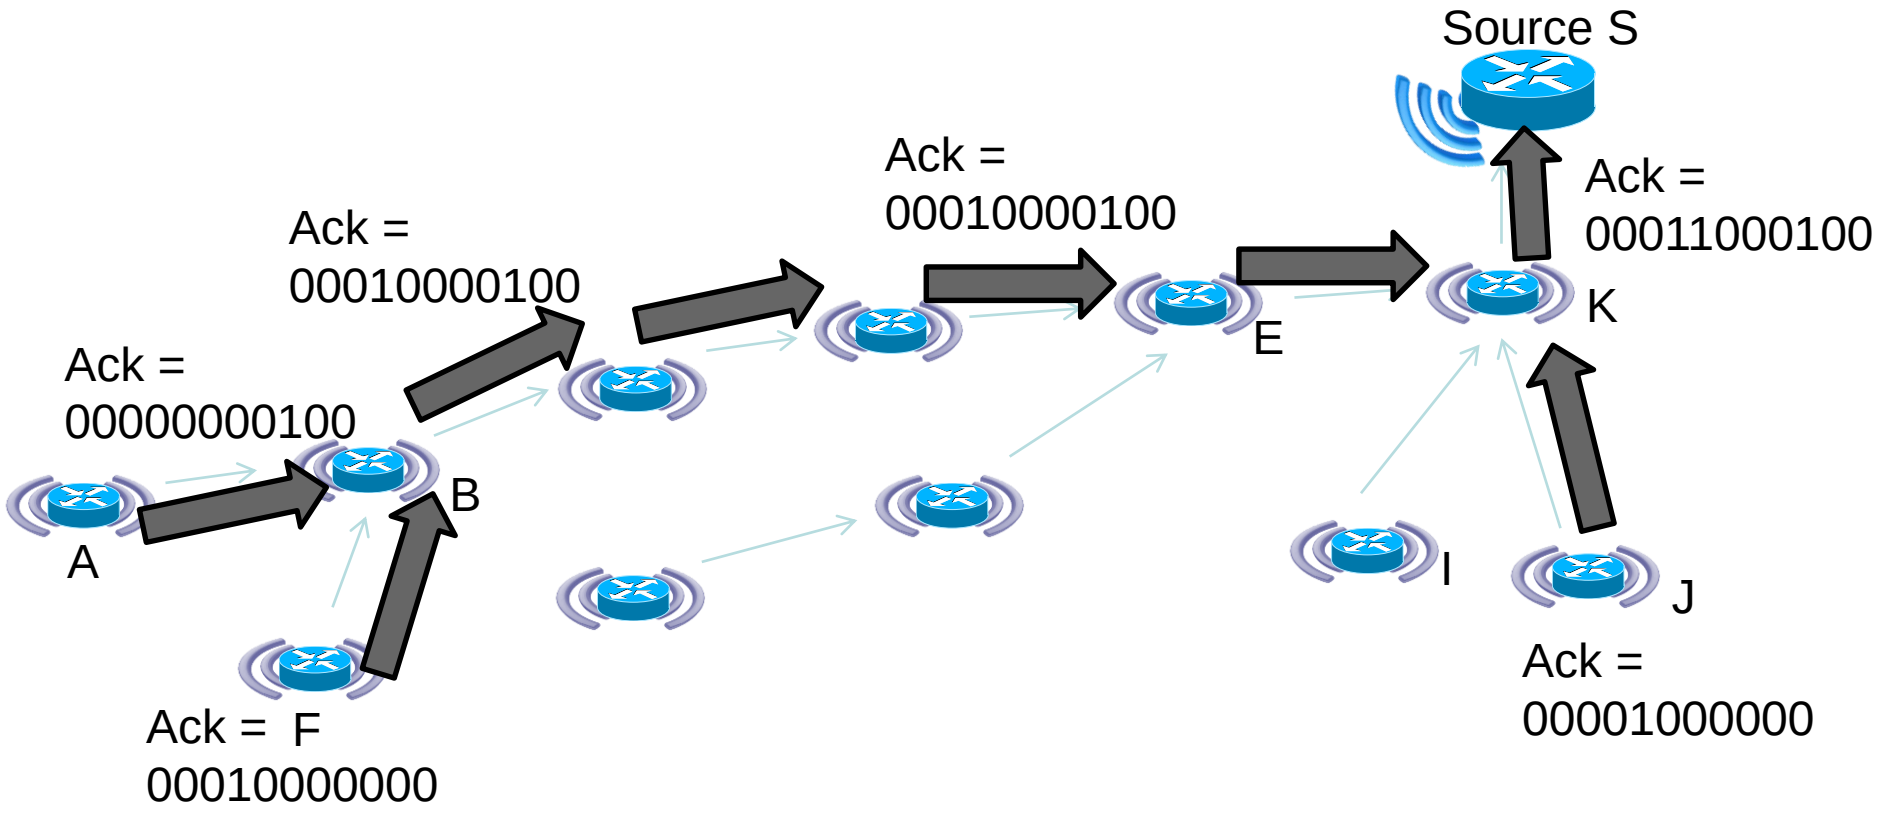


Acks are aggregated on the return pa

# Transmission failure down a branch

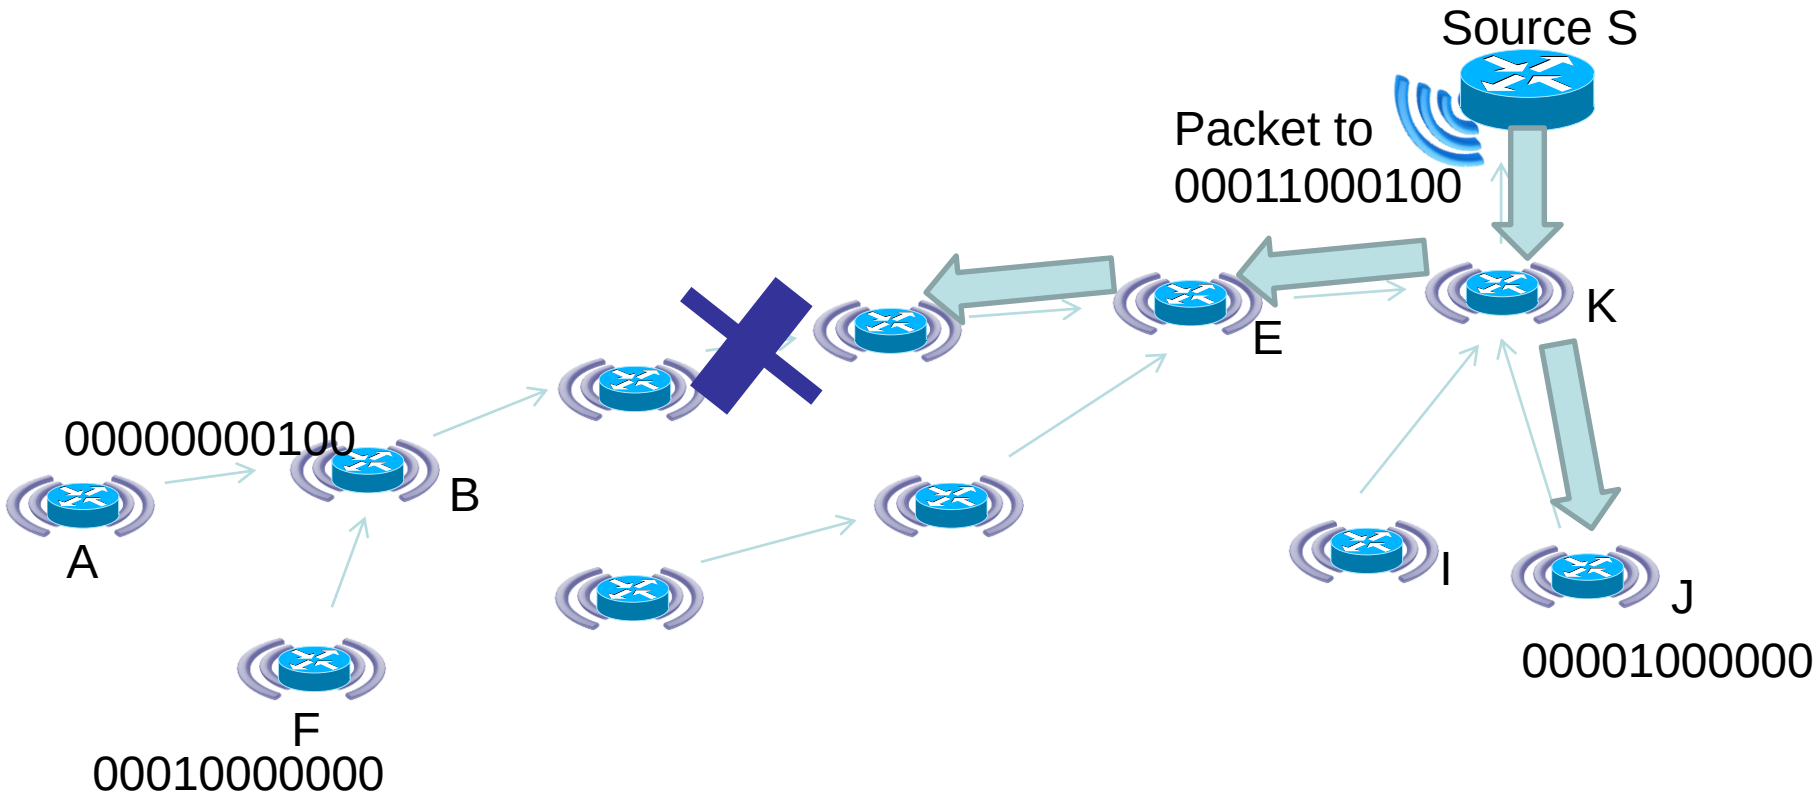


loss on the way in the branch that leads to A and F

# Ack bitmap reports only J reception

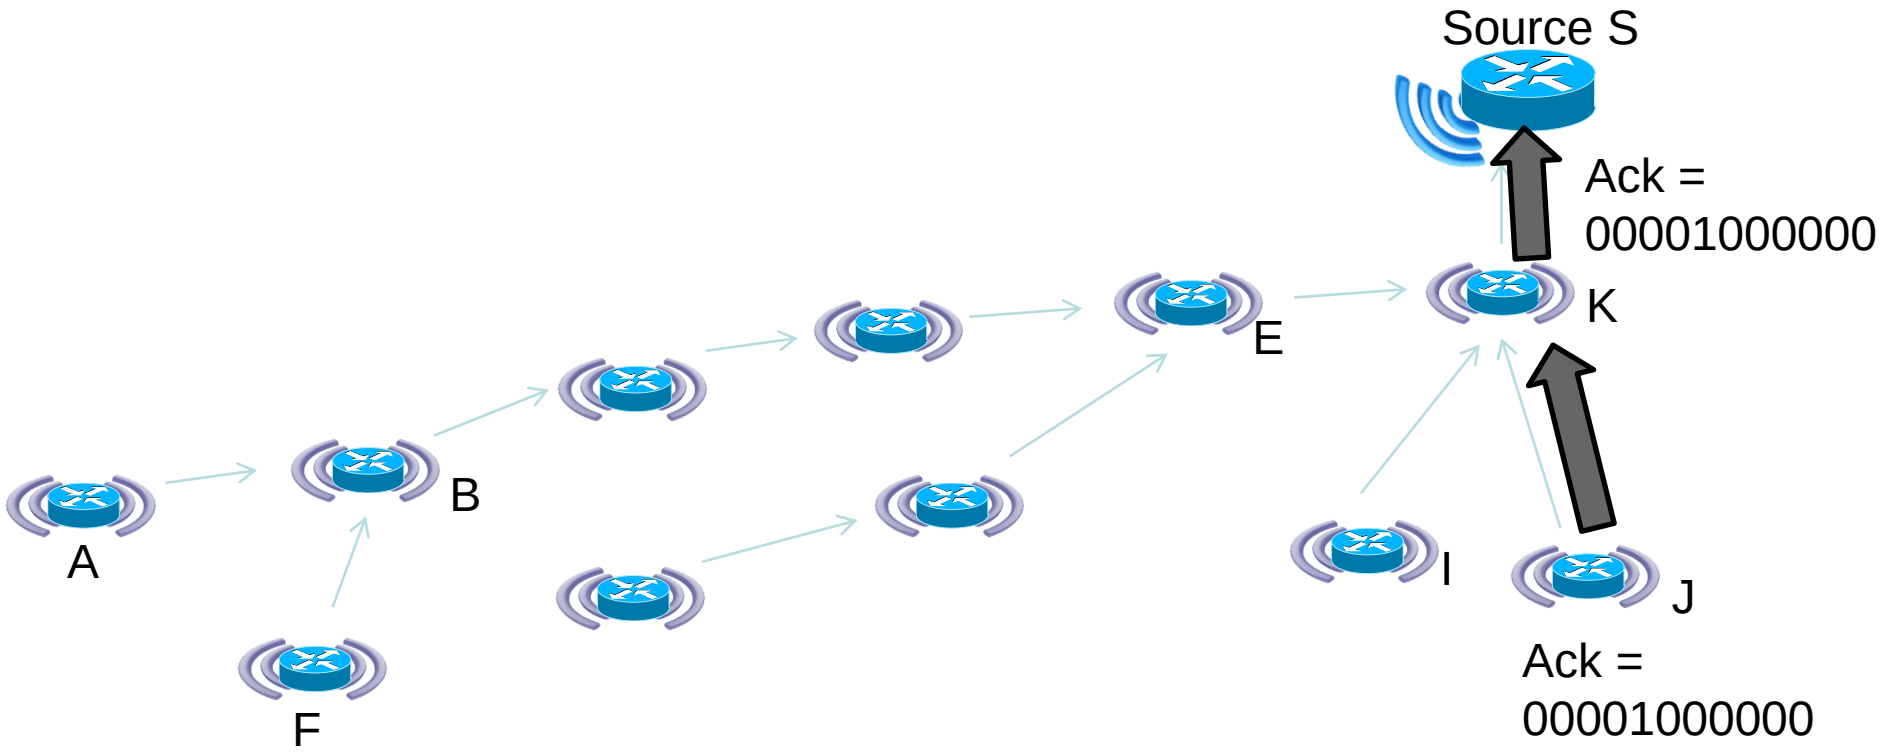


Root computes destination bitmap – ack bitpmap
Retrans bitmap = Dest bitmap - Ack bitmap
= 00011000100 - 00001000000
= 00010000100

# e.g. reliable mcast message to A, F and J

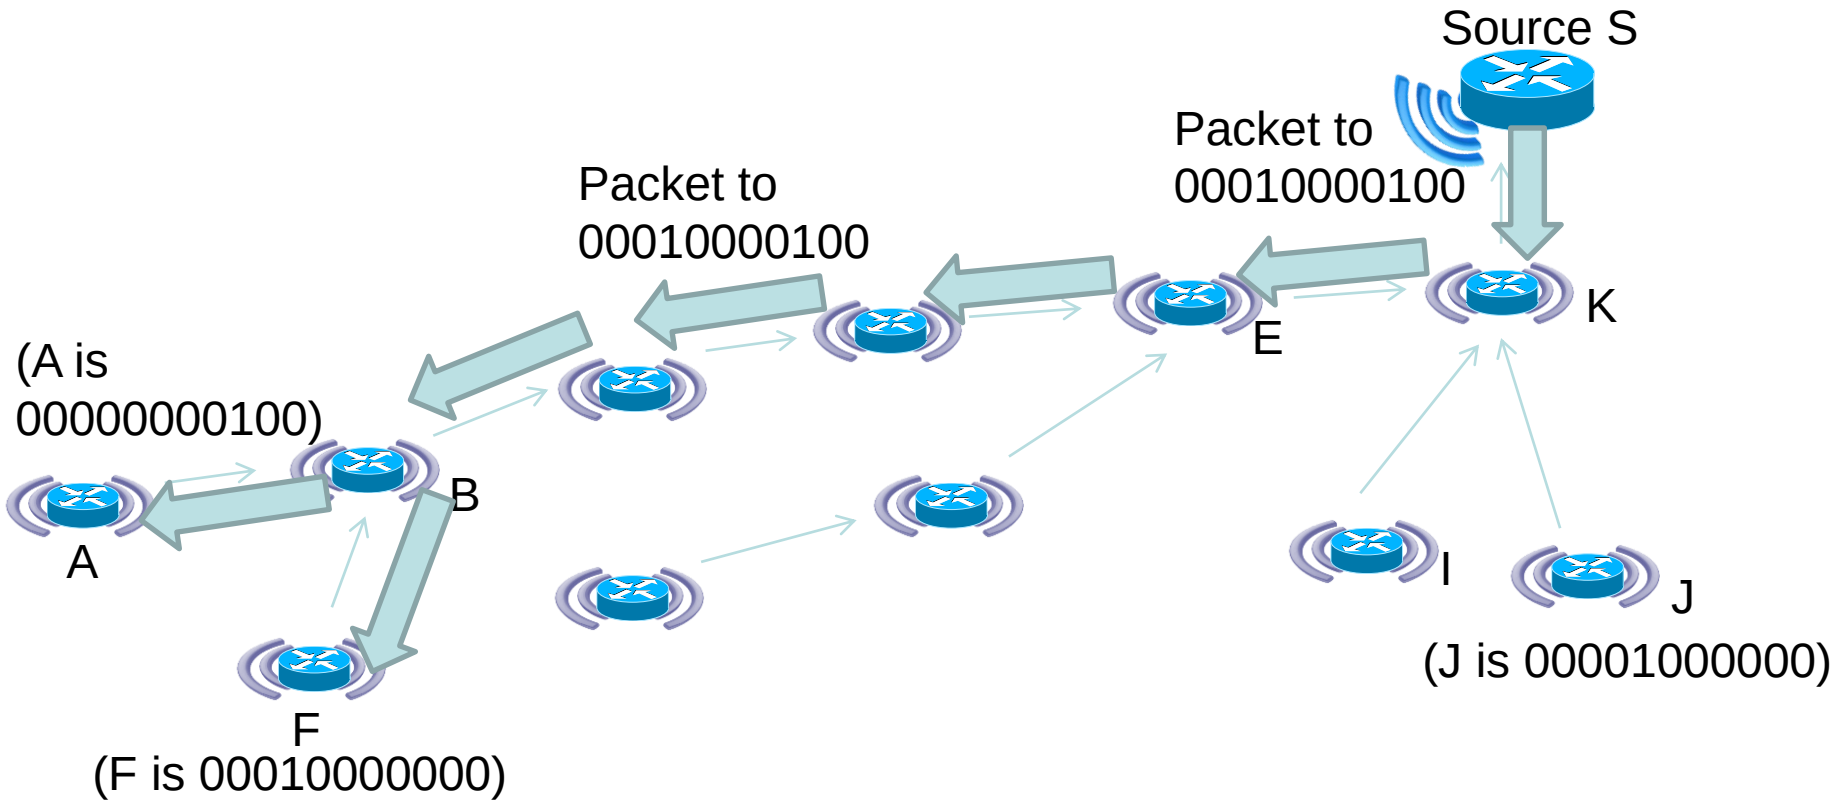


Retransmission bitmap indicates only along failed branch(es)

Forwarding expected to follow

# Ack from A and F

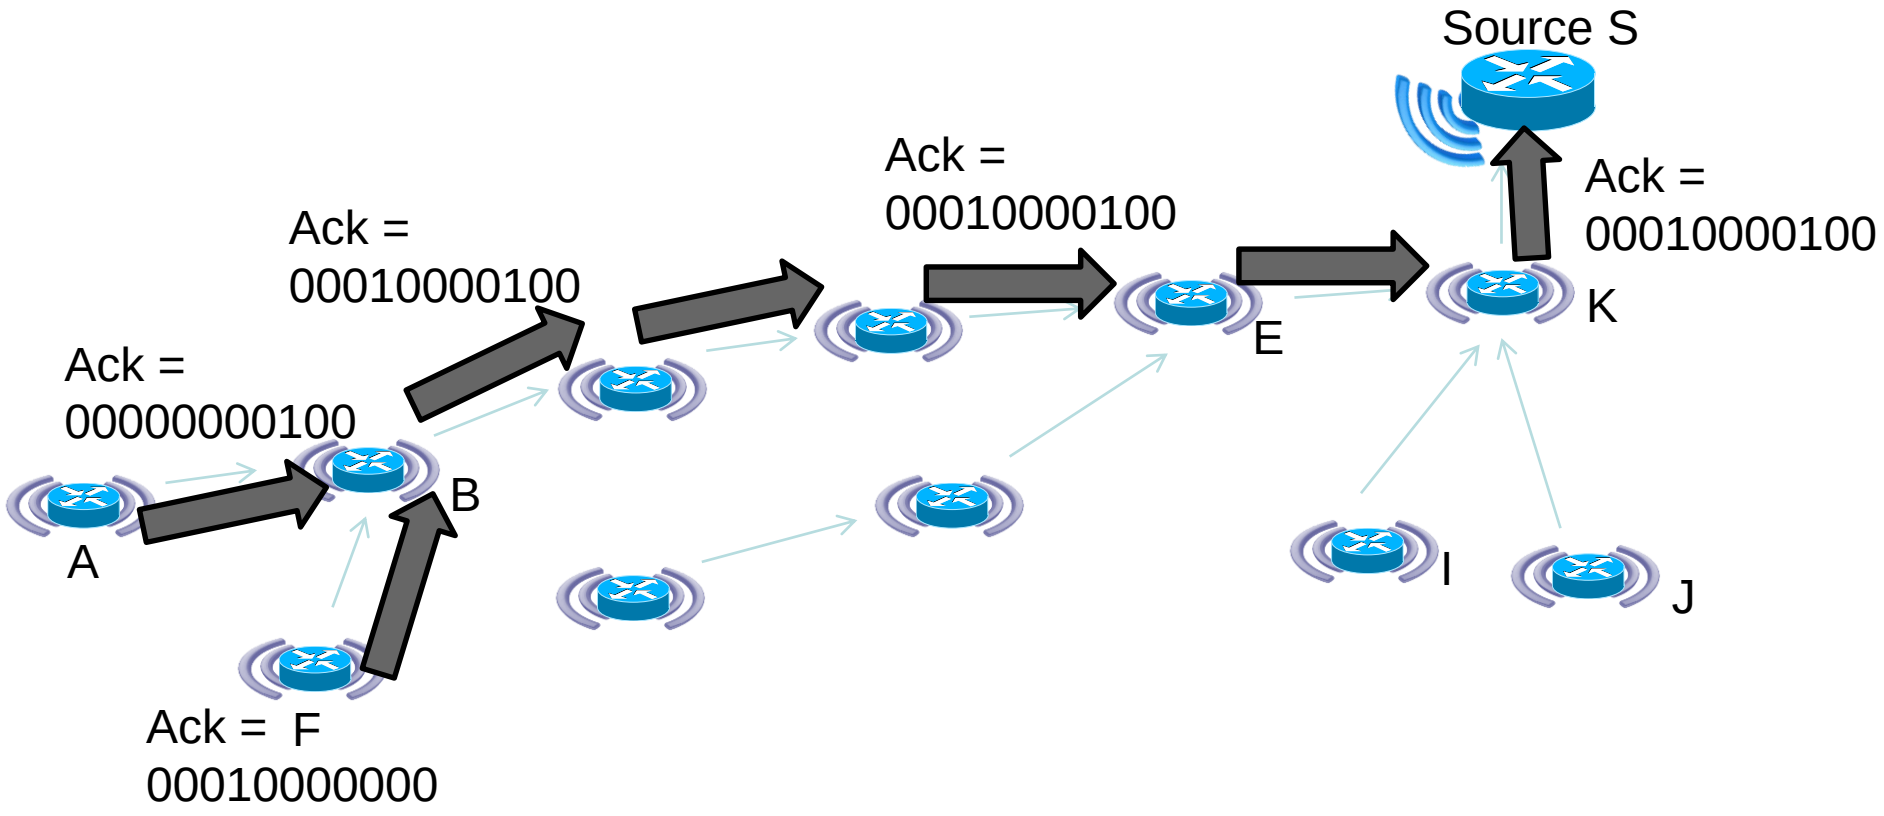


Acks are aggregated on the return pa

# Enhanced Beacon

## Metrics and Values (version 2)

Michael Richardson
mcr+ietf@sandelman.ca

https://www.sandelman.ca/SSW/ietf/meeting/ietf101/ietf101_6tisch_roll_beacon_info

# Overview

- What’s the problem?
- What’s the 6tisch part?
- What’s the ROLL part?
- What’s the problem?
- Discussion and Questions.

Same slides for 6tisch and for ROLL, but different discussion!

# What's the problem? Network Selection

- A (new!) device (pledge!) will not know which network it should enroll in.
- A single network will be visible multiple times.


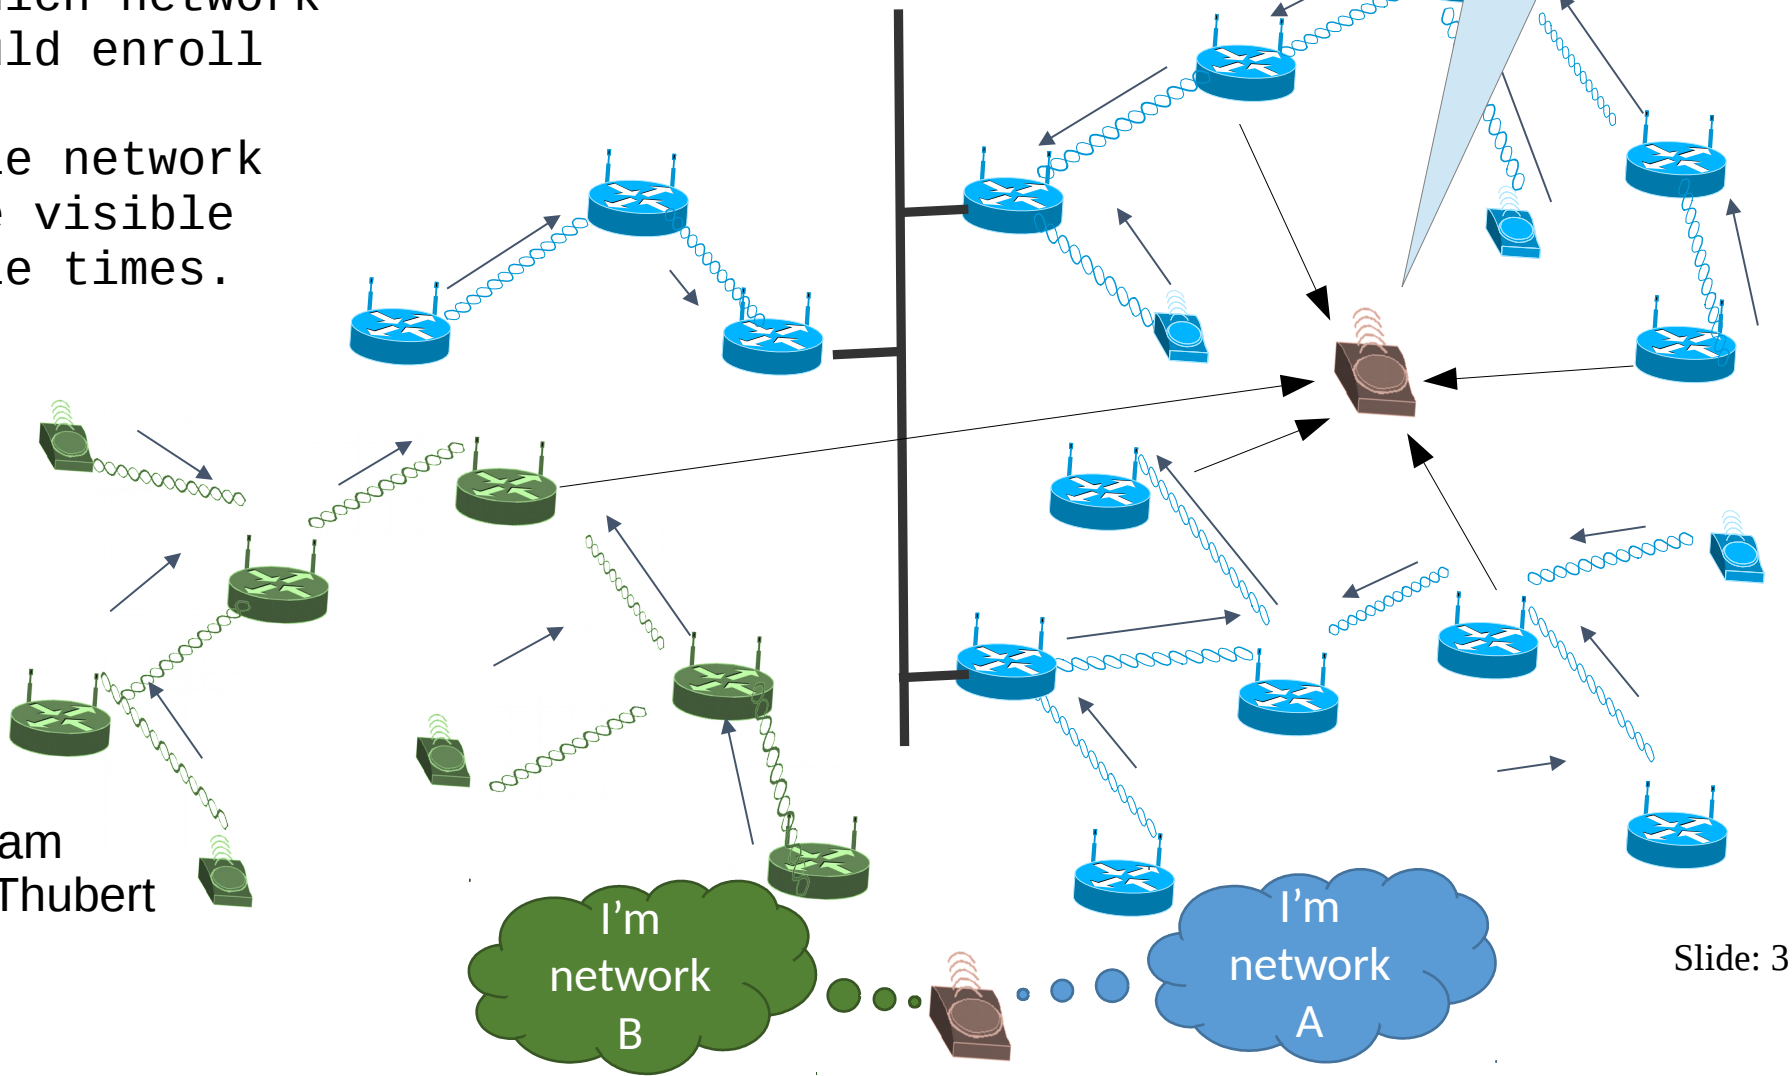


Diagram
By P.Thubert

# Different meaning of Join

- getting the network keys/ credentials
  - ENROLLMENT
- JOINing a DODAG
  - Parent Selection

# Observations

- Clearly, Blue network and Green network are different.
- Blue Network may have three PANids, and therefore MUST have three different keys.
- Different PANids leads to different IPv6 derived addresses.
- In order to balance the load on the network, the ENROLLMENT decision must be informed by the network load.
- We prefer intelligent end nodes, so express some of the network information into Beacon.

# What's the 6tisch part?

IEEE802.15.4 Informational Element encapsulation of 6tisch Join and Enrollment Information
draft-richardson-6tisch-enrollment-enhanced-beacon-00

Creates a new 802.15.4 **Informational Element**, using the IETF allocation in rfc8137

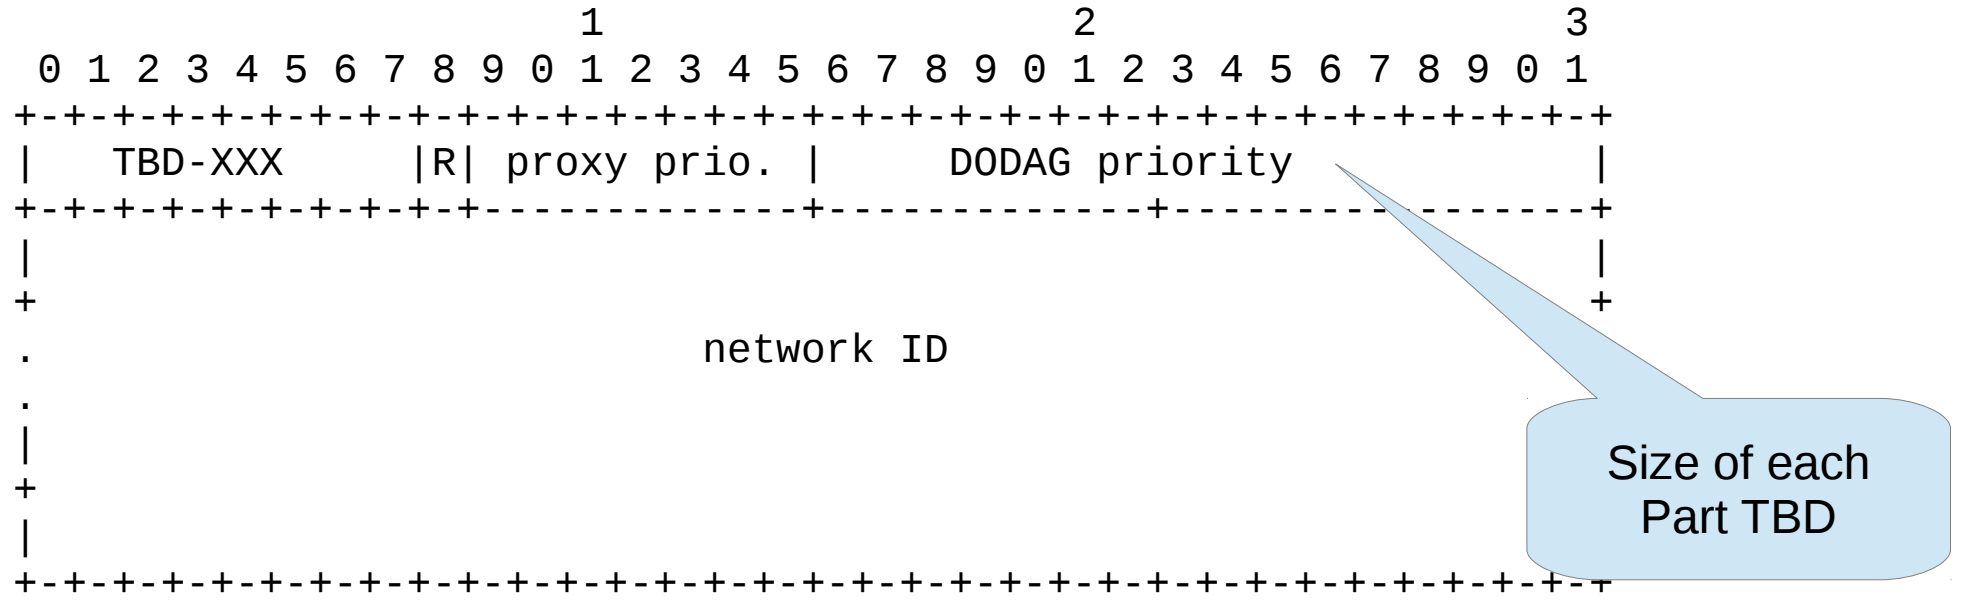


- R a flag to indicate device will answer unicast Router Solicitations
- Network ID is variable length nonce, probably derived from PIO.
- DODAG Priority: how desirable is this **network**
- Proxy prio: how desirable is this **proxy**
- Rank Priority: already part of another IE

# Different preferences

- Network ID
  - Shows which networks are the same

    (blue networks vs green networks)
- DODAG Preference ID (needs to be sent in DIO)
  - Less desirable as network gets more busy.
- Proxy Preference ID (calculated locally by Proxy)
  - Determines which proxy has more capacity.
- Join proxy rank
  - Tie breaker between different proxy, a local property.

# What's the ROLL part?

Enabling secure network join in RPL networks
draft-richardson-6tisch-roll-enrollment-priority-02

Defines a new metric Option:

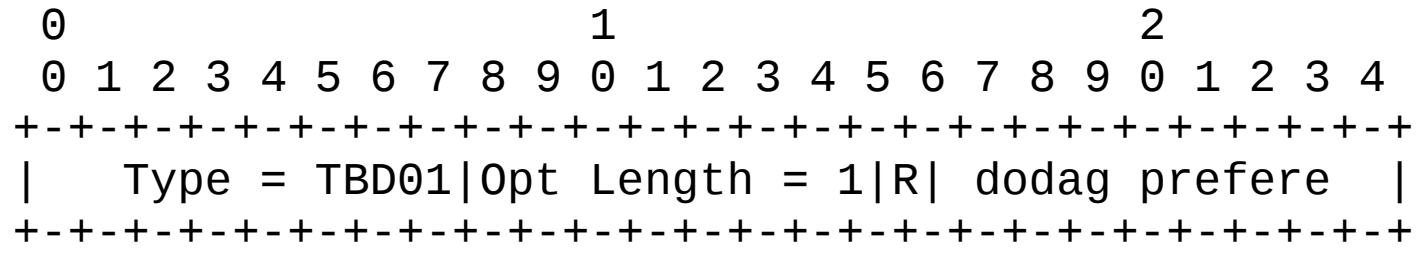


DODAG preference: a 7 bit field which provides a base value for the Enhanced Beacon Join priority. A value of 0x7f (127) disables the Join Proxy function entirely.

R a reserved bit that SHOULD be set to 0 by senders, and MUST be ignored by receivers. The reserved bit SHOULD be copied to options created.

The Minimum Priority influences the Proxy Priority that is announced in the Enhanced Beacon. The local node will apply additional criteria (such as number of neighbor cache entries it can allocate for untrusted nodes).

# What's the ROLL part? (2)

Enabling secure network join in RPL networks
draft-richardson-6tisch-roll-enrollment-priority-02

Defines a new DIO configuration Option:

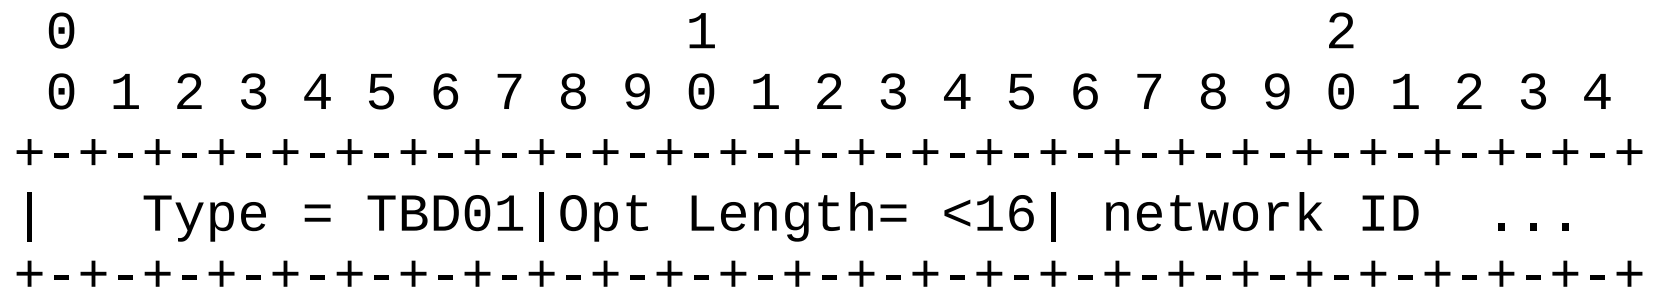


- Network ID: a 1 to 16 byte identifier which is set by the operator. Suggestion, is SHA256 hash of PIO, sent by DODAG root. Maybe created any other way.

# Goals in 6tisch

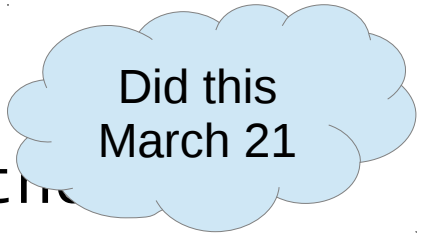


- Decide what set of things we want in the Enhanced Beacon.
  - Write this down somewhere, and ask ROLL to document how those numbers are derived, creating any new metrics or configuration containers needed.
- Document the security risk of exposure of these values.

ADOPT
richardson-6tisch-enrollment-enhanced-beacon

# Goals in ROLL

- Determine how the newly exposed metrics interact with or are derived from DIO things.
  - A value in an enhanced beacon vs a value in a subsequent DIO.
- There are two additional things related to Enrollment Priority and also the Parent Selection:
  - Number of children
    - Multiple drafts about balancing children
      - Children require (privileged) neighbour cache entries.
      - Enrollment requires unprivileged neighbor cache entries
  - Availability of bandwidth for Enrollment
    - Turn off enrollment when there are issues.

# Questions/Discussion

?

# Open Mic

Questions, Comments, Suggestions, …. :-)

Friday 23/3 9h30-11h30

# IETF 101 ROLL

## Routing over Low-Power And Lossy Networks


**Chairs:**
Peter van der Stok
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

- Friday 23/3 9h30-11h30: 120 minutes
- Remote Participation
  - Jabber Room: roll@jabber.ietf.org
  - Meetecho: http://www.meetecho.com/ietf101/roll
- Etherpad:
  - http://tools.ietf.org/wg/roll/minutes
- Audio Streaming:
- Minutes taker:
- Jabber Scribe:
- **Please sign blue sheets :-)**

# AGENDA - Friday

| Friday 23/3 9h30am-11h30am: 120 minutes | | |
|---|---|---|
| 9:30 - 9:35 (5 min.) | Introduction | Ines + Peter |
| 9:35 - 9:50 (15 min.) | Asymmetric AODV-P2P-RPL in LLNs<br>draft-ietf-roll-aodv-rpl-03 | Charlie |
| 9:50 - 10:30 (40 min.) | Routing for RPL Leaves<br>draft-thubert-roll-unaware-leaves-04<br>ND only device connects to RPL DODAG | Pascal |
| 10:30-10:45 (15 min.) | No-Path DAO modifications - Performance -<br>draft-ietf-roll-efficient-npdao-01 | Rahul |
| 10:45 -11:00 (15 min.) | Root initiated routing state in RPL<br>draft-ietf-roll-dao-projection-03 | Pascal |
| 11:00 - 11:05 (5 min.) | Status of draft-ietf-roll-dis-modifications-00 | Ines/Peter/Authors |
| 11:05 -11:30 (25 min.) | Discussion/Open Mic | Everyone |

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

draft-ietf-roll-aodv-rpl-03

IETF 101, London

Satish Anamalamudi <satishnaidu80@gmail.com>
Mingui Zhang <zhangmingui@huawei.com>
AR. Sangi <sangi_bahrian@yahoo.com>
Charlie Perkins <charles.perkins@earthlink.net>
S.V.R Anand <anand@ece.iisc.ernet.in>
Remy Liubing <remy.liubing@huawei.com>

# AODV-RPL: Overview

- Differences with P2P-RPL
  - Two DODAGs rooted separately in the Source and the Target
  - Support symmetric/asymmetric routes for upward and downward
  - Higher route diversity in asymmetric thanks to decoupling constraints on two directions
  - Encapsulate RREQ and RREP of AODV into RPL Options
  - Enable gratuitous RREP
- Note: Bi-directional asymmetric link
  - Can be used in both directions for DIOs but the two directions may have different values for, e. g. bandwidth, latency

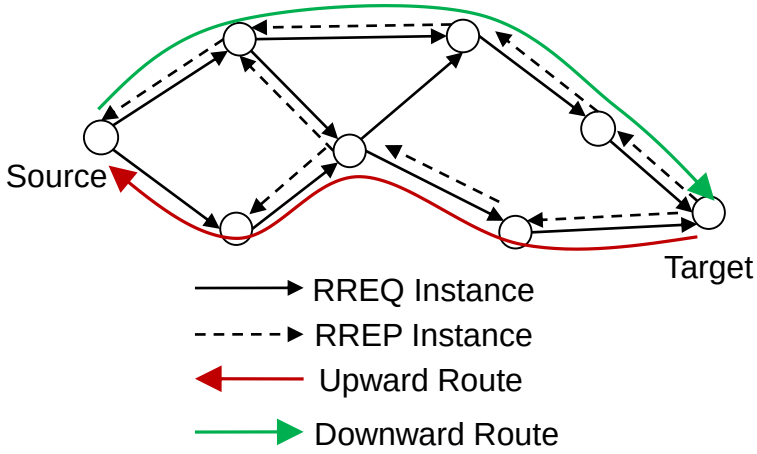


- Some comments during the last call
  - Inappropriate to have constraints on ETX without priori knowledge on link quality
  - Losing features compared to P2P-RPL: multiple targets in one DIO, DIO transmission scope limit
  - Inappropriate InstanceID setting in RREQ instance
  - Use DSTN for sequence number?
  - Why not use AODV's method for black-holing unidirectional links?

# Clarifying the scope of AODV-RPL

- AODV-RPL has not proposed methods for determining link asymmetry. But it is presumed that an intermediate node has that knowledge.
- The determination of whether or not a one-hop link is symmetric is implementation specific.
- The values of the metrics may be acquired by
  - Local information of intermediate routers
  - A priori knowledge
  - Estimation methods, such as averaging techniques

# New RREQ/RREP Option

RREQ Option

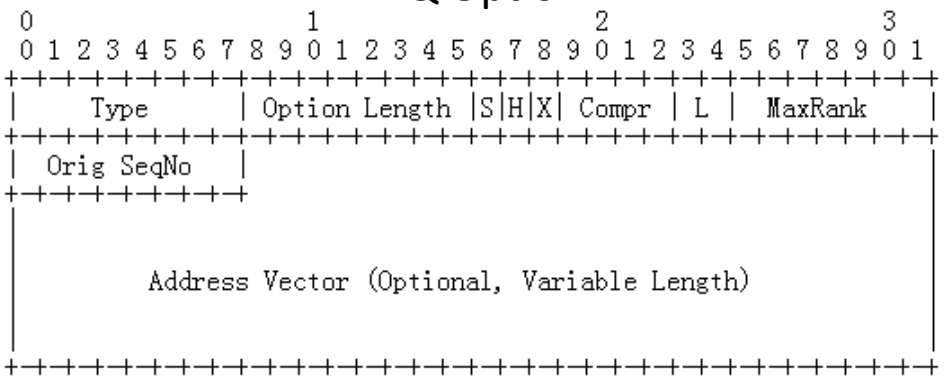


RREP Option

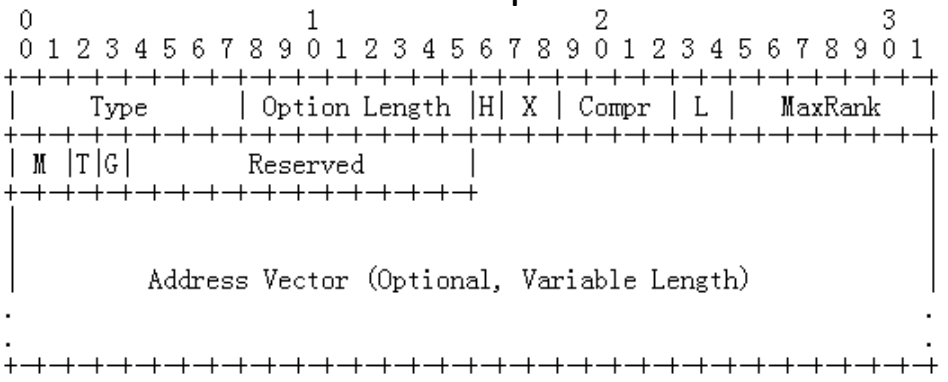


- "Option Length" added
- Relocate the 'S' bit (symmetric) from the base DIO to RREQ option
- 'H' bit: source routing or hop-by-hop
- Adoption of features in P2P-RPL (RFC 6997)
  - Compr: Address Vector compression
  - L: Router's residence duration in the DODAG
  - MaxRank: maximum acceptable cost for a route, limit the transmission scope of DIO
- Address Vector for source routing, not present when hop-by-hop

# New defined AODV-RPL Target Option

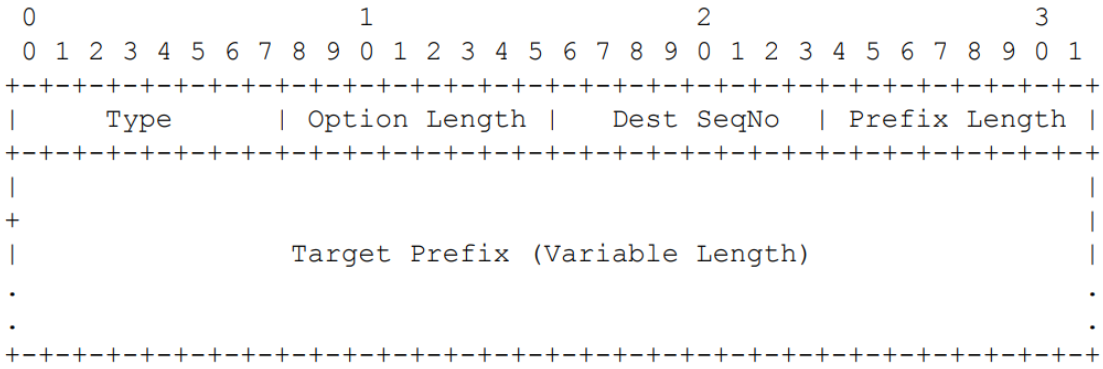


- Add the Destination Sequence Number to Target Option defined in RFC 6550
- Carry the TargNode/OrigNode addresses which are originally in RREQ/RREP options
- RREQ-DIO MUST carry at least one AODV-RPL Target Options
- Route discovery for multiple targets in the price of building one DODAG (in case of same RD requirements)

# RPLInstance ID Settings

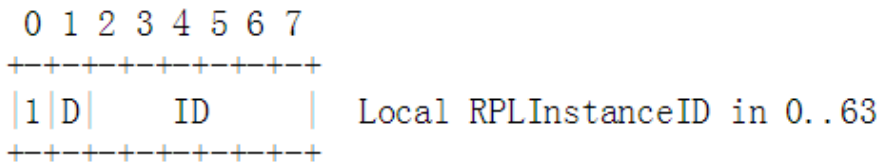


- Local Instance ID assigned by the OrigNode
- Paired Instance IDs
  - Same value in the 6 right most bits for RREQ and RREP instances
  - D=0 for RREQ-instance, D=1 for RREP-instance
  - Pairing the RREQ-instance and RREP-instance, especially when there are more than one RD processes between a pair of OrigNode & TargNode.
- When the Instance ID to be used is occupied
  - Shift it to another number (still between 0 and 63)
  - Recover it to the original one according to the SHIFT field in RREP option

# Next Steps

- Eliminate normative references to RFC 6997
- Make sure that the the Target Option can correctly handle multiple targets along with their sequence numbers.
- Analyze and resolve the corner case with multiple OrigNodes using the same RPLInstanceID
- Resolve whether or not to import black-hole link detection from AODV
- Other items that are identified in the WG meeting
- Projection for submission of next revision (proposal: mid-April)

# Routing for RPL Leaves

**draft-thubert-roll-unaware-leaves-03**

Pascal Thubert

IETF 101

London

# Unmet expectations

- Connectivity for a Non-RPL aware node in a RPL domain
  - Forwarding is described but not the control plane
- Integration of the EDA Exchange (EDAR/EDAC) used as keep-alive with the RPL signaling to avoid duplication
  - At the moment both are needed periodicallThis spec uses a common lifetime and the EDA exchange is proxied
- Separation of the RPL Root and the 6LBR and proxy registration to the 6BBR
  - The RPL root proxies the EDA with the 6LBR and the NS(EARO) with the 6BBR

# RFC 6775 Update

P.Thubert, E. Nordmark, S. Chakrabarti, C. Perkins

# What are the 6LoWPAN ND extensions?

- draft-ietf-6lo-rfc6775-update
  - Simplifies the protocol (no DAR/DAC for LL, no secondary NC)
  - Enables proxy registration
- draft-ietf-6lo-ap-nd
  - Protects addresses against theft (Crypto ID in registration)
- draft-ietf-6lo-backbone-router
  - Federates 6lo meshes over a high speed backbone
  - ND proxy that mimics 802.11 association but at Layer 3

# RFC 6775 update new features: the Length

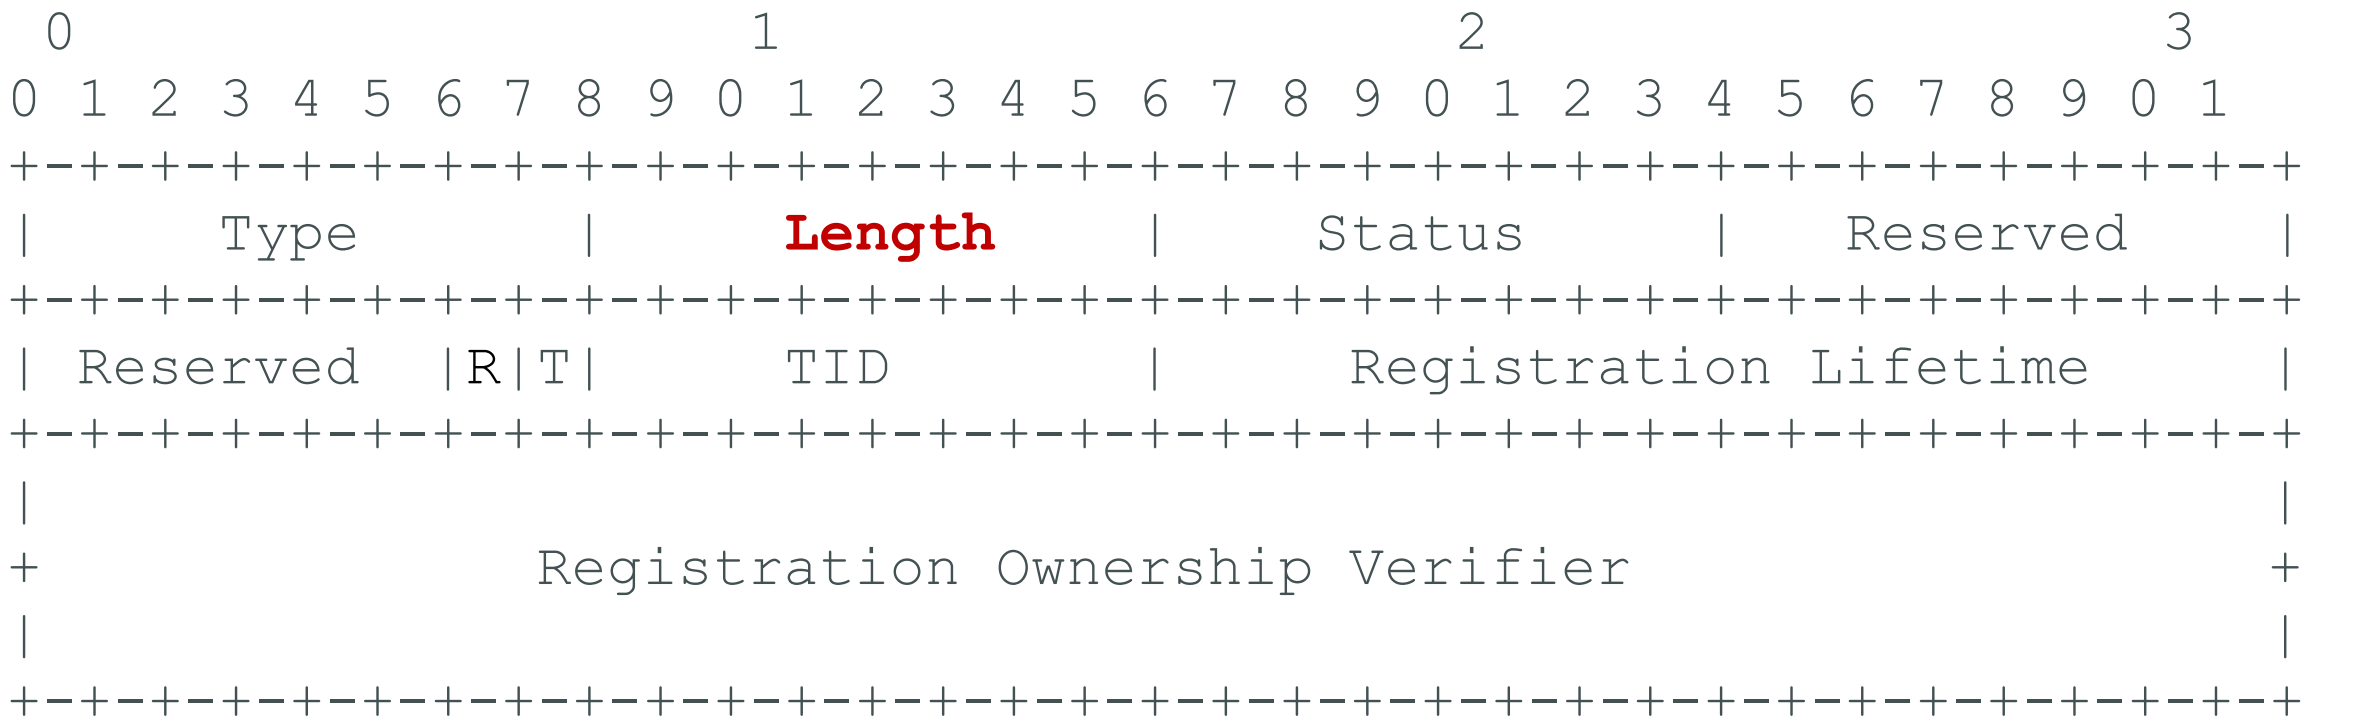


Length: 8-bit unsigned integer. The length of the option in units of 8 bytes. **It MUST be 2 when operating in backward-compatible mode**. It MAY be 3, 4 or 5, denoting a **ROVR size of 128, 192 and 256 bits respectively**.

