

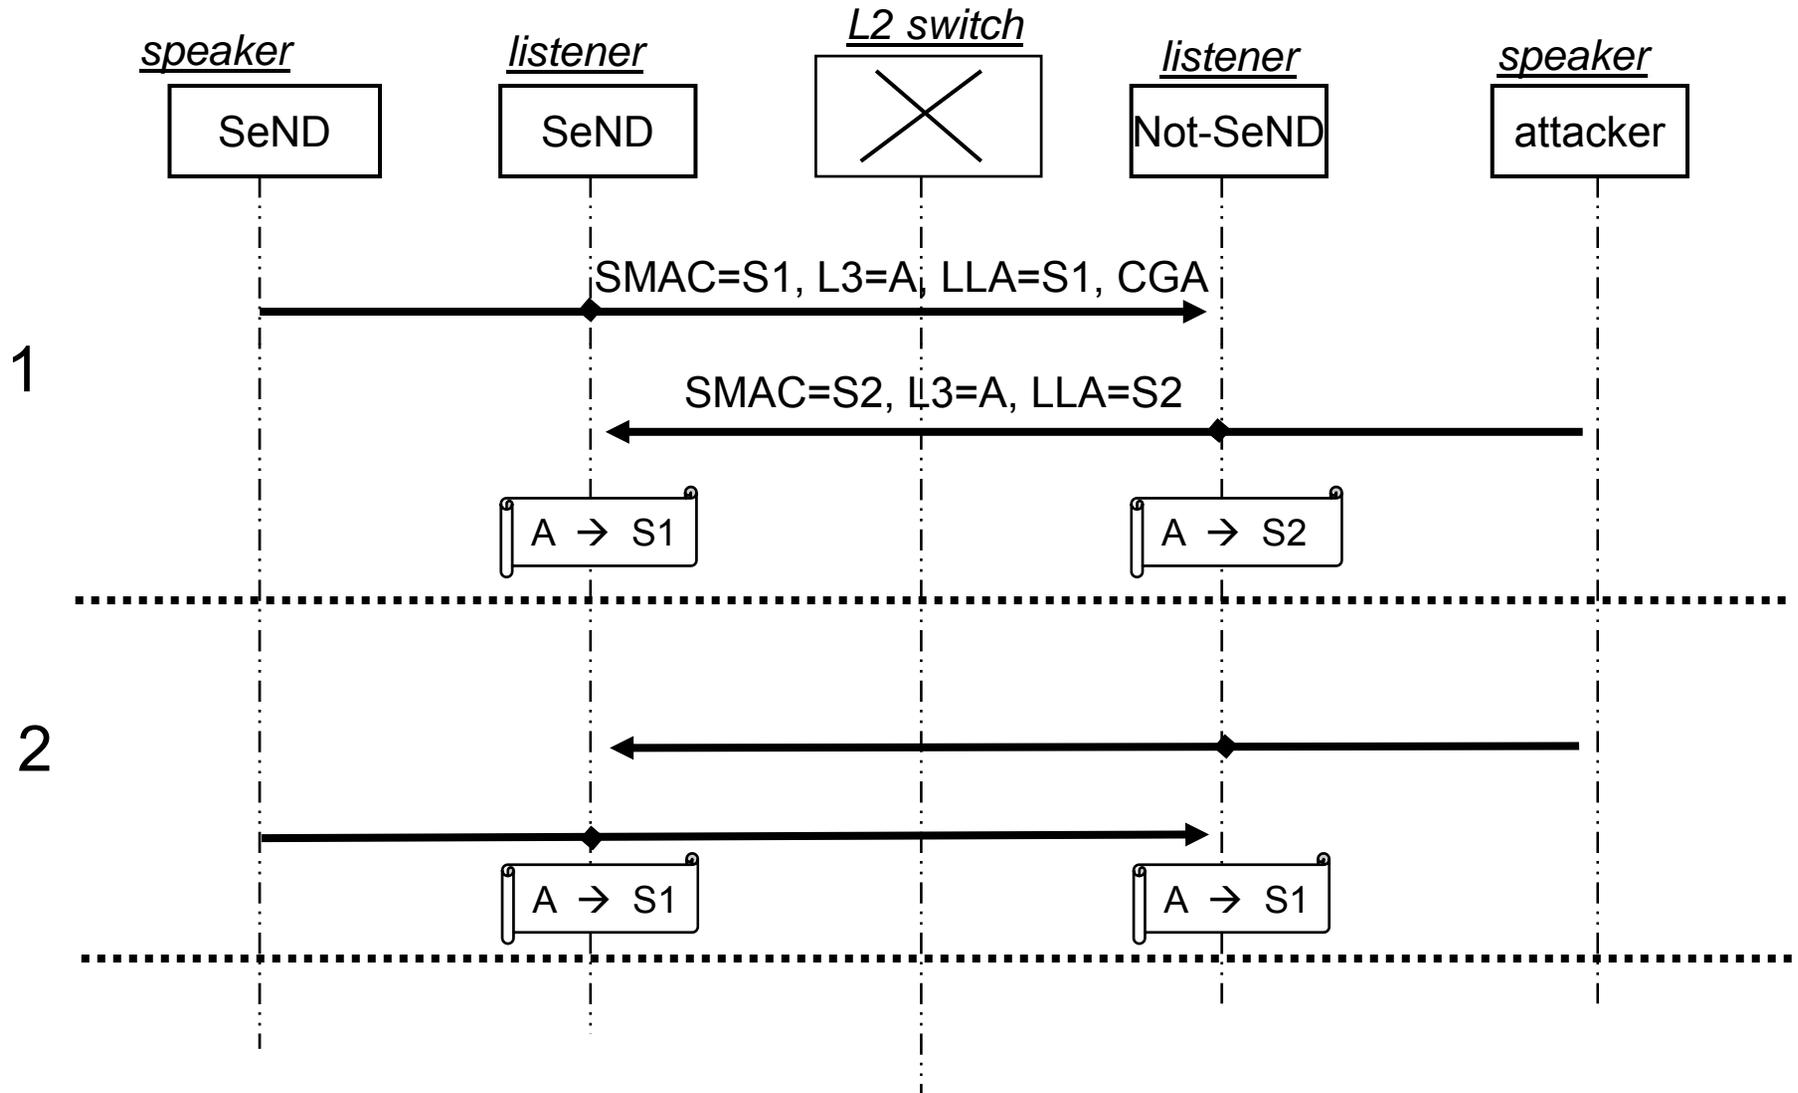
**IETF-75**

[draft-levy-abegnoli-savi-plbt-00](#)