# RFC 6775 update new features: the ‘R’ flag

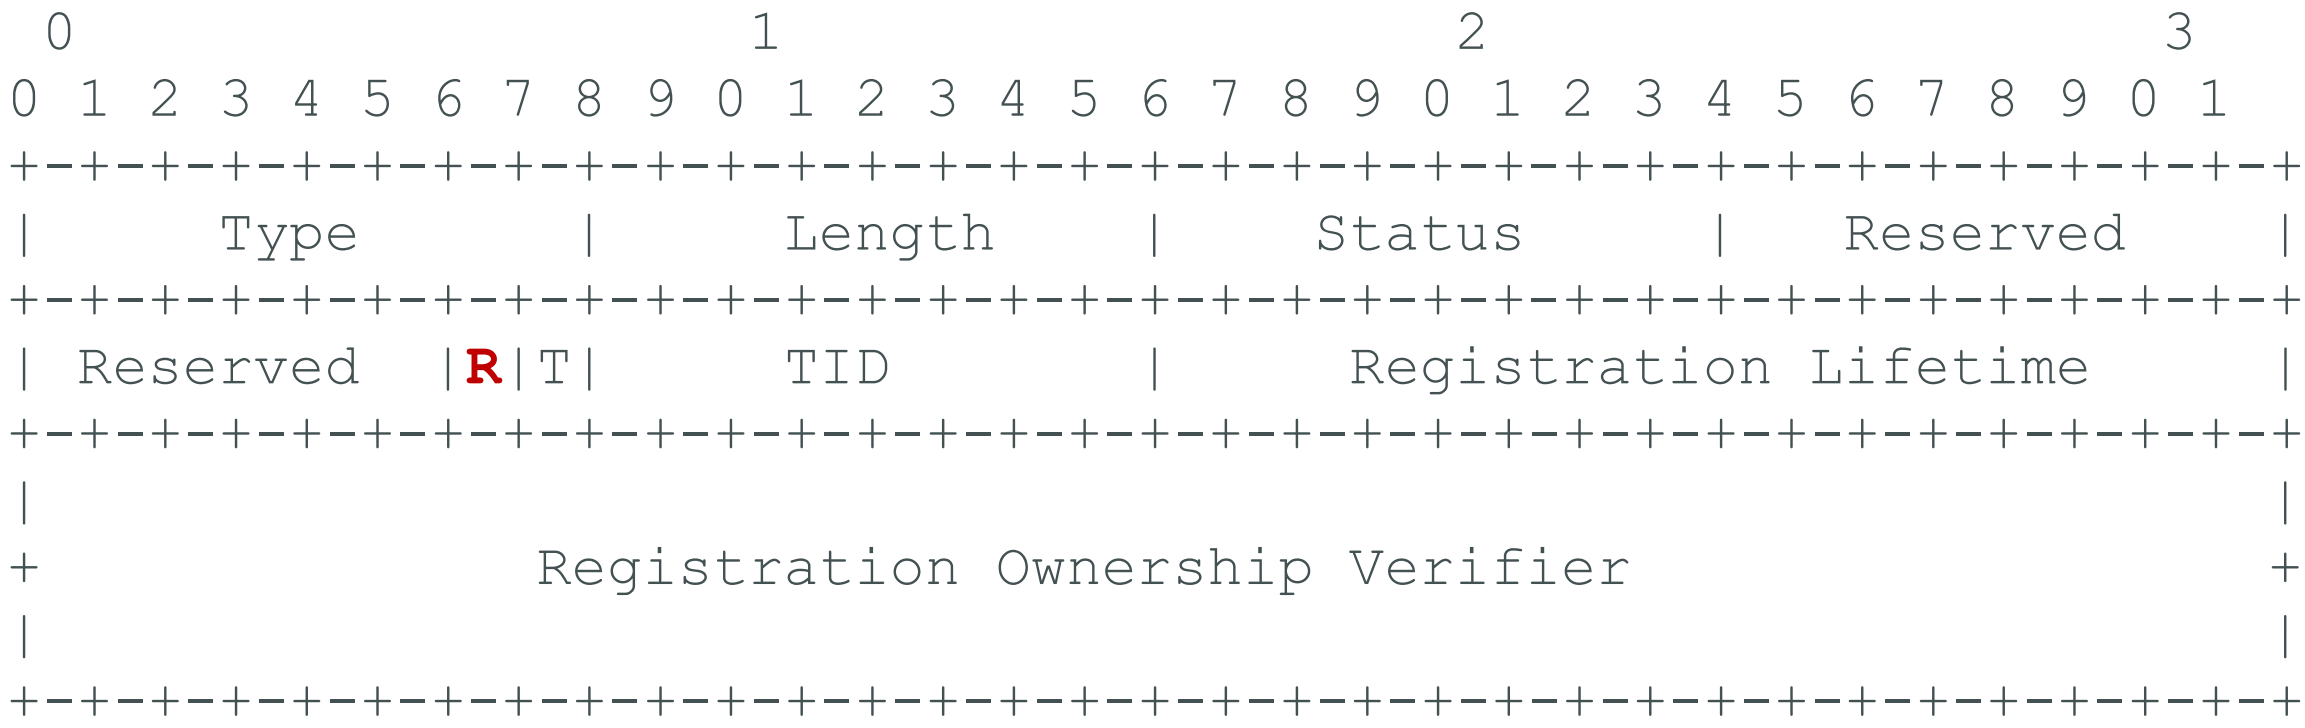


R: One-bit flag. If the 'R' flag is **set**, the registering node expects that **the 6LR ensures reachability for the registered address**, e.g., by injecting the address in a Route-Over routing protocol or proxying ND over a Backbone Link.

# RFC 6775 update new features: the Transaction ID

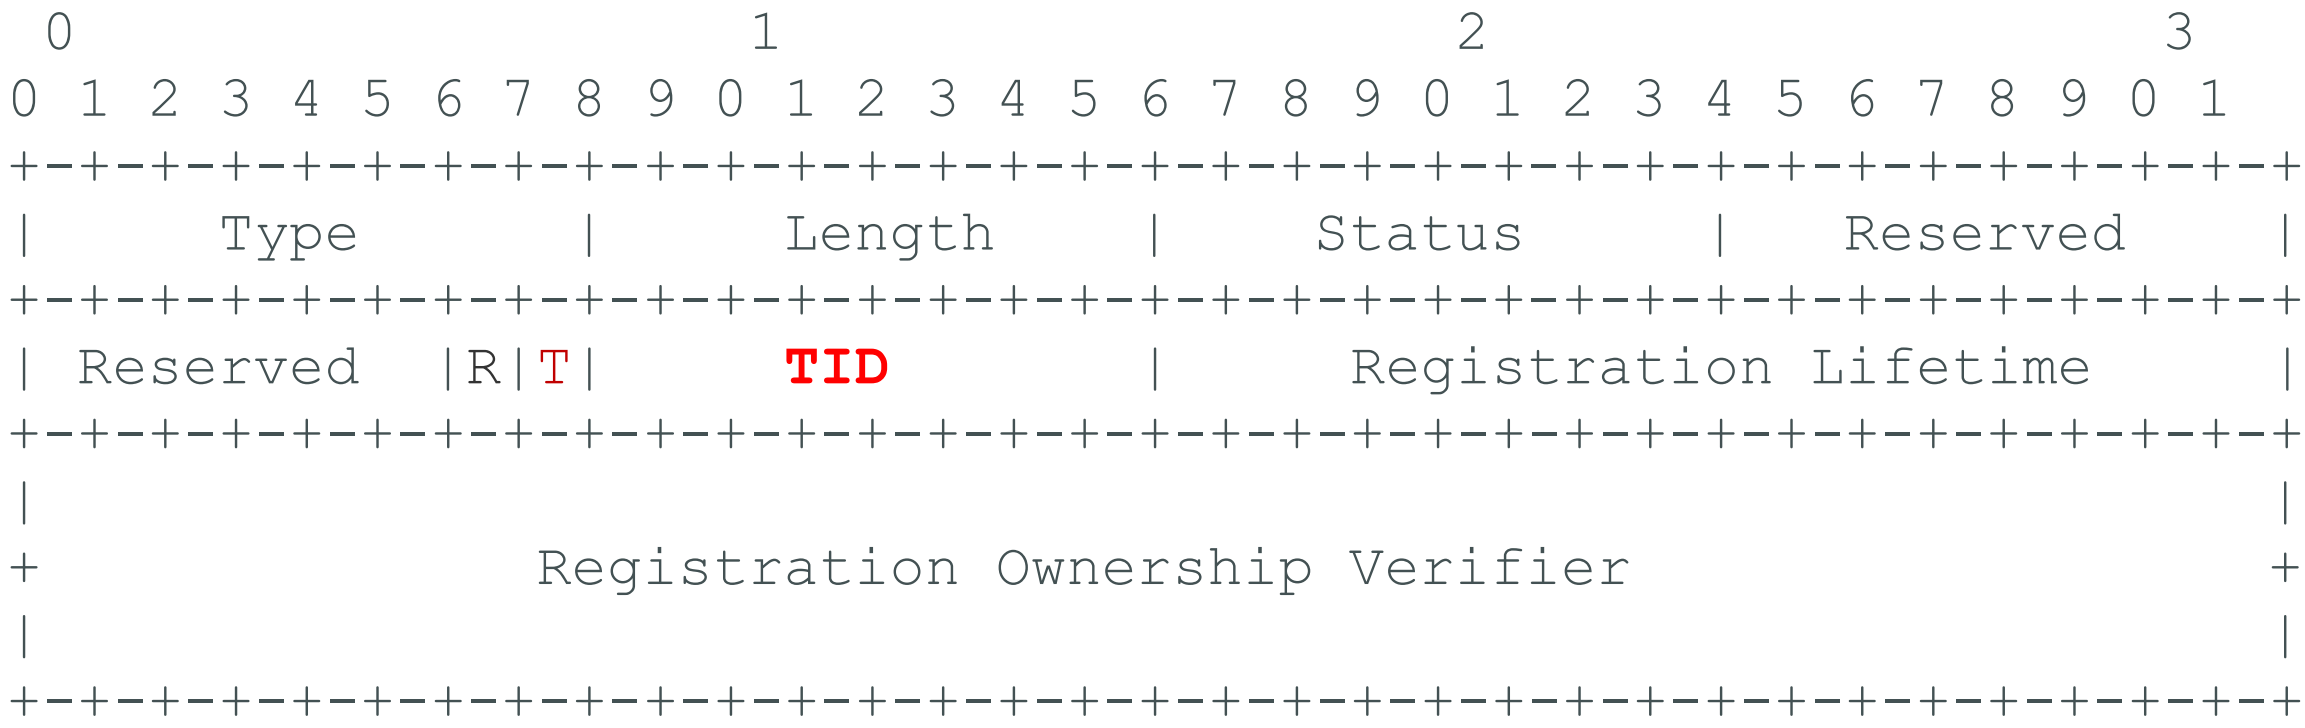


4.2.1. Comparing TID values: **The TID is a sequence counter and its operation is the exact match of the path sequence specified in RPL**, the IPv6 Routing Protocol for Low-Power and Lossy Networks [RFC6550] specification.

# RFC 6775 update new features: Registration Lifetime

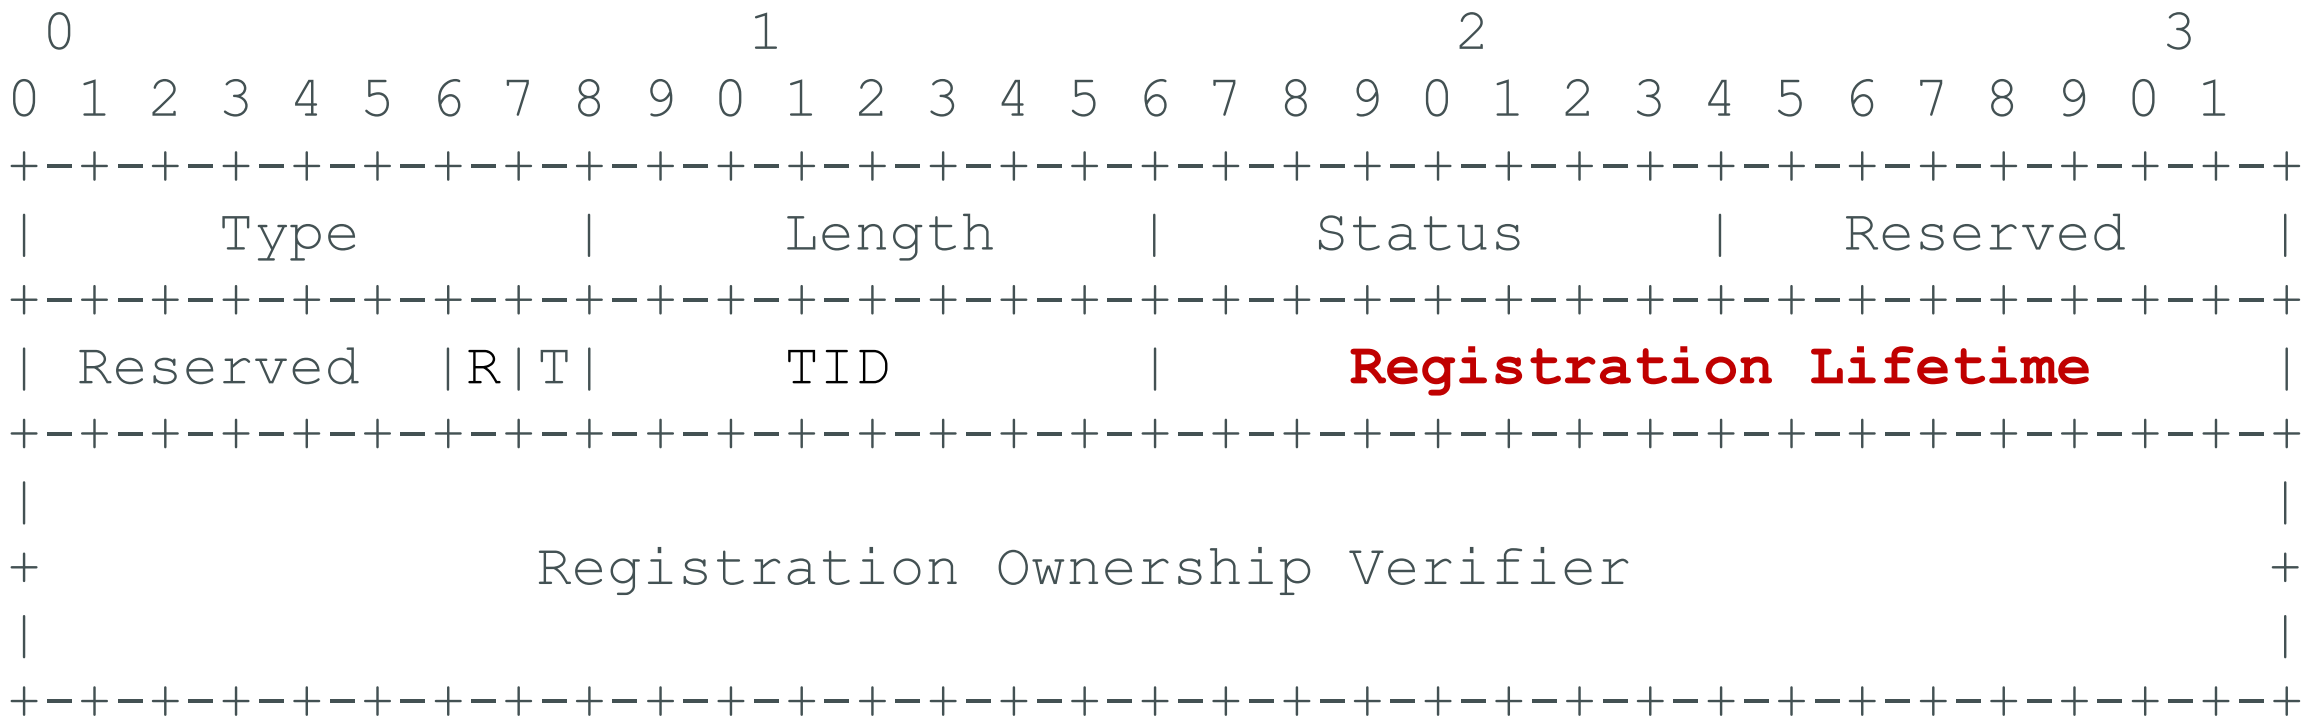


Registration Lifetime: 16-bit integer; expressed in minutes.

0 Registration Lifetime: 16-bit integer; expressed in minutes. 0 means that the registration has ended and the associated state MUST be removed.

# RFC 6775 update new features: ROVR

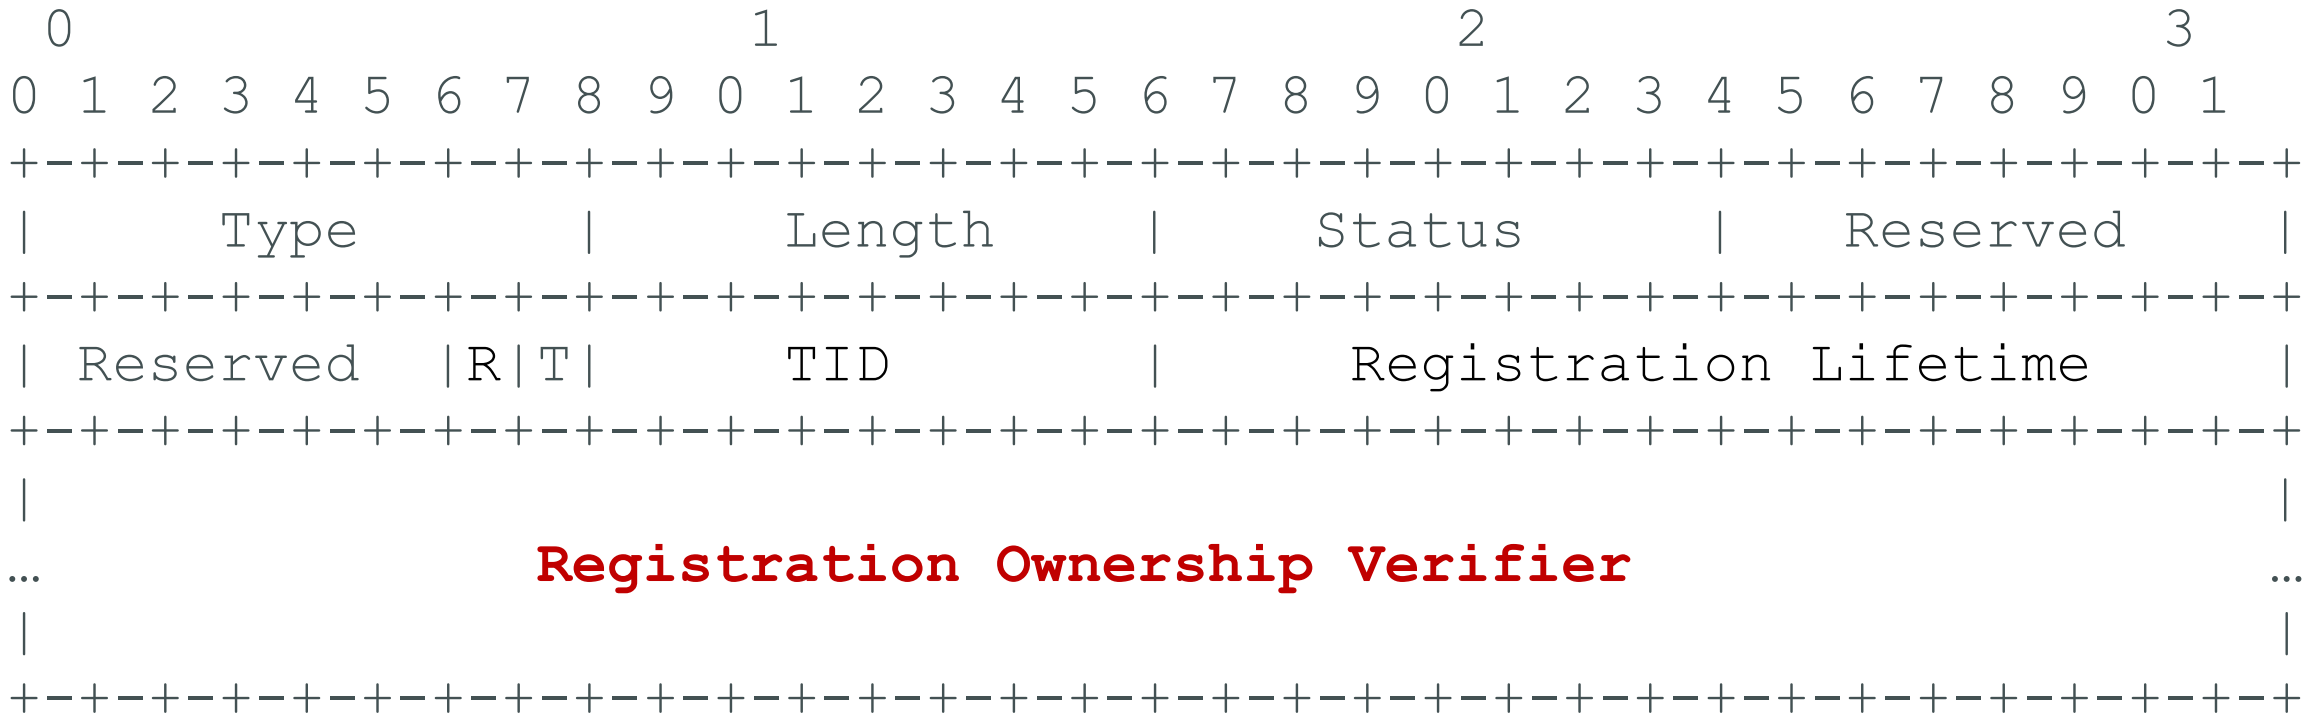


Registration Ownership Verifier (ROVR): Enables the correlation between multiple attempts to register a same IPv6 Address. This can be a unique ID of the Registering Node, such as the EUI-64 address of an interface. This can also be a token obtained with methods and used as proof of ownership of the registration.

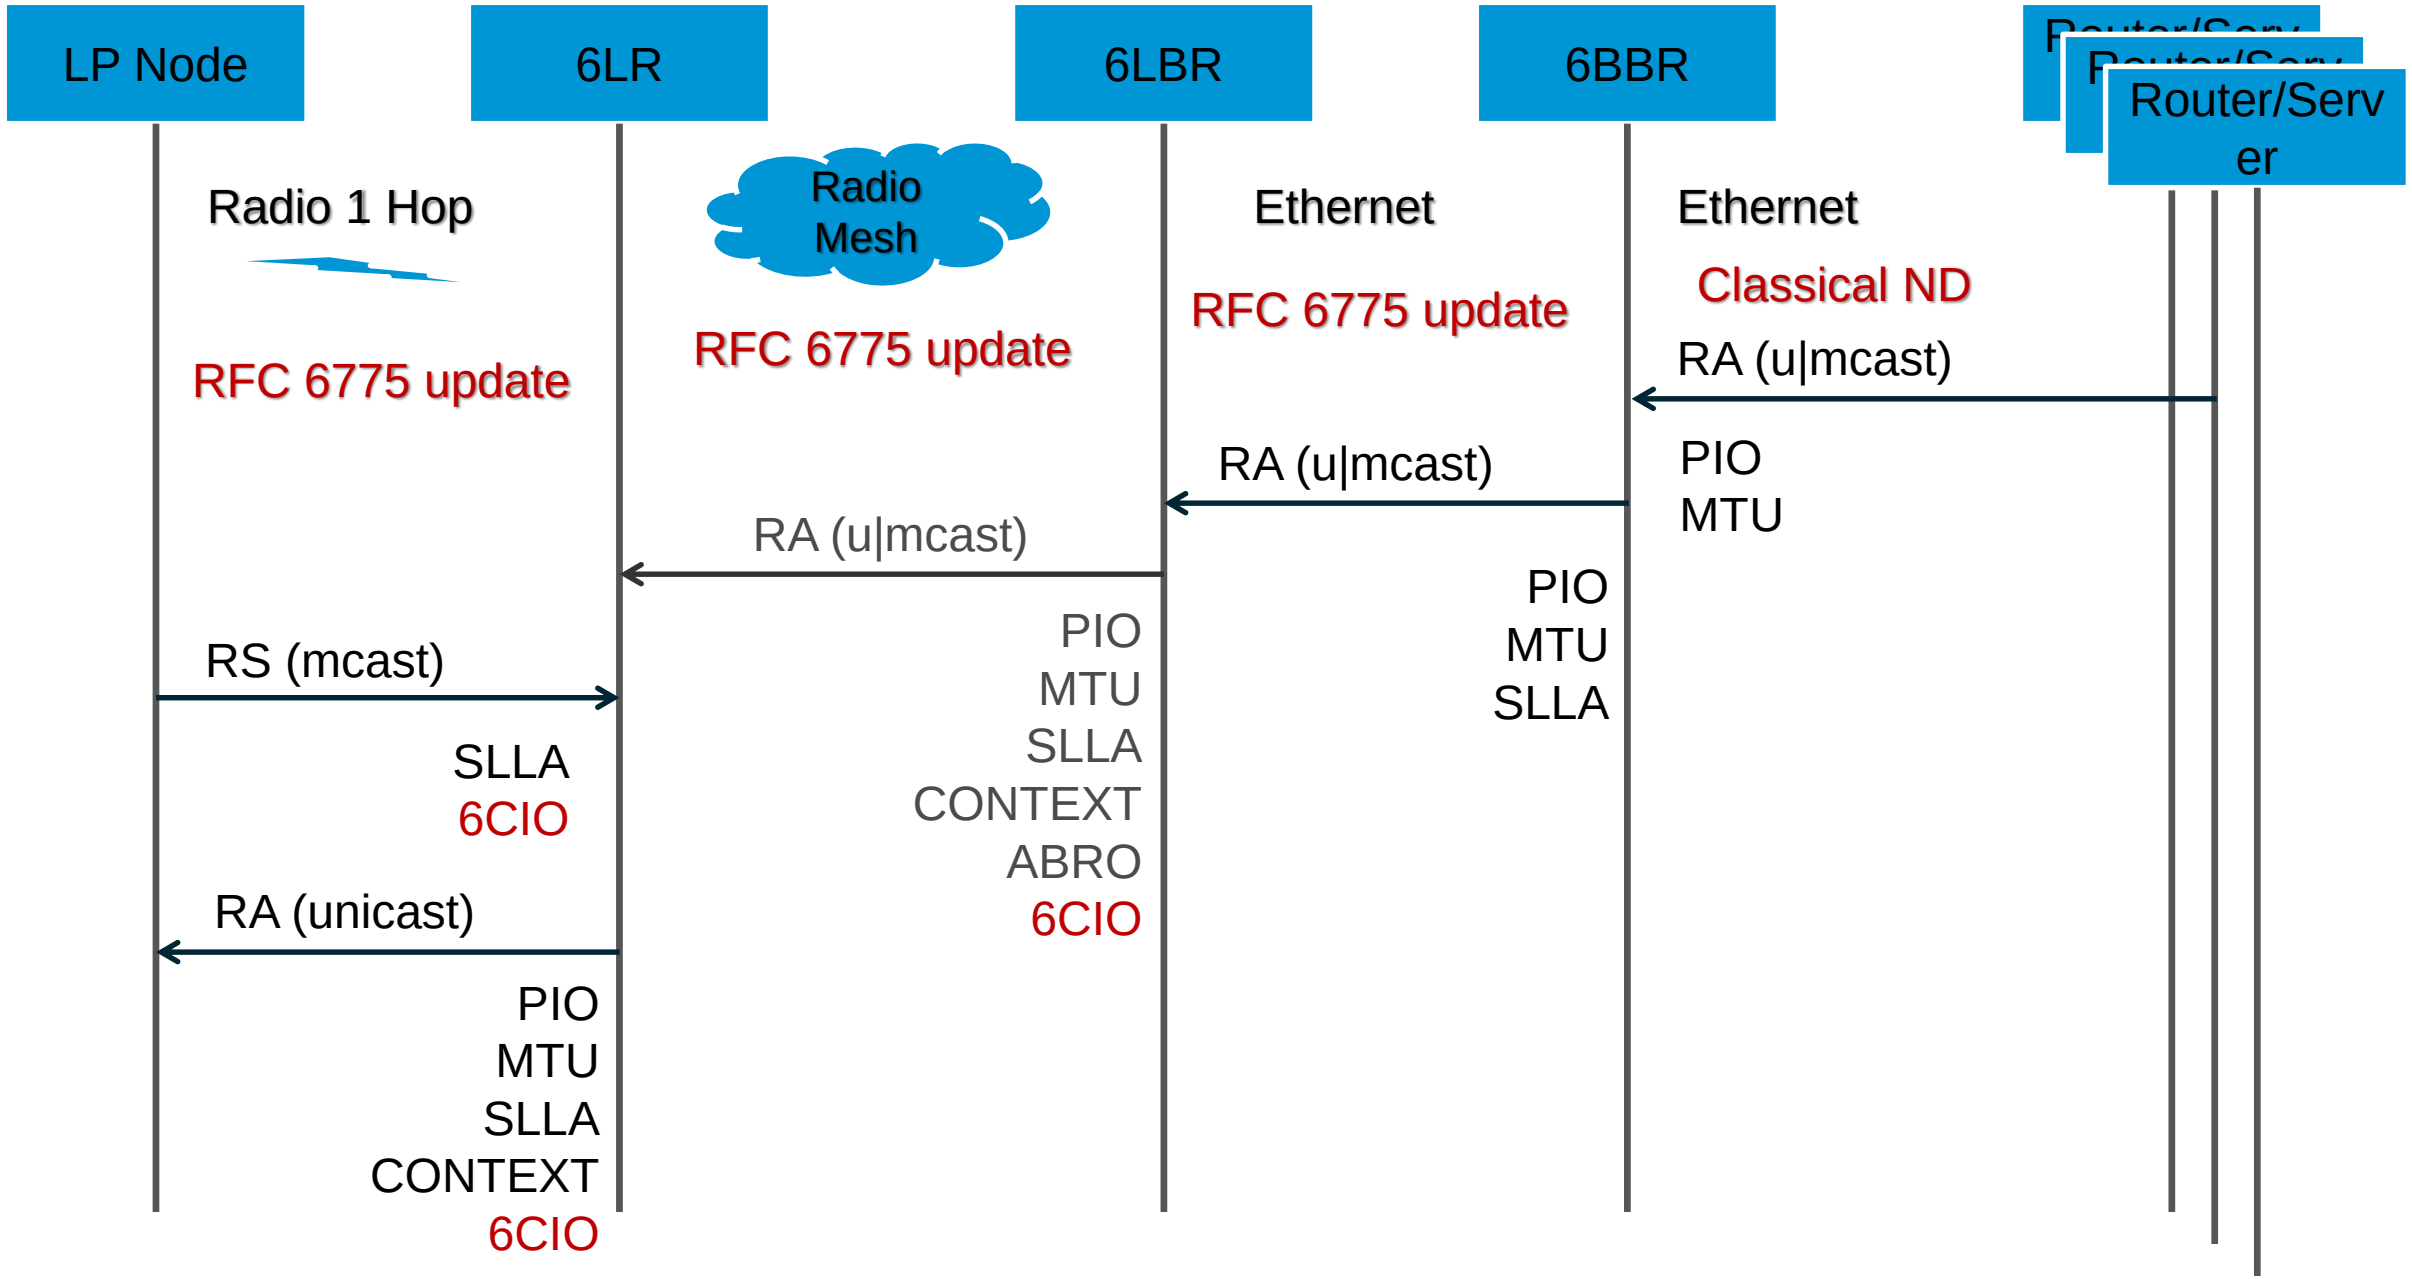
LP Node
6LR
6LBR
6BBR
Router/Server
Radio 1 Hop
RFC 6775 update
Radio Mesh
RFC 6775 update
Ethernet
RFC 6775 update
Ethernet
Classical ND
RA (u|mcast)
PIO
MTU
RA (u|mcast)
RA (u|mcast)
PIO
MTU
SLLA
PIO
MTU
SLLA
CONTEXT
ABRO
6CIO
RS (mcast)
SLLA
6CIO
RA (unicast)
PIO
MTU
SLLA
CONTEXT
6CIO

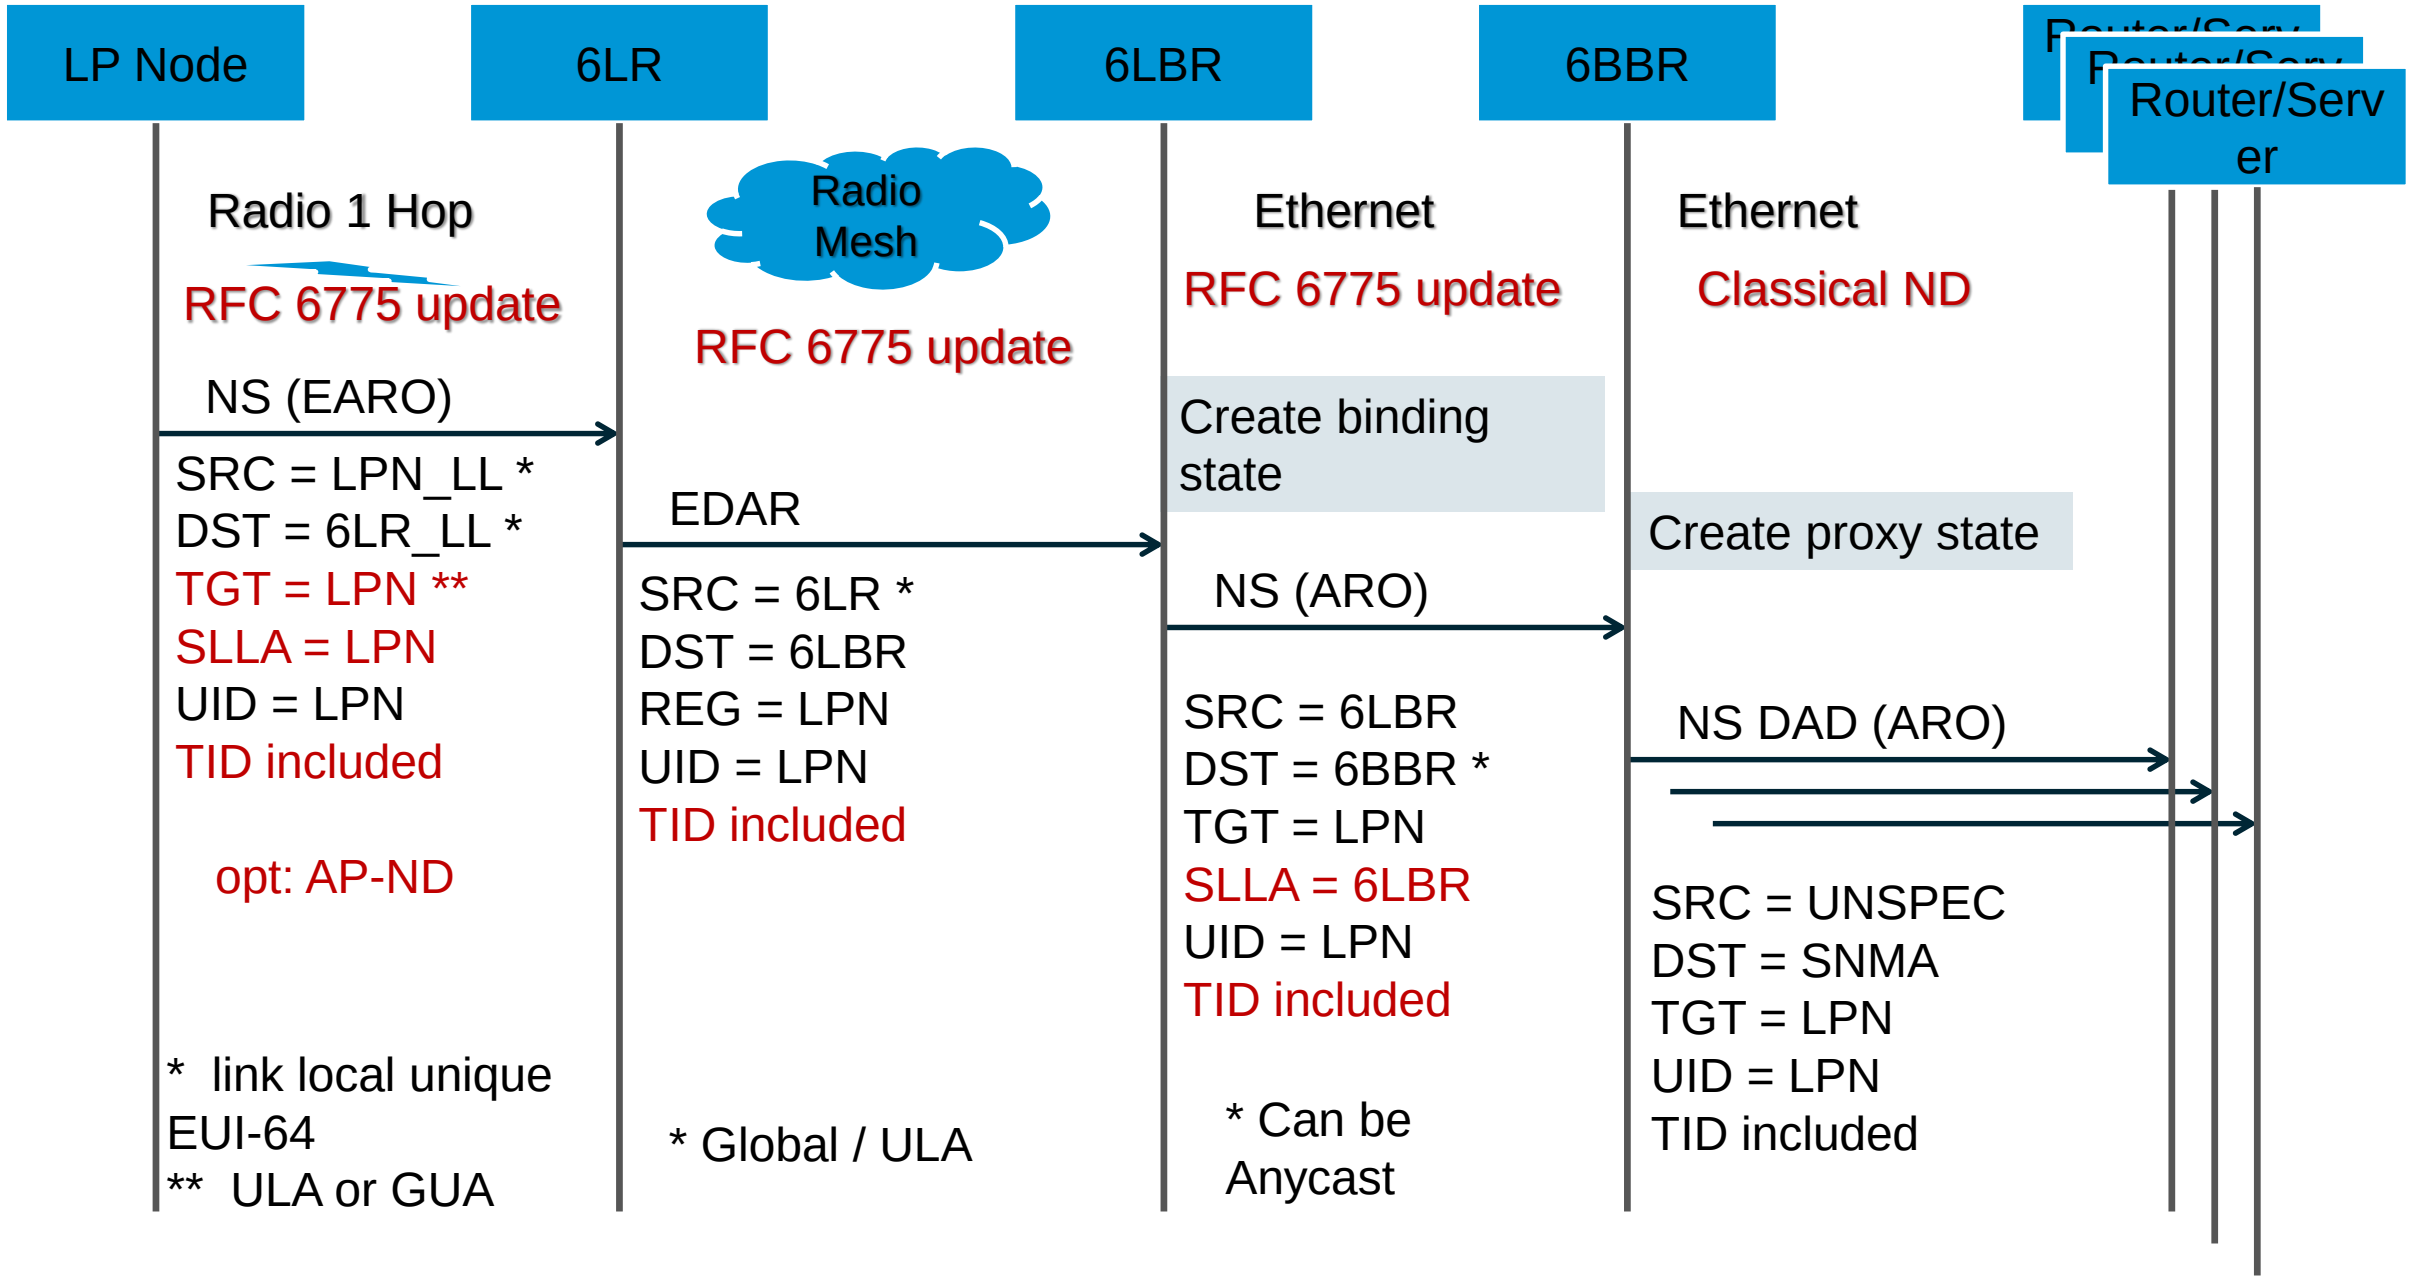
LP Node
6LR
6LBR
6BBR
Router/Server
Radio 1 Hop
RFC 6775 update
Radio Mesh
RFC 6775 update
Ethernet
RFC 6775 update
Ethernet
Classical ND
NS (EARO)
SRC = LPN_LL *
DST = 6LR_LL *
TGT = LPN **
SLLA = LPN
UID = LPN
TID included
opt: AP-ND
EDAR
SRC = 6LR *
DST = 6LBR
REG = LPN
UID = LPN
TID included
Create binding state
NS (ARO)
SRC = 6LBR
DST = 6BBR *
TGT = LPN
SLLA = 6LBR
UID = LPN
TID included
Create proxy state
NS DAD (ARO)
SRC = UNSPEC
DST = SNMA
TGT = LPN
UID = LPN
TID included
* link local unique EUI-64
** ULA or GUA
* Global / ULA
* Can be Anycast

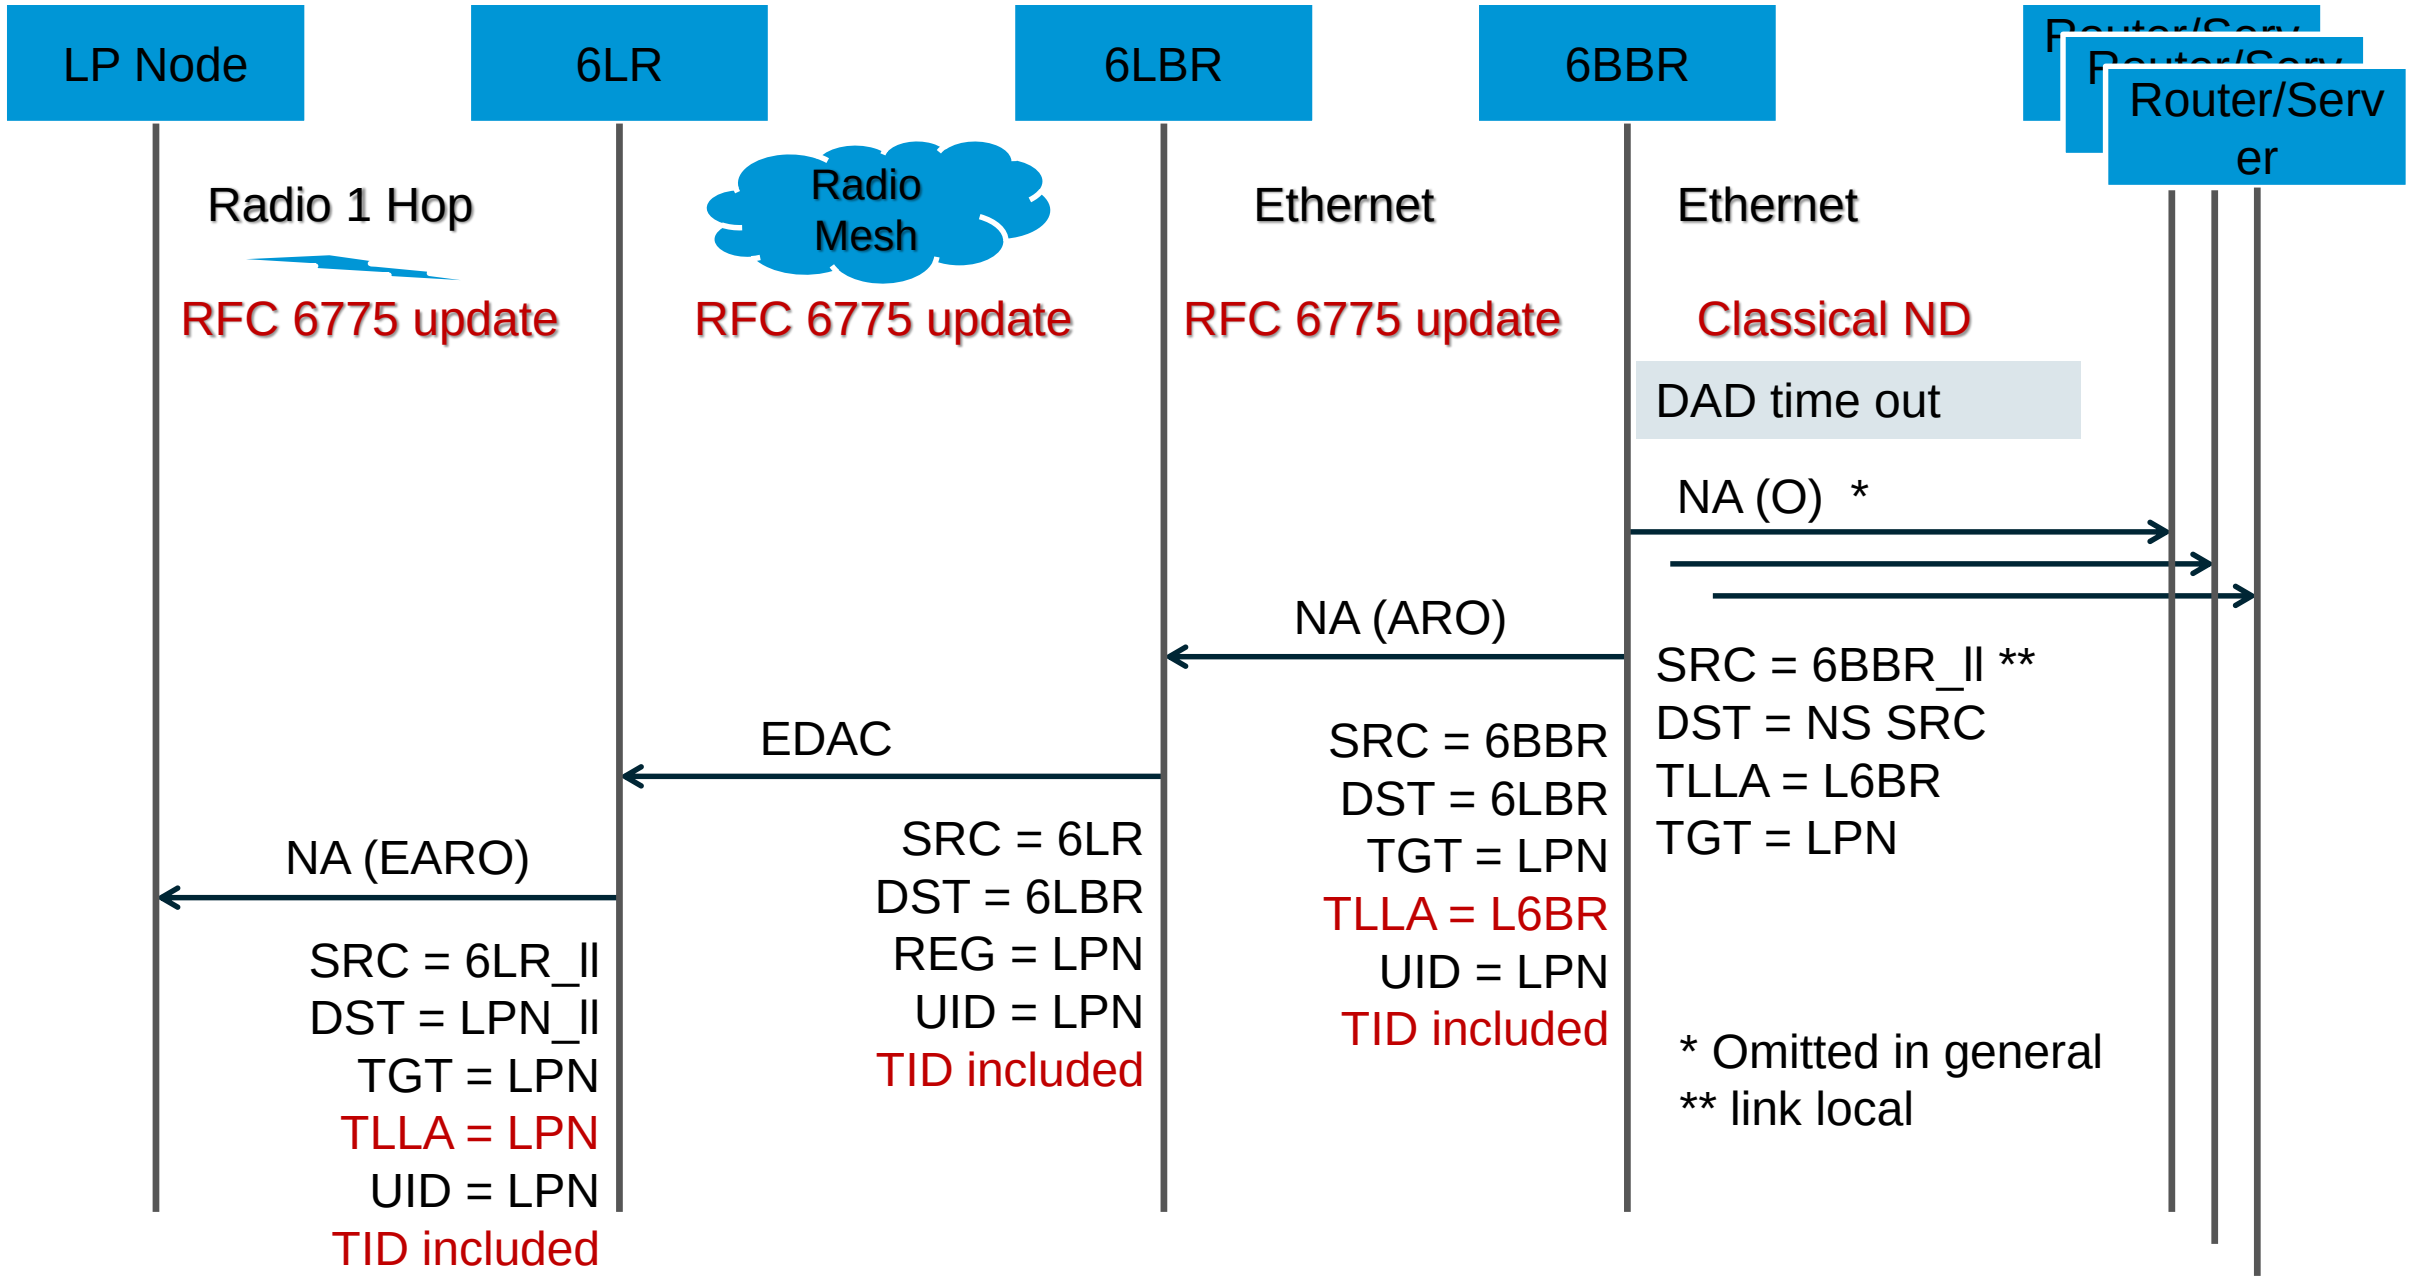
LP Node
6LR
6LBR
6BBR
Router/Server
Radio 1 Hop
Radio Mesh
Ethernet
Ethernet
RFC 6775 update
RFC 6775 update
RFC 6775 update
Classical ND
DAD time out
NA (O) *
NA (ARO)
SRC = 6BBR_ll **
DST = NS SRC
TLLA = L6BR
TGT = LPN
EDAC
SRC = 6BBR
DST = 6LBR
TGT = LPN
TLLA = L6BR
UID = LPN
TID included
NA (EARO)
SRC = 6LR
DST = 6LBR
REG = LPN
UID = LPN
TID included
SRC = 6LR_ll
DST = LPN_ll
TGT = LPN
TLLA = LPN
UID = LPN
TID included
* Omitted in general
** link local

# IESG Review

RFC 6775 Update

Draft-…-12 to -16

# INT-DIR (Tim Chown) => v-12

- Use of EUI-64, should it be deprecated (for privacy reasons) ?
- Clarifications on privacy addresses
- Added a matrix matching specs and requirements in appendix
- Added a glossary
- Suggestion to ask 6MAN about the need for a LRU algorithm

# OPS-DIR (Jürgen Schönwälder) + SEC-DIR (Chris Lonvick) => v-13

- Moved terminology up for readability
- Changed “legacy” to “RFC6775-only” refering to RFC 6775
- Changed OUI field to RUID
- Added Appendix B.7.

  “Requirements Related to Operations and Management”

# IOT-DIR (Dave Thaler) => v-14

- Reworded Intro (and many other things)
- Introduced the ‘R’ flag based on parallel discussion with ROLL
- Reworded RUID description
- Limiting the number of addresses => What is the minimum?
- Clarification on address duplication over backbone

# RTG-DIR (Adrian Farrel) => v-15

- RUID => ROVR  Registration Ownership Verifier ; new text on ROVR functionality and collision scope and consequences
- 6CIO now the only way to discover 6LR capabilities. New flag for 6LBR capability to support extended DA messages
- Use of ICMP code: non-NULL code => Extended DA message
- EARO Length extended due to side discussion on AP-ND

# GEN-ART (Peter Yee) => v-16

- Clarifications, e.g., RECOMMENDED for implementations
- How properties are discovered (completing Adrian's review)
- Review of the requirements and security section
- Clarified / fixed IEEE references
- A lot of editorials, syntax corrections

# draft-ietf-6lo-ap-nd

P.Thubert, B. Sarikaya, M Sethi, (and expecting R. Struik but not there yet)

# Unmet expectations

- First come first Serve address registration
  - First registration for an address owns that address till it releases it
  - The network prevents hijacking
- Source address validation
  - Address must be topologically correct
  - Source of the packet owns the source address
- First Hop Security only?
  - Proxy ownership and routing advertisements not protected yet

# AP-ND new features: ‘C’ flag

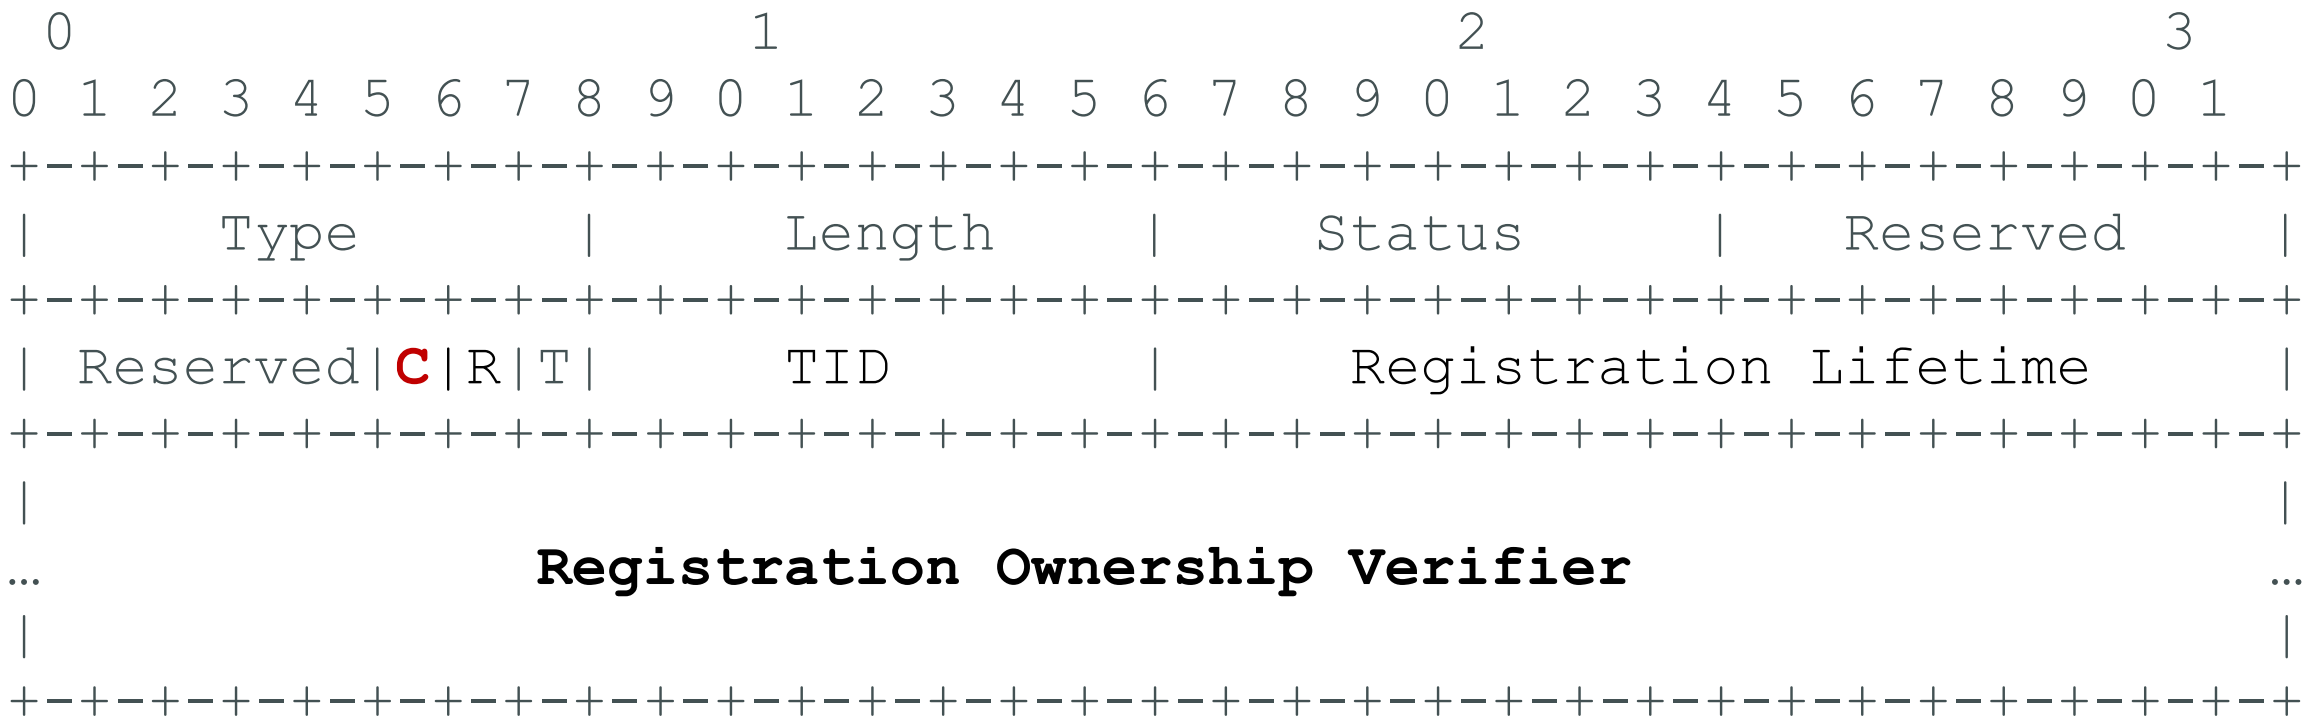


C: The "C" flag is set to indicate that the Registration Ownership Verifier field contains a Crypto-ID and that the 6LN MAY be challenged for ownership as specified in this document.

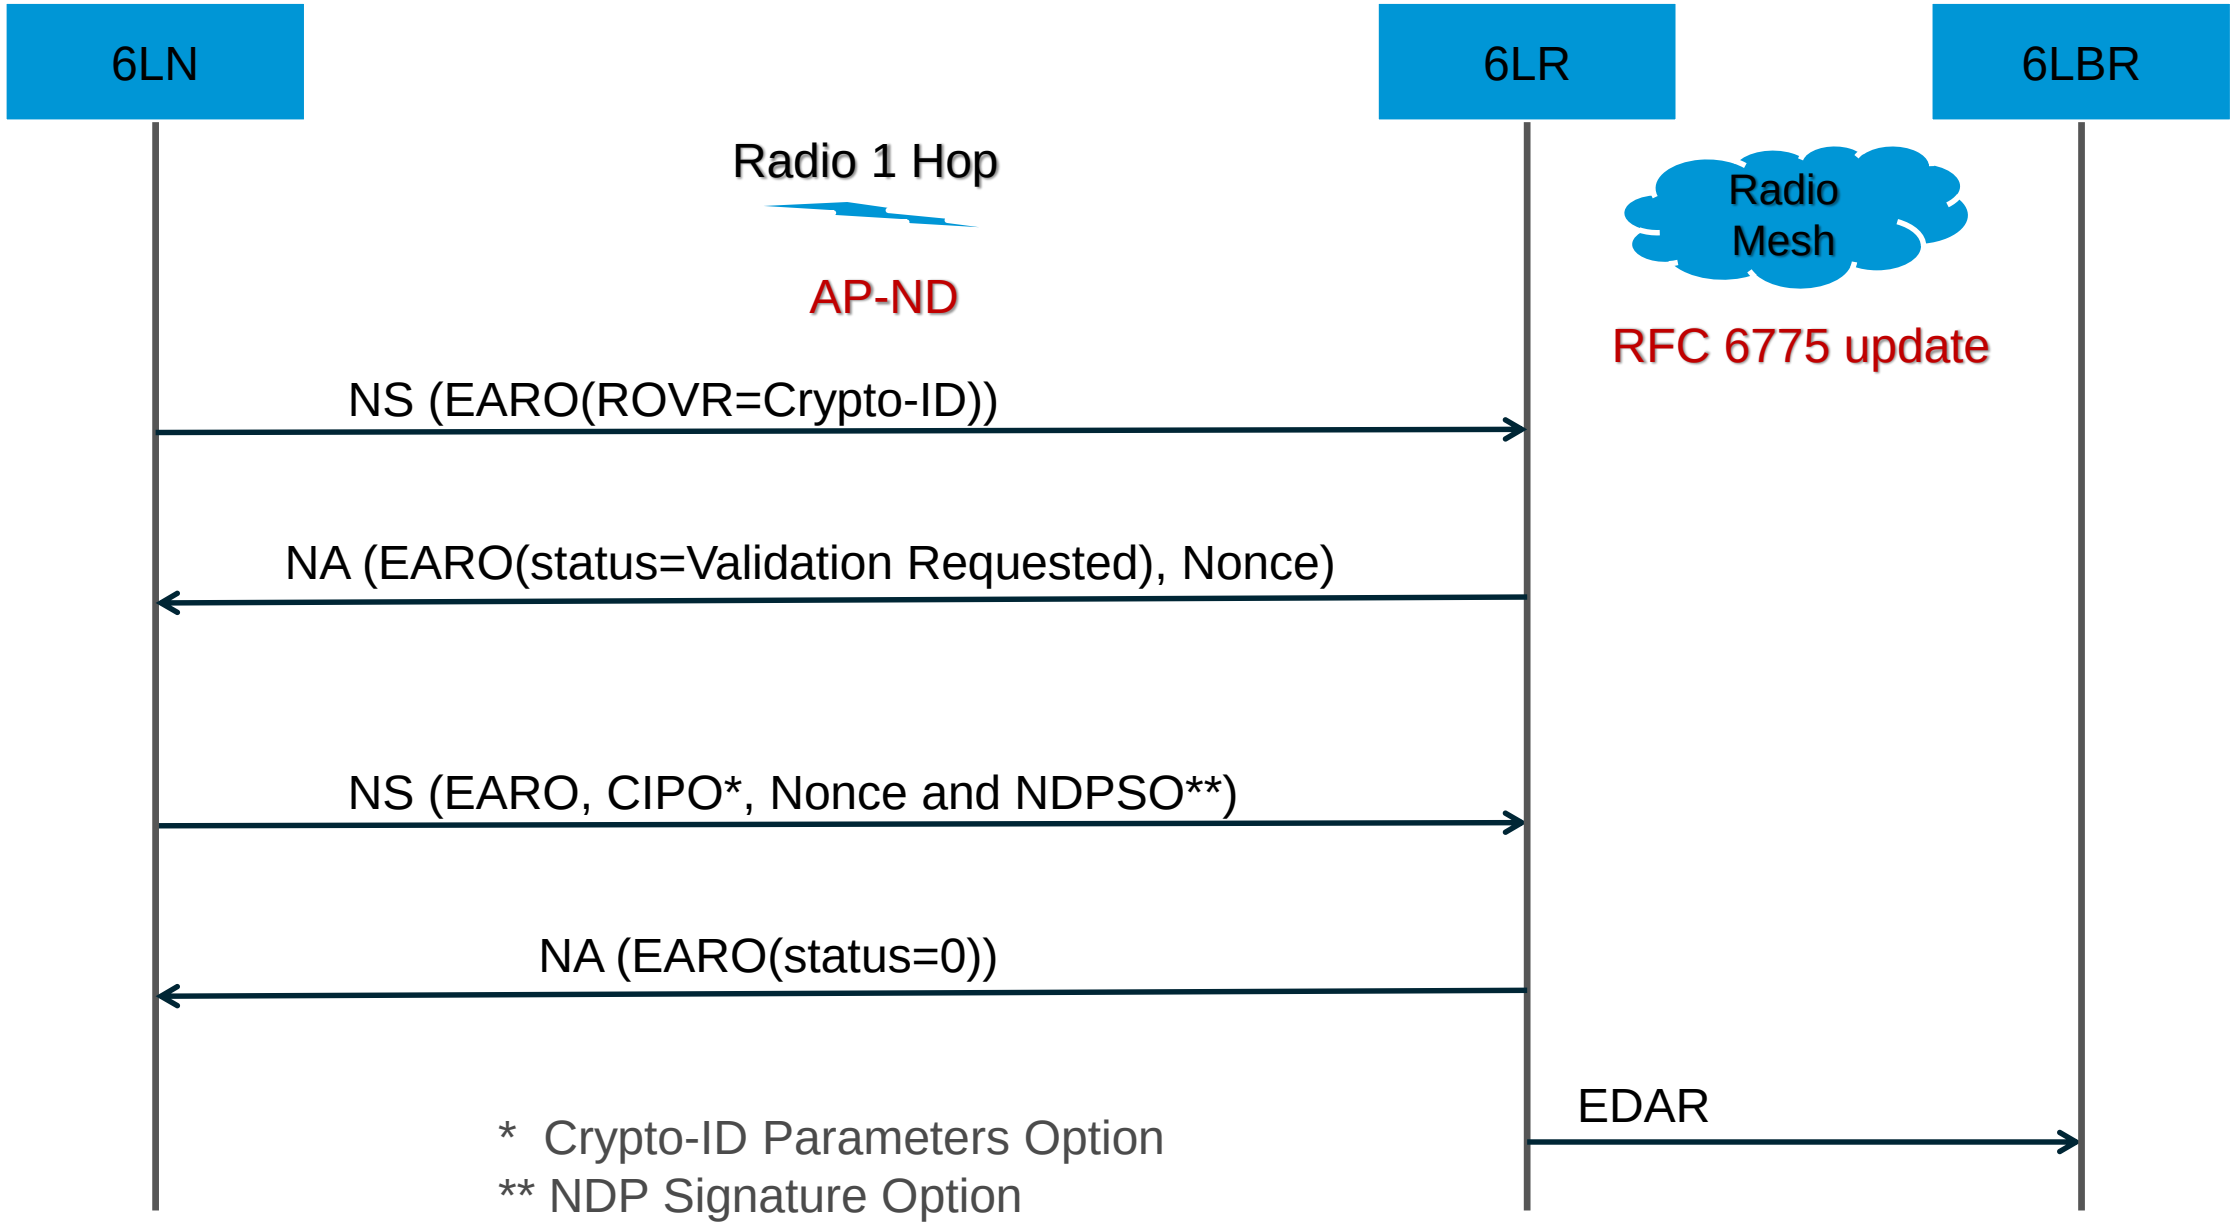
6LN
6LR
6LBR
Radio 1 Hop
AP-ND
Radio Mesh
RFC 6775 update
NS (EARO(ROVR=Crypto-ID))
NA (EARO(status=Validation Requested), Nonce)
NS (EARO, CIPO*, Nonce and NDPSO**)
NA (EARO(status=0))
EDAR
* Crypto-ID Parameters Option
** NDP Signature Option

# Recent changes

- Simplified the computation of the Crypto-ID
  - Digital signature (SHA-256 then either NIST P-256 or EdDSA) is executed on the concatenation of short modifier and public key
  - Modifier not used to make computation complex as opposed to CGA. This simplifies the operation of a constrained node
  - But 64 bits ROVR might not suffice for adequate protection => Longer ROVR
- Reuse options defined in RFC 3971 for SEND
  - Crypto-ID Parameters Option, a variation of the CGA Option
  - Nonce Option
  - NDP Signature Option, a variation of the RSA Signature Option
    - the option is extended for non-RSA Signatures
    - this specification defines an alias to avoid the confusion.

# Security properties

- We made the size of the ROVR tunable so we can get high security
- At the moment a joining 6LN is challenge from the 6LR
  - The 6LBR MUST trust the 6LR
  - A rogue 6LR may pretend that it represents a 6LN that passed the challenge
  - Should we challenge all the way from the 6LBR?
  - Can the Crypto-ID be used in routing protocols, how?

# draft-ietf-6lo-backbone-router

P.Thubert

# Unmet expectations

- Scale an IOT subnet to the tens of thousands
  - With device mobility (no renumbering)
  - Controlled Latency and higher Reliability using a backbone
- Deterministic Address presence
  - Route towards the latest location of an address
  - Remove stale addresses

# Recent changes

- Uses of the ‘R’ flag
  - Indicates the need for proxy operation
- Clarifications
- TBD : RPL Root / 6LBR separation

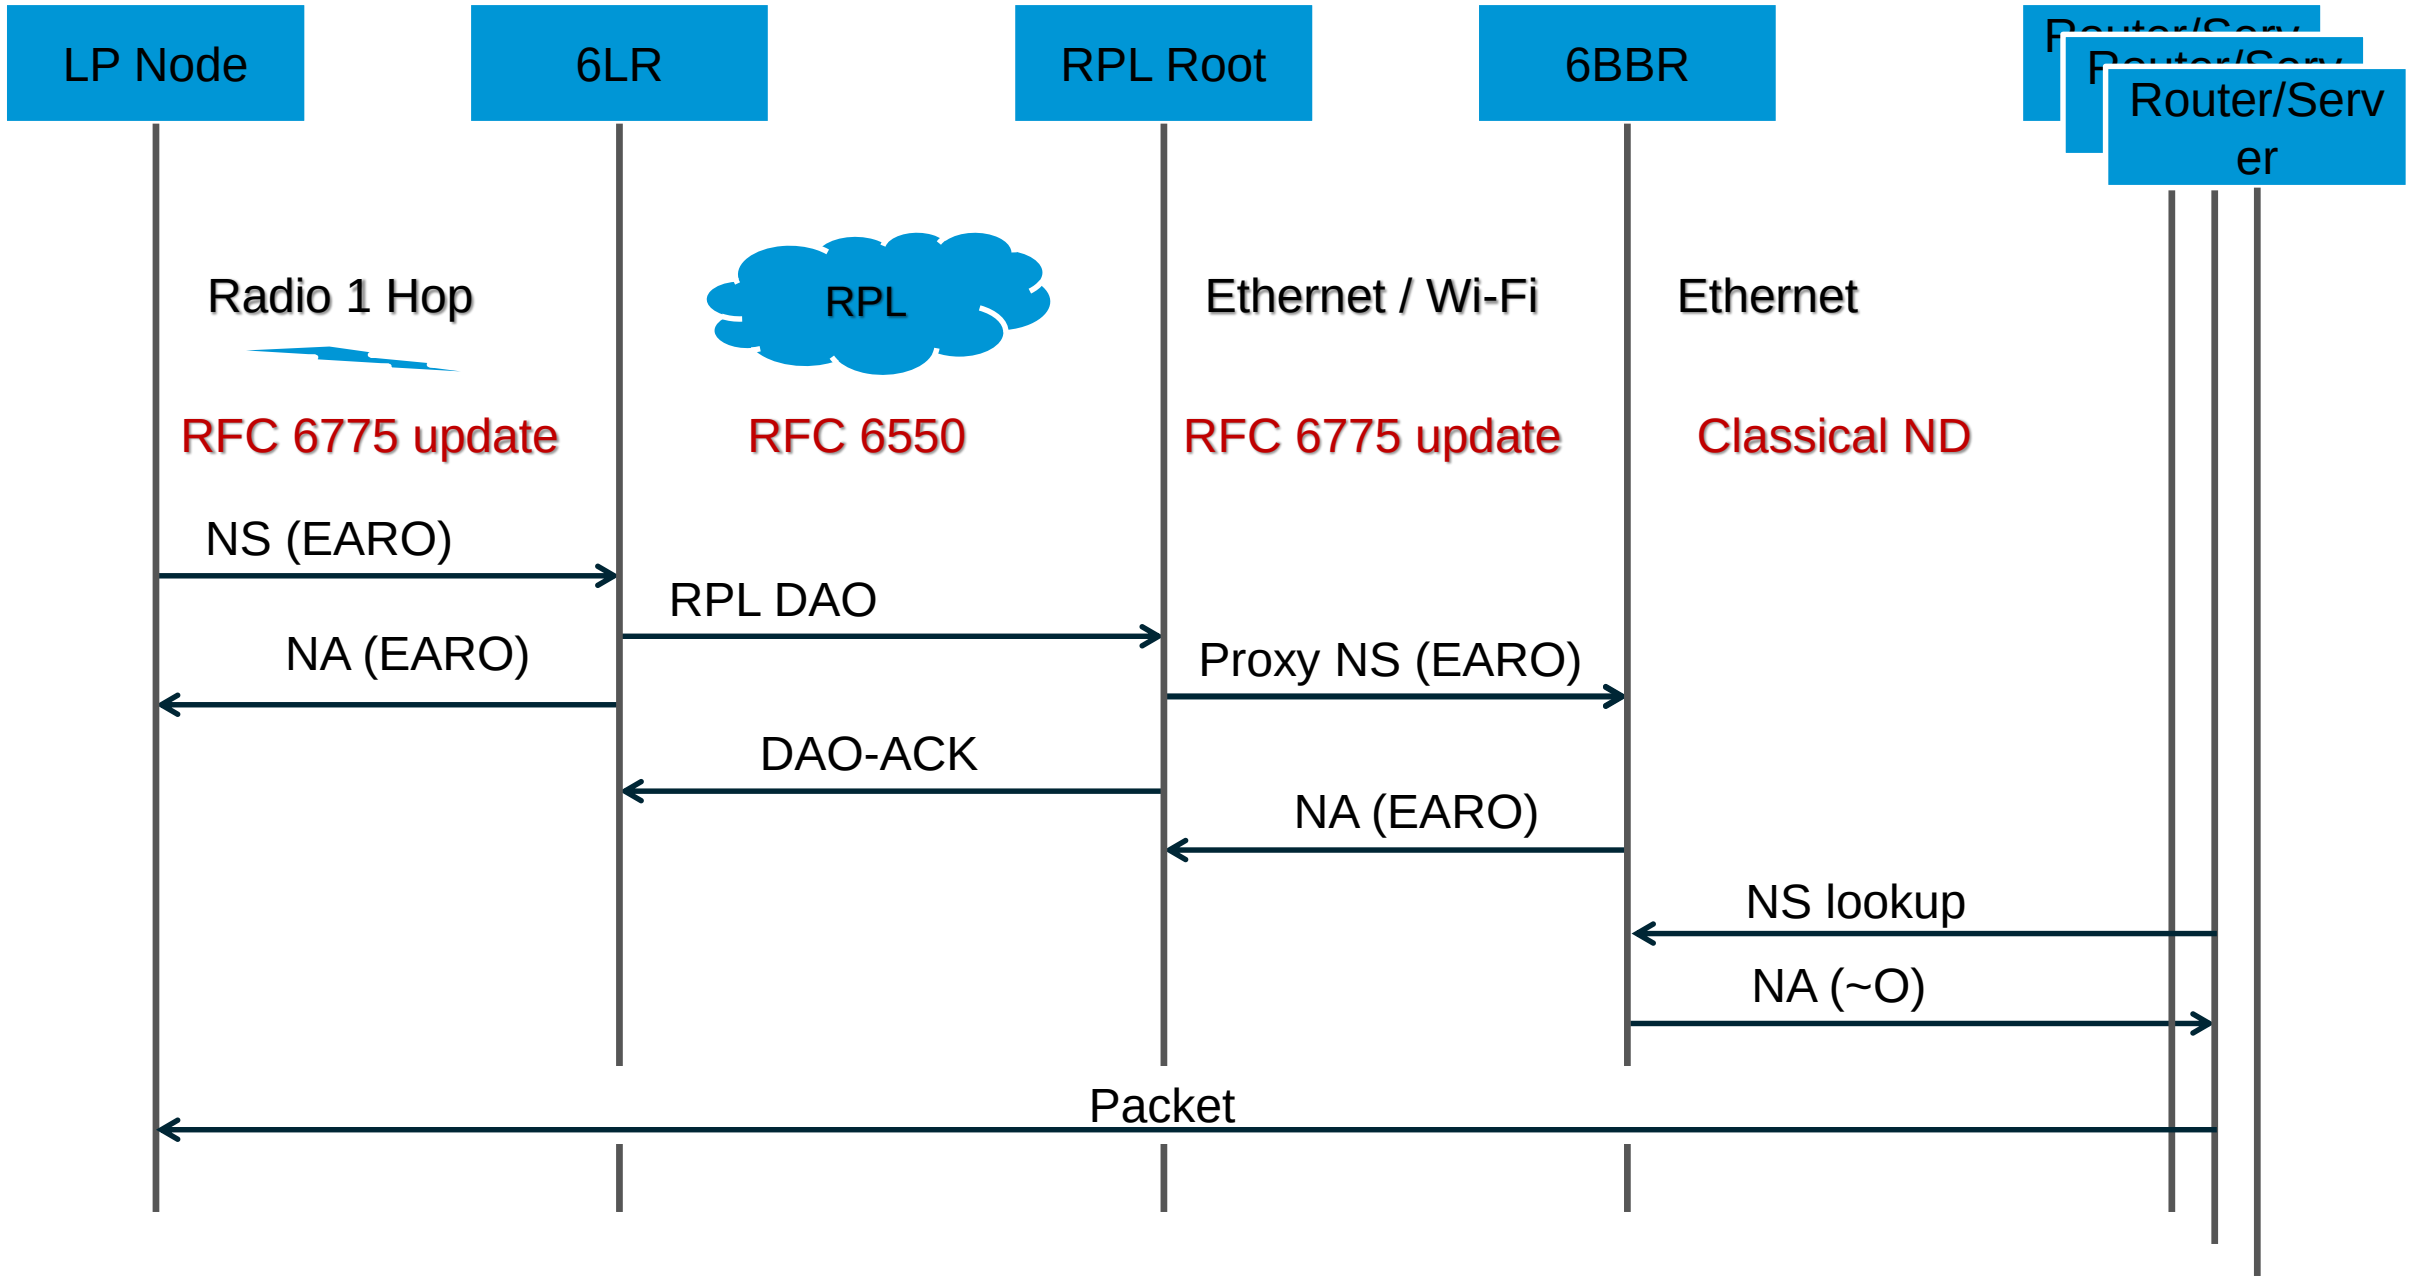
LP Node
6LR
RPL Root
6BBR
Router/Server
Radio 1 Hop
RPL
Ethernet / Wi-Fi
Ethernet
RFC 6775 update
RFC 6550
RFC 6775 update
Classical ND
NS (EARO)
RPL DAO
NA (EARO)
Proxy NS (EARO)
DAO-ACK
NA (EARO)
NS lookup
NA (~O)
Packet

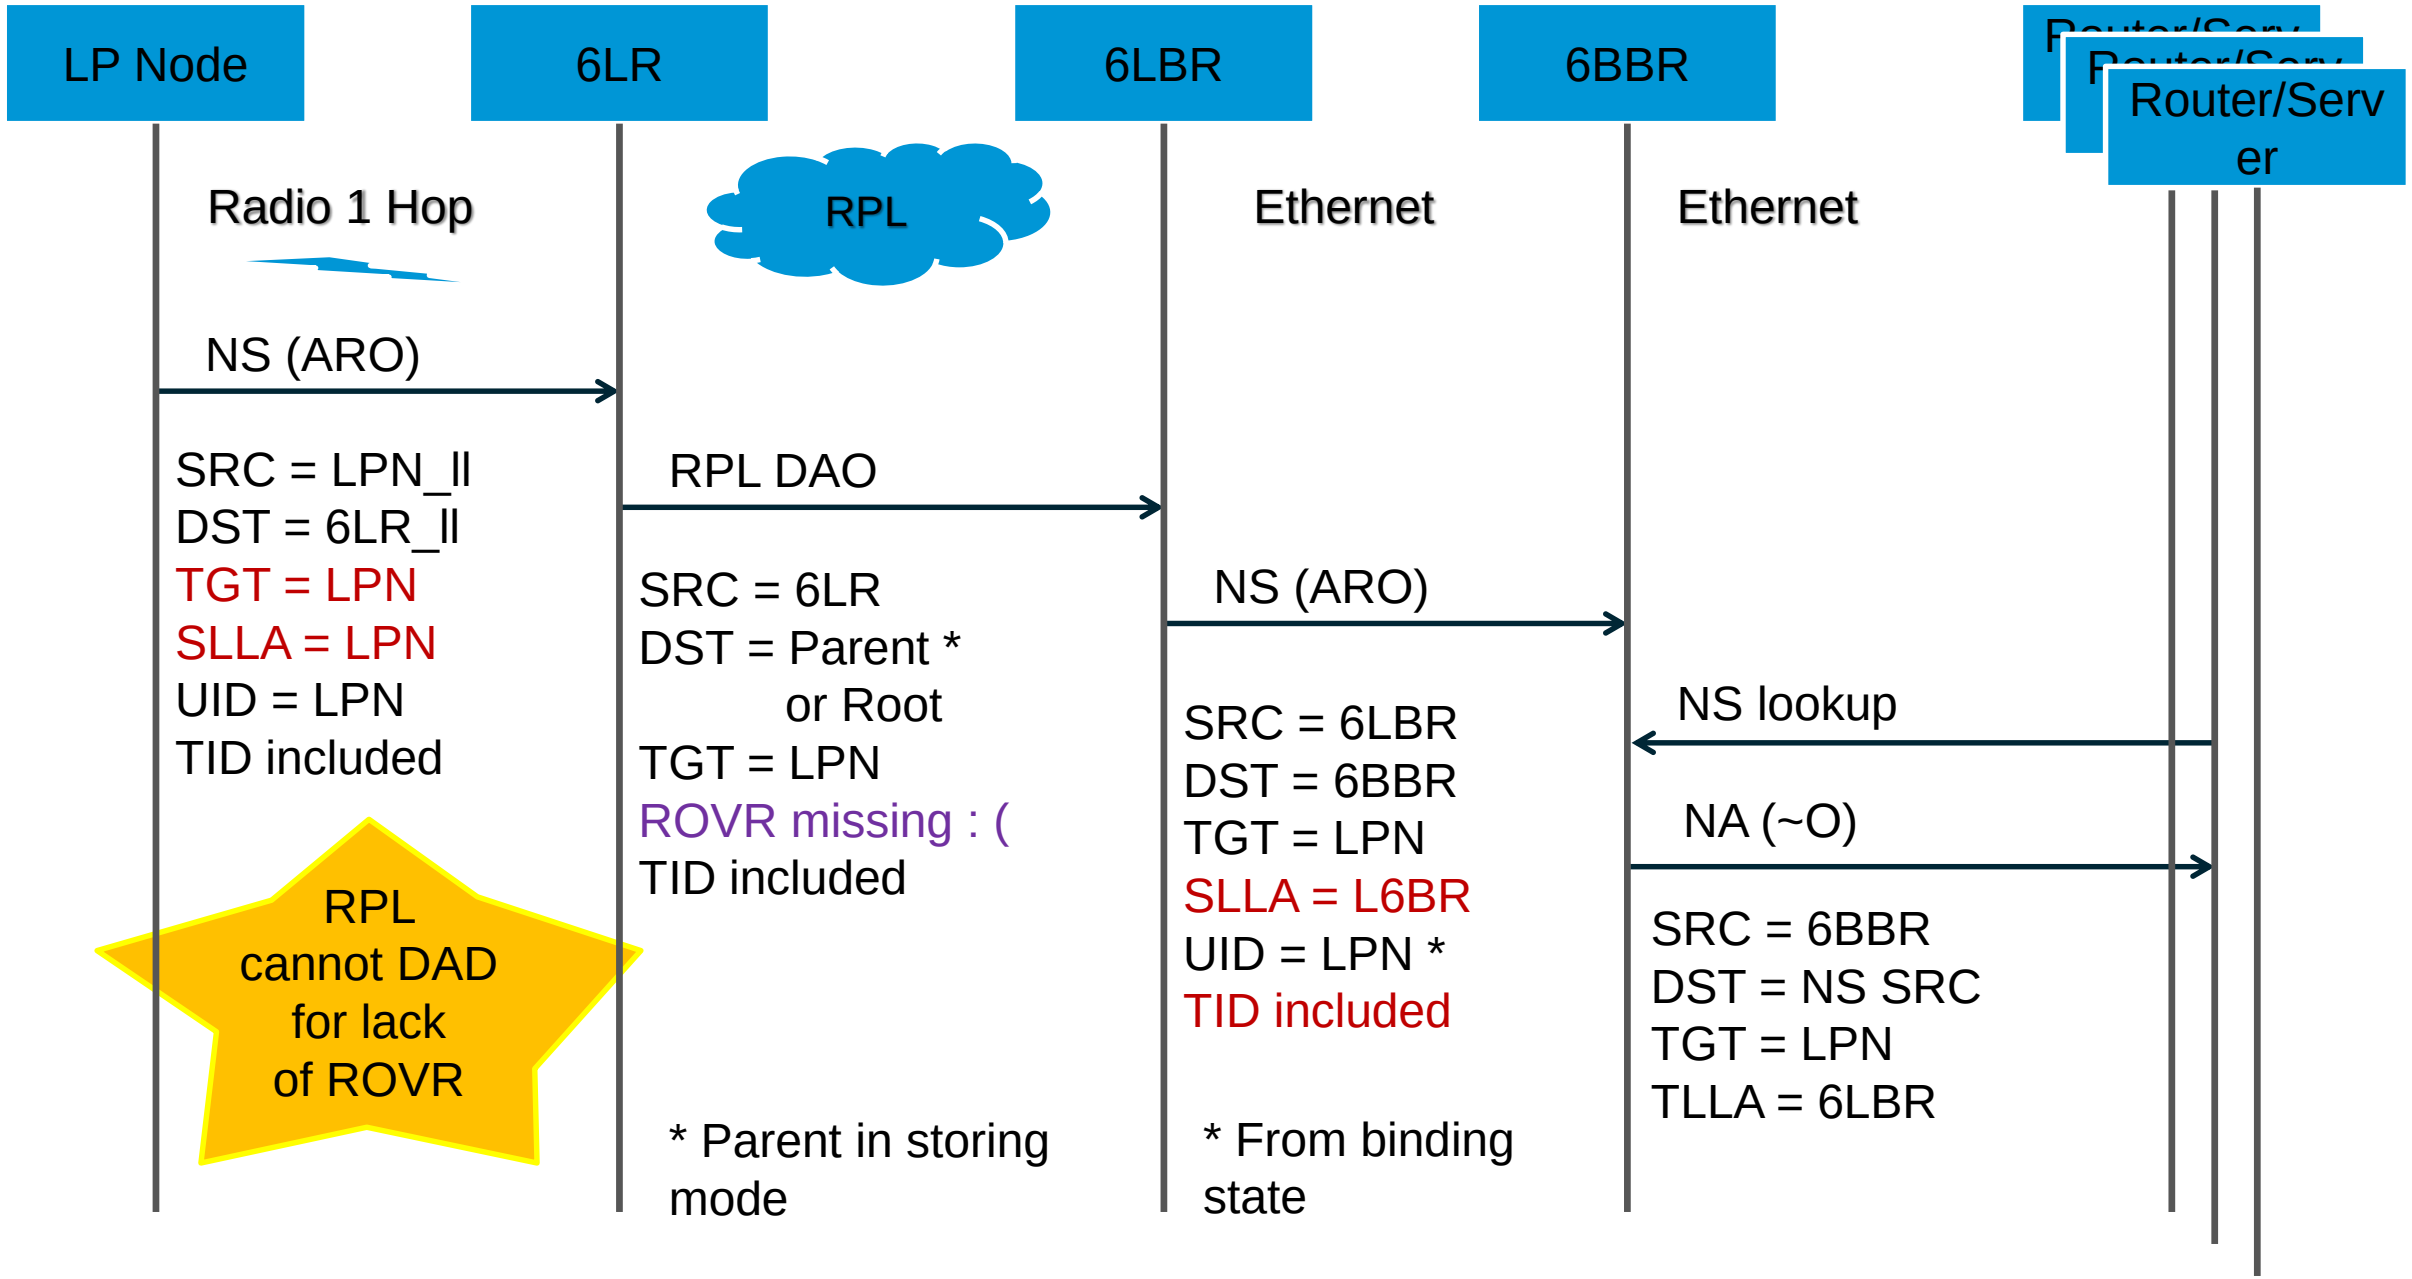
LP Node
6LR
6LBR
6BBR
Router/Server
Radio 1 Hop
RPL
Ethernet
Ethernet
NS (ARO)
SRC = LPN_ll
DST = 6LR_ll
TGT = LPN
SLLA = LPN
UID = LPN
TID included
RPL DAO
SRC = 6LR
DST = Parent *
or Root
TGT = LPN
ROVR missing : (
TID included
NS (ARO)
SRC = 6LBR
DST = 6BBR
TGT = LPN
SLLA = L6BR
UID = LPN *
TID included
NS lookup
NA (~O)
SRC = 6BBR
DST = NS SRC
TGT = LPN
TLLA = 6LBR
RPL cannot DAD for lack of ROVR
* Parent in storing mode
* From binding state

# WGLC?

# draft-thubert-roll-unaware-leaves

P.Thubert

IETF 101

London

# Terminology

- RFC 6550:
  - A RPL leaf may understands RPL
  - But does not Act as a router
- This draft: A RPL-unaware leaf does not implement anything specific to RPL, but it MUST support draft-rfc6775-update

# Notes on the ‘R’ flag (defined in draft-rfc-6775-update)

- A RPL Unaware Leaf does not know that there is routing in place and that the routing is RPL; draft-thubert-roll-unaware-leaves does not require anything from the Leaf.
- draft-rfc-6775-update specifies a new flag in the EARO, the 'R' flag.
- If the 'R' flag is set, the Registering Node expects that the 6LR ensures reachability for the Registered Address, e.g., by means of routing or proxying ND.
- Conversely, when it is not set, the 'R' flag indicates that the Registering Node is a router, which for instance participates to RPL and that it will take care of injecting its Address over the routing protocol by itself.
- A 6LN that acts only as a host, when registering, MUST set the 'R'  to indicate that it is not a router and that it will not handle its own reachability.
- A 6LR that manages its reachability SHOULD NOT set the 'R' flag; if it does, routes towards this router may be installed on its behalf and may interfere with those it injects.

# RPL Unaware Leaf (RUL) operation

- Note: The RUL does not know that there is routing in place and that the routing is RPL; draft-thubert-roll-unaware-leaves does not require anything from the Leaf Node. The ‘R’ flag is defined in draft-rfc-6775-update and plain 6LNs MUST set it.
- A RPL-Unaware Leaf (RUL) sets the 'R' flag in the EARO to declare itself as a host with the expectation that the 6LR that accepts the registration injects routing information for the Registered Address in the RPL domain as described in draft-rfc-6775-update.
- The packet forwarding operation by the 6LR serving a Leaf 6LN is described in draft-ietf-roll-useofrplinfo.
- This doc draft-thubert-roll-unaware-leaves adds the capability by a 6LR to advertise the IPv6 address(es) of the 6LN in the RPL protocol.
- Examples of routing-agnostic 6LN may include lightly-powered sensors such as window smash sensor (alarm system), or the kinetically powered light switch.

# First registration

- Upon the first registration, the EDAR / EDAC populates a state in the 6LBR including the ROVR field and the 6LR sends a first DAO message.
- The RPL Root acts as a proxy on behalf of the 6LR upon the reception of the DAO propagation initiated at the 6LR. Should we allow splitting from the 6LBR, e.g.:

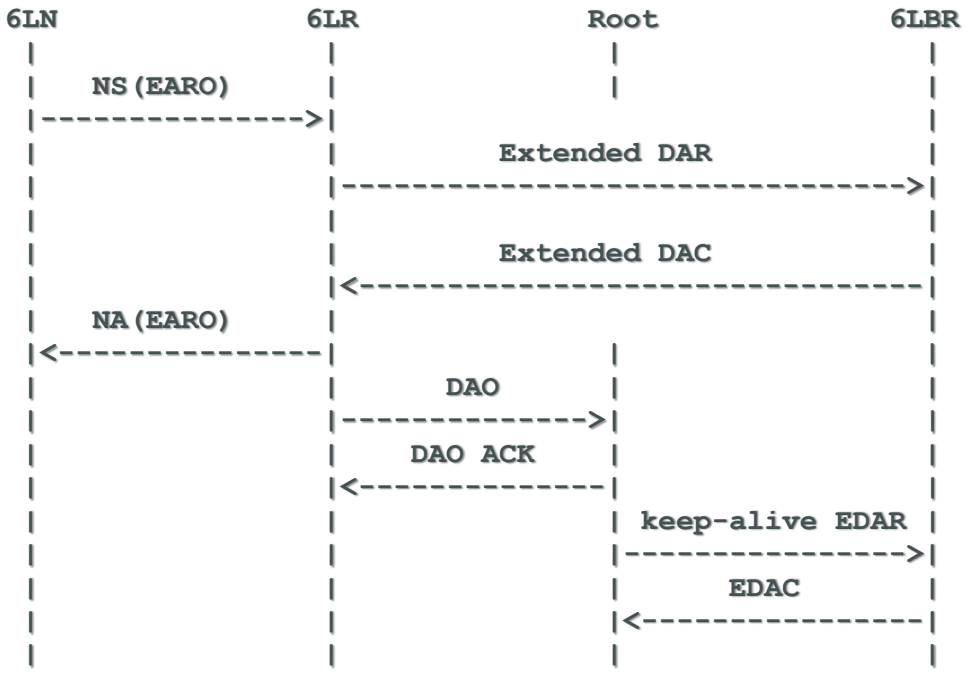

# EDA (DAR, DAC) message Proxying

- Upon the renewal of a 6lowPAN ND registration: if the 'R' flag is set, the 6LR injects a DAO targeting the Registered Address, and refrains from sending a DAR message.
- With a Root/6LBR split that could give:

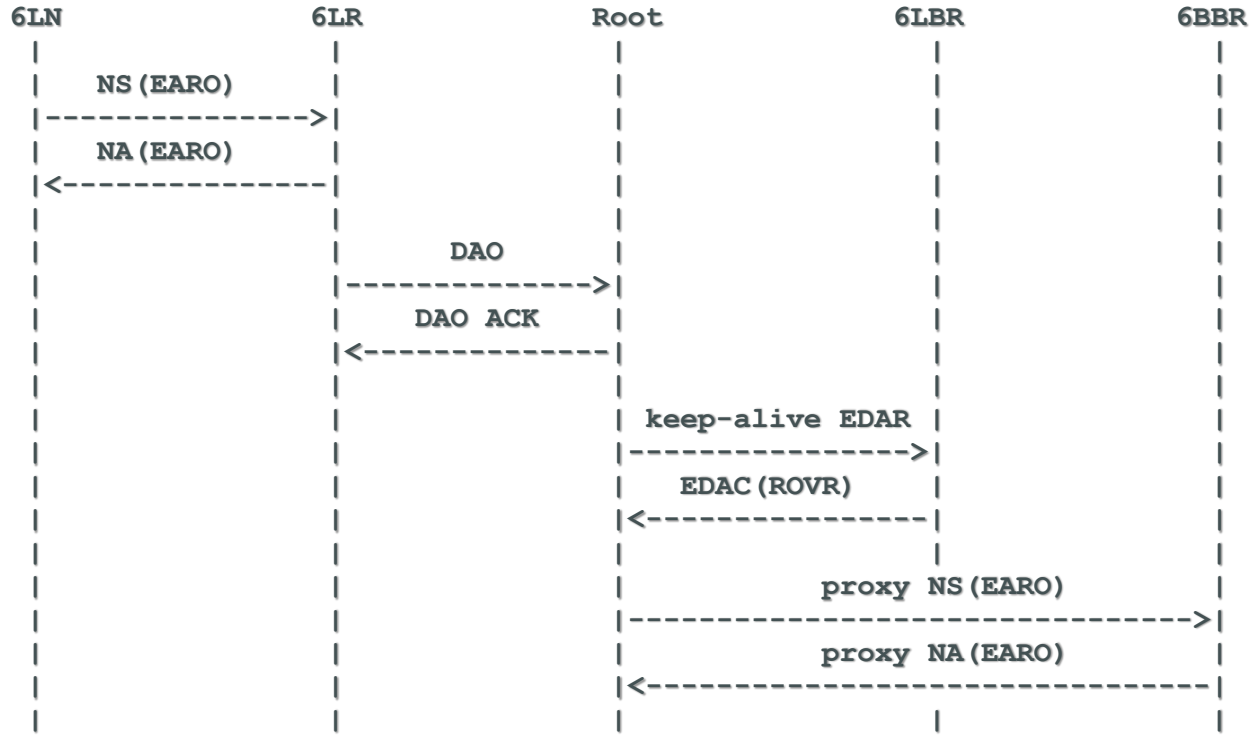

# Mapping Fields from RPL DAO to NS(EARO) and EDA

- The Registered Address in a RPL Target Option is a direct match to the Registered Address field of the EDAR message and in the Target field of the NS, respectively
- EARO's TID is a direct match to Path Sequence in Transit Information option (TIO)
- EARO's Lifetime unit is 60s. RPL uses Lifetime Units that is passed in the DODAG Configuration Option. Converting EARO to DAO and back requires mapping of units.
- The Registration Ownership Verifier (ROVR) field in keep-alive EDAR messages by the Root is set to 64-bits of all ones to indicate that it is not provided. It is obtained in the EDAC from the 6LBR and used in proxy registration.

  Q: Should we carry it in a RPL option in DAO messages?

# Efficient route invalidation for RPL - Performance Report

draft-ietf-roll-efficient-npdao-01

Rahul, Rabi, Zhen@ Huawei

IETF101, London

History:

IETF95 - Presented the problem statement

IETF96 - Presented existing solutions based on comments rcvd and why those fall short

IETF98 – Presented new solution for improving route invalidation

IETF99 – adopted as WG document , thank you for the review

IETF100 – Changes to message codes

IETF101 – Implementation Performance report

# Recap: the problem and the solution

## Current RPL NPDAO

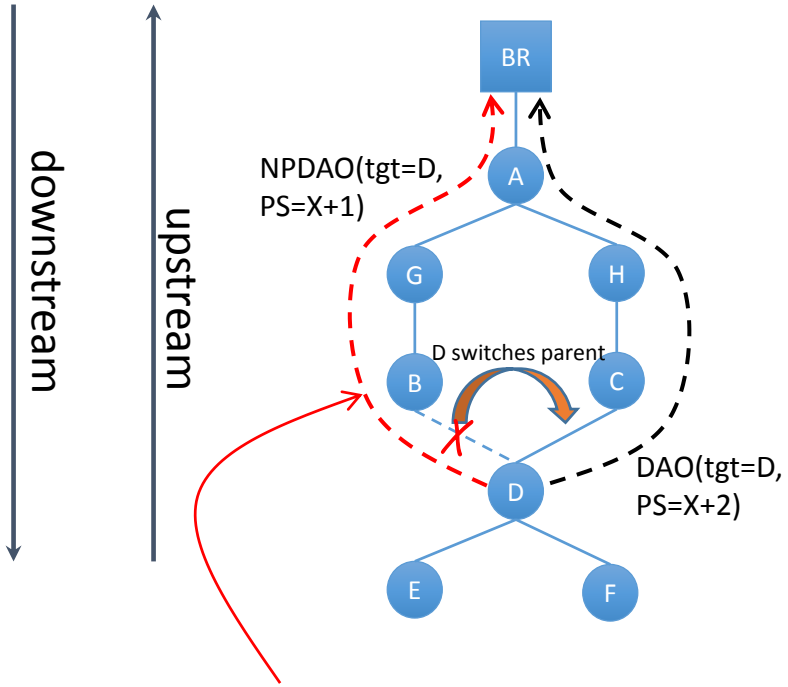


NP-DAO via broken links will cause many problems such as reachability and efficiency

## Proposed Change

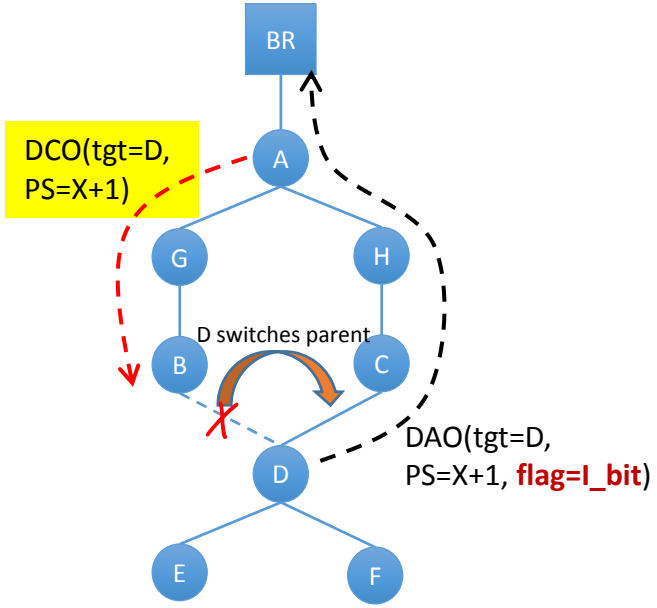


- Send the DAO via the new parent;
- Common parent to trigger the DCO to invalidate the previous path

PS = PathSequence
Tgt = Target

# Implementation Report

- Network stack: Contiki
- LOC: ~120 (link)
- δ RAM: 0 bytes
- δ flash: 840 bytes for CC2538

| system | mode | text | data | bss |
|---|---|---|---|---|
| x86_64 | without DCO | 137008 | 1916 | 26796 |
| | with DCO | 139020 | 1916 | 26796 |
| cc2538 | without DCO | 58477 | 2671 | 21551 |
| | with DCO | 59317 | 2671 | 21551 |

- Implementation Notes:
  - Send DCO on parent switching
  - Continue using NPDAO in other cases such as route lifetime expiry

- Further scope for improvisation…
  - Refactor dao_input/output and dco_input/output code
  - Refactor dao/dco ack handling code

# Performance Report

https://github.com/nyrahul/IETF_npdao_optimize/blob/master/DCO_performance_report.md

Two Scenarios:

1. Metric Deterioration (regular parent switching)
2. Connectivity Impairment

Two Tools:

1. Contiki
2. Whitefield-Framework (internally using NS3 lr-wpan)

Two Topologies:

1. n50_udp30
2. n100_udp30

# Topology

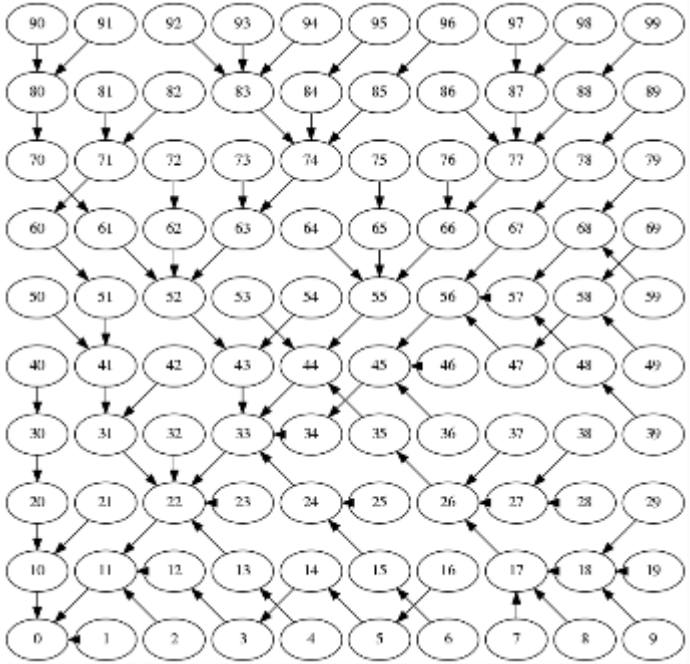


n100_udp30

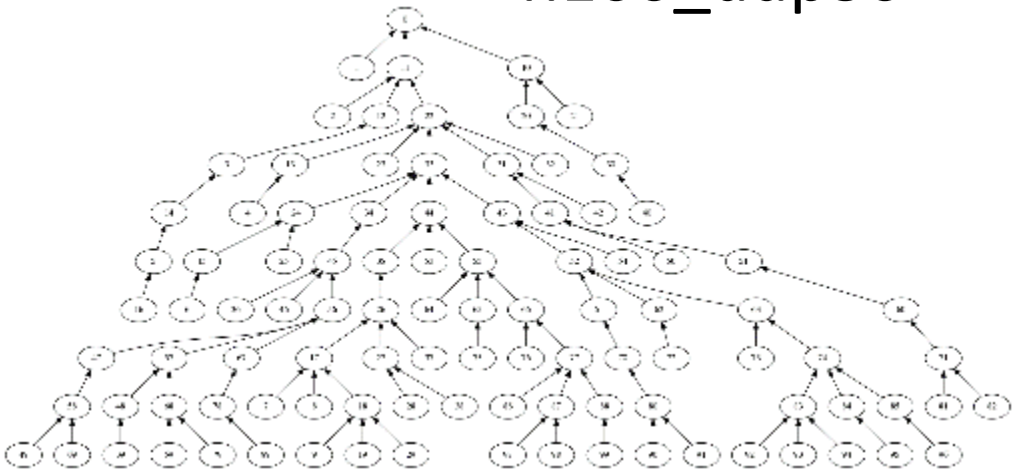

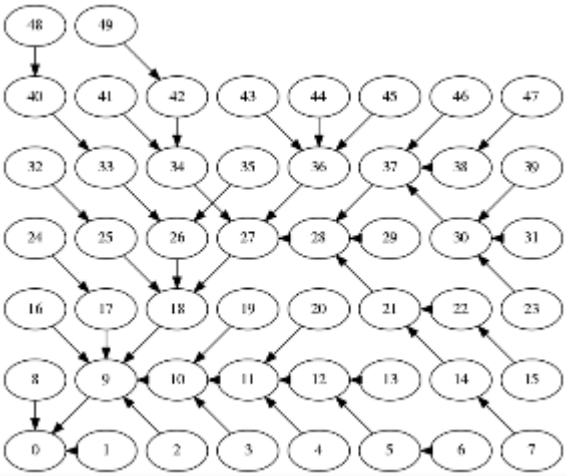


n50_udp30

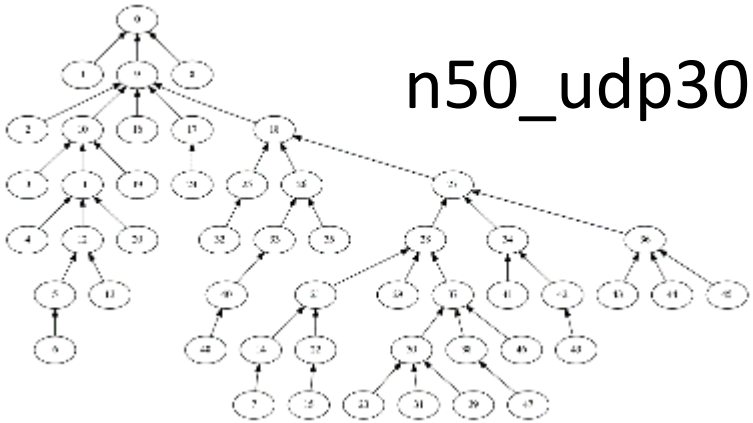

# Scenario 1: Regular Case

- No deliberately introduced impairment
- During initial network formation, the parent switching is high
- Measure:
  - Impact on control overhead
  - Impact on stale entries
- Observations
  - Control overhead decrease
    - Contiki NPDAO sub-optimal
  - Stale entries reduction
    - Steep gradient for DCO

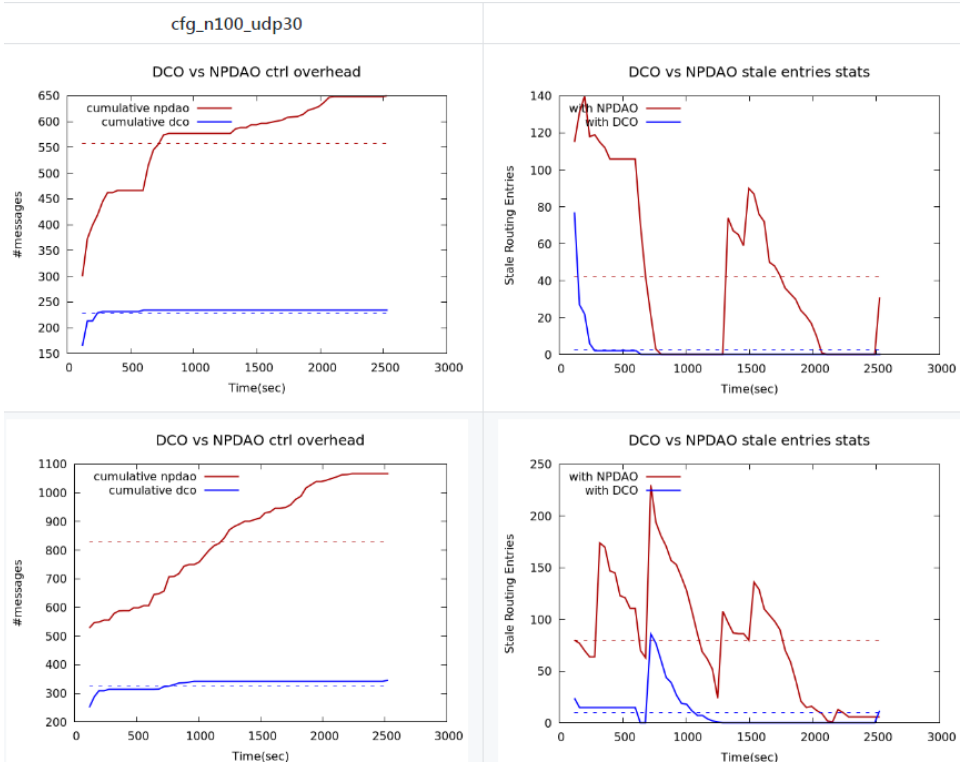

# Scenario 2: Connectivity impairment

- Deliberate impairment
- Thus more #parent_switches
- Higher impact on route invalidation control overhead
- Stale entries still linger
  - Because for dependent nodes we wait for the child nodes to send DAO without triggering a subDODAG DAO update by incrementing DTSN increase

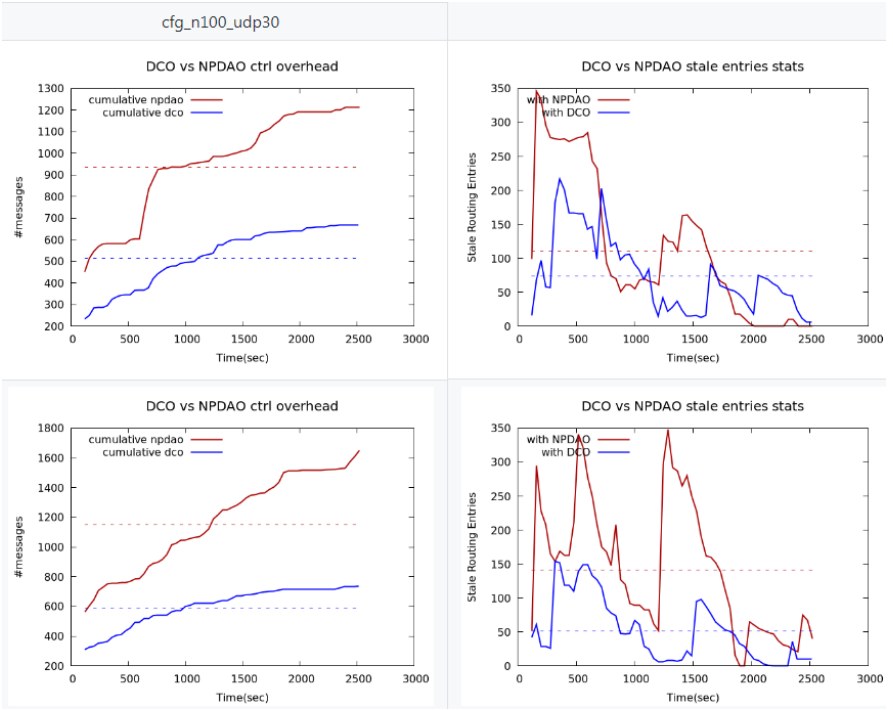

# Important observation

- How NPDAO is initiated?
  - Before DAO or after DAO?
  - Usual flow:

```
onParentSwitch(oldparent, newparent) {
    invalidateRoute(oldparent); //NPDAO gets sent now
    scheduleDAO(newparent); //DAO gets sent after random delay
}
```

  - Problem with this:
    - NPDAO reaches all the way to the BR, causing ‘significantly’ higher control overhead
    - Results in route downtime because of async DAO/NPDAO operation…
- Contiki and RIOT follows the same implementation flow!
- Hence DCO control overhead is significantly low
  - It flows only within sub-DODAG rooted at common ancestor

# Implementation ease with DCO

- Since the target node does not send DCO by itself

```
onParentSwitch(oldparent, newparent) {
    //No sending of NPDAO/DCO here
    scheduleDAO(newparent); //DAO gets sent after random delay
}
```

- Thus synchronized operation of DAO and DCO
  - Routes are cleared only when new path is established

# Next Step

- Quite stable for some time
- We think the document is ready for the next step (WGLC)
  - More reviews will certainly help.

# Thank you

# Root initiated routing state in RPL

## draft-ietf-dao-projection

Pascal Thubert

IETF 101

London, July 2017

# Root initiated routing state in RPL

P.Thubert

IETF 101

London

# Changes Highlights

- No major change
- Need to revisit the MOP, 3 bits, gets saturated

# Differentiating Storing vs. non-storing

- The non-storing mode P-DAO discussed in section Section 4.1 has a single VIO with one or more Via Addresses in it, the list of Via Addresses indicating the source-routed path to the target to be installed in the router that receives the message, which replies to the root directly with a DAO-ACK message.
- The storing mode P-DAO discussed in section Section 4.2 has at least two Via Information options with one Via Address each, for the ingress and the egress of the path, and more if there are intermediate routers. In normal operations, the P-DAO is propagated along the chain of Via Routers from the egress router of the path till the ingress one, which confirms the installation to the root with a DAO-ACK message.

# Discussions

How is the topology known to the root?

How are the node capabilities known to the root?

Complexity of mixed modes

MOP saturation

Compression of the Via Info option (so far full addresses)

Loop avoidance

- in particular for loose and not end to end route
- Recommend Setting the ‘O’ bit

<RFC6550>: “Down 'O': 1-bit flag indicating whether the packet is expected to progress Up or Down. A router sets the 'O' flag when the packet is expected to progress Down (using DAO routes), and clears it when forwarding toward the DODAG root to a node with a lower Rank). A host or RPL leaf node MUST set the 'O' flag to 0.”

# Backup

# New: non-storing mode transversal route

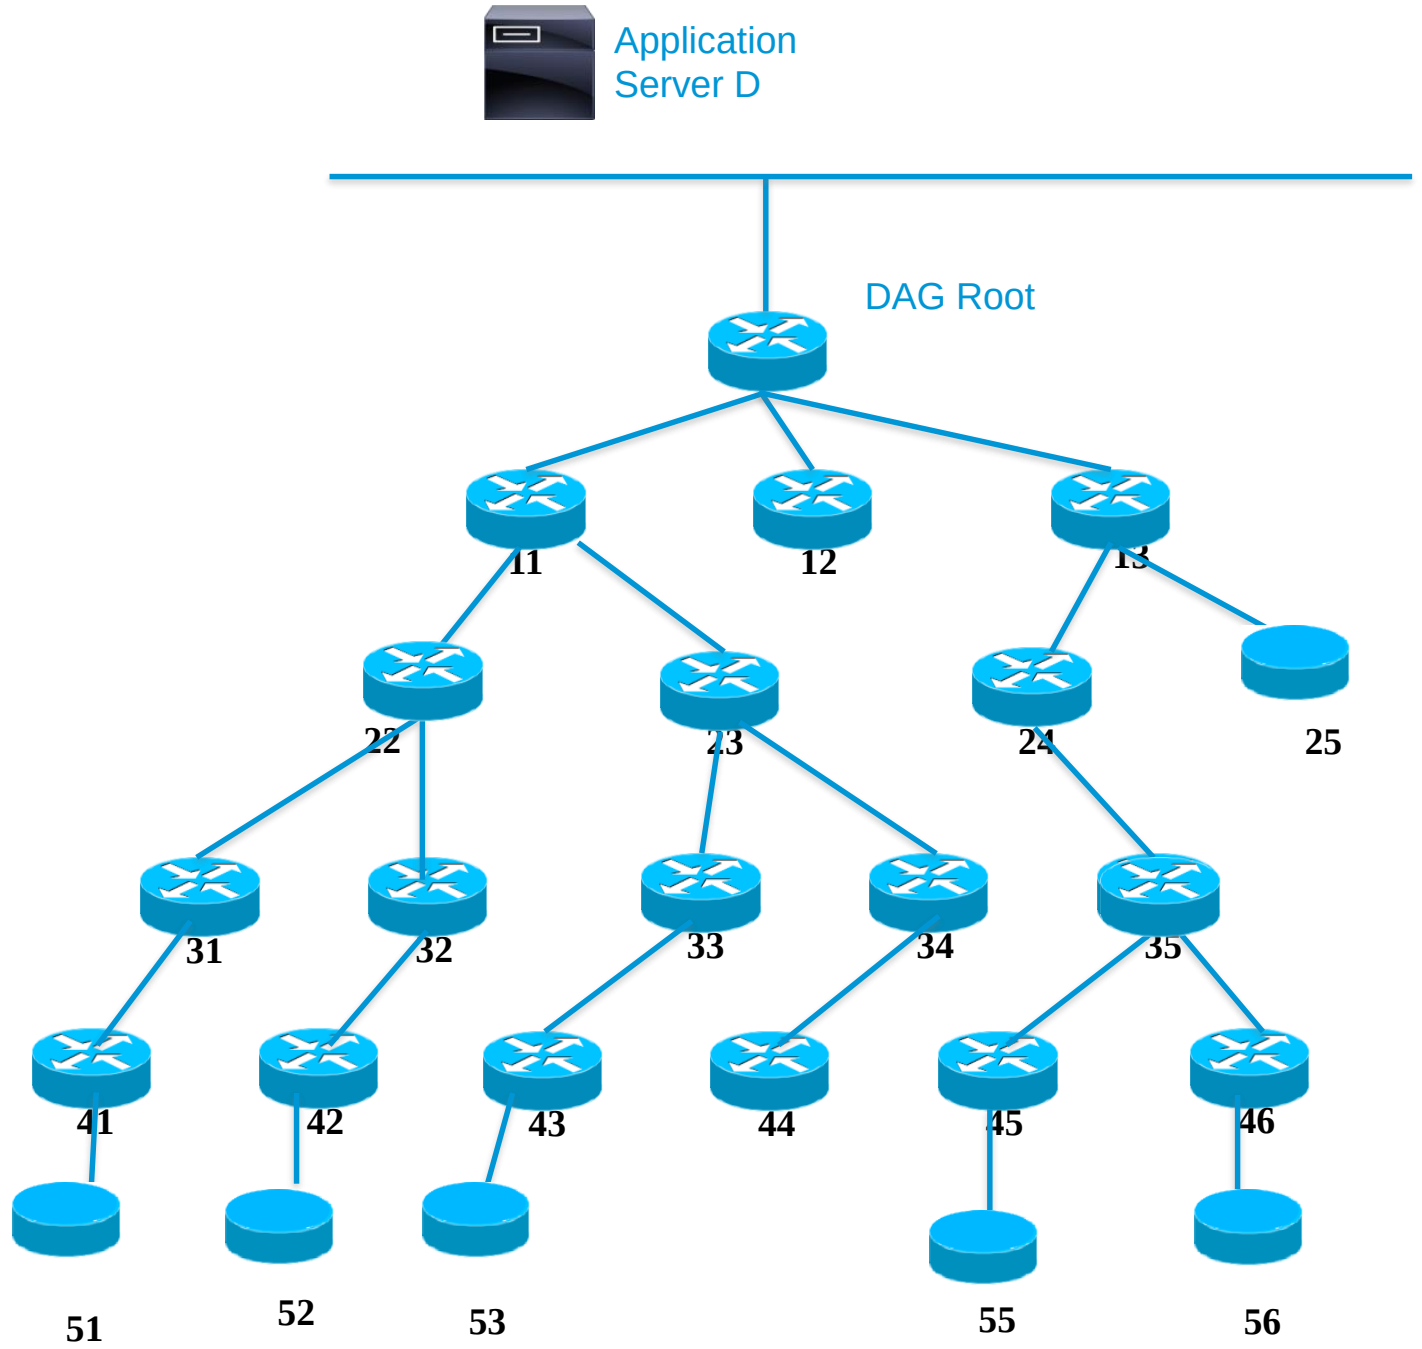
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

# Stretch in non-storing mode

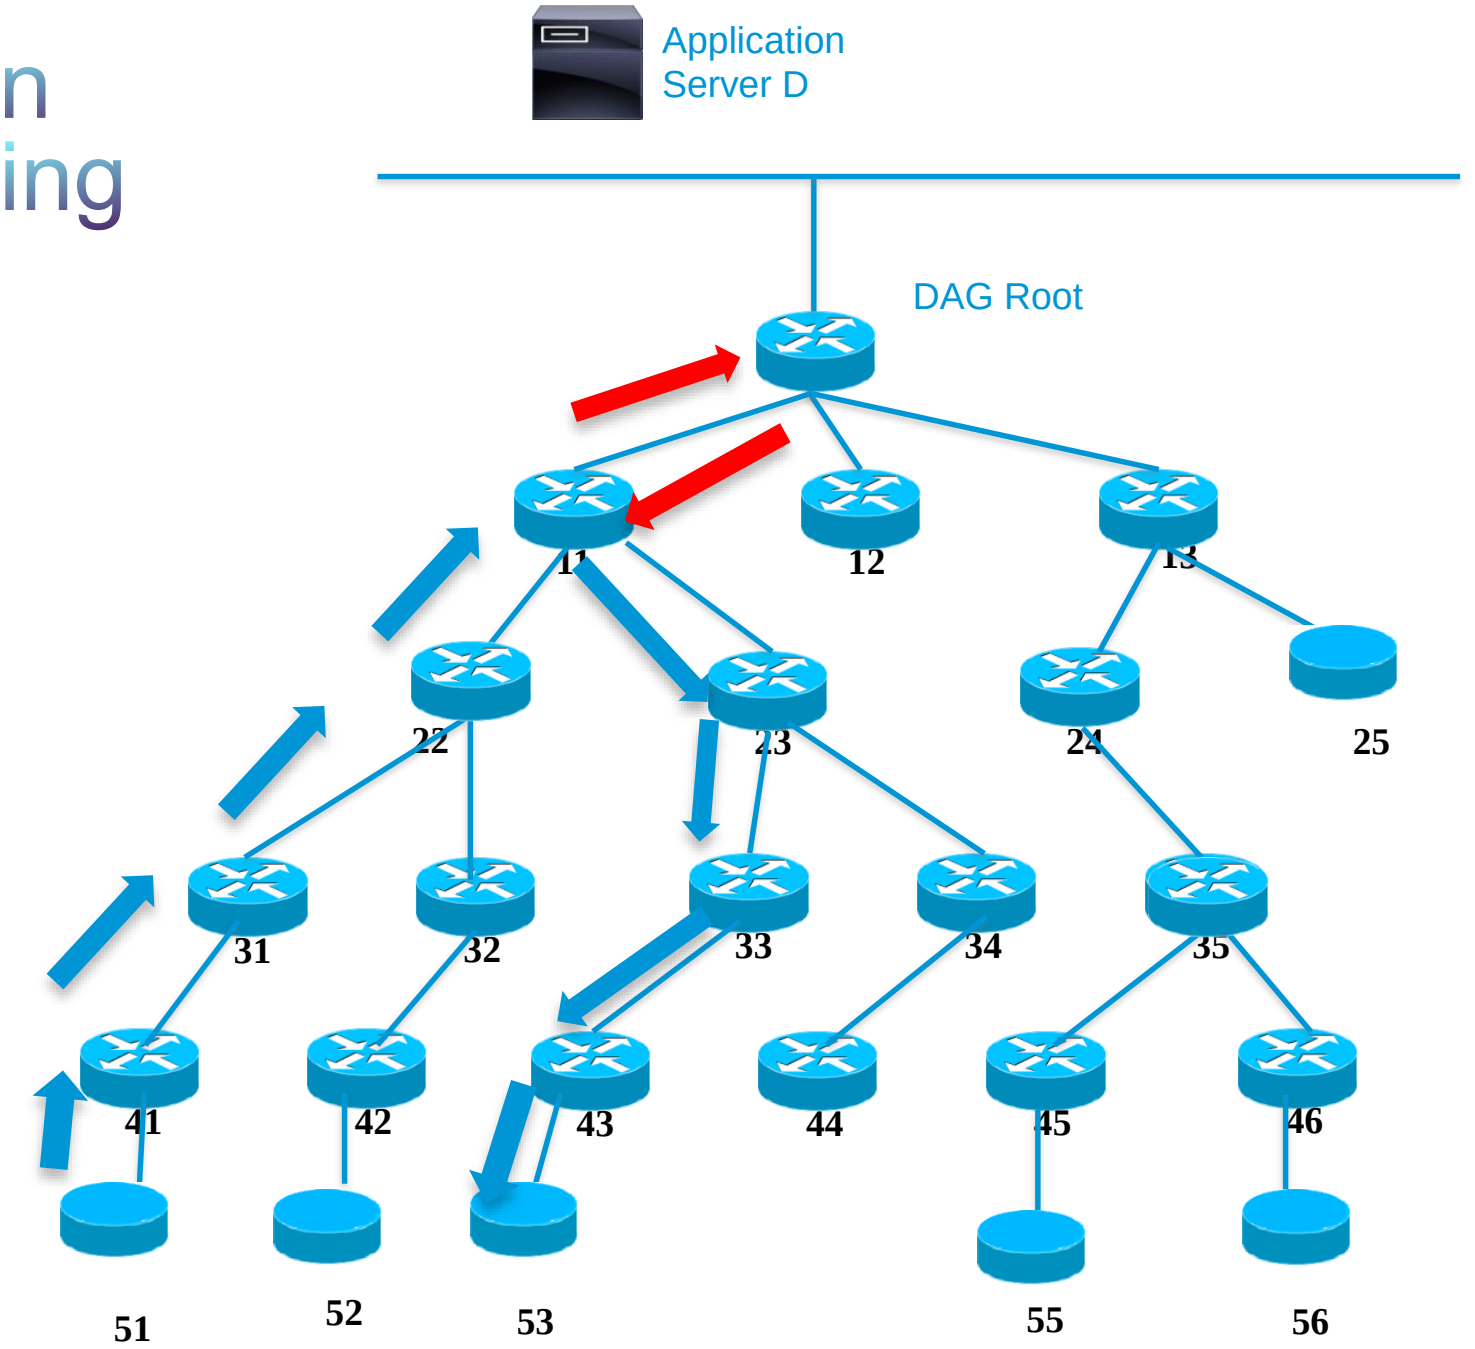

# Stretch in storing mode

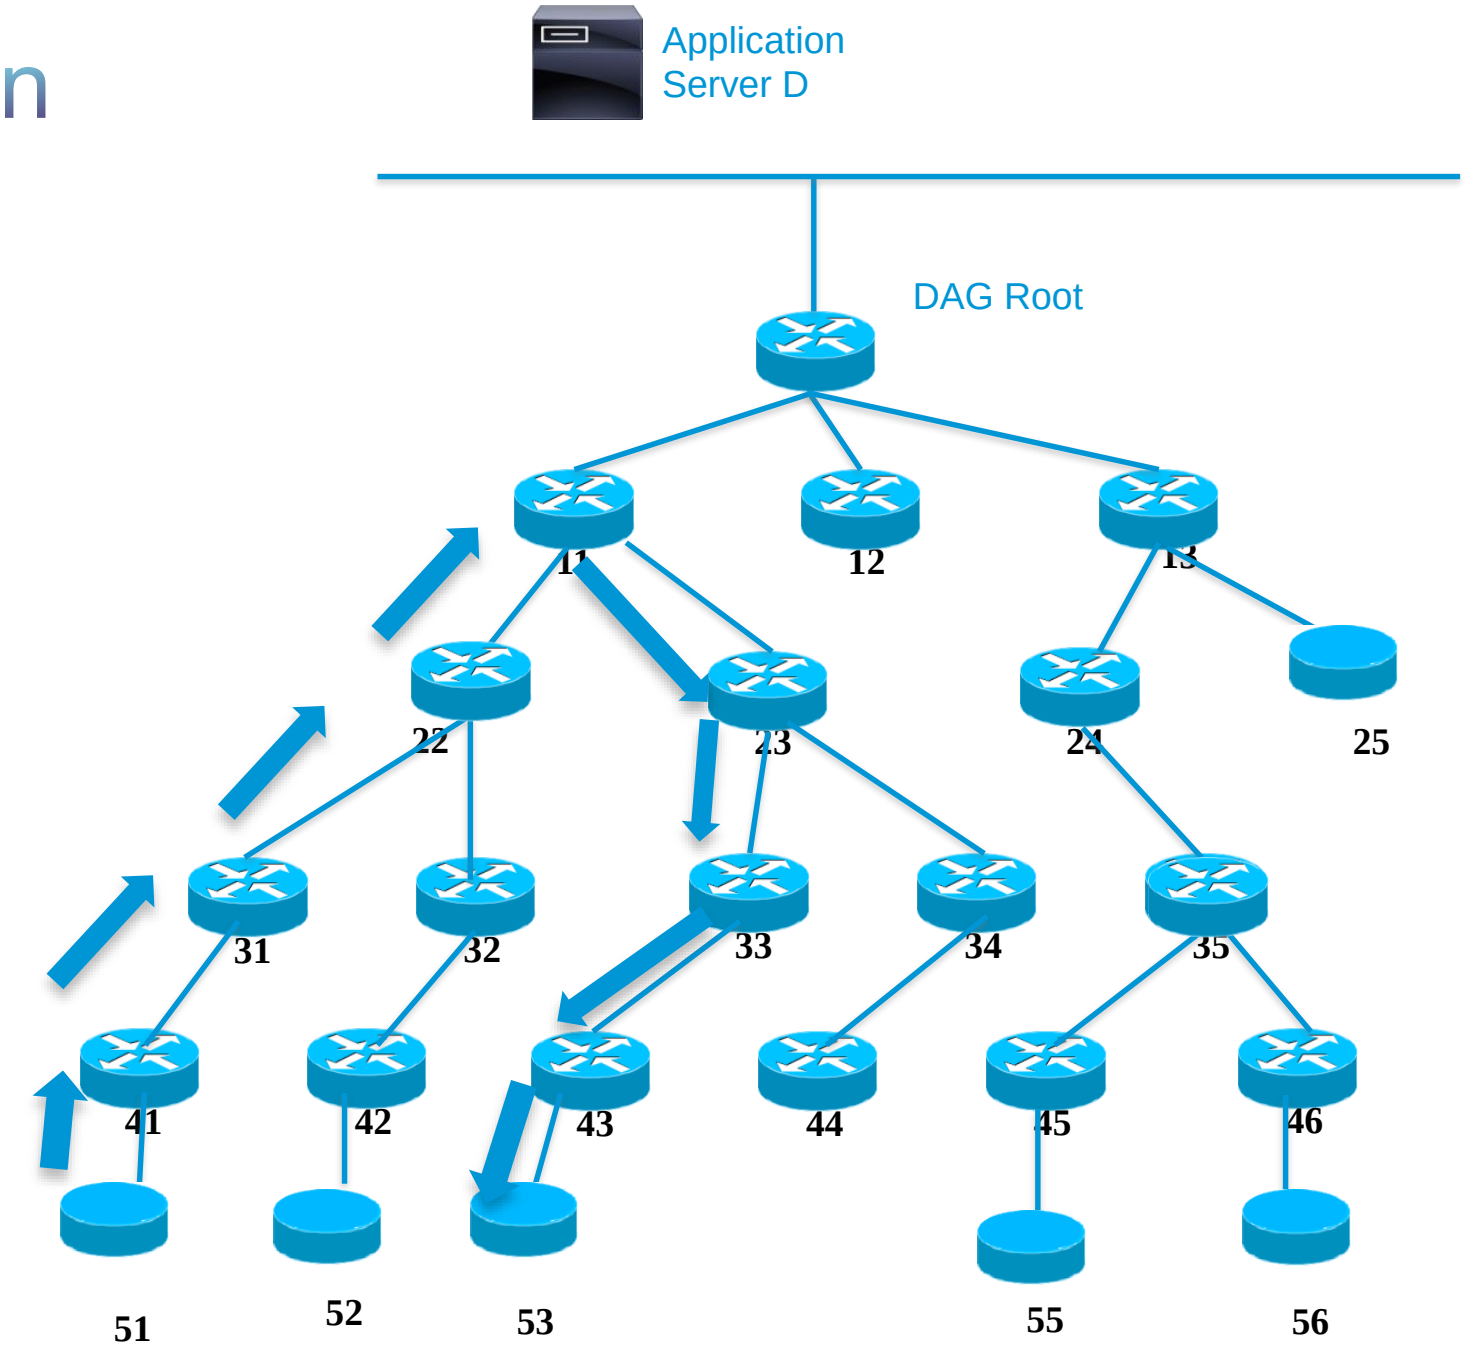

# DAO projection

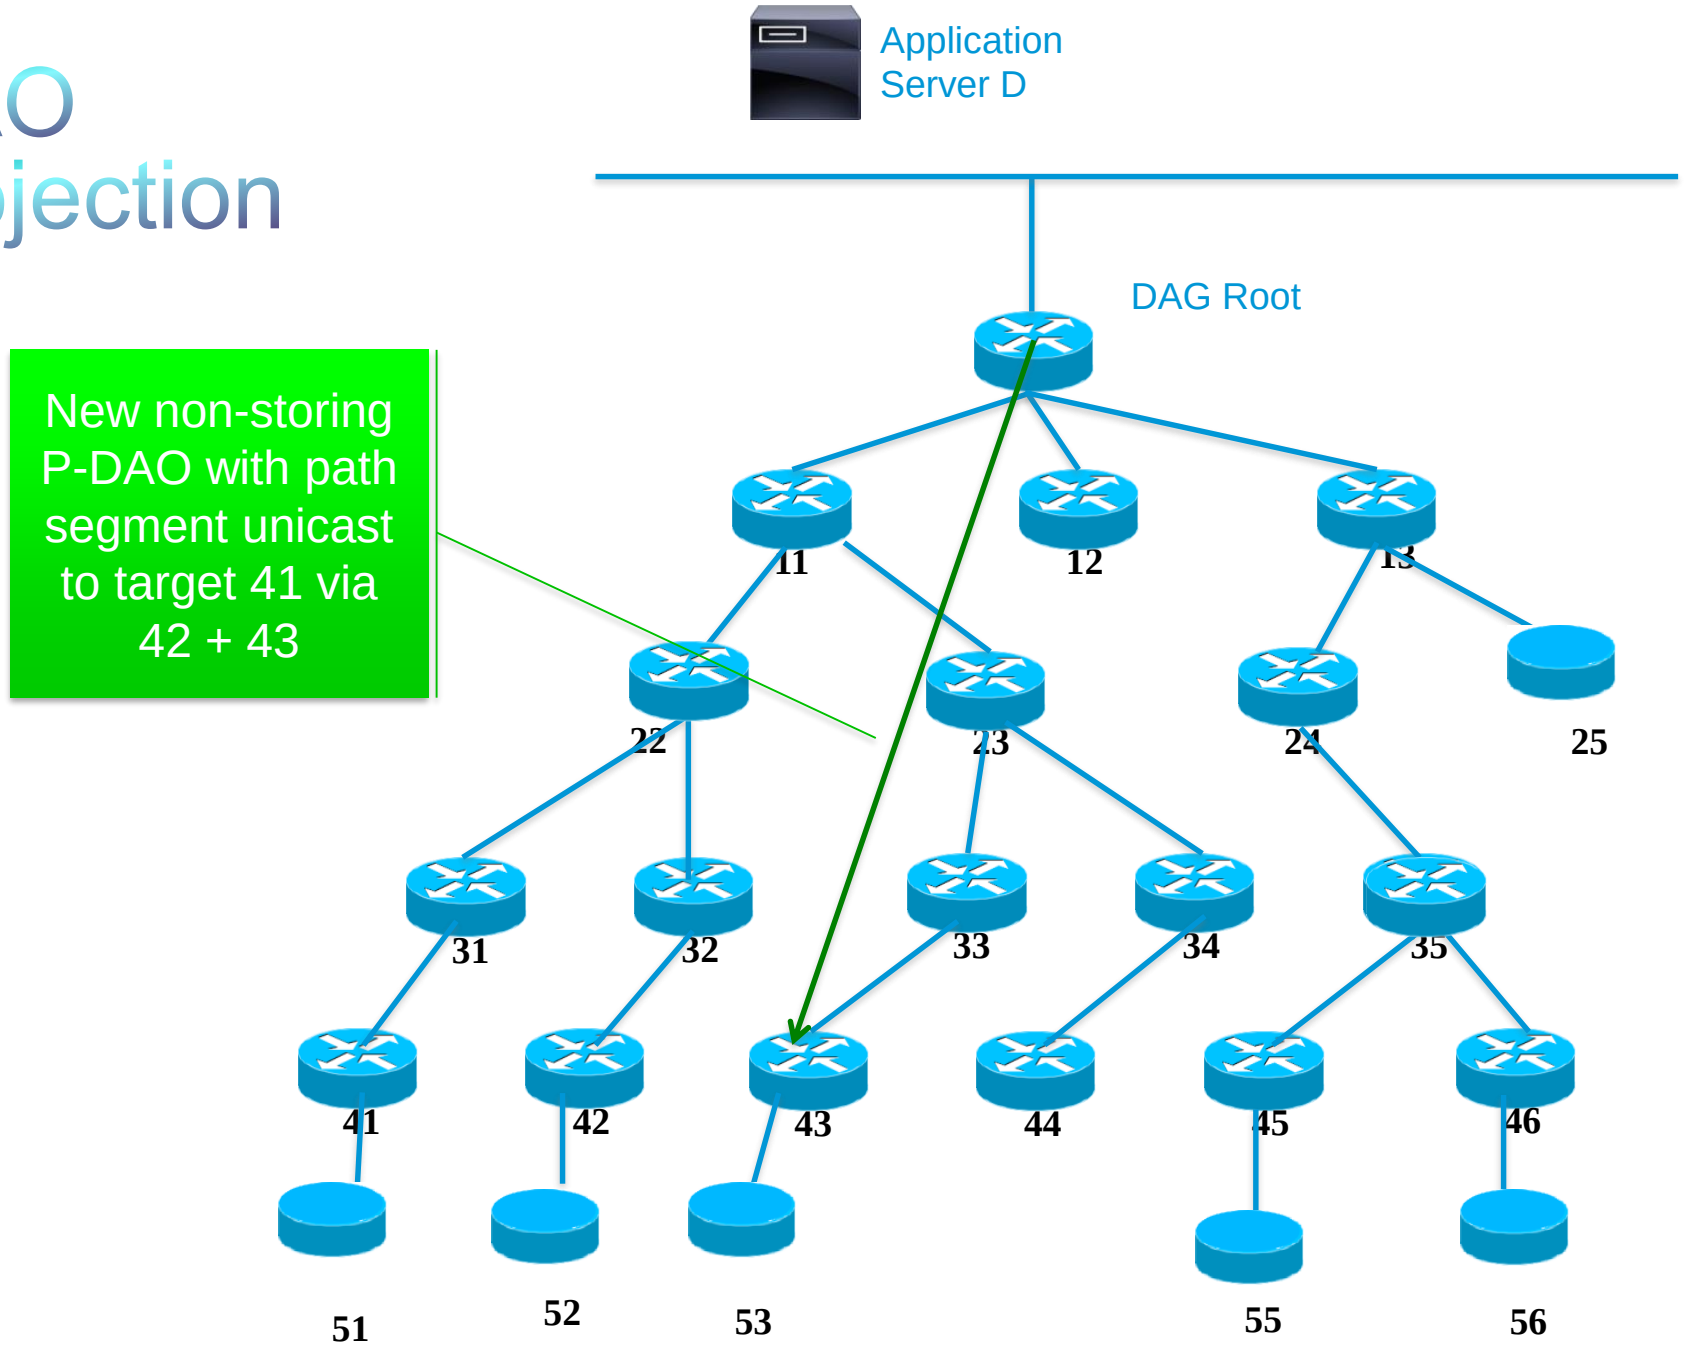

# Source routed DAO projection

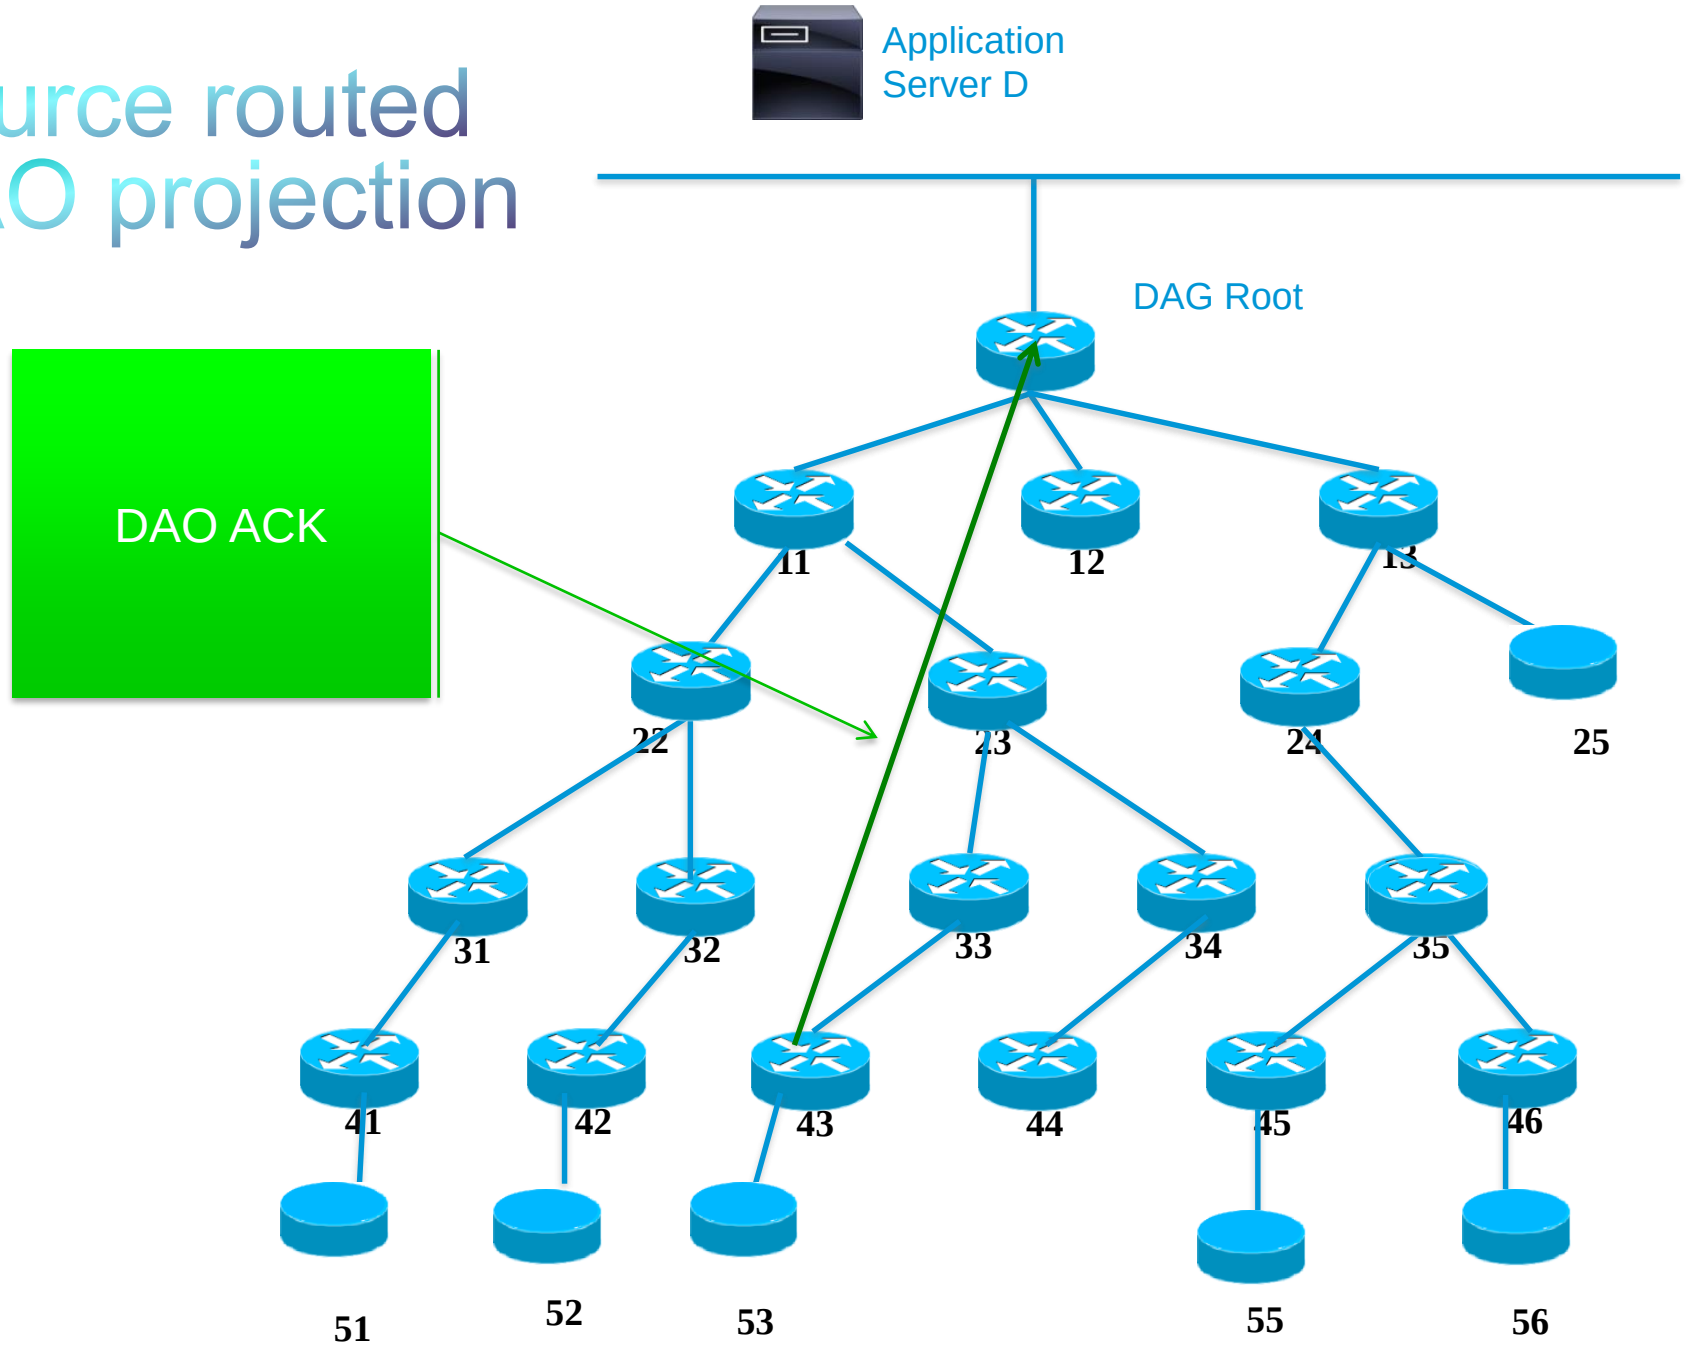

# DAO projection

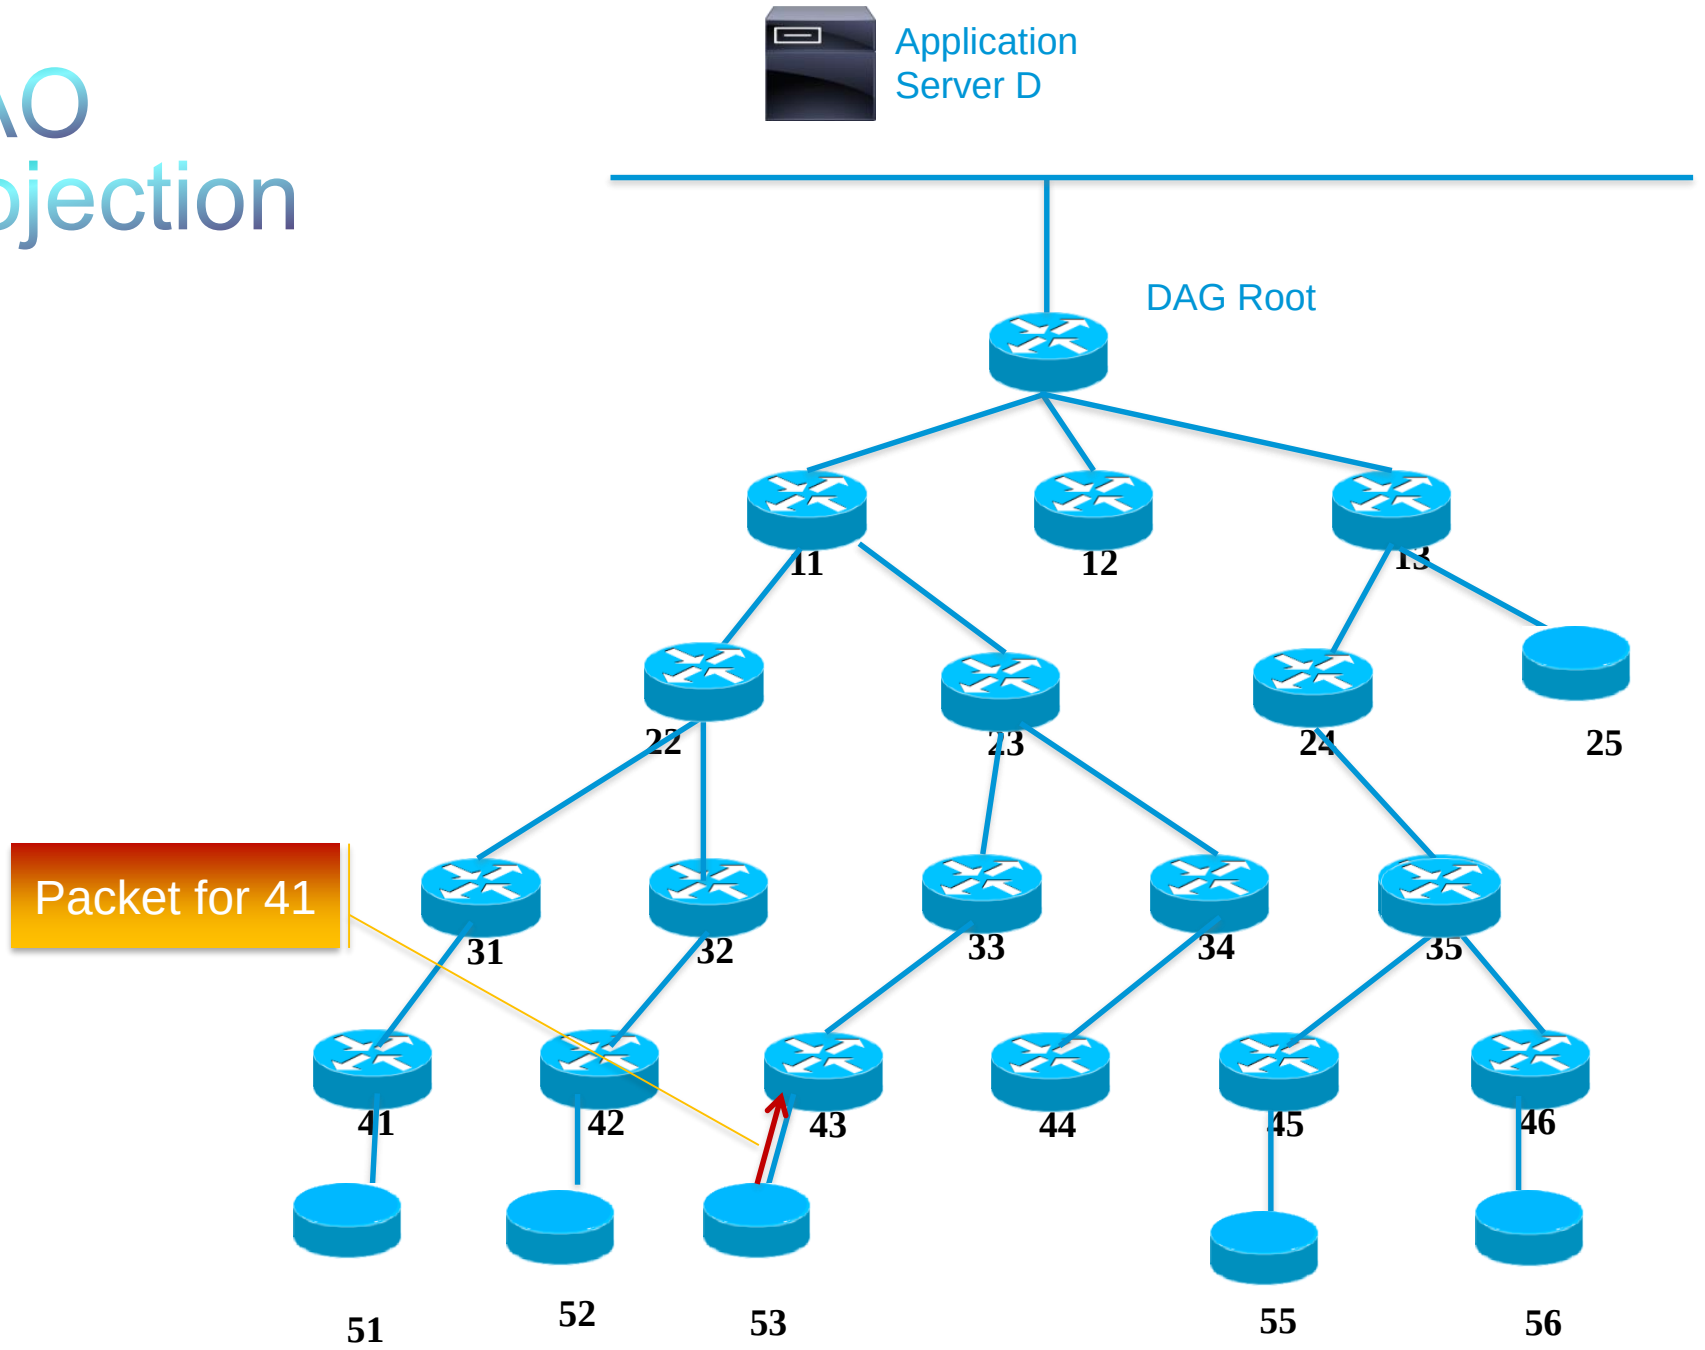

# DAO projection

Application Server D

DAG Root

IP in IP encapsulation with SRH 42, 41

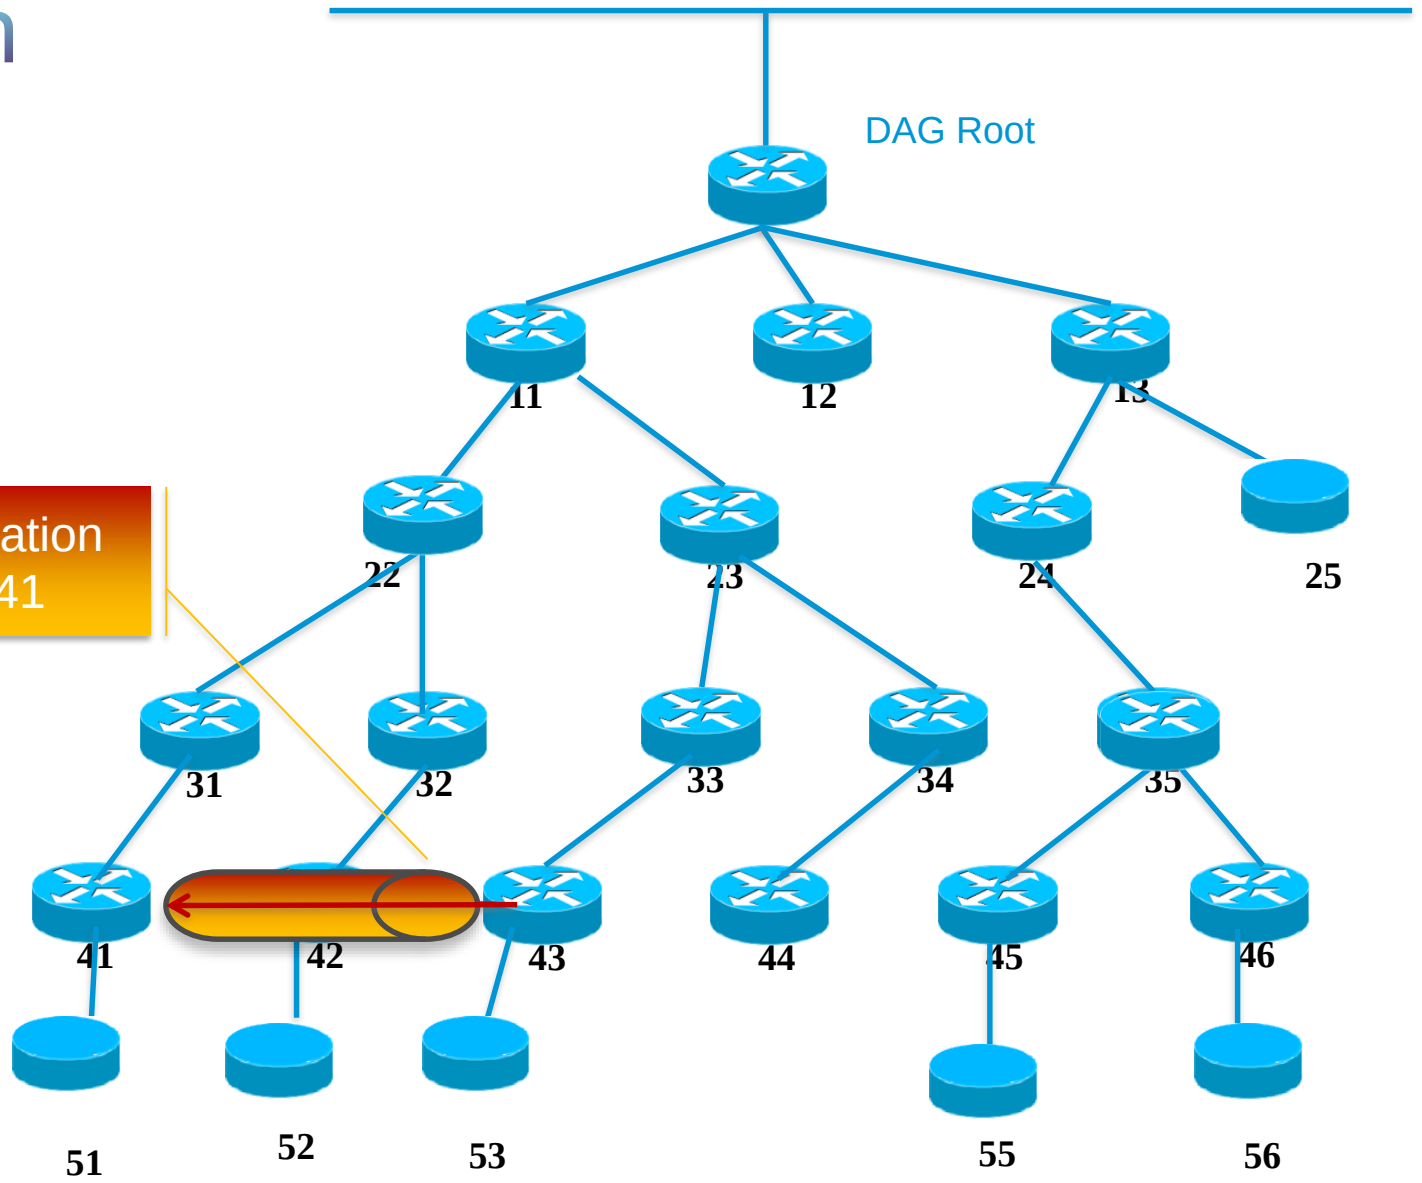

# Not done yet
# DAO projection

Application
Server D

DAG Root

storing P-DAO with path segment unicast to target 51 via 43==41 (egress)

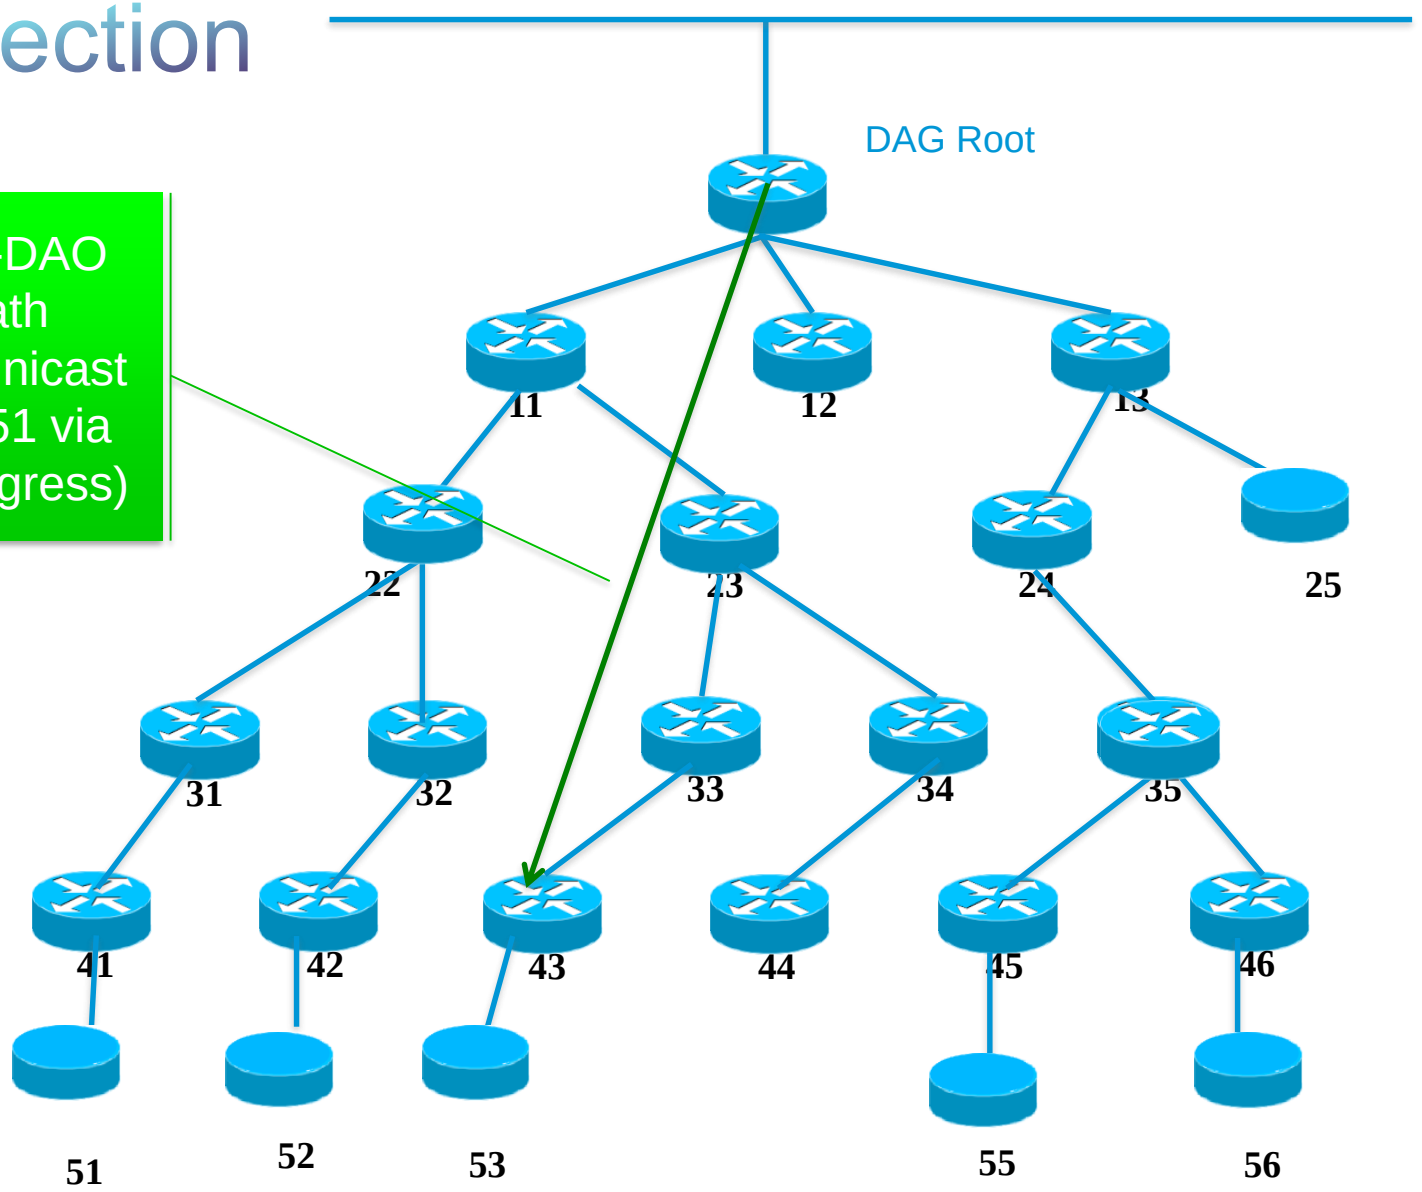

# DAO projection

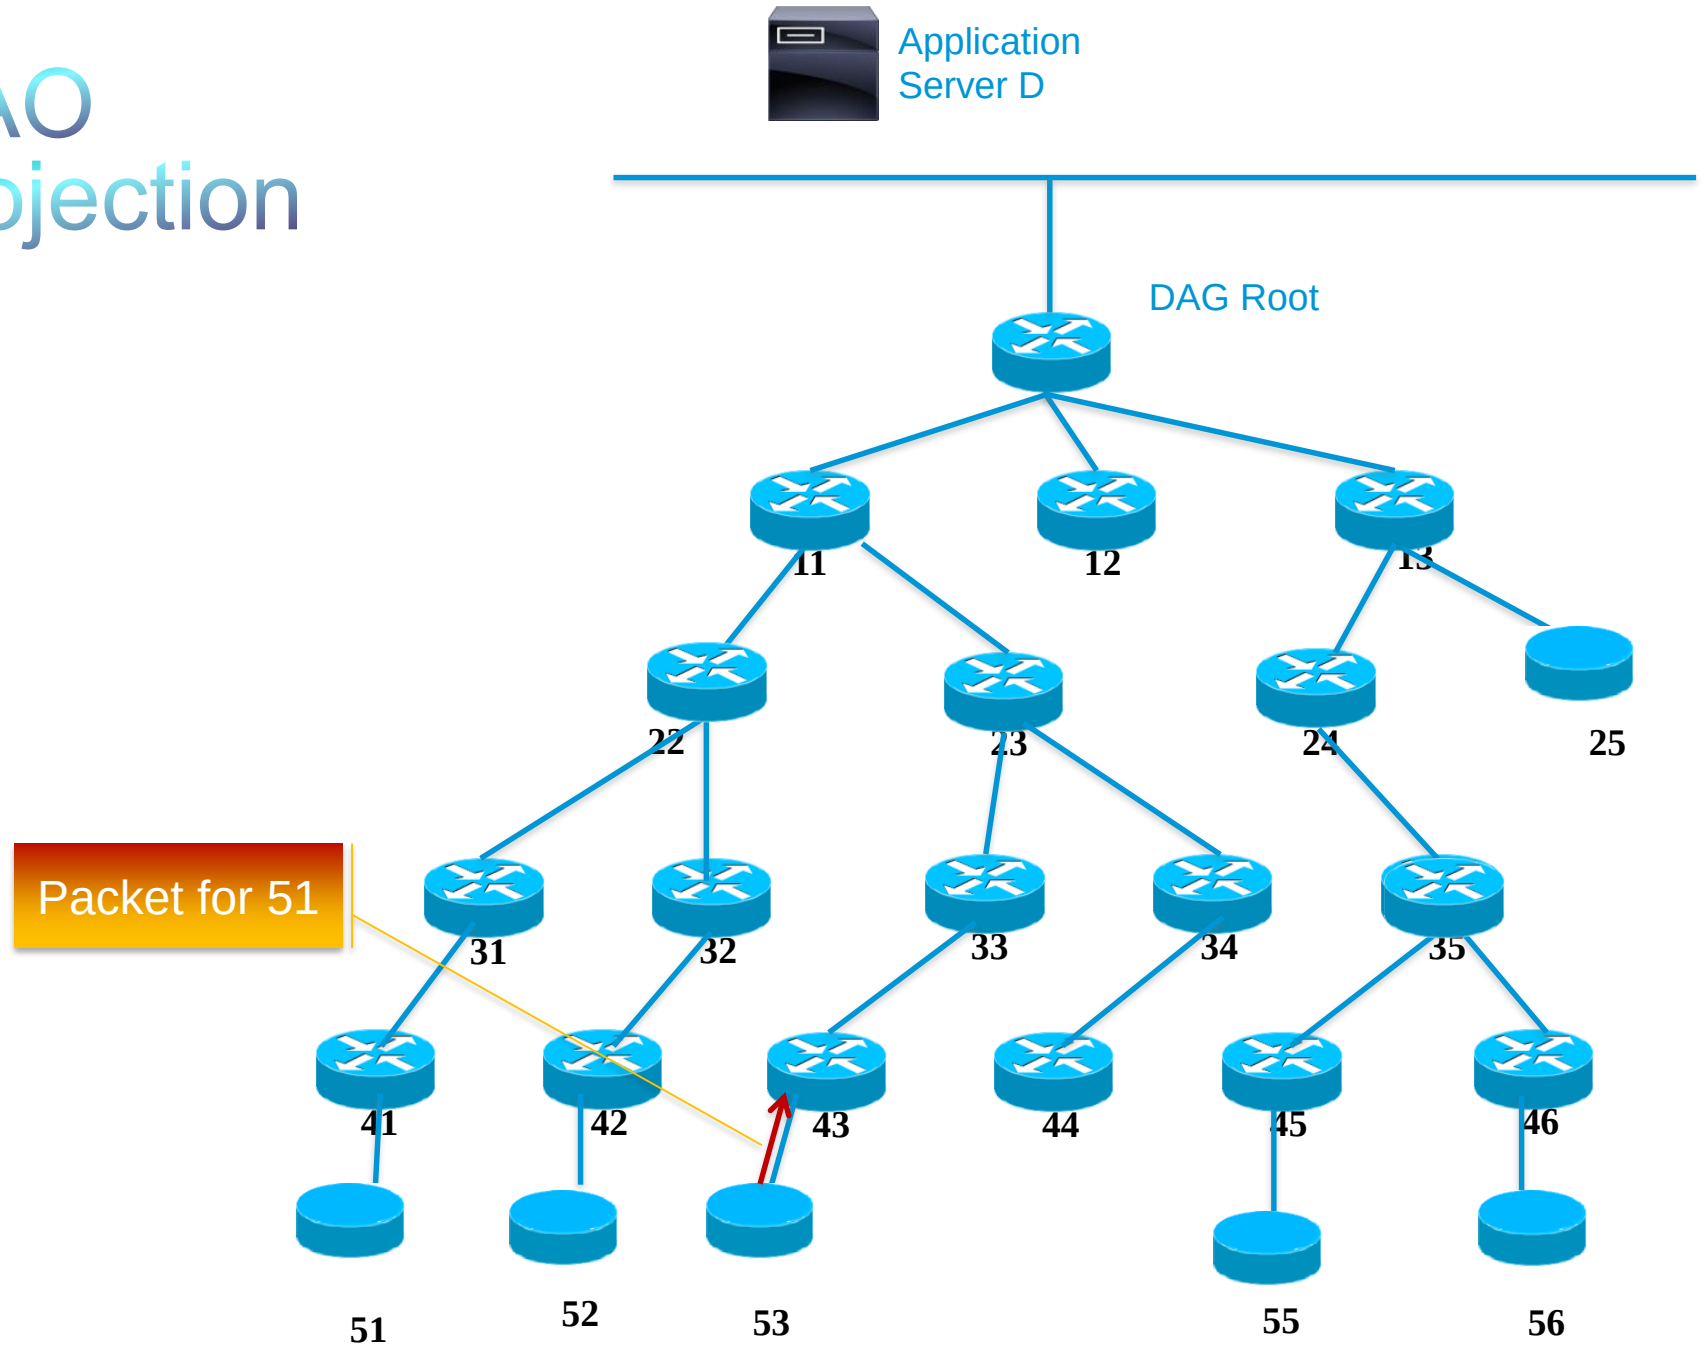

# DAO projection

Application Server D

DAG Root

IP in IP encapsulation with SRH 42, 41

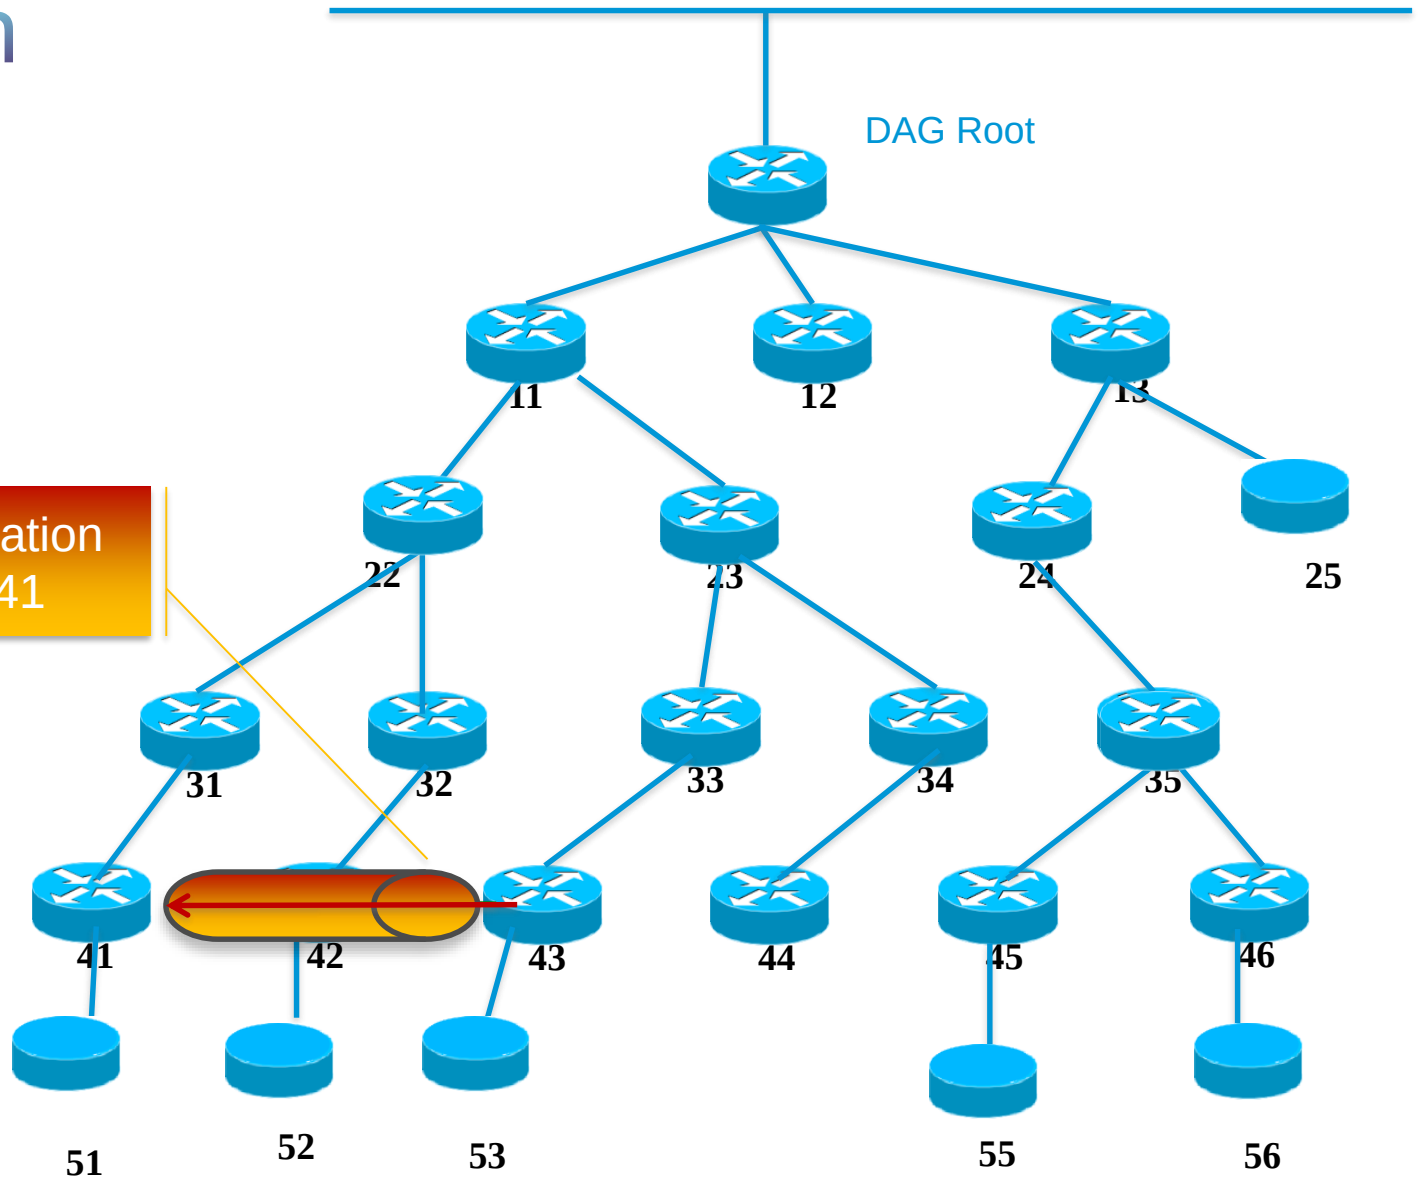

# DAO projection

Application Server D

DAG Root

11 12 13

22 23 24 25

31 32 33 34 35

41 42 43 44 45 46

51 52 53 55 56

Packet for 51

# Storing mode transversal route

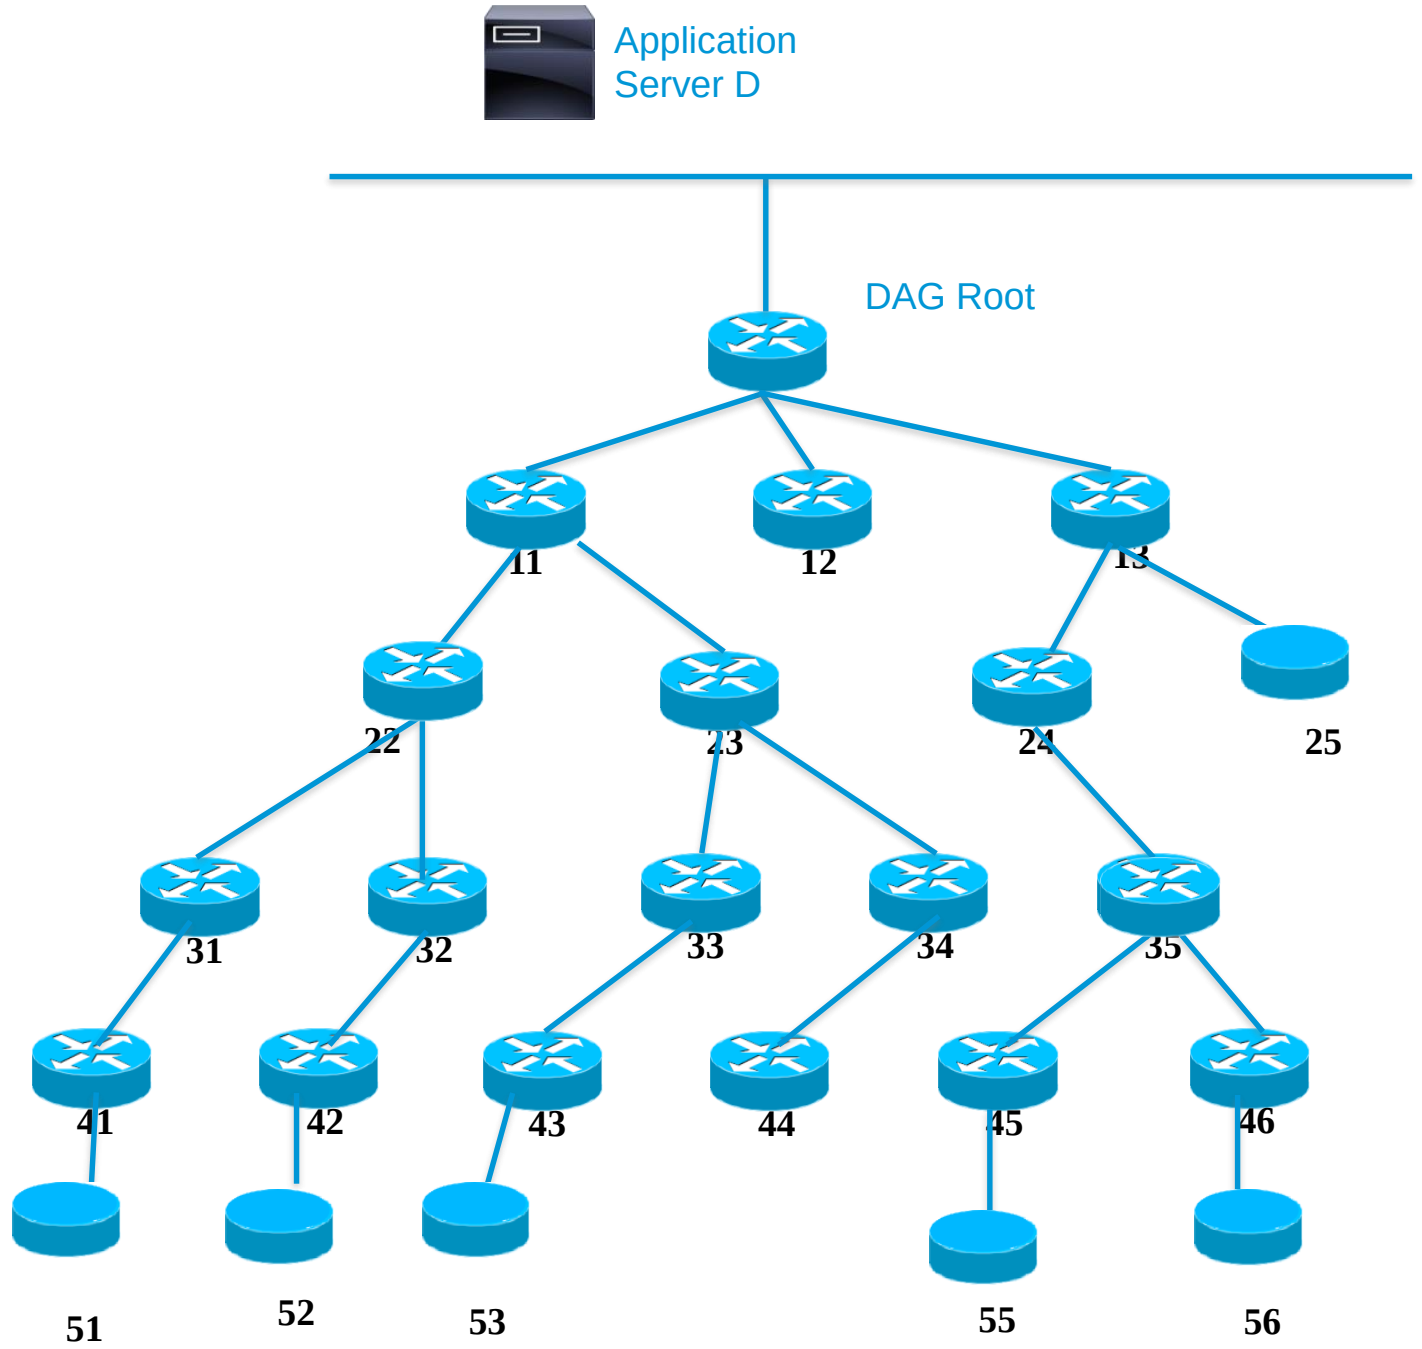
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

# Stretch in non-storing mode

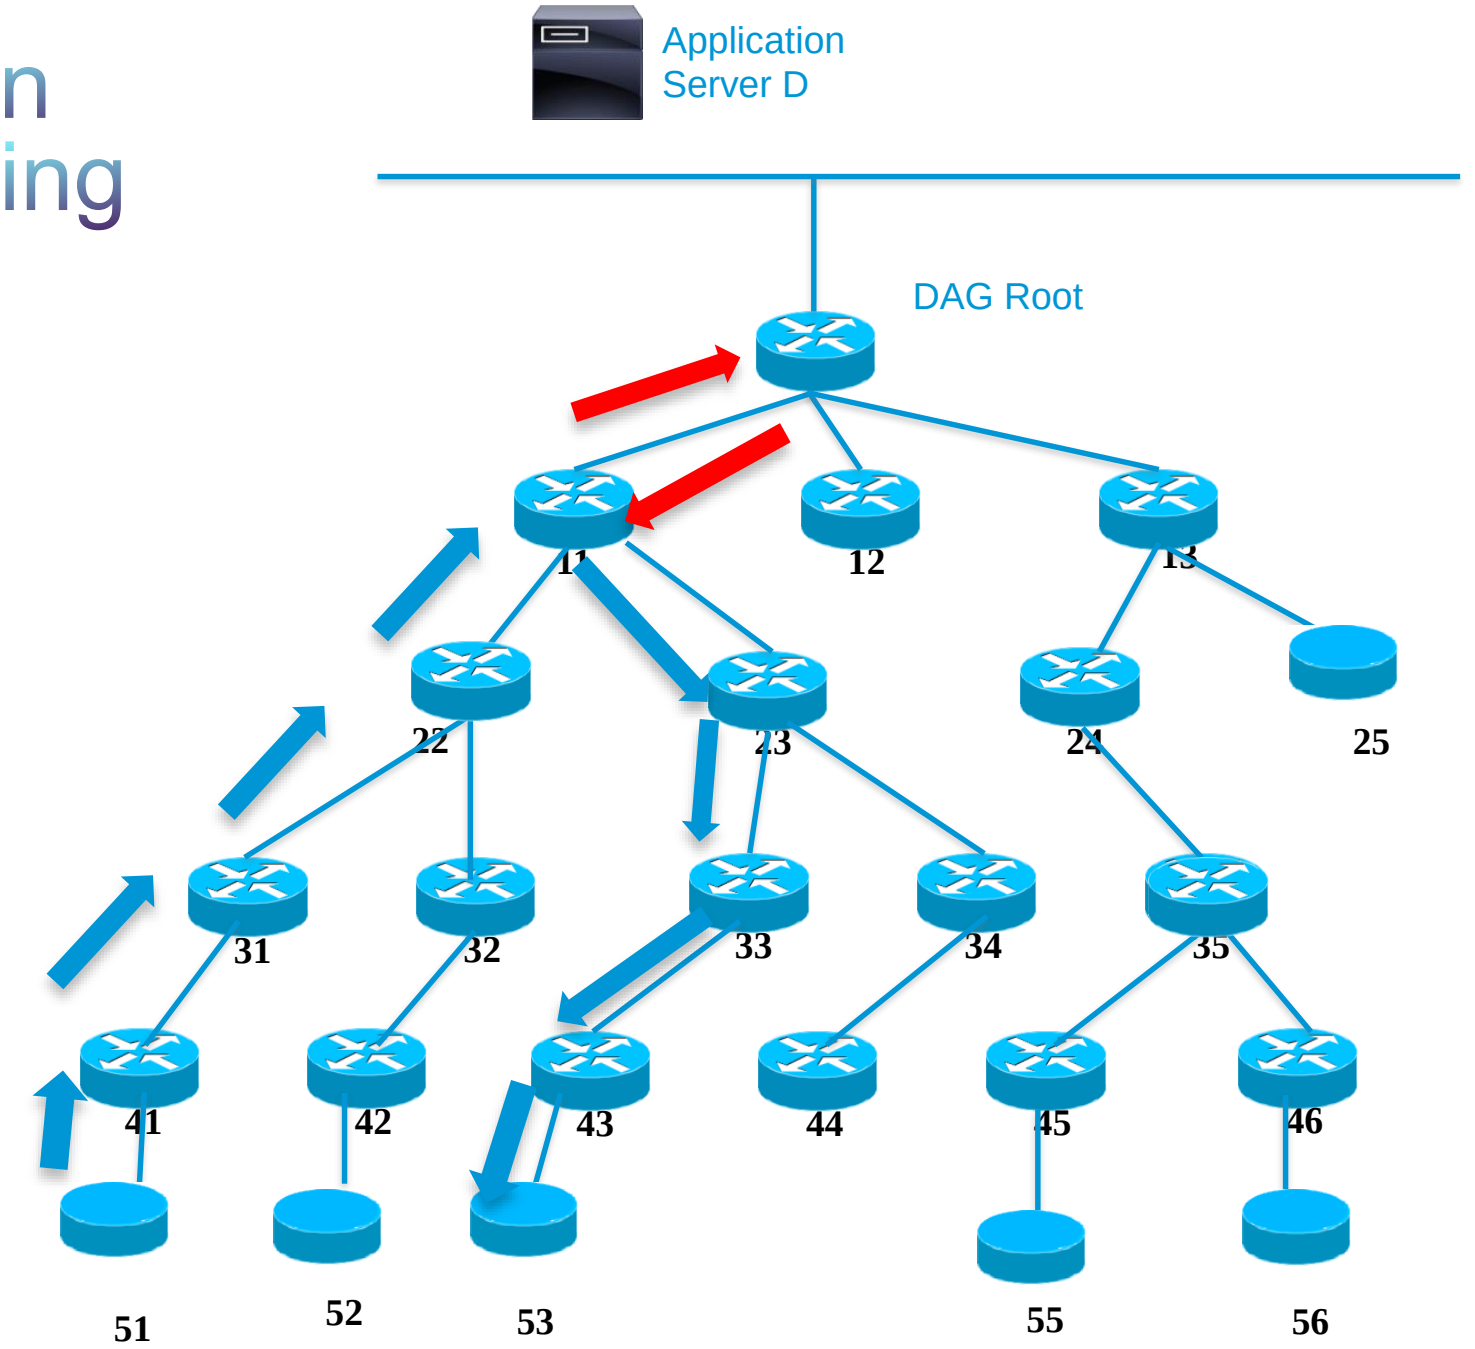

# Stretch in storing mode

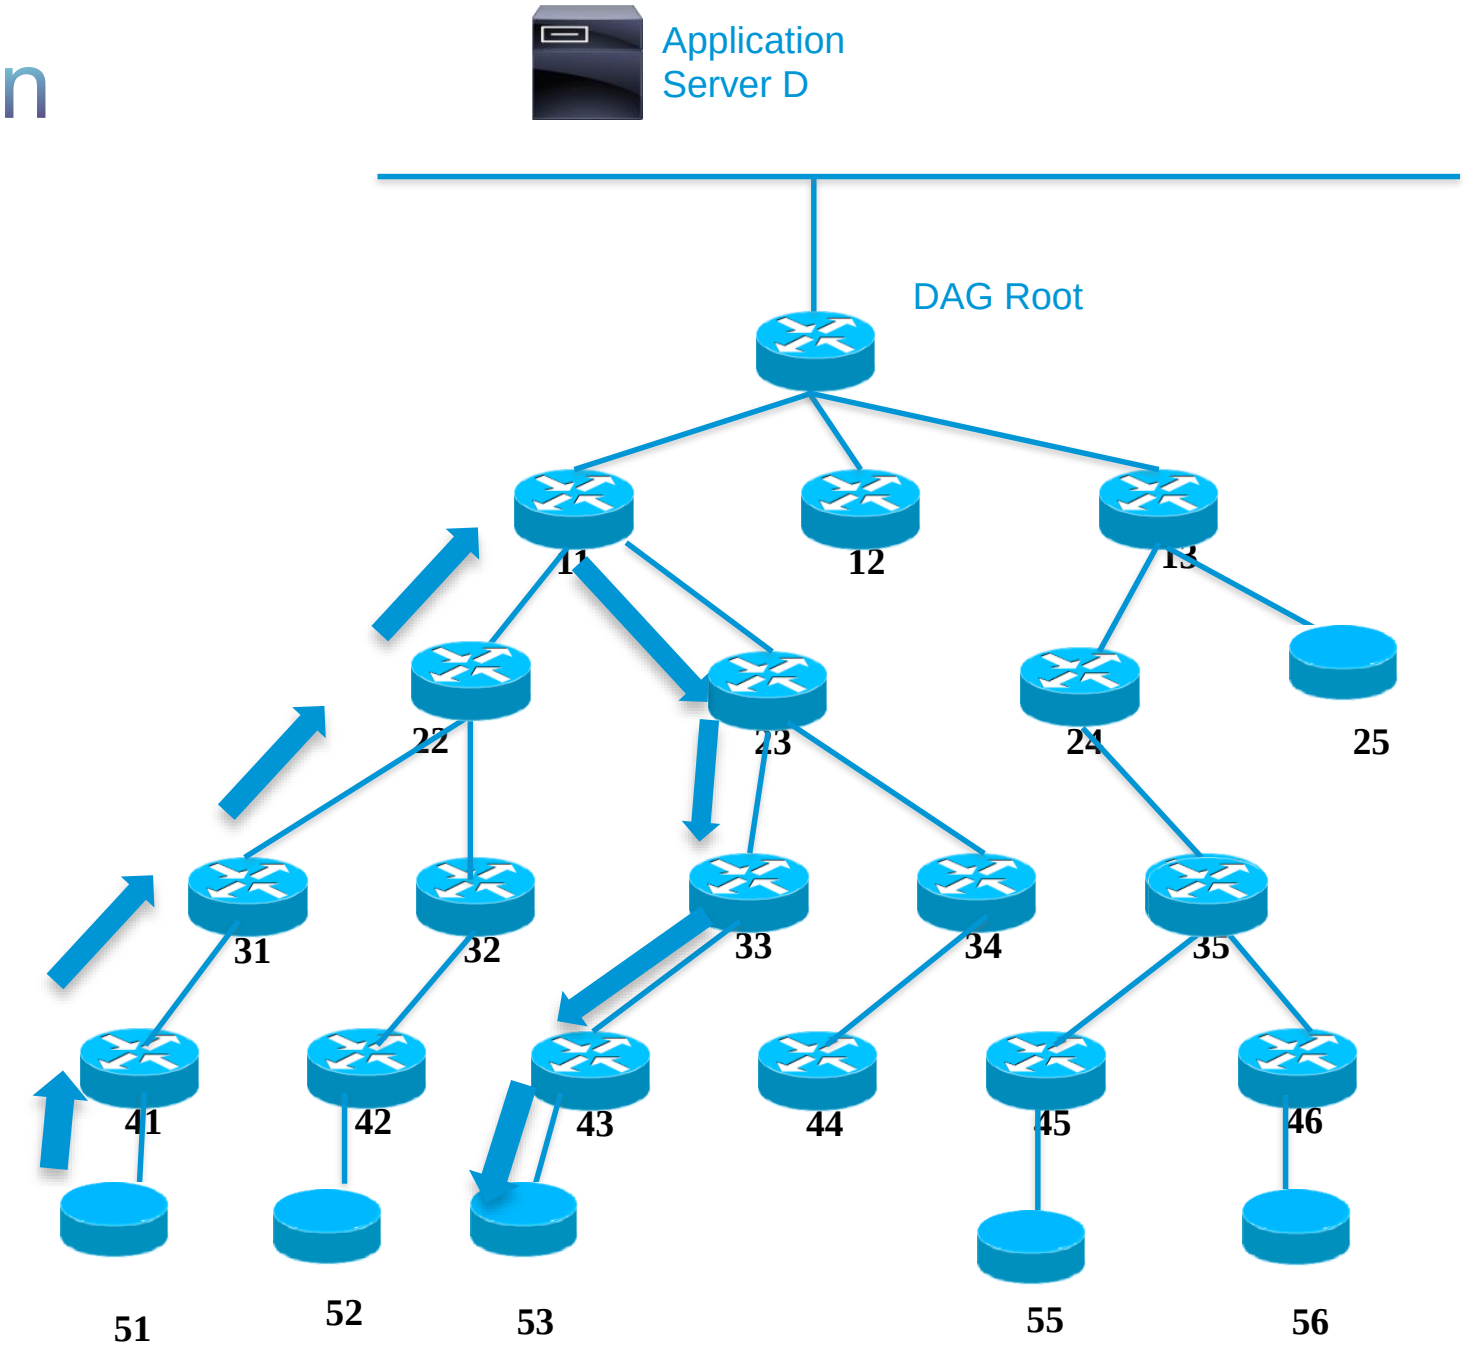

# DAO projection

Application Server D

New (projected) DAO with path segment unicast to target 53 via 41 (ingress), 42 and 43 (egress)

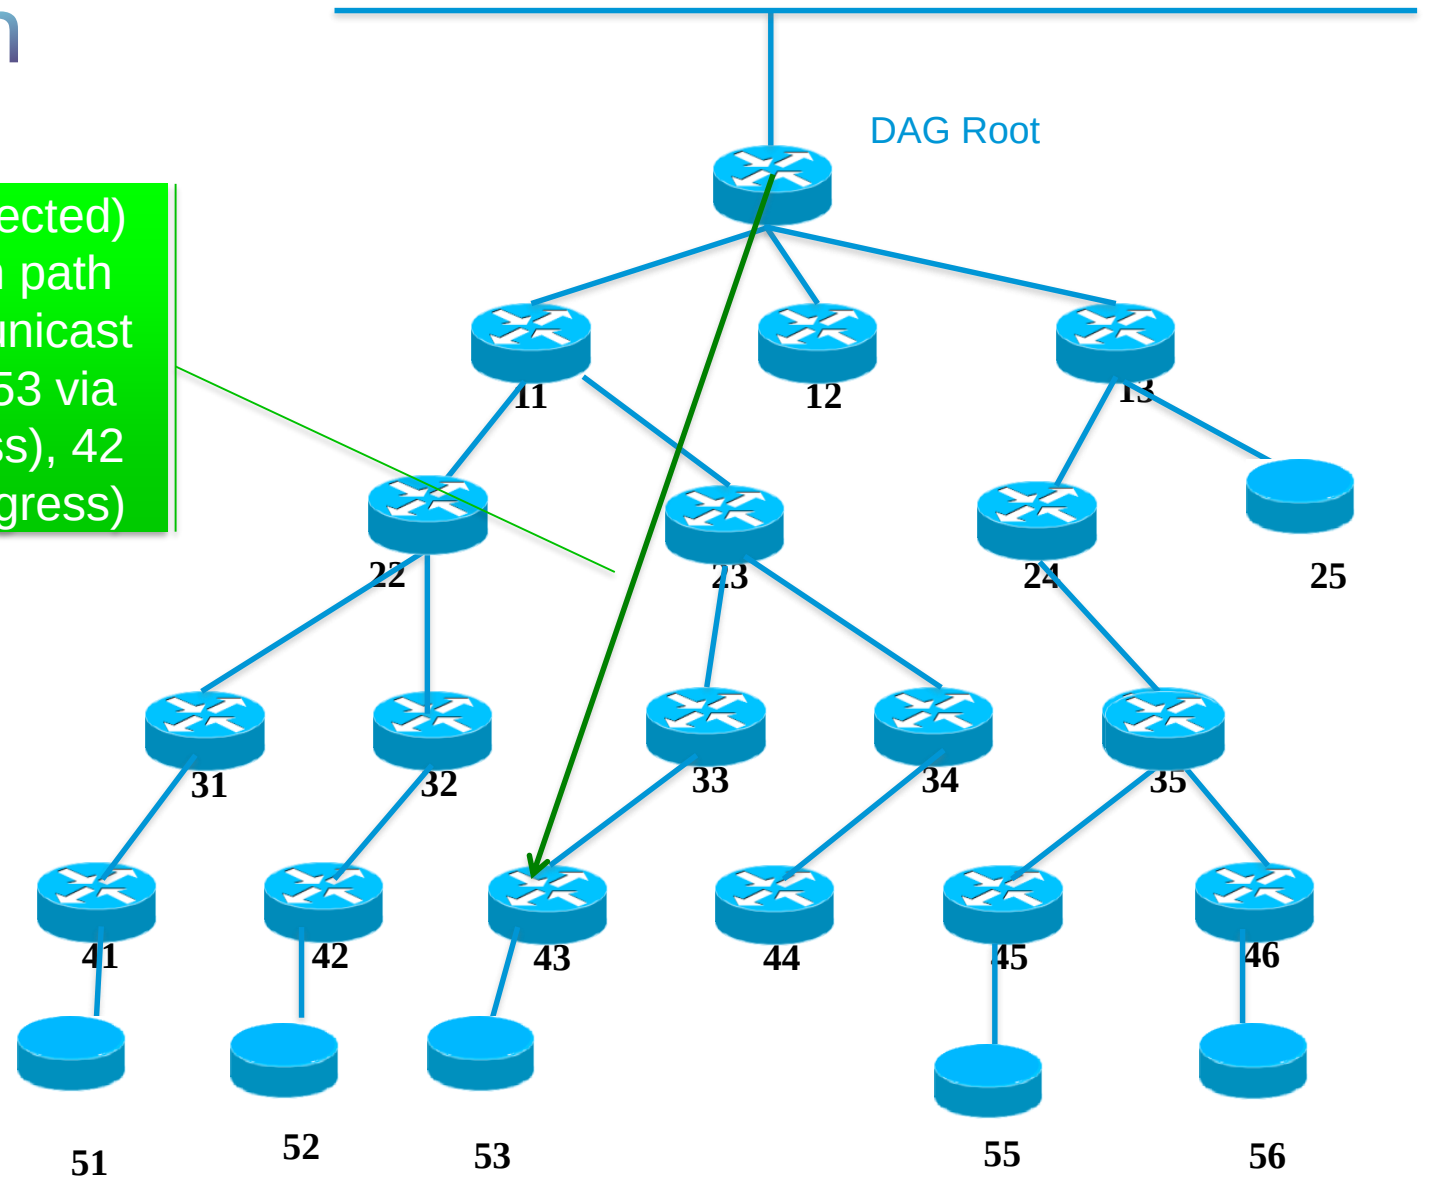

# DAO projection

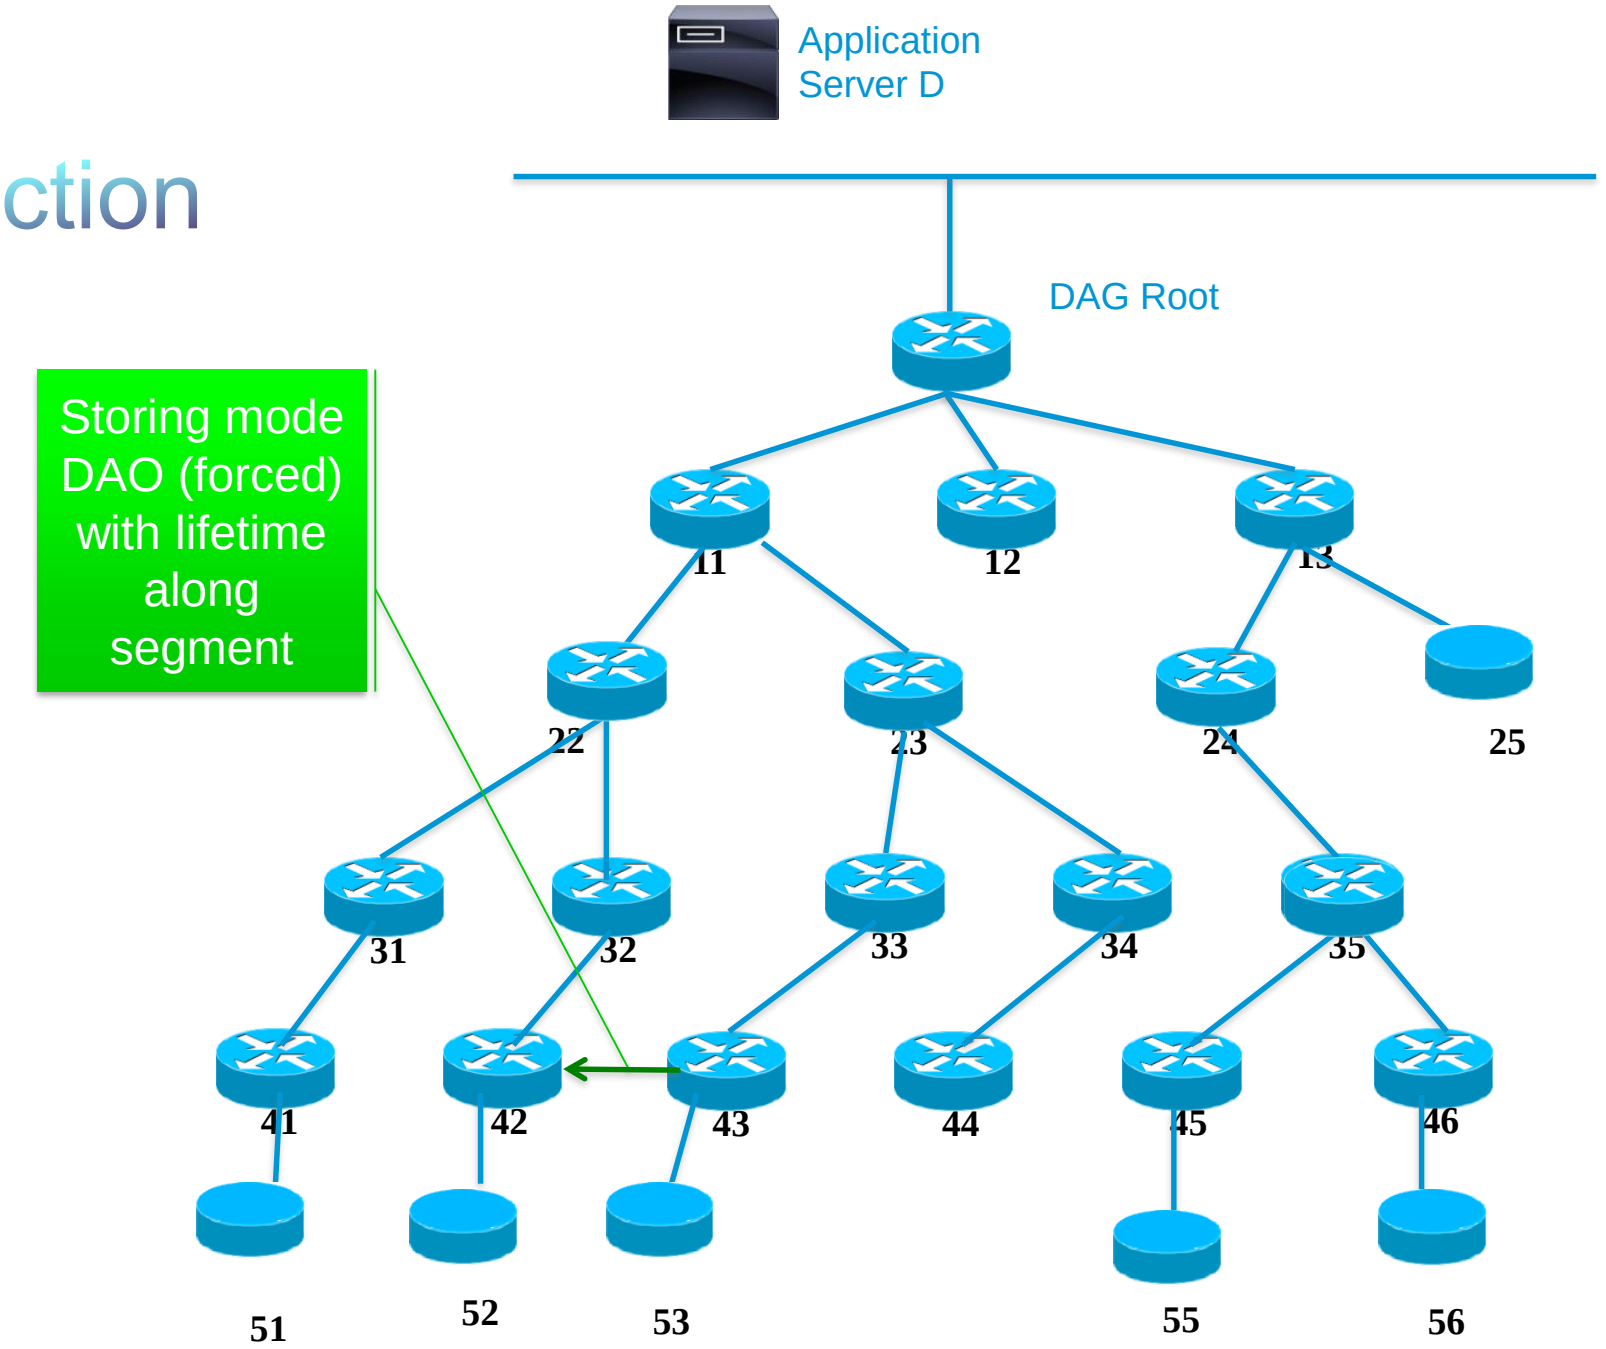

# DAO projection

Application Server D

DAG Root

Storing mode DAO (forced) with lifetime along segment

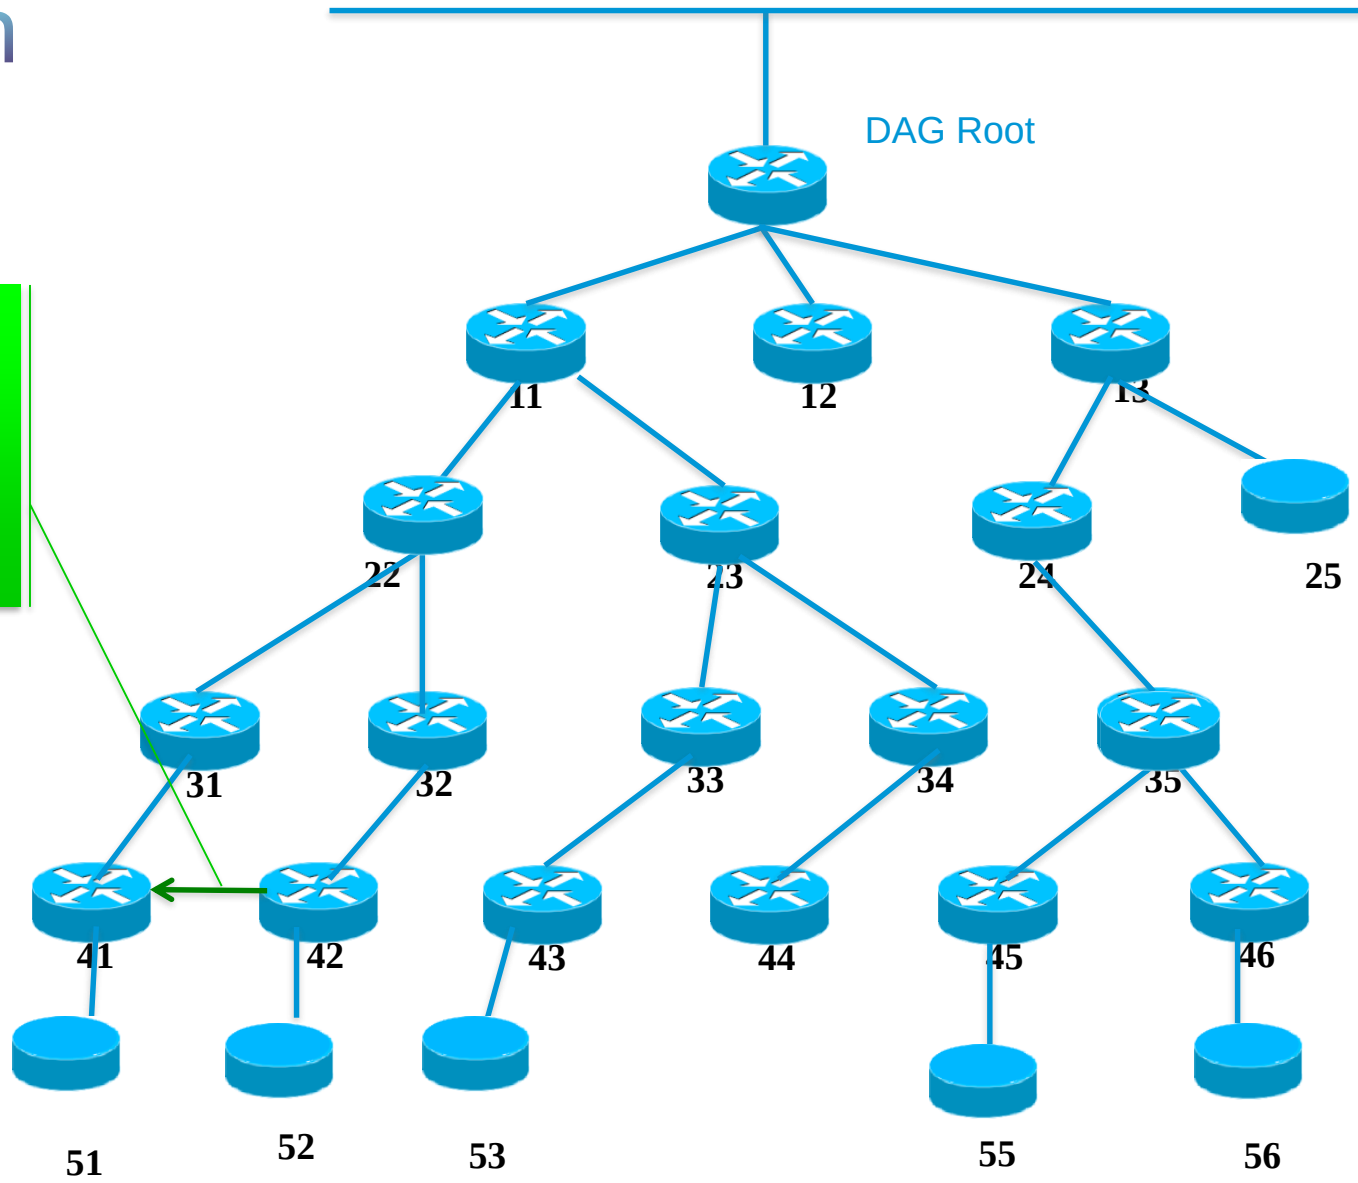

# DAO projection

Application Server D

DAG Root

DAO-ACK unicast, self 41 as parent, final destination 53 as target

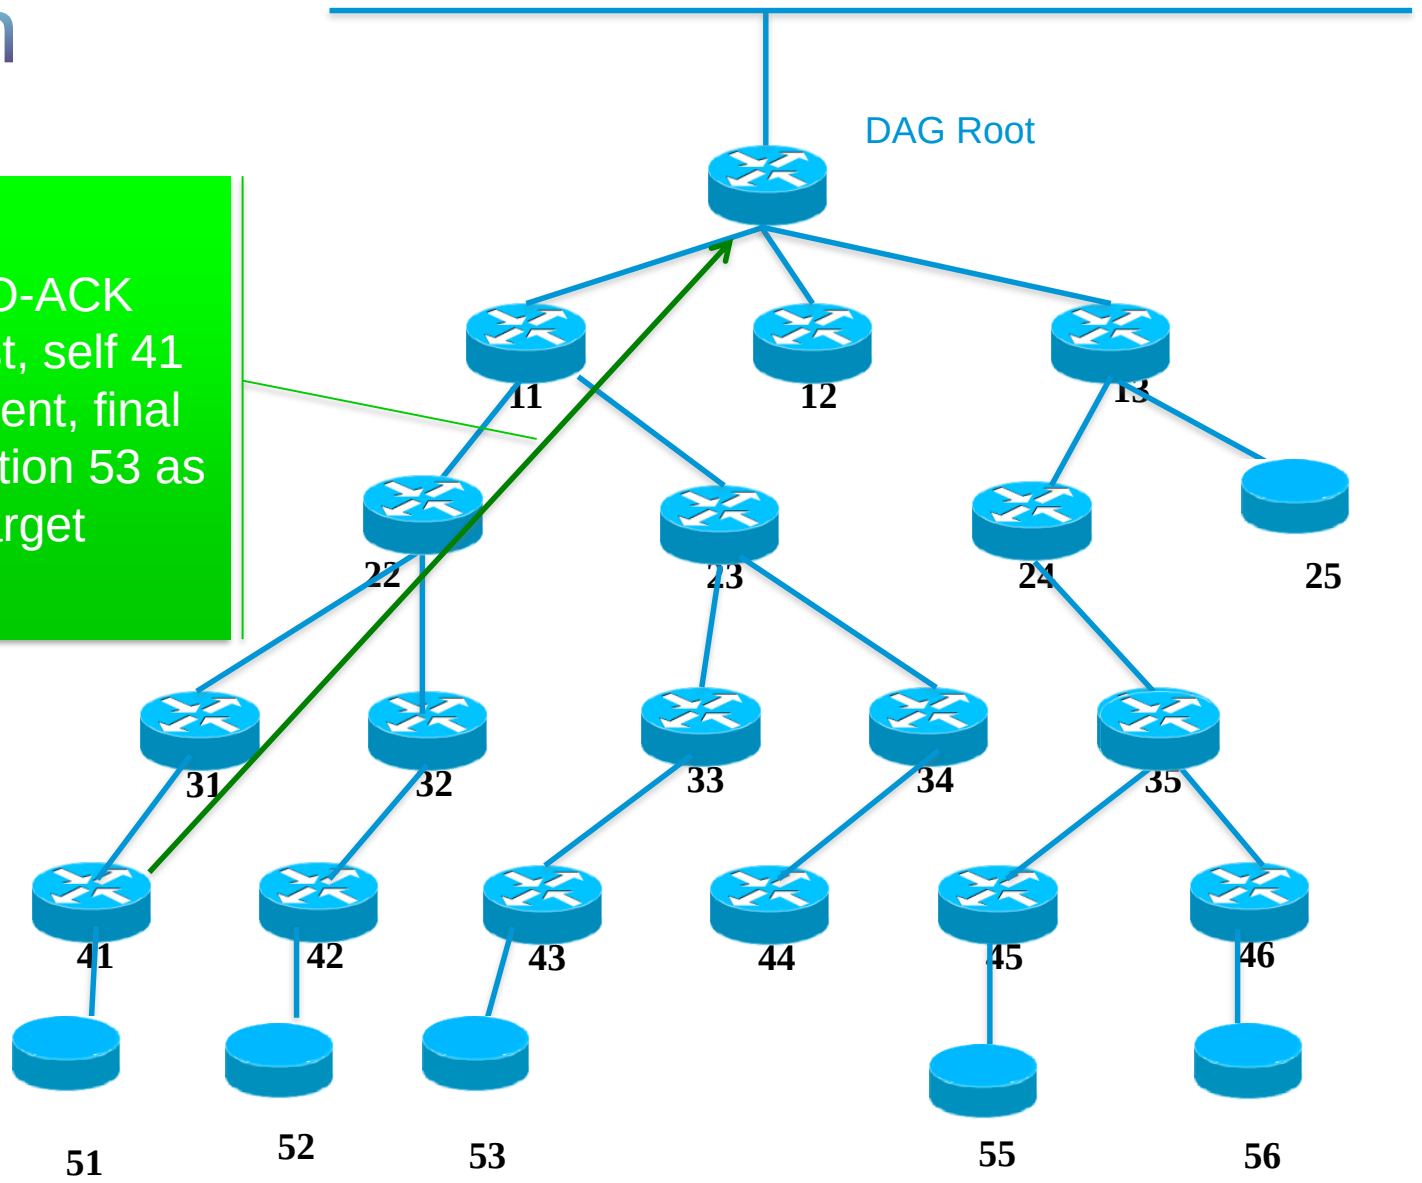

# Optimized Path

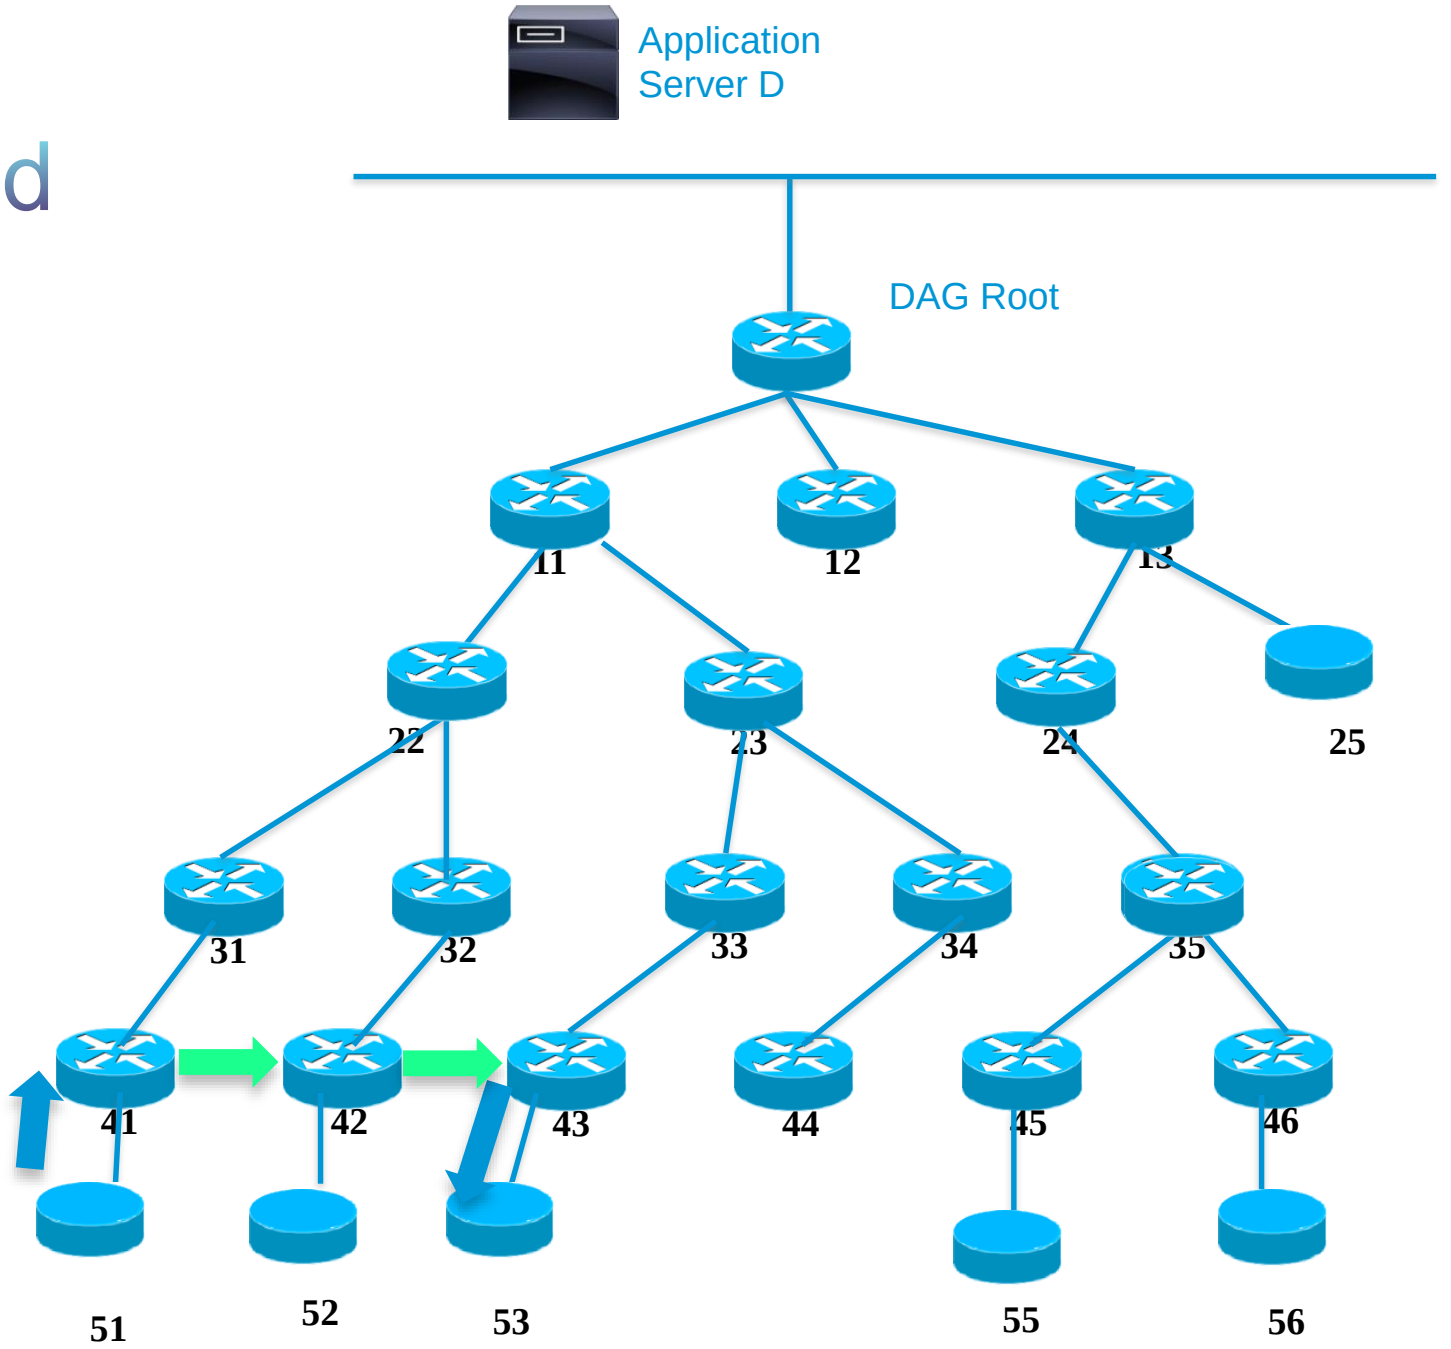

# Existing non storing optimization

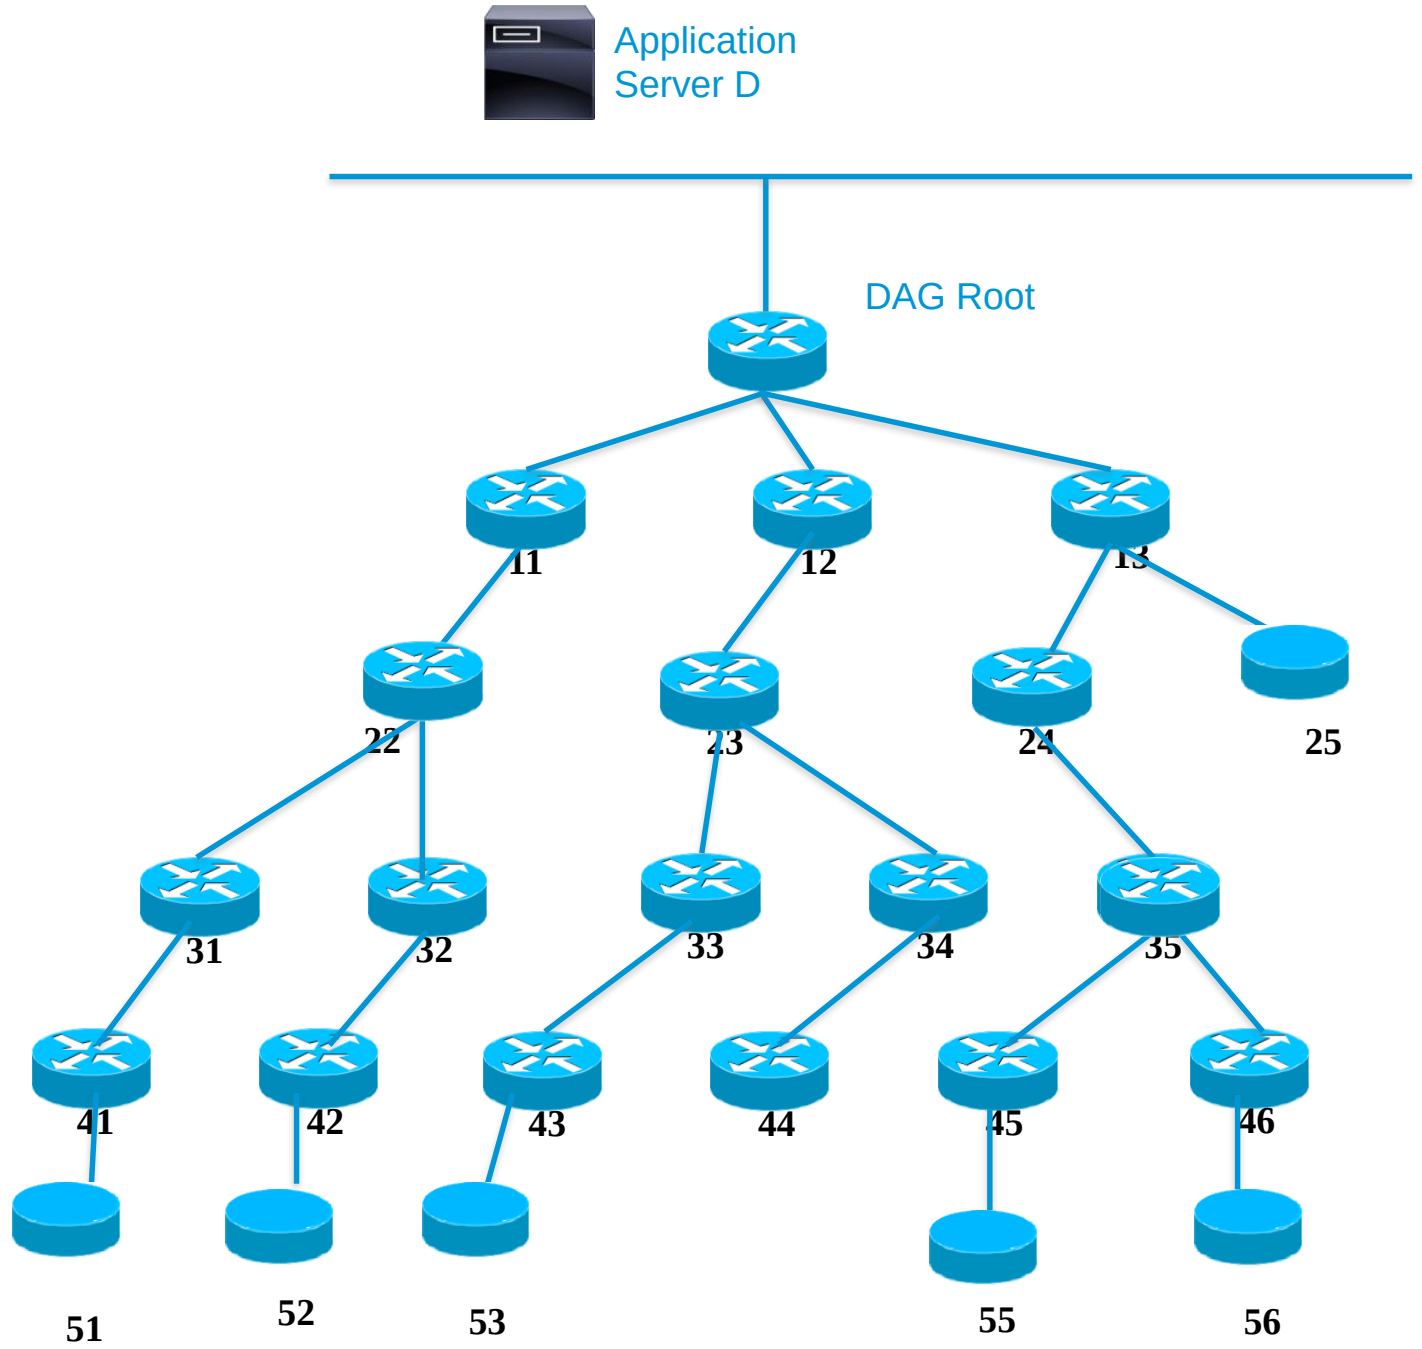
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

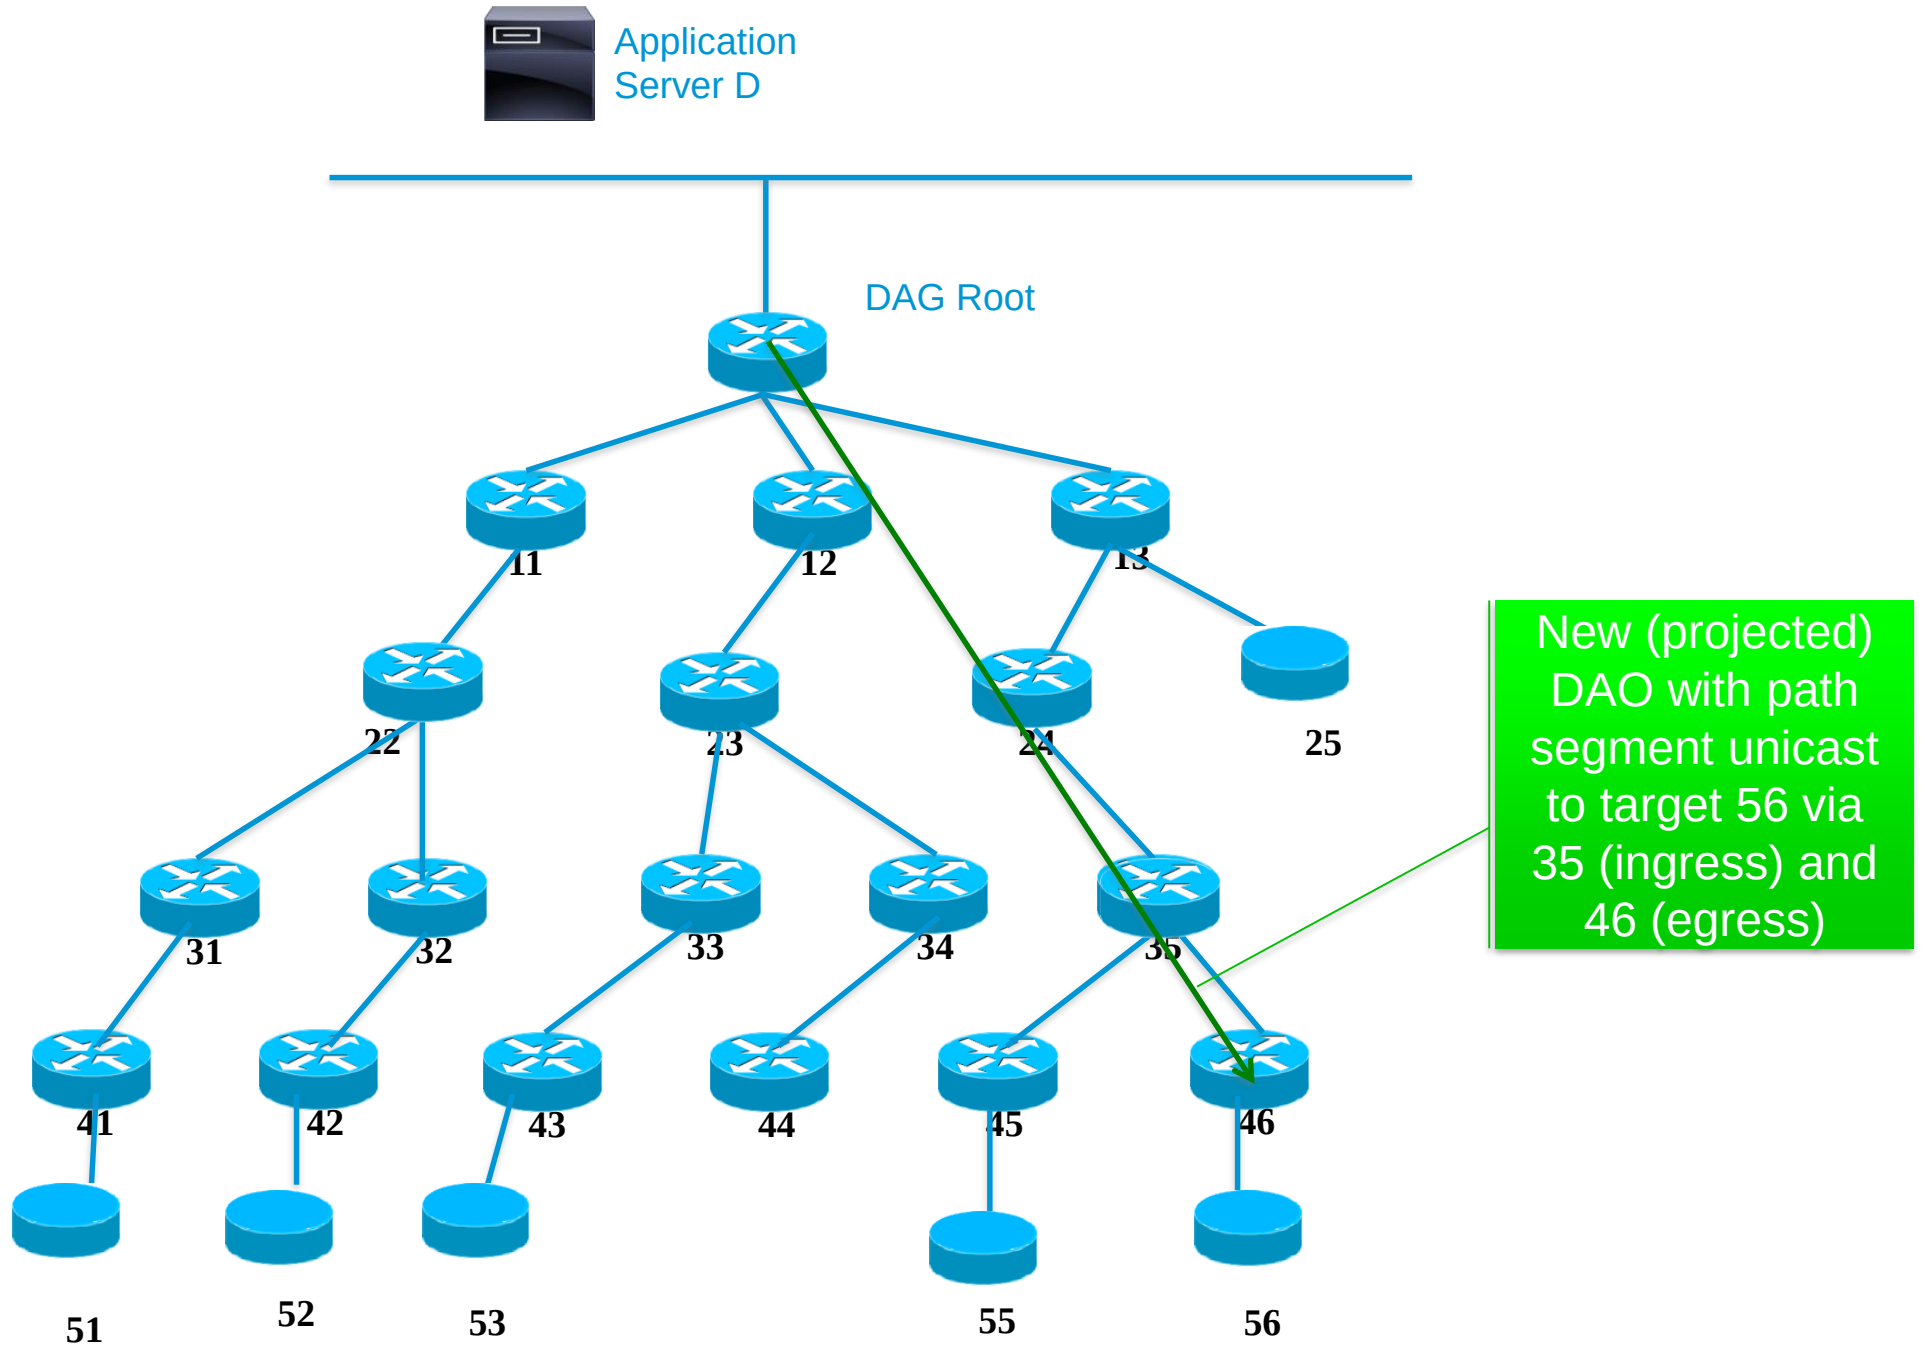
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56
New (projected) DAO with path segment unicast to target 56 via 35 (ingress) and 46 (egress)

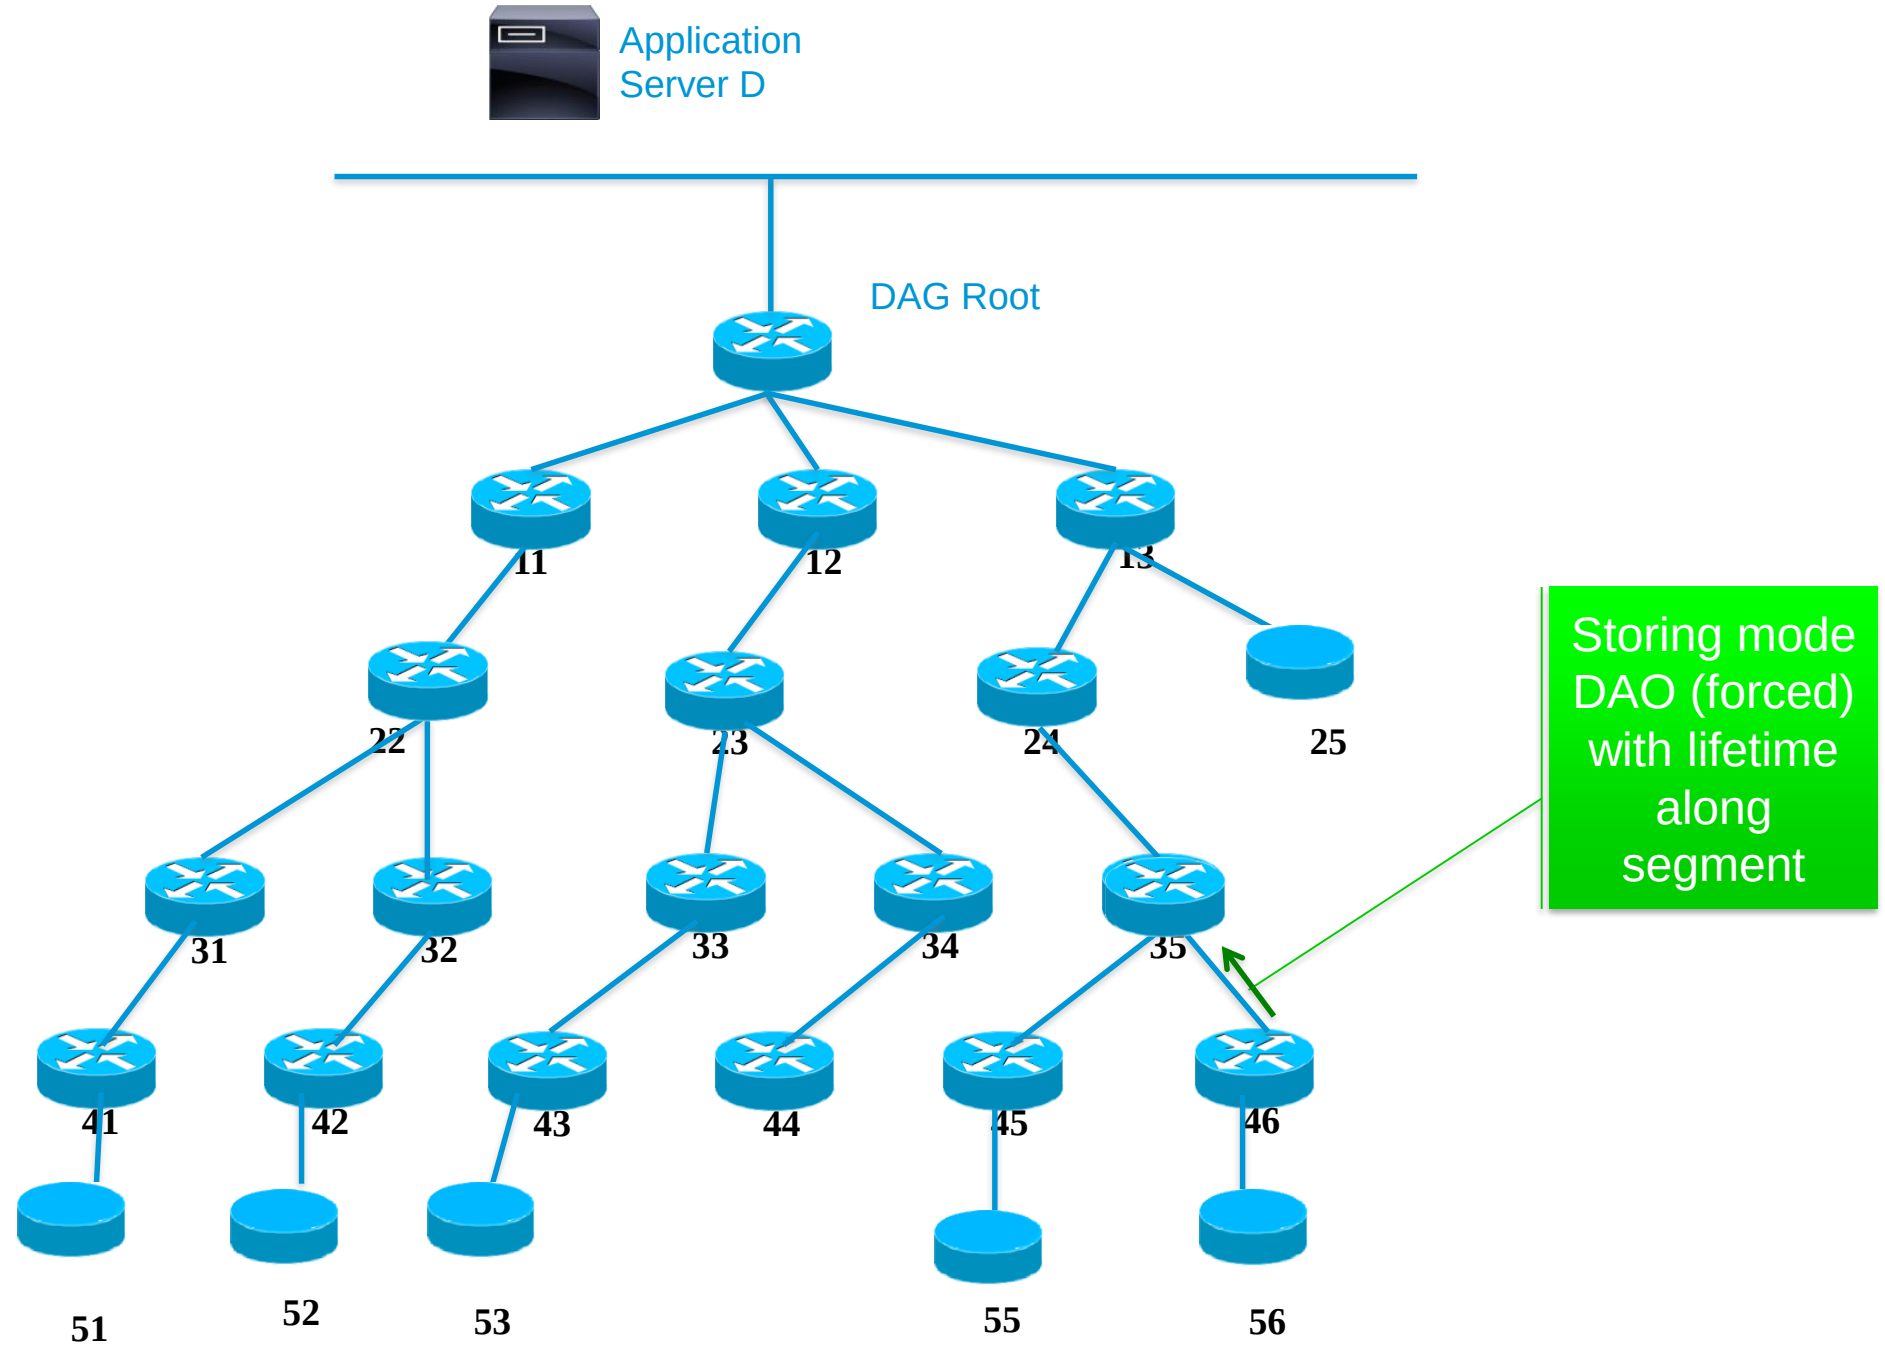
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56
Storing mode
DAO (forced)
with lifetime
along
segment

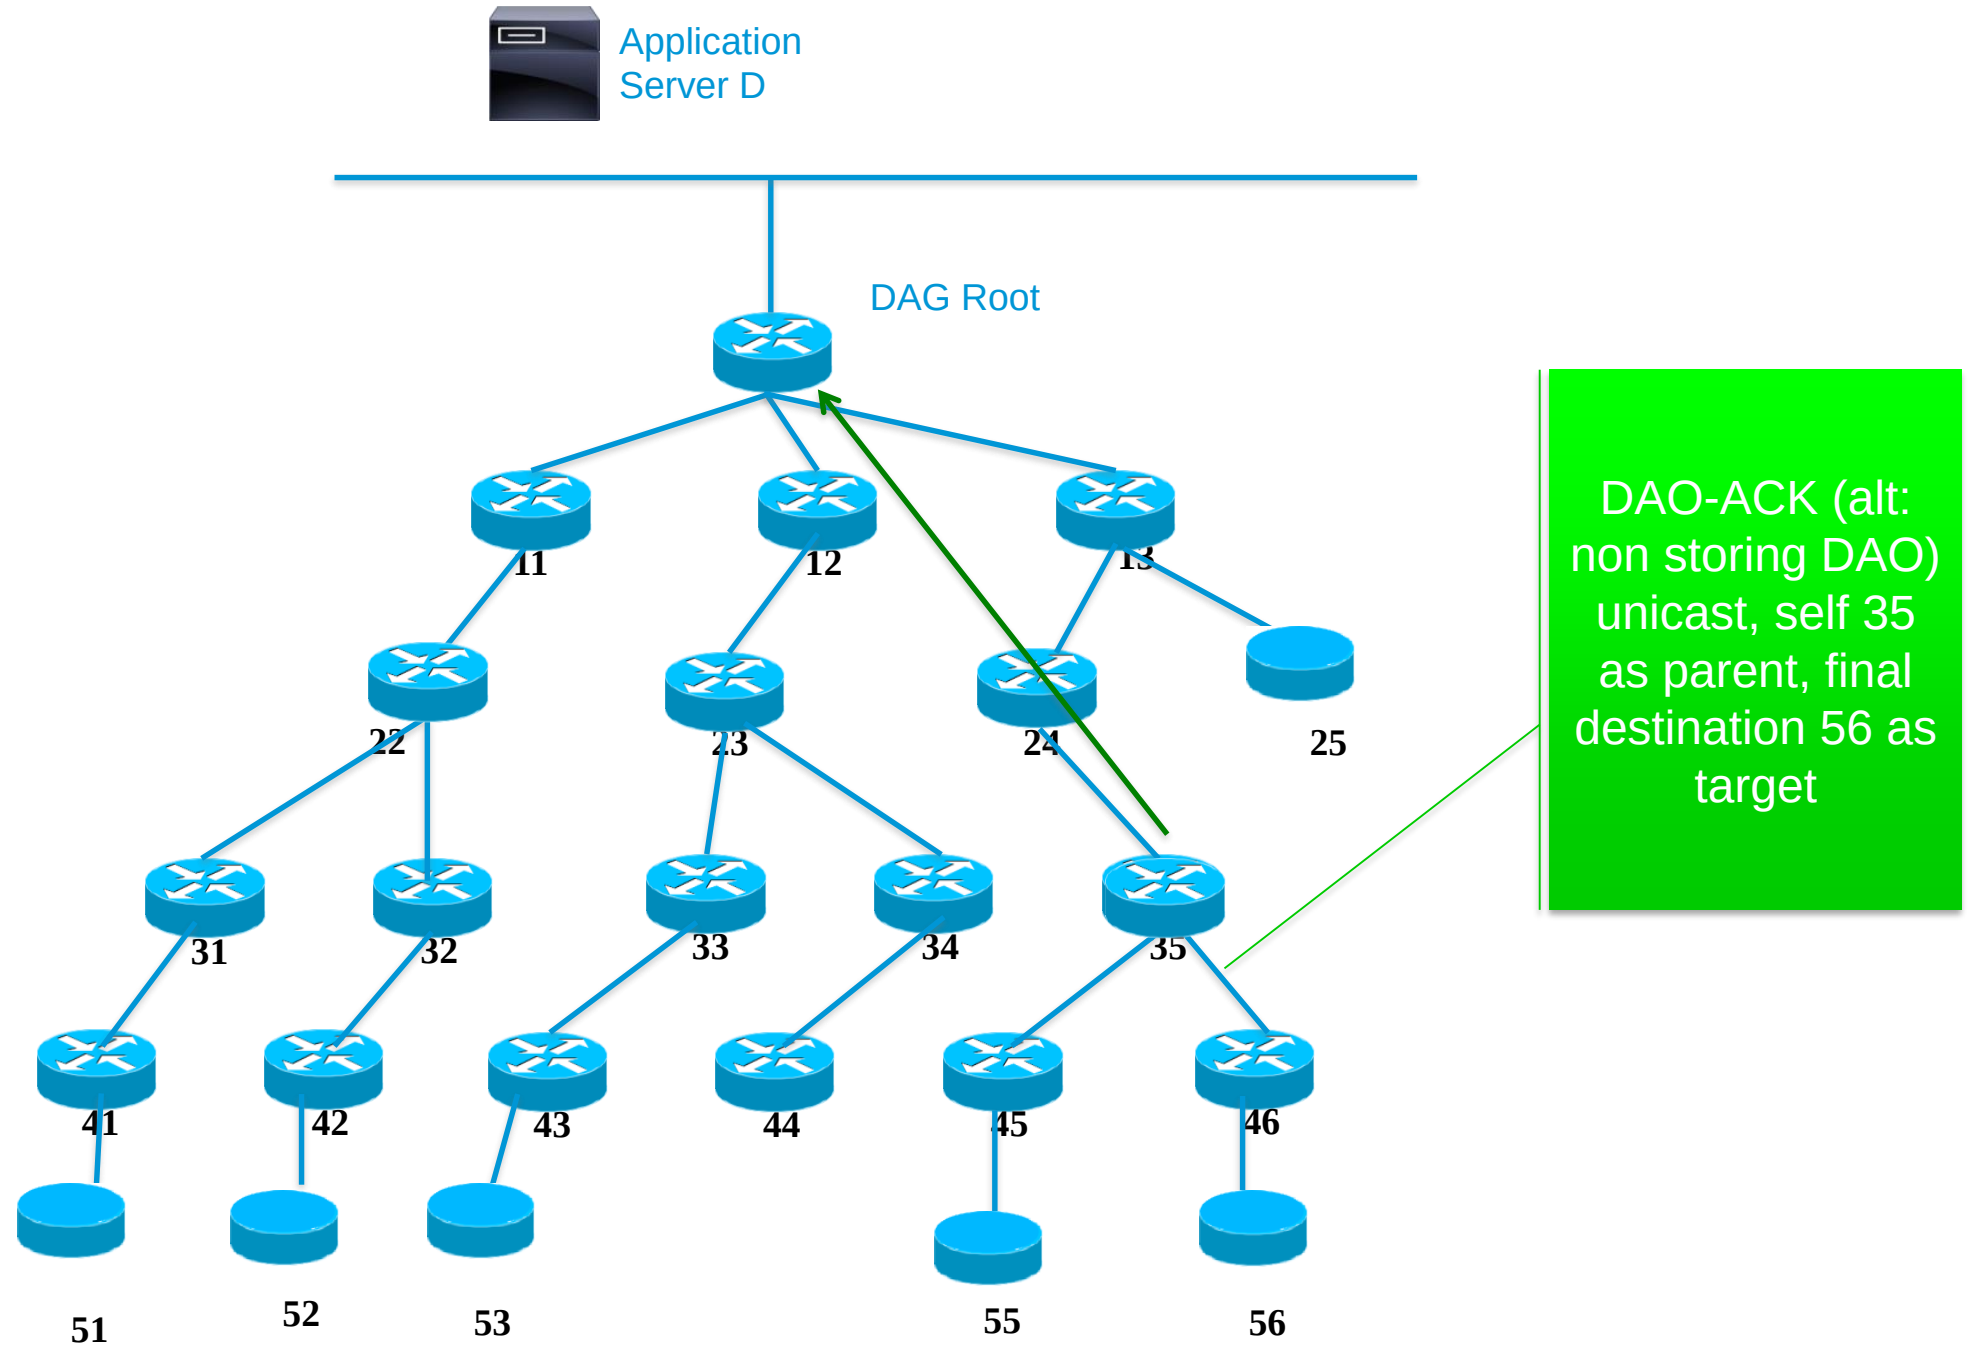
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56
DAO-ACK (alt: non storing DAO) unicast, self 35 as parent, final destination 56 as target

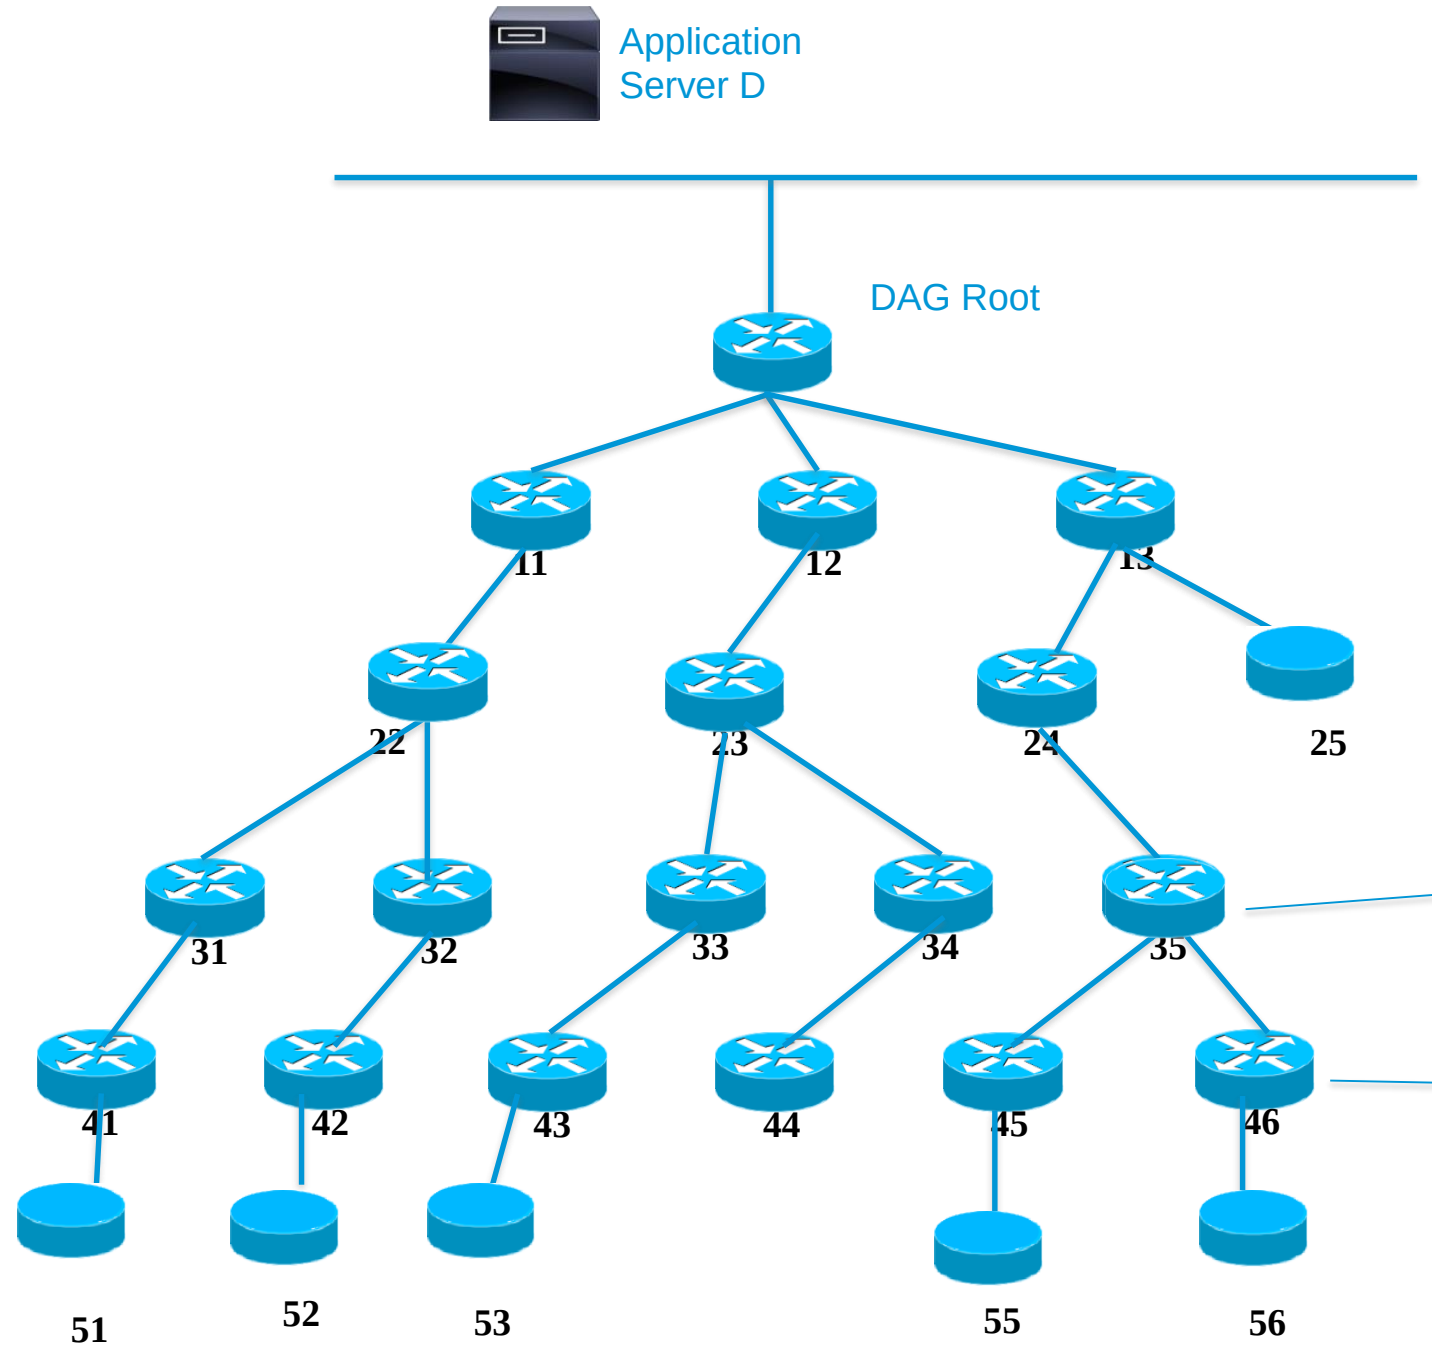


DAO from 46 installs a route to 56 in 35
(all nodes in projected route from ingress included to egress excluded)
=> egress should already have a route to target

56 via 46

Preexisting connected route to 56

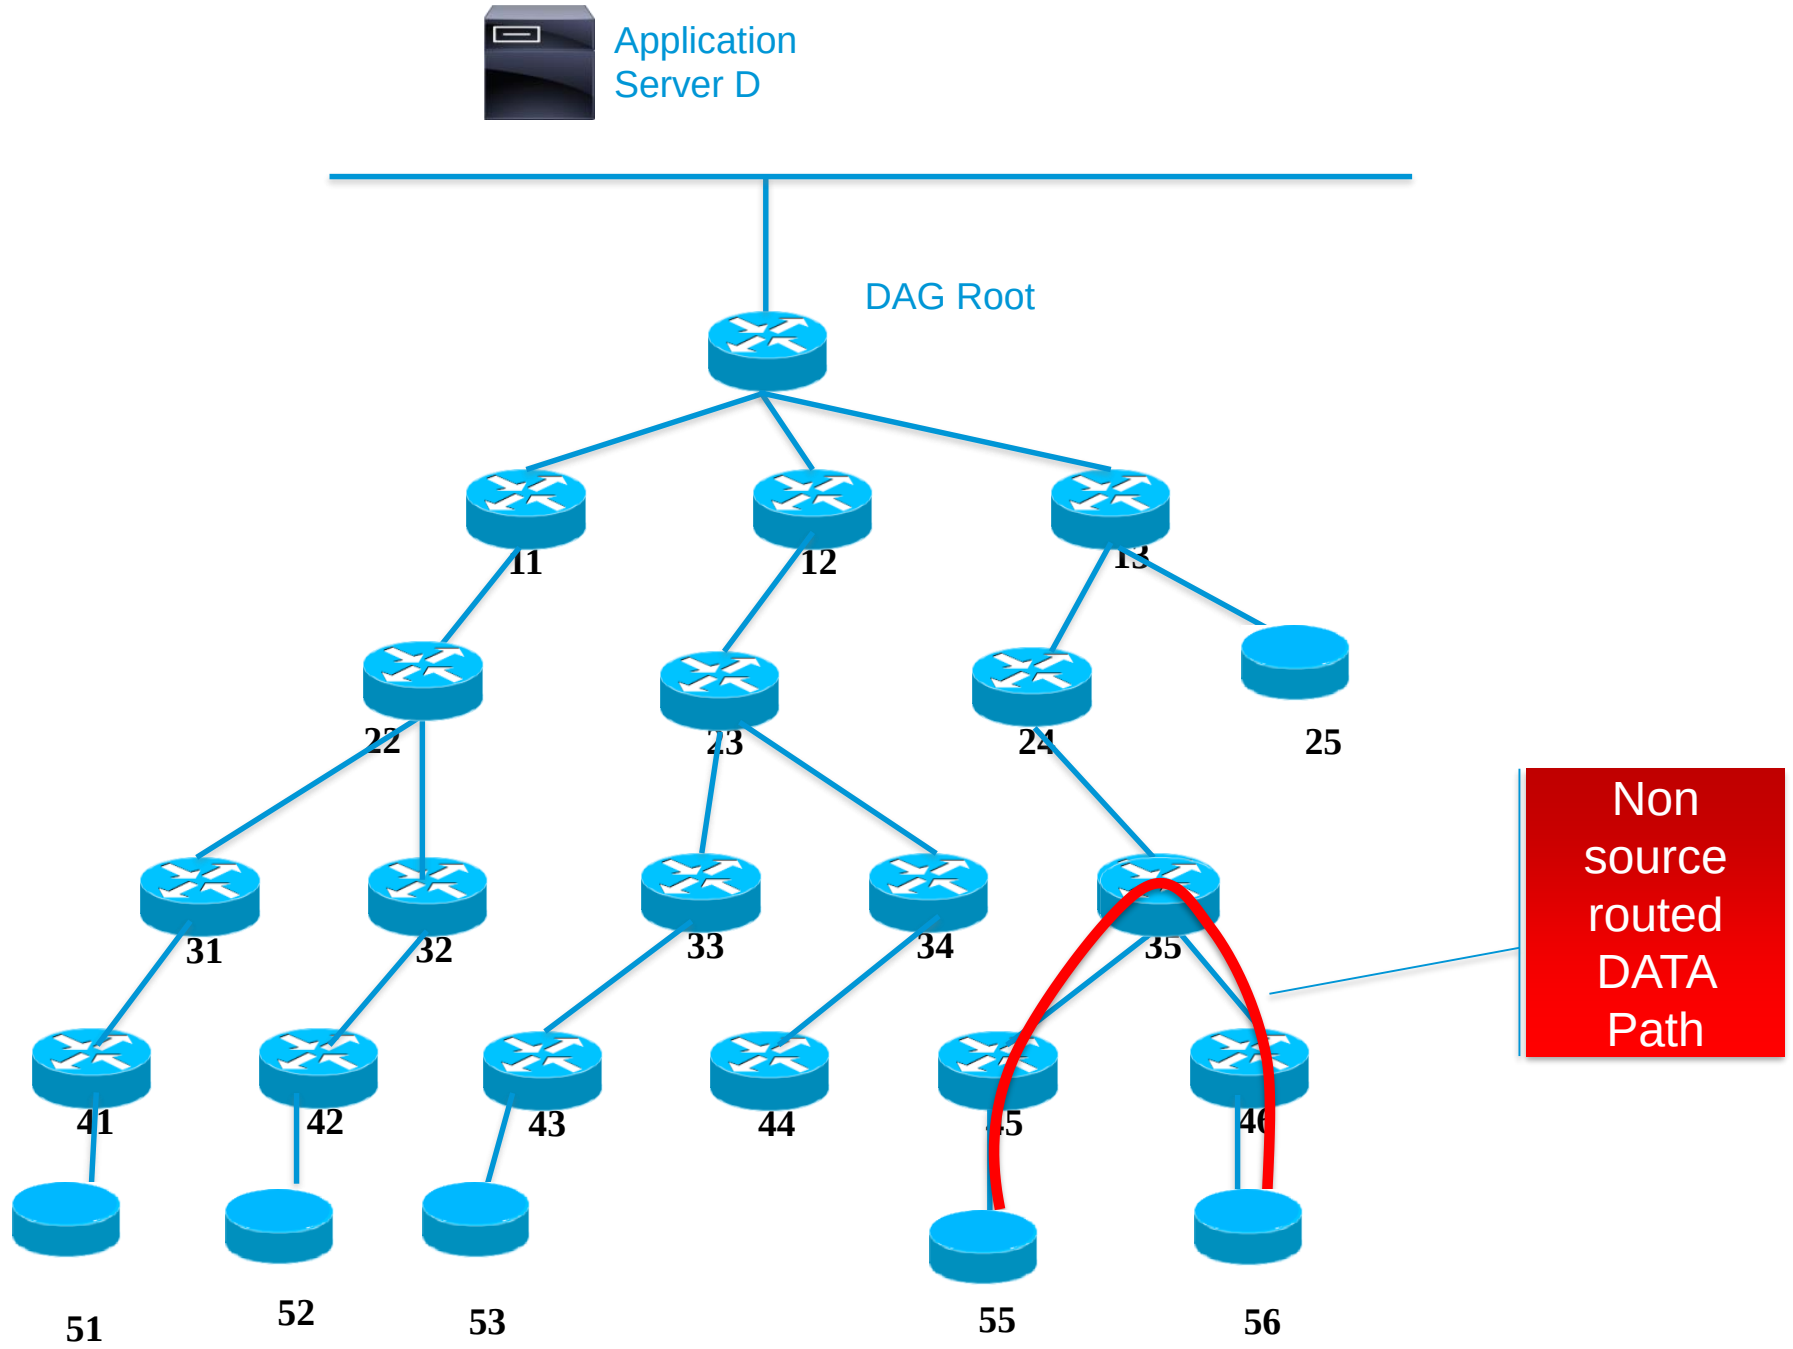
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
46
51
52
53
55
56
Non source routed DATA Path

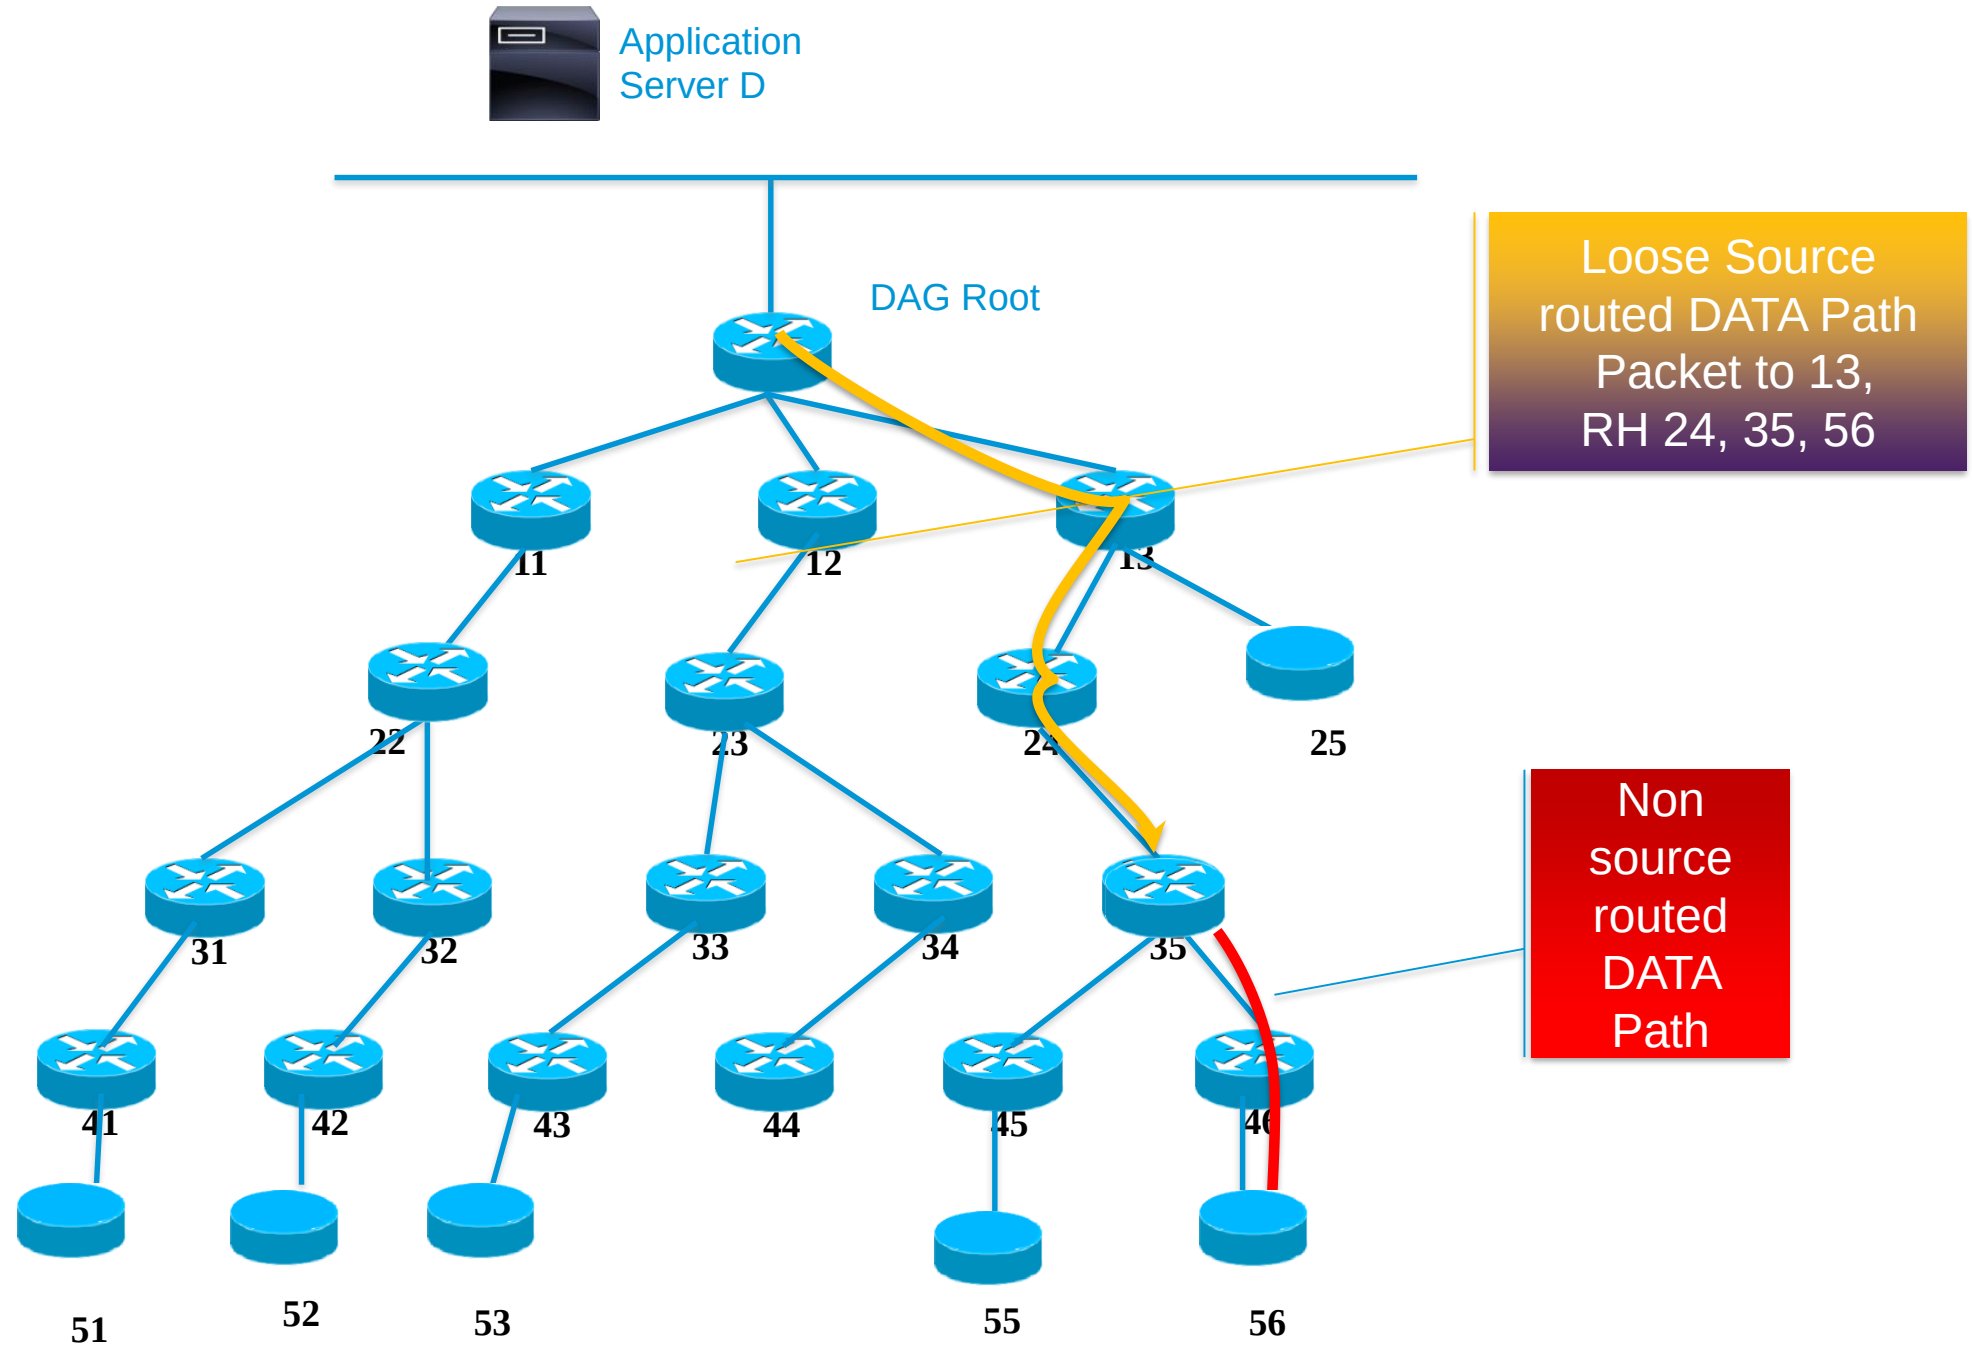
Application
Server D
DAG Root
Loose Source routed DATA Path Packet to 13, RH 24, 35, 56
Non source routed DATA Path
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

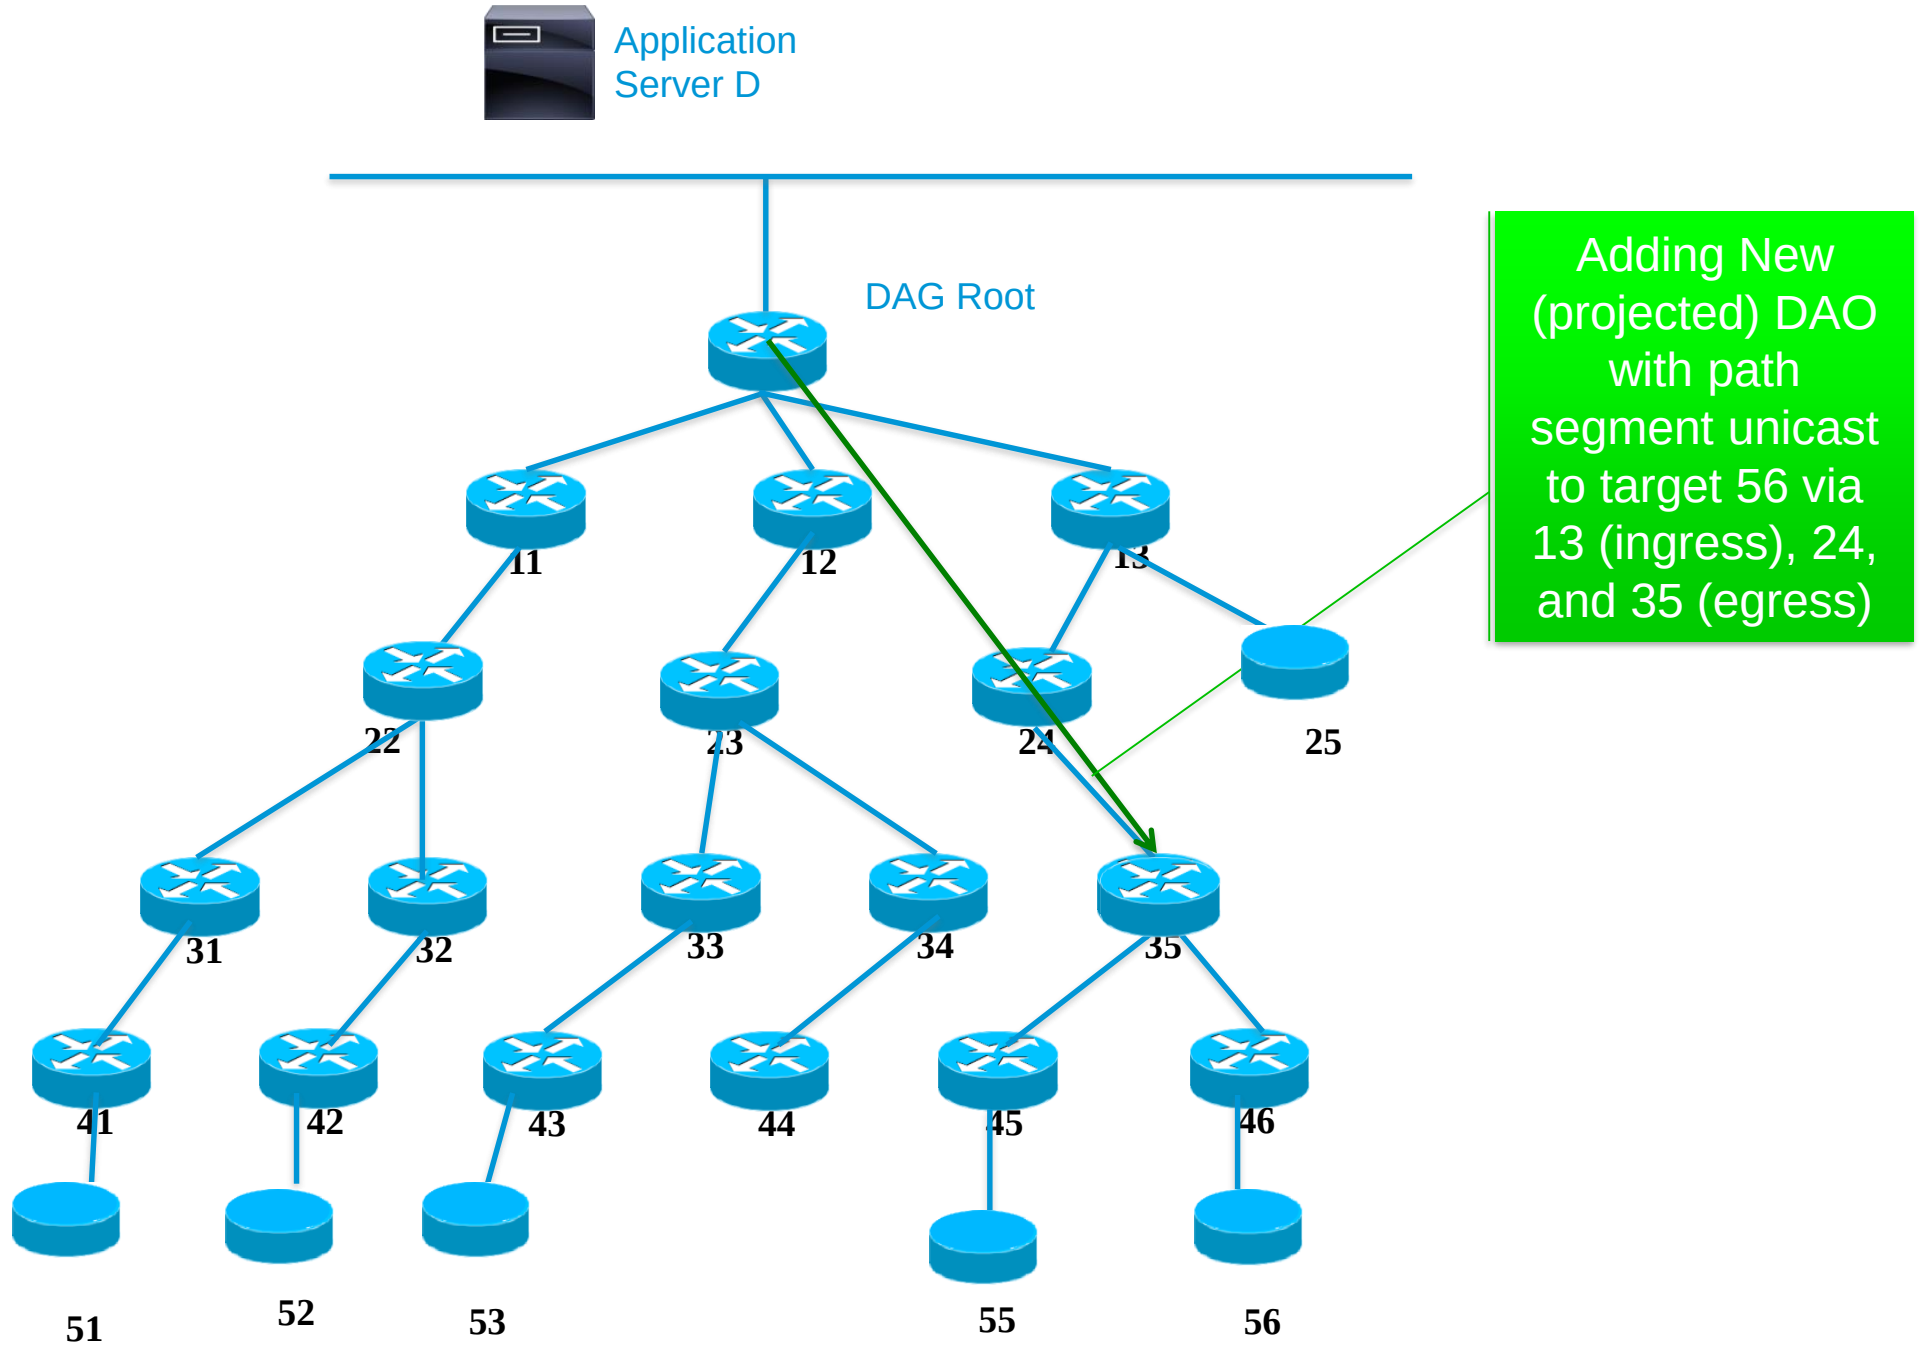
Application
Server D
DAG Root
Adding New (projected) DAO with path segment unicast to target 56 via 13 (ingress), 24, and 35 (egress)
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

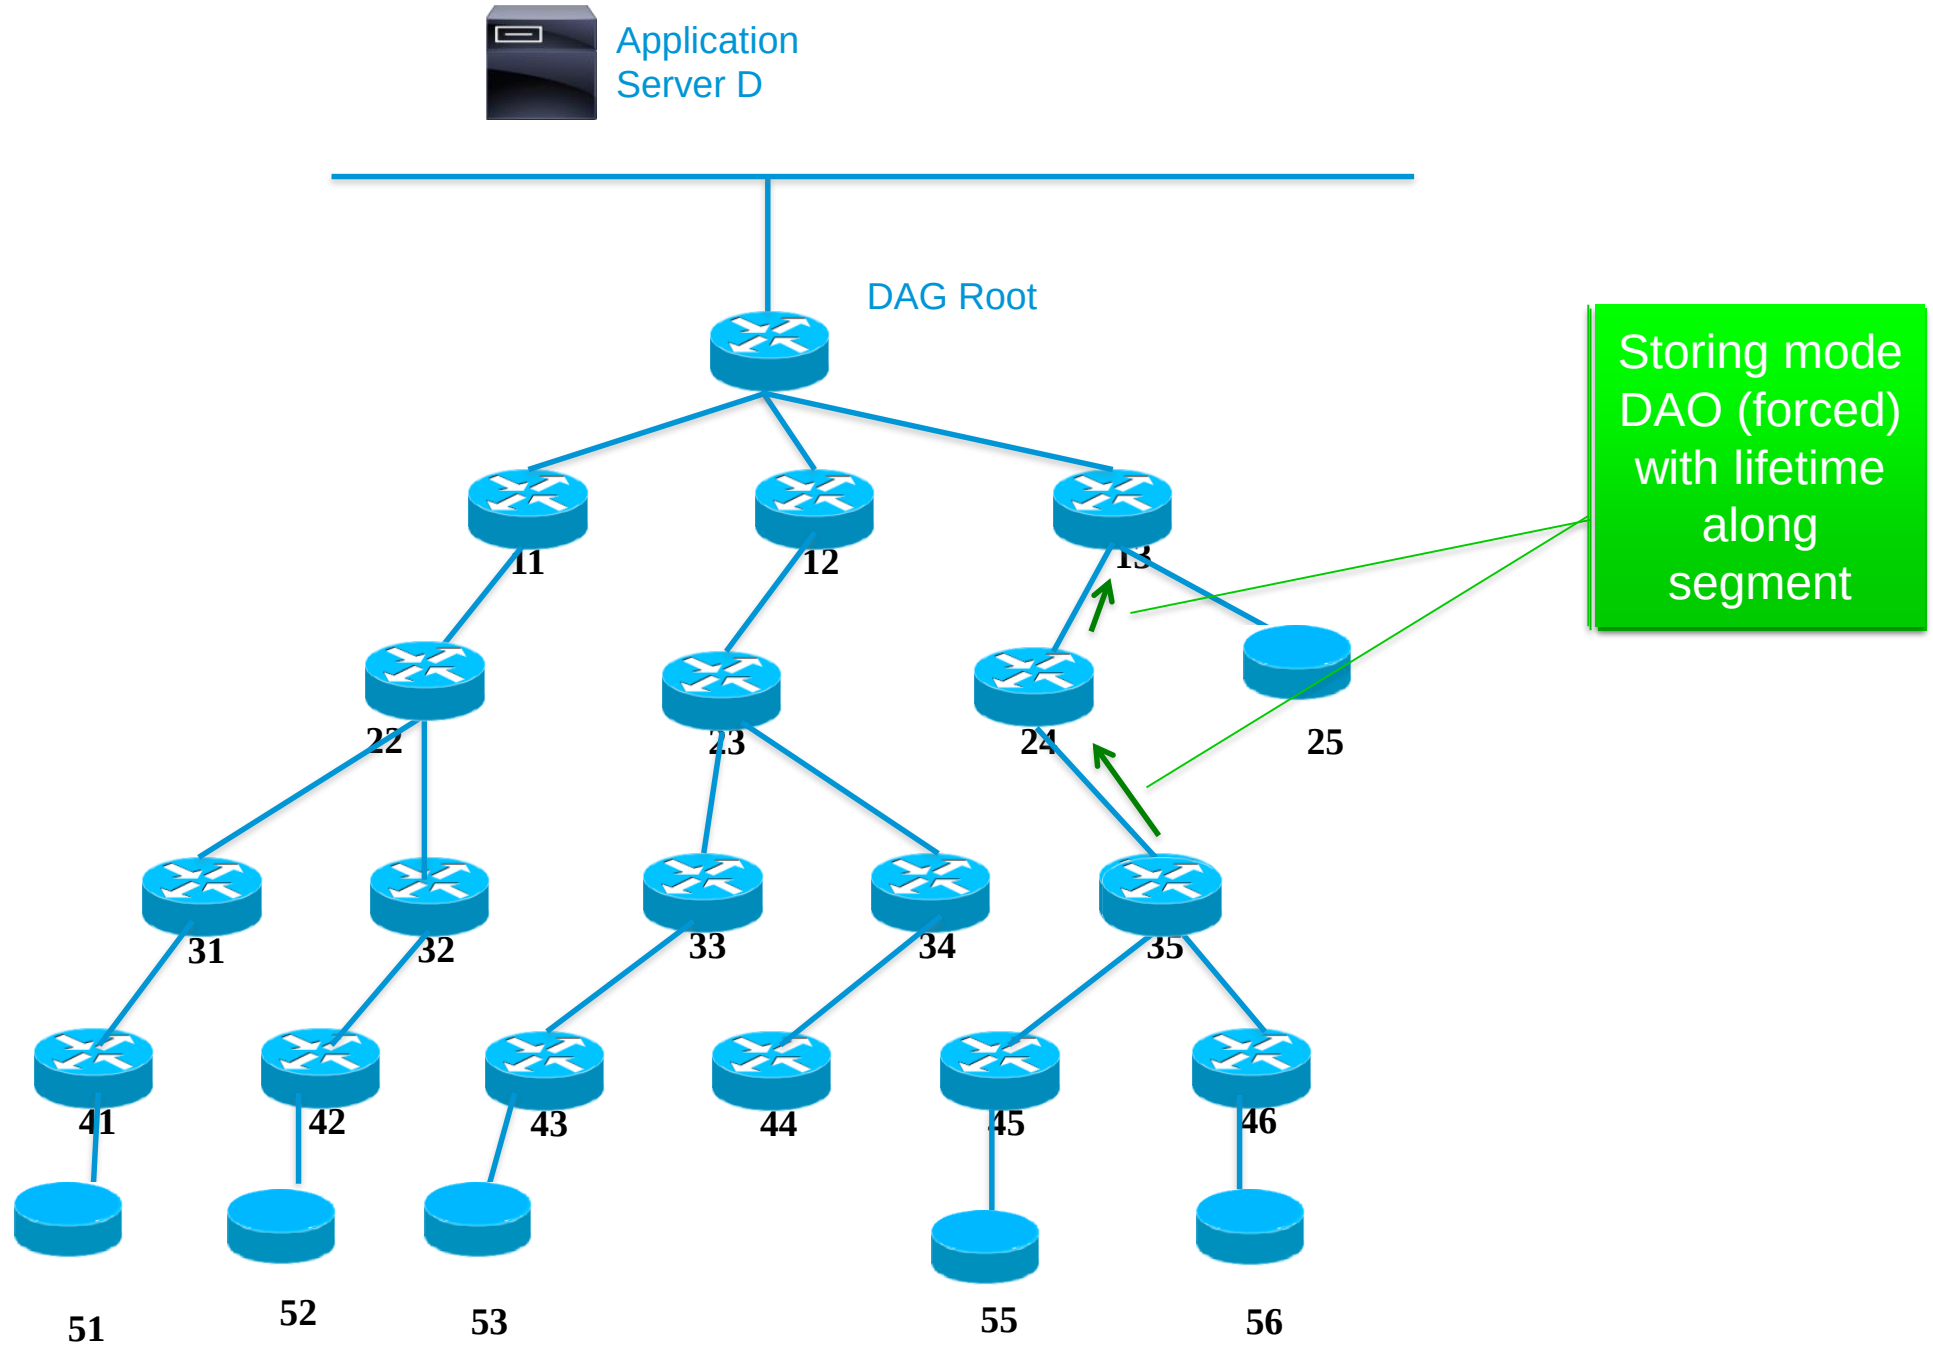
Application
Server D
DAG Root
Storing mode
DAO (forced)
with lifetime
along
segment
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

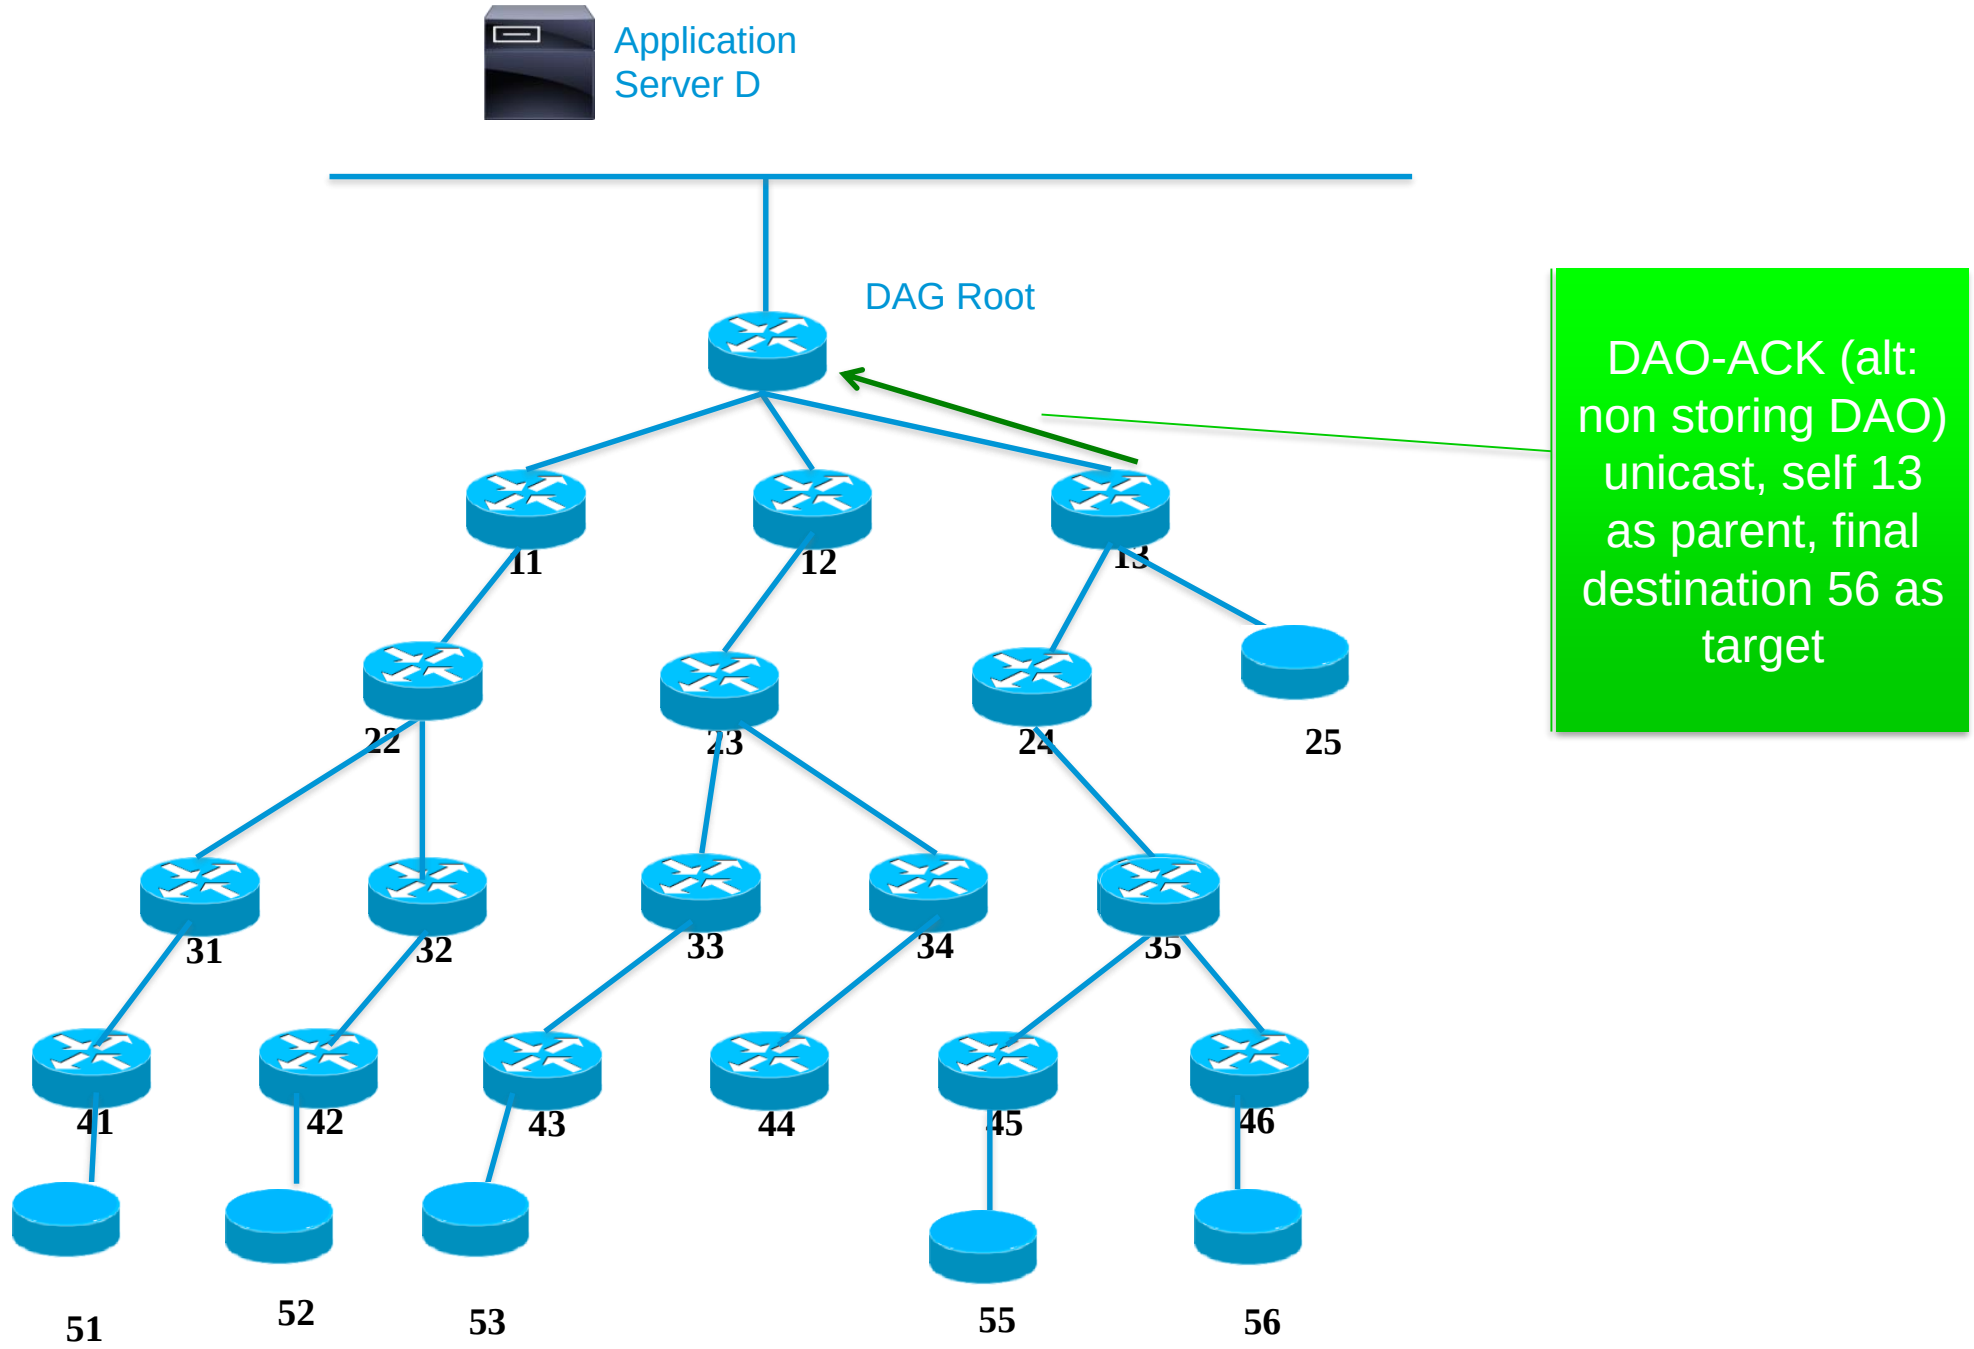
Application
Server D
DAG Root
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56
DAO-ACK (alt: non storing DAO) unicast, self 13 as parent, final destination 56 as target

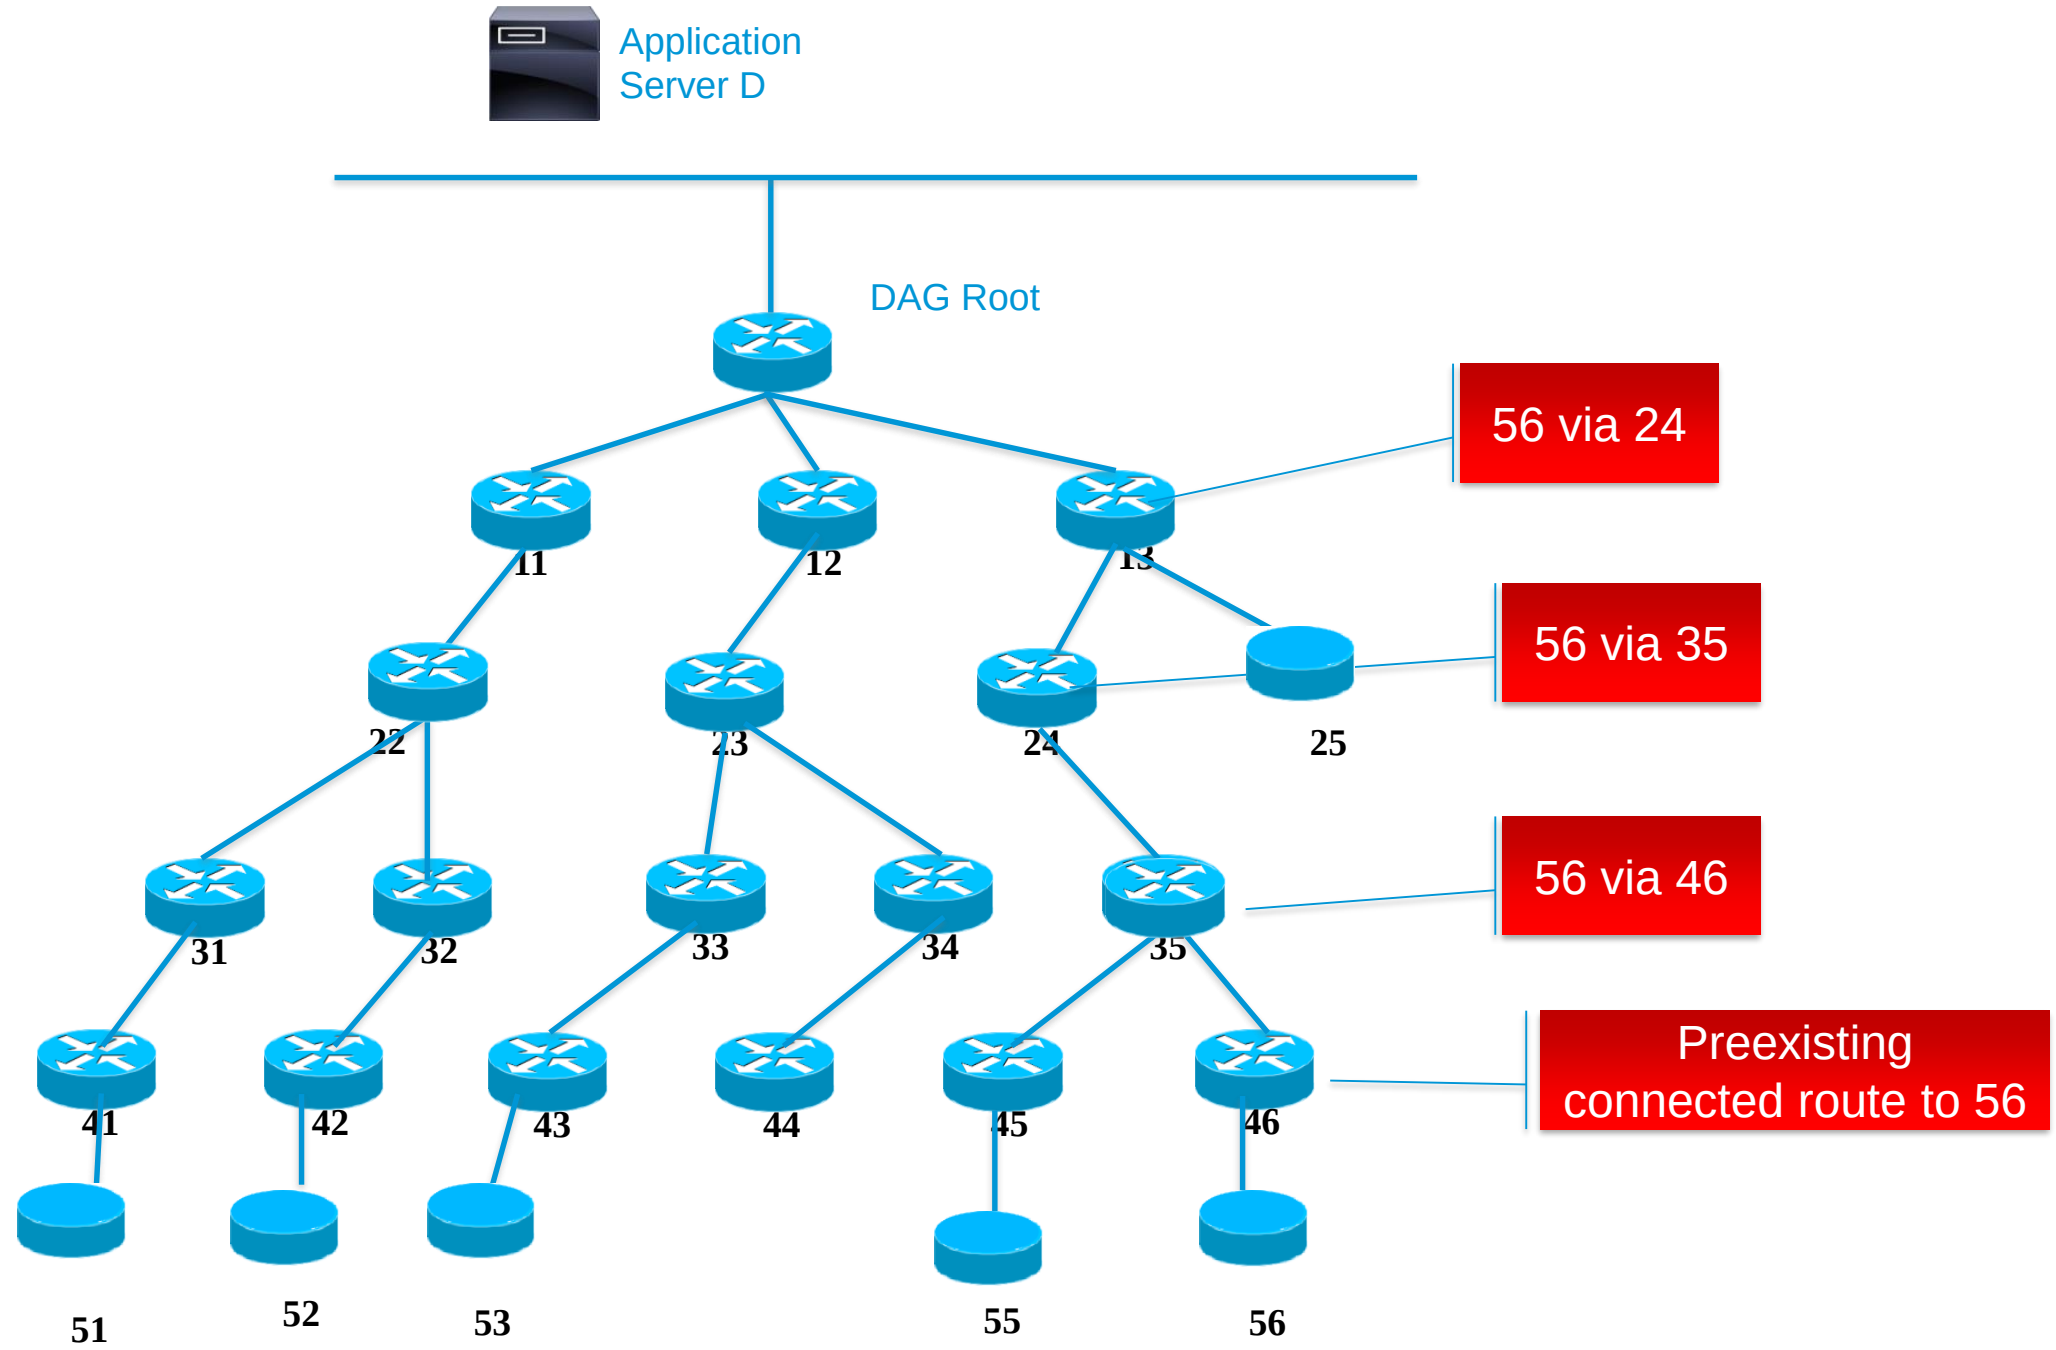
Application
Server D
DAG Root
56 via 24
56 via 35
56 via 46
Preexisting
connected route to 56
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

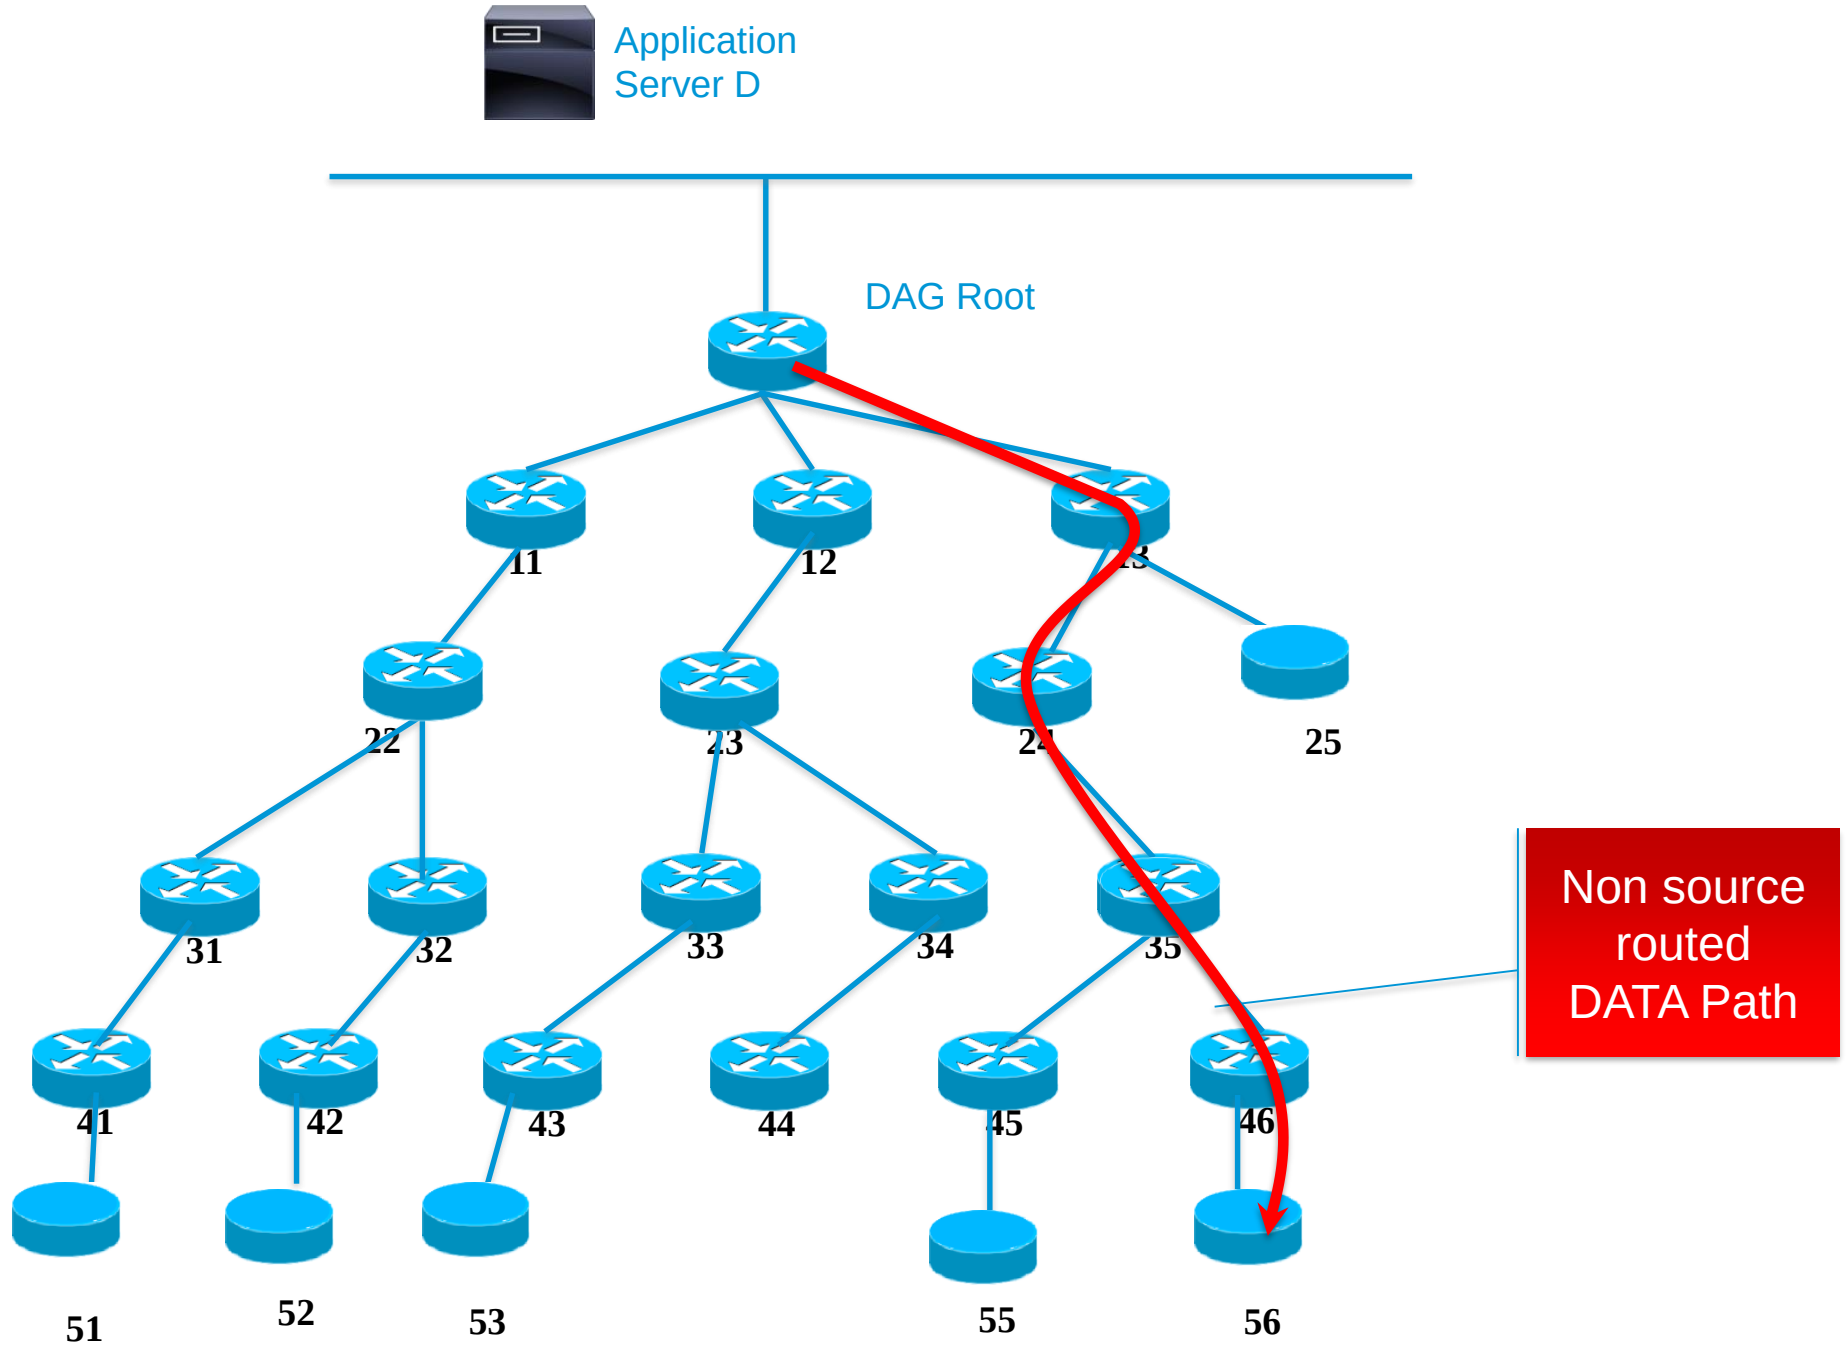
Application
Server D
DAG Root
11
12
22
23
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56
Non source
routed
DATA Path

# Alternate Programming By the root (Michael)

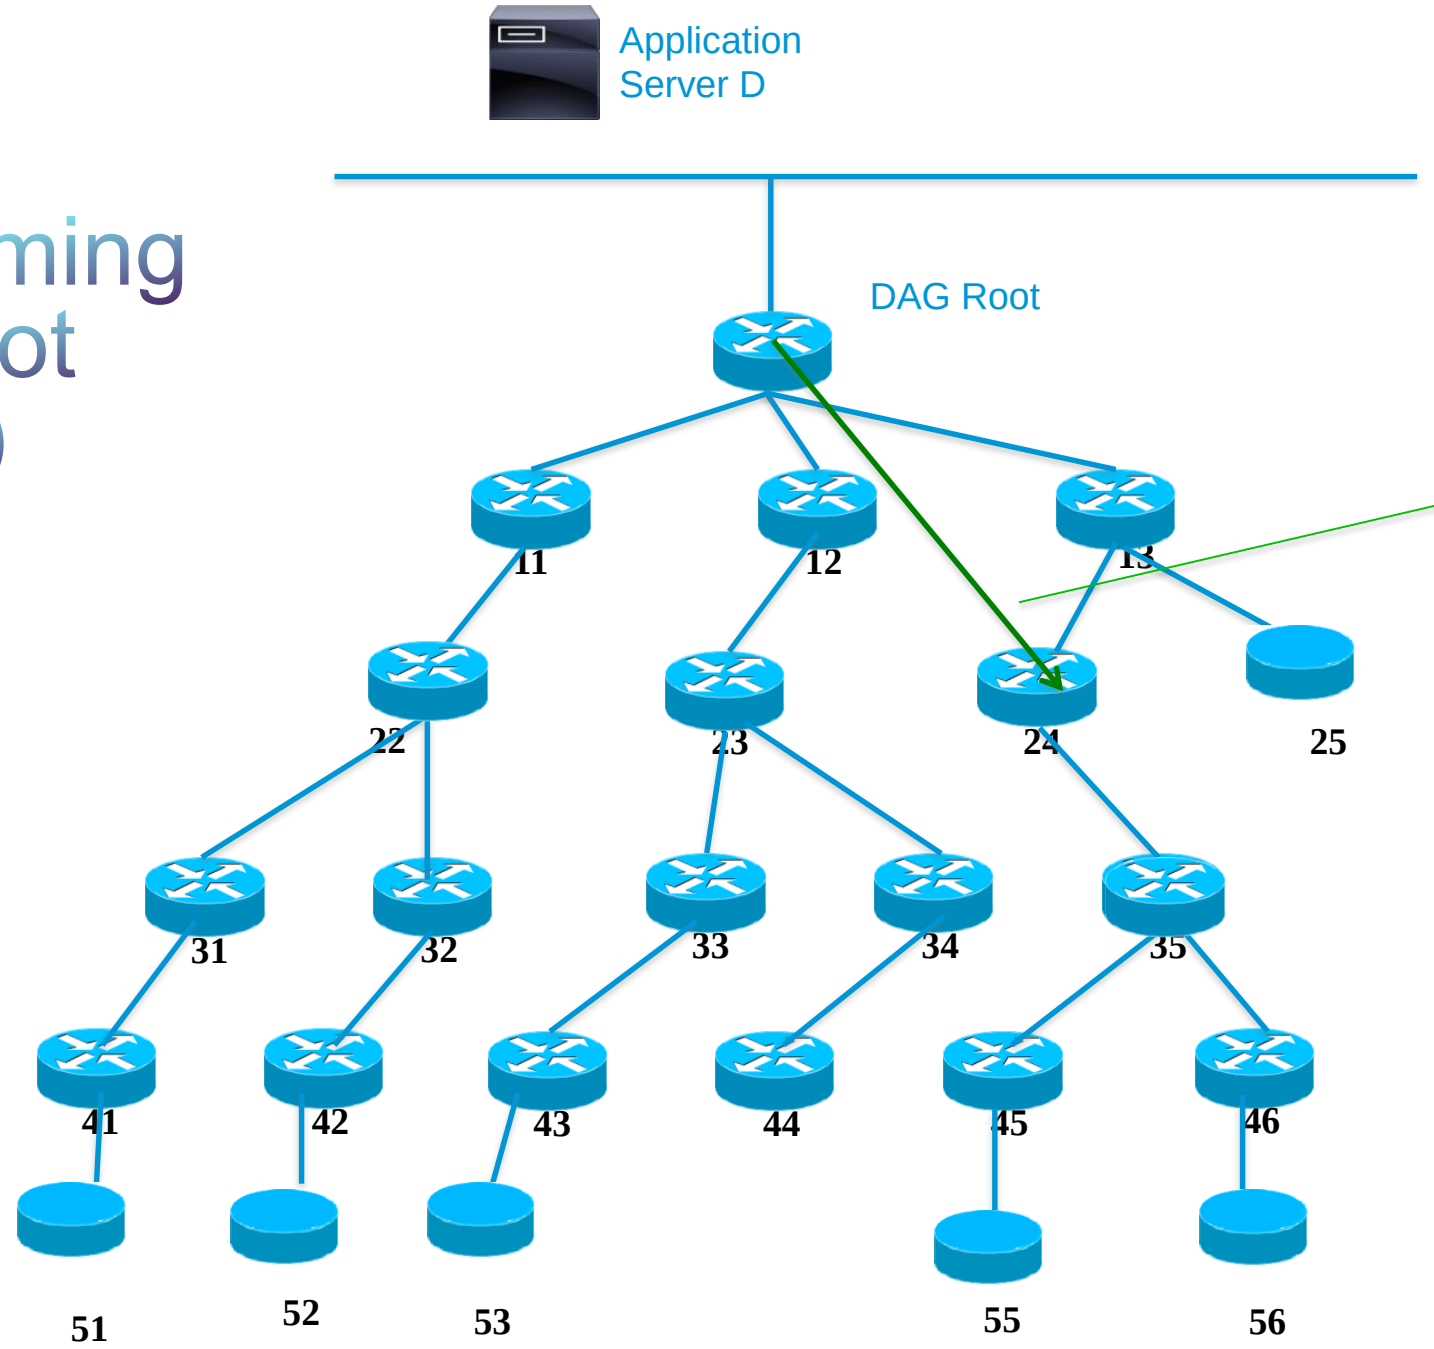


ALT: Adding New (projected) DAO with path segment unicast to target 35 via 13 (ingress) and 24 (egress)

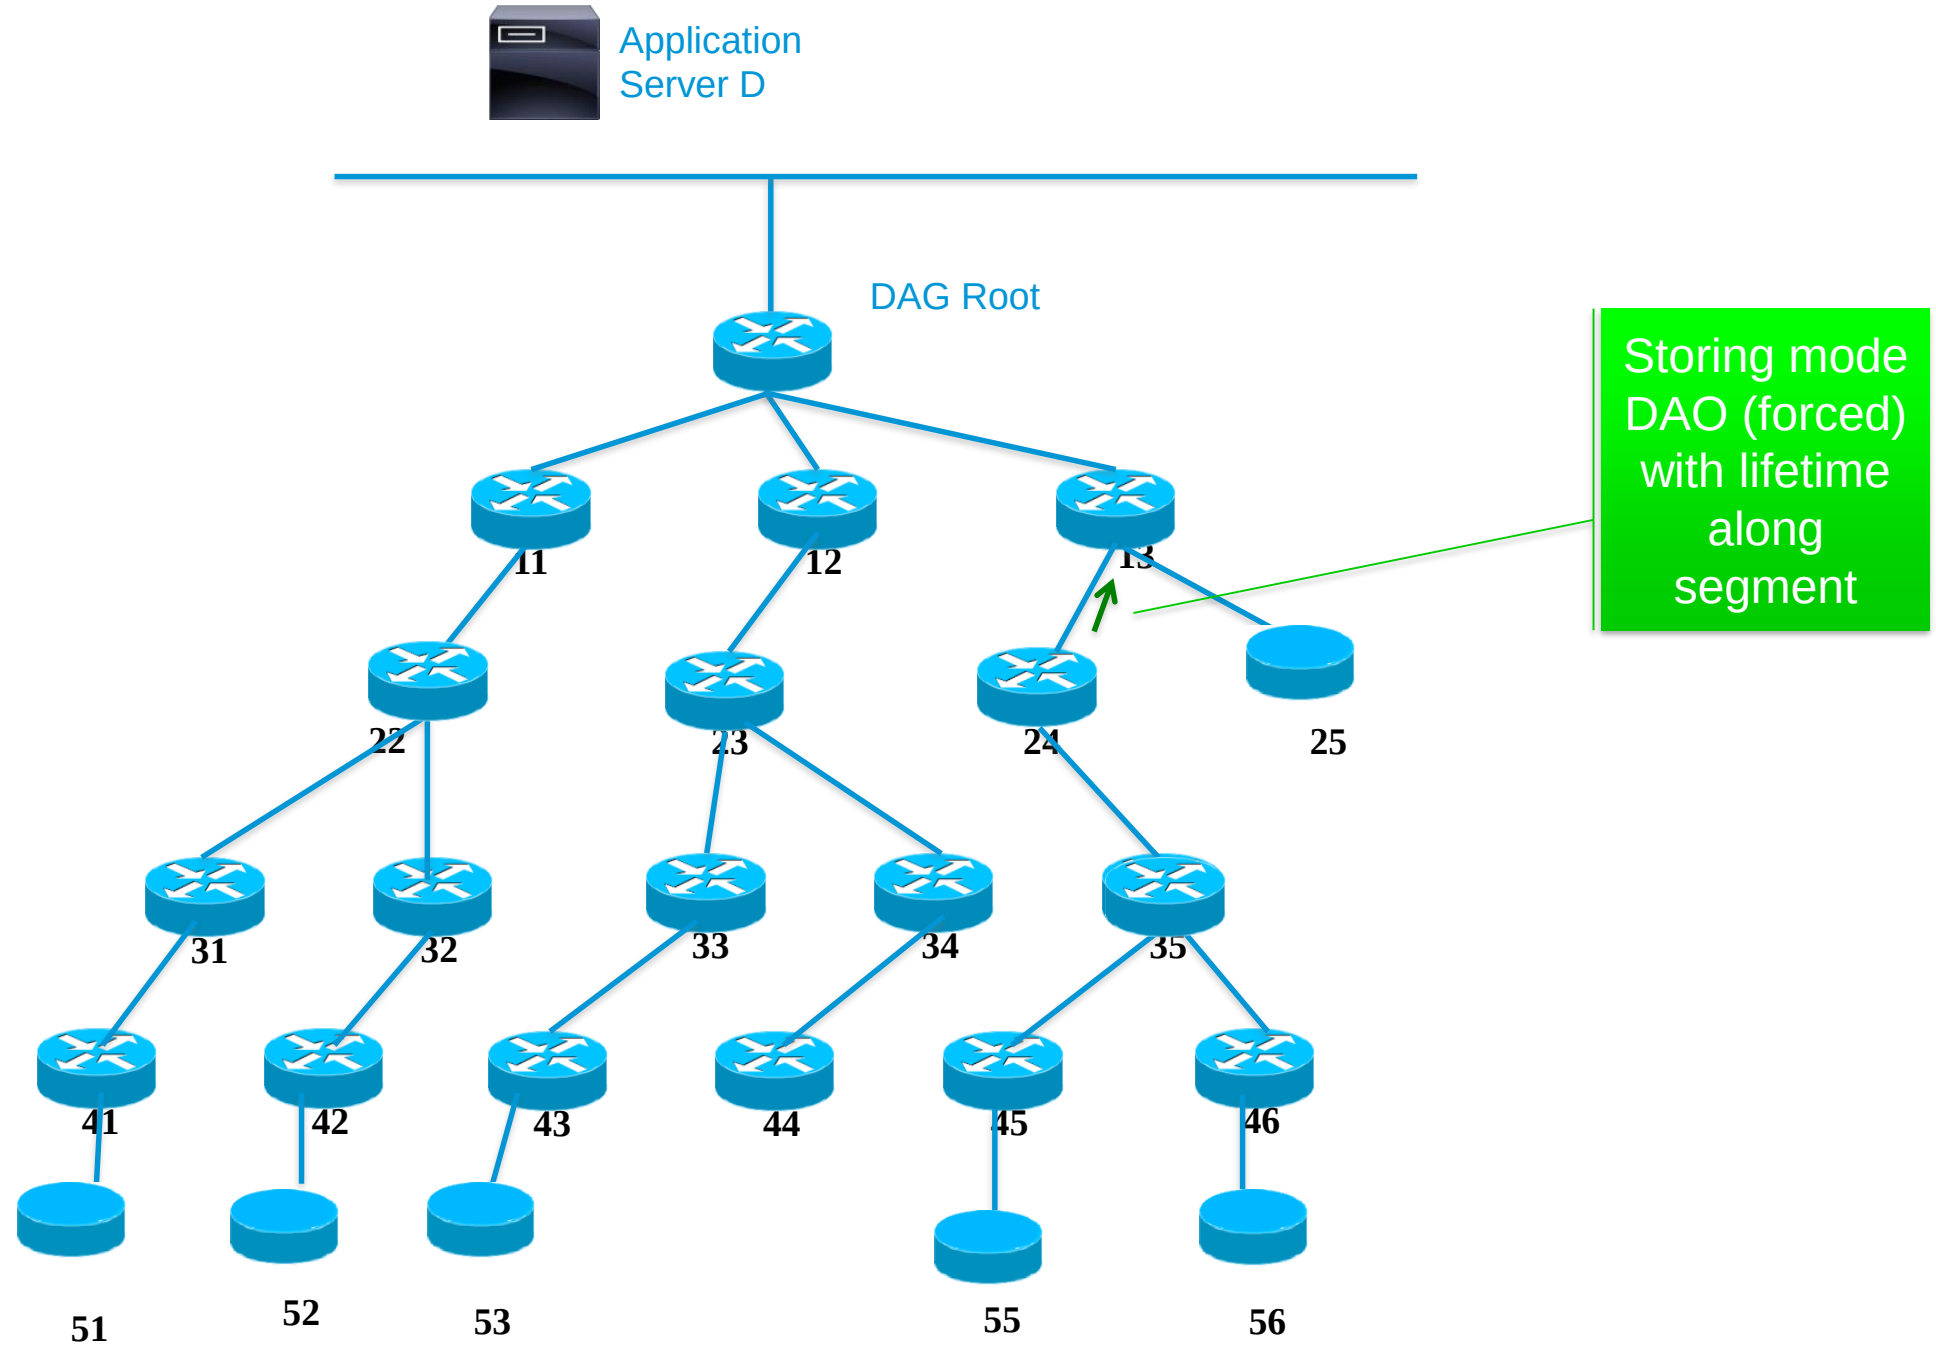
Application
Server D
DAG Root
Storing mode
DAO (forced)
with lifetime
along
segment
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

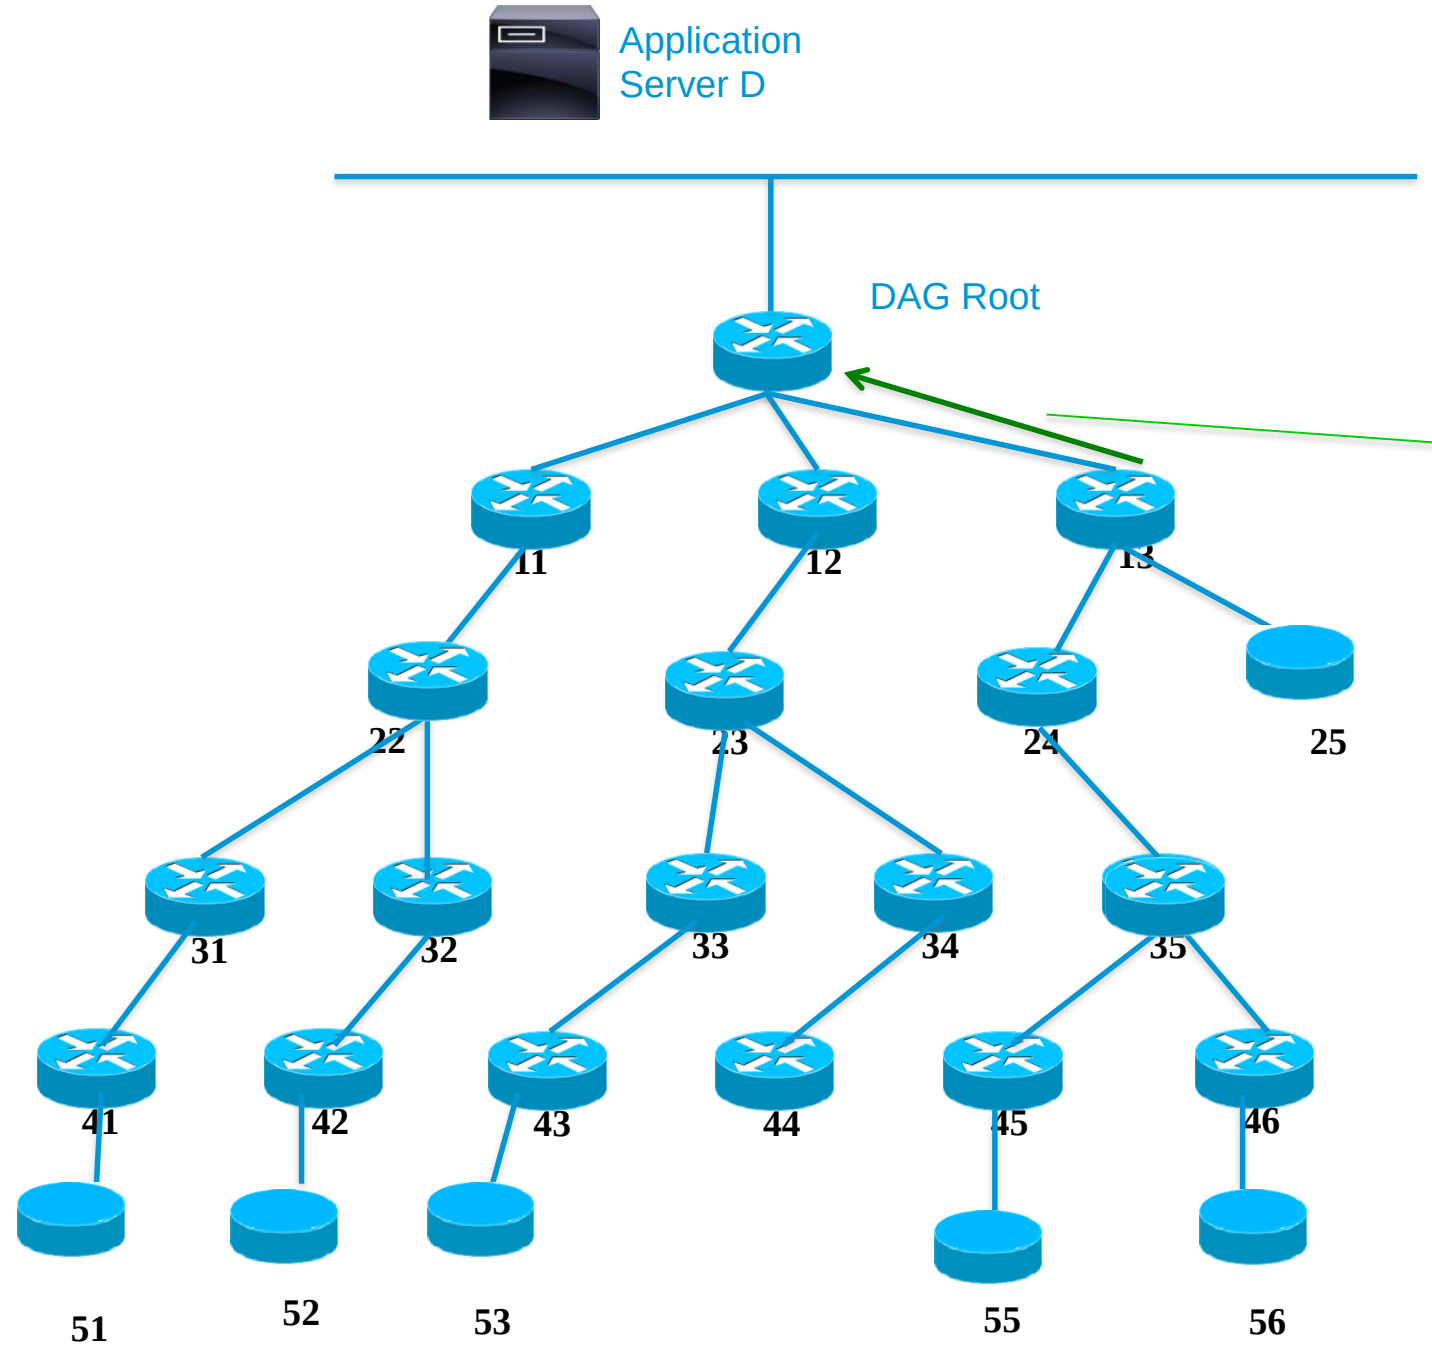


DAO-ACK (alt: non storing DAO) unicast, self 13 as parent, final destination 56 as target

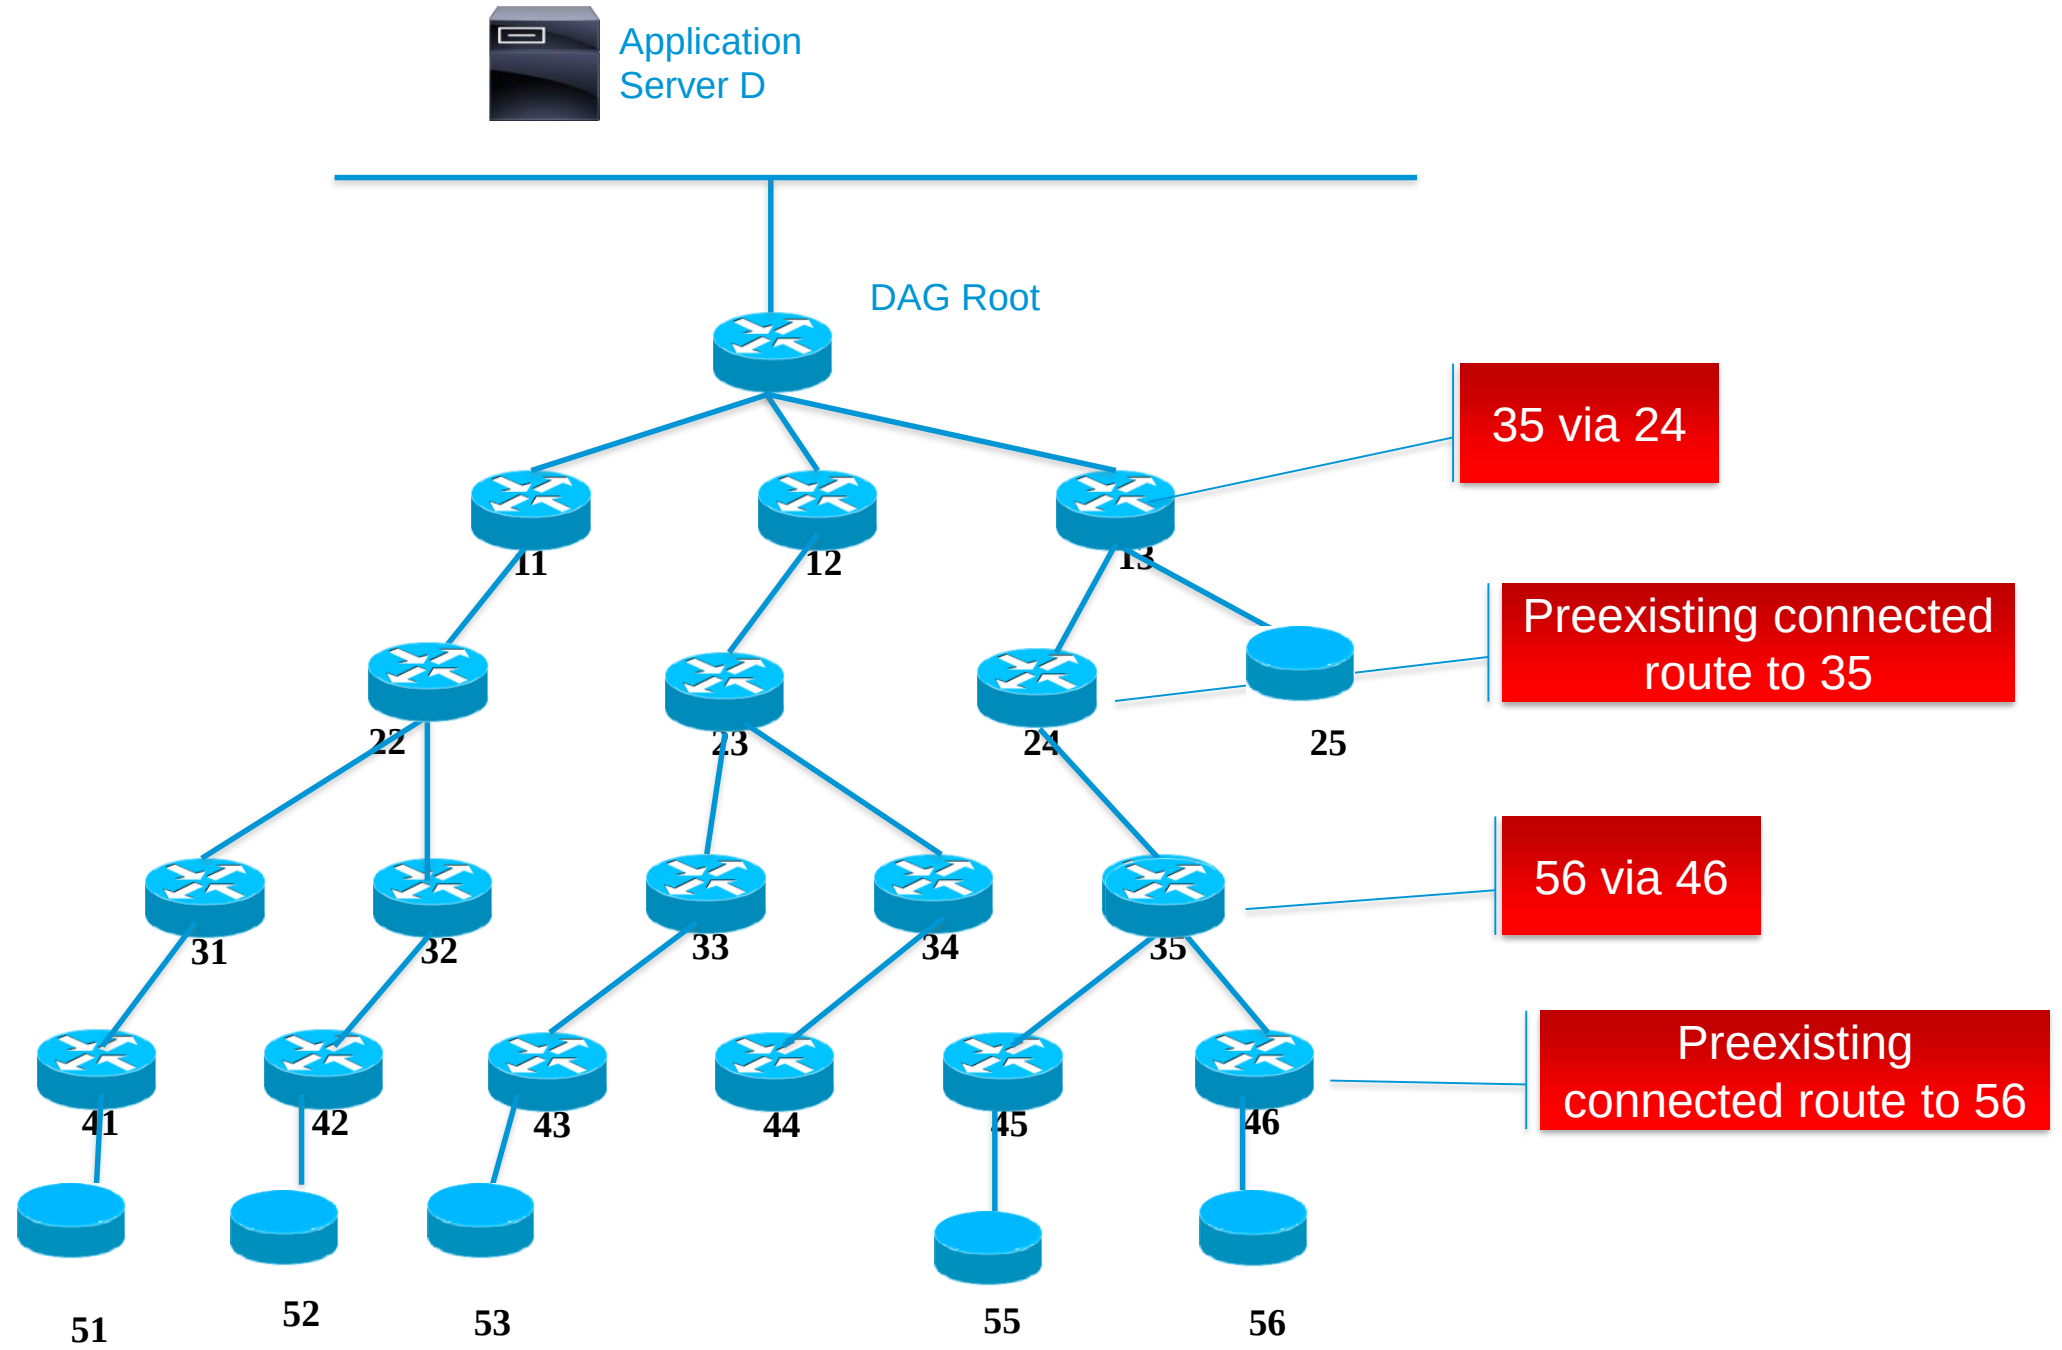
Application
Server D
DAG Root
35 via 24
Preexisting connected
route to 35
56 via 46
Preexisting
connected route to 56
11
12
13
22
23
24
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

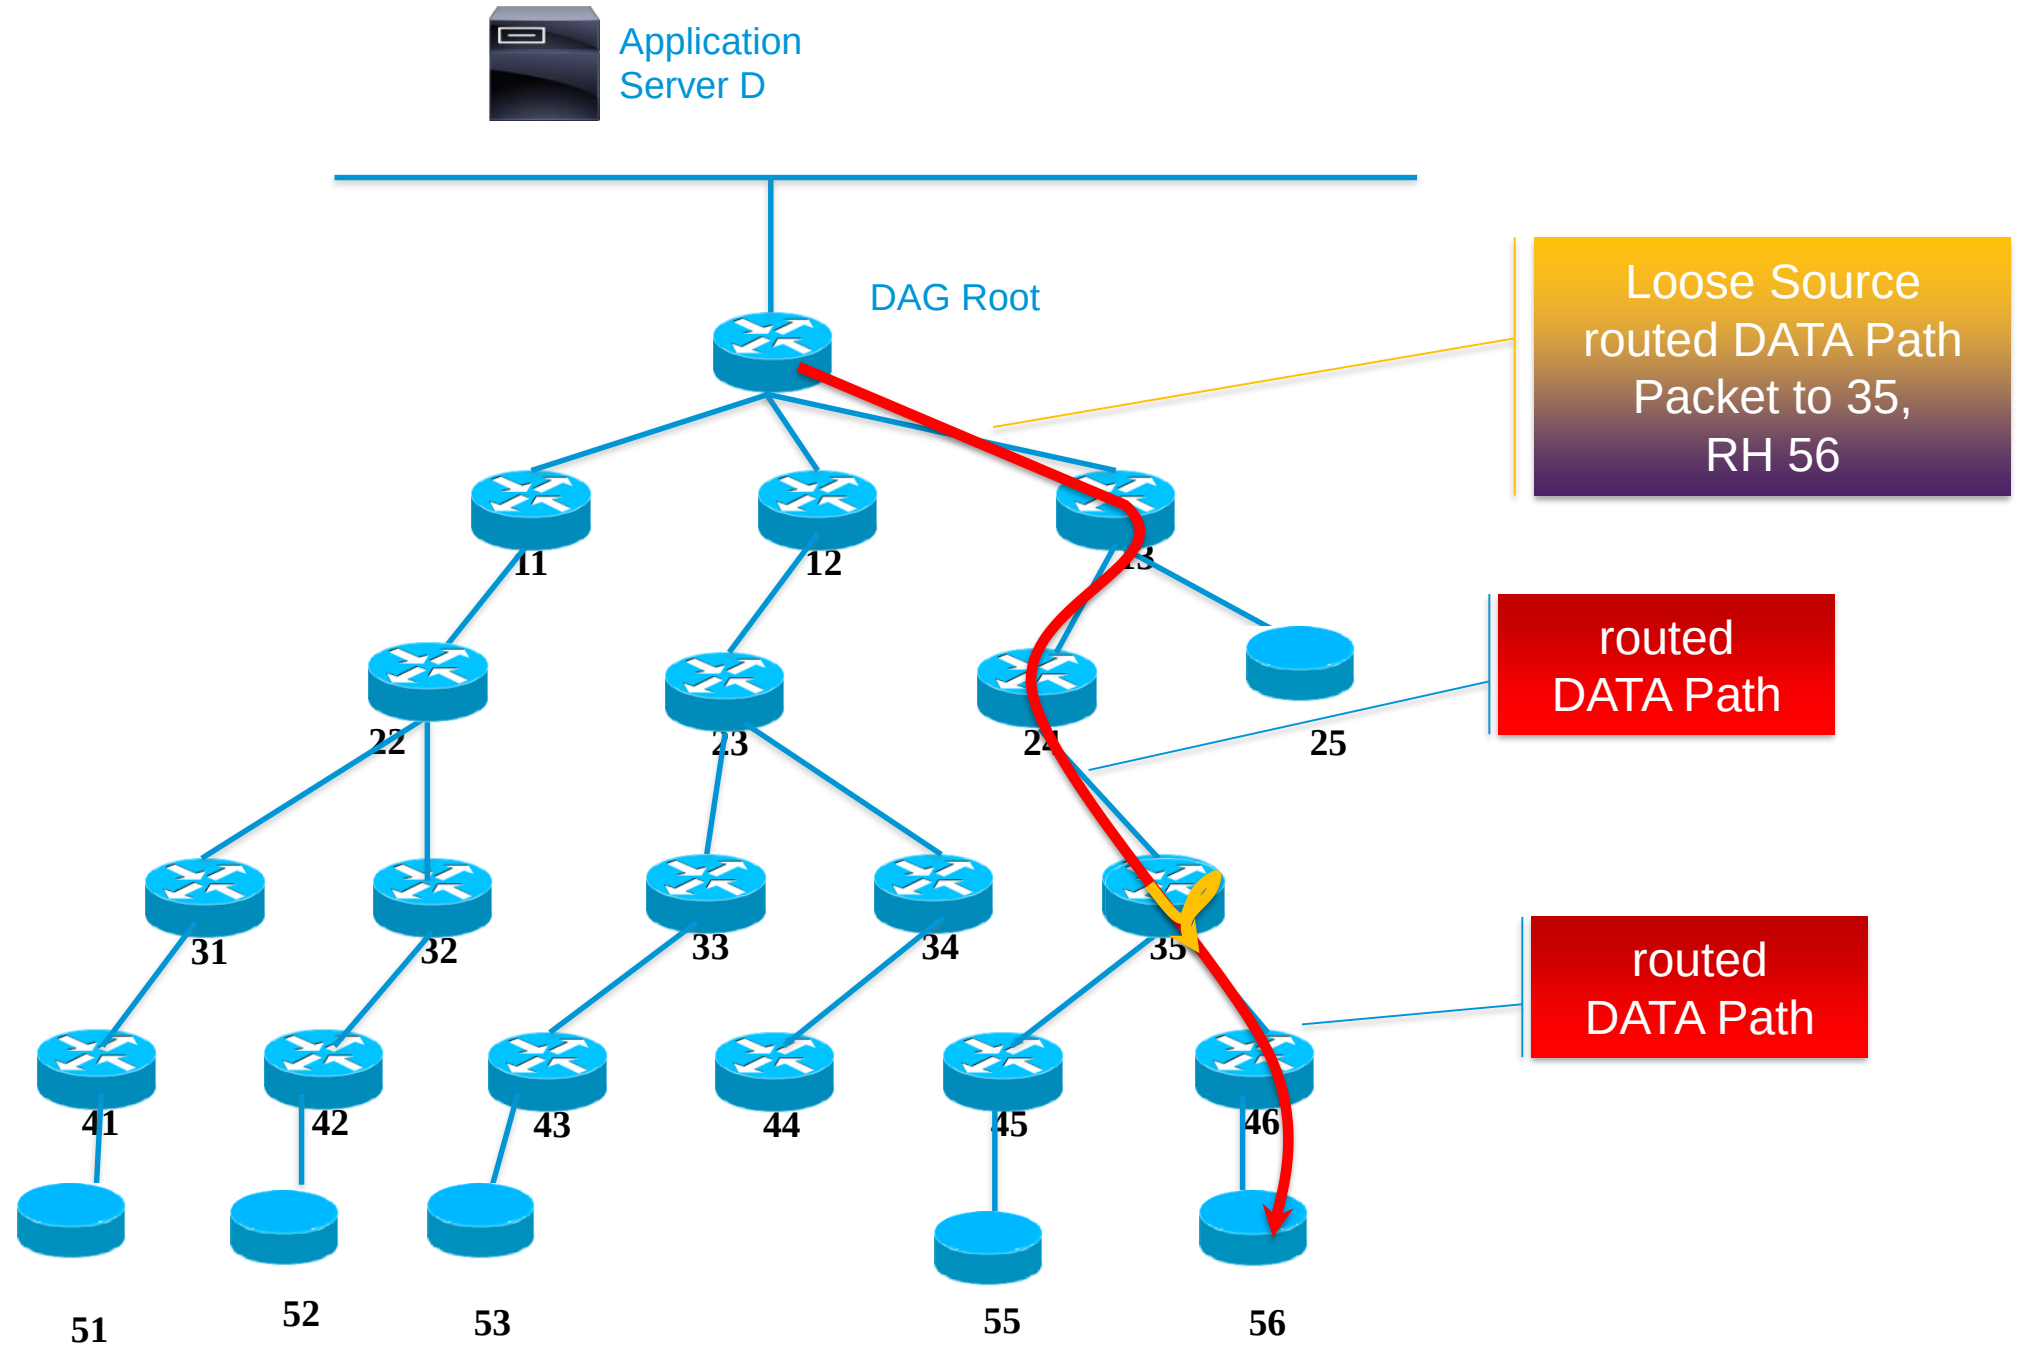
Application
Server D
DAG Root
Loose Source
routed DATA Path
Packet to 35,
RH 56
routed
DATA Path
routed
DATA Path
11
12
22
23
25
31
32
33
34
35
41
42
43
44
45
46
51
52
53
55
56

# Status of draft-ietf-roll-dis-modifications-00

New version: June/July 2018 :-)

# Open Mic

Questions, Comments, Suggestions, …. :-)