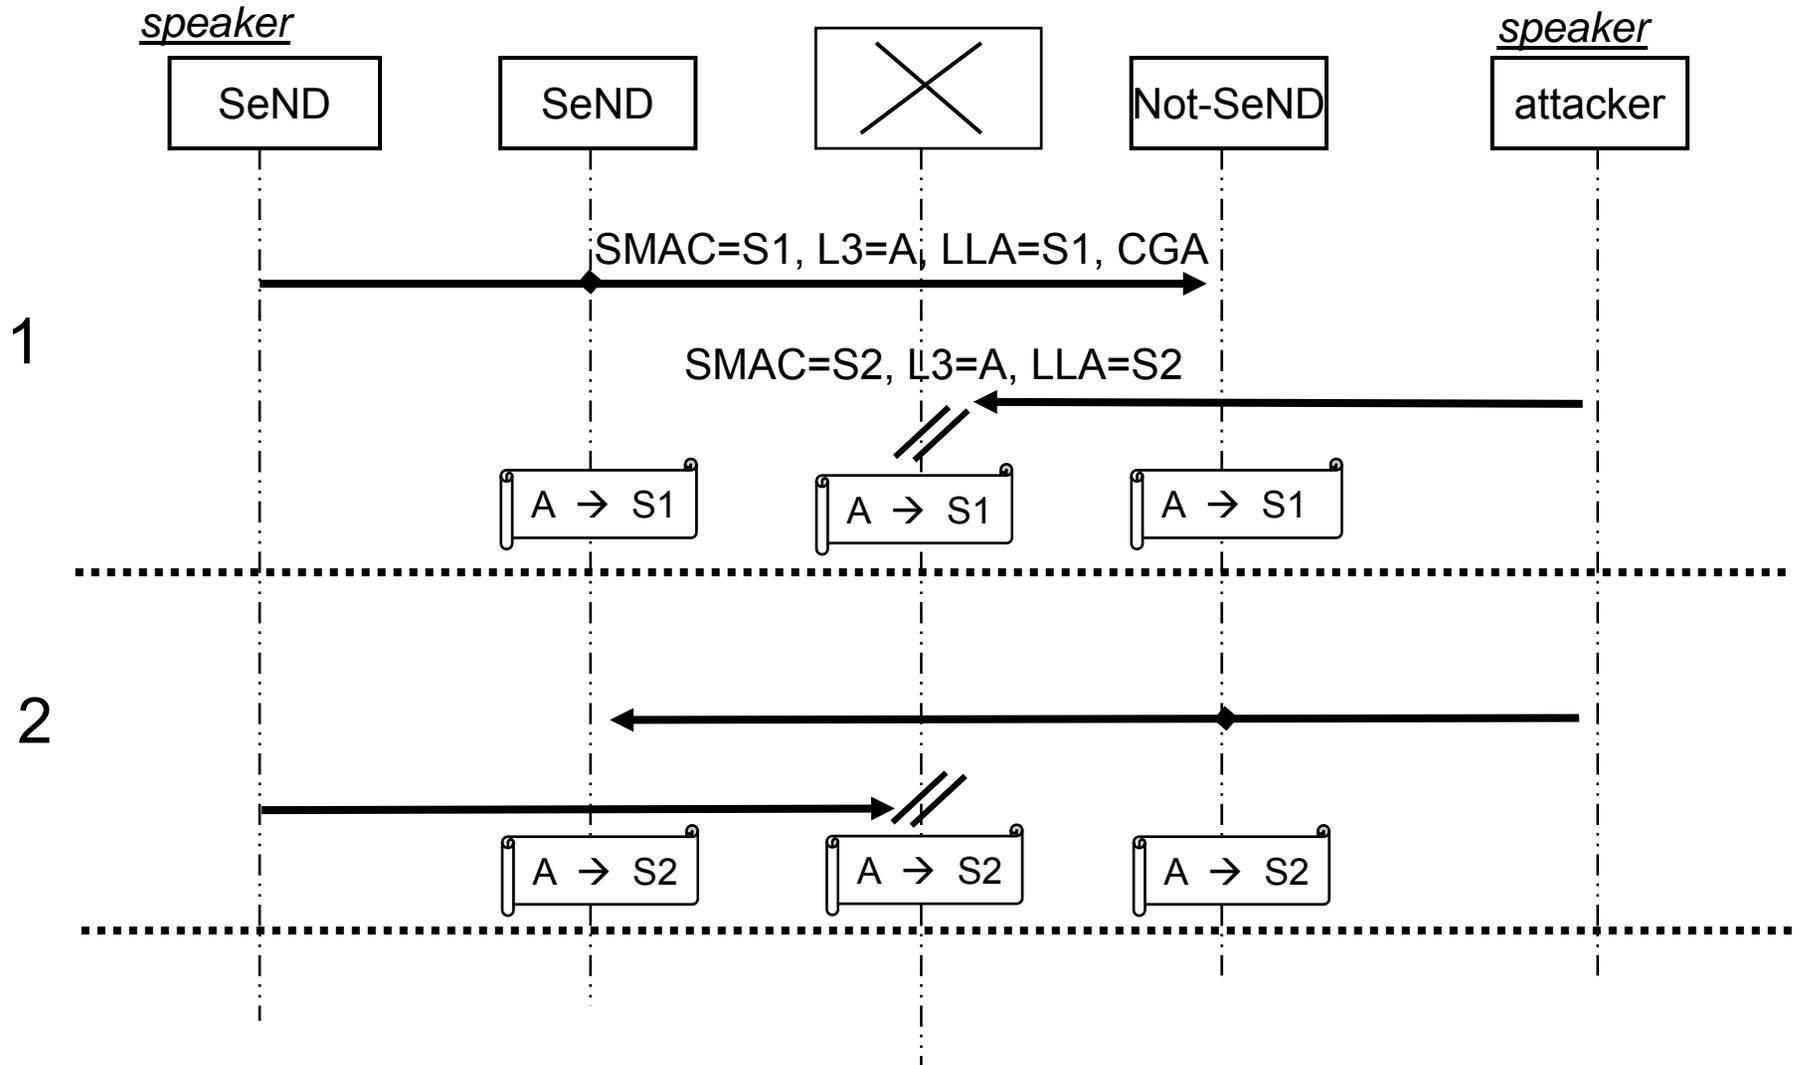
# Goals

- Build and maintain a L3/L2 binding table on the L2 switch as *accurate as possible (definition of “accurate”?)*
- Focus on SLACC address assignment
- Separate control and data plane
  - Bindings established with control plane protocols
  - DAD and data packets can be used as “hints”
  - Simple rustic data plane decisions: bridge, drop, count, punt to cpu
- Analyses and possibly handle “hard” cases:
  - SLACC
  - Multiple switches, with mix of savi & non-savi
  - Mix of trusted and untrusted switch ports
  - Mix of cryptographic, static, dhcp, slacc addresses
  - NDP-LLA and SMAC can be different
  - Heterogeneous movement requirements

# Bindings accuracy



# Bindings accuracy



# Approach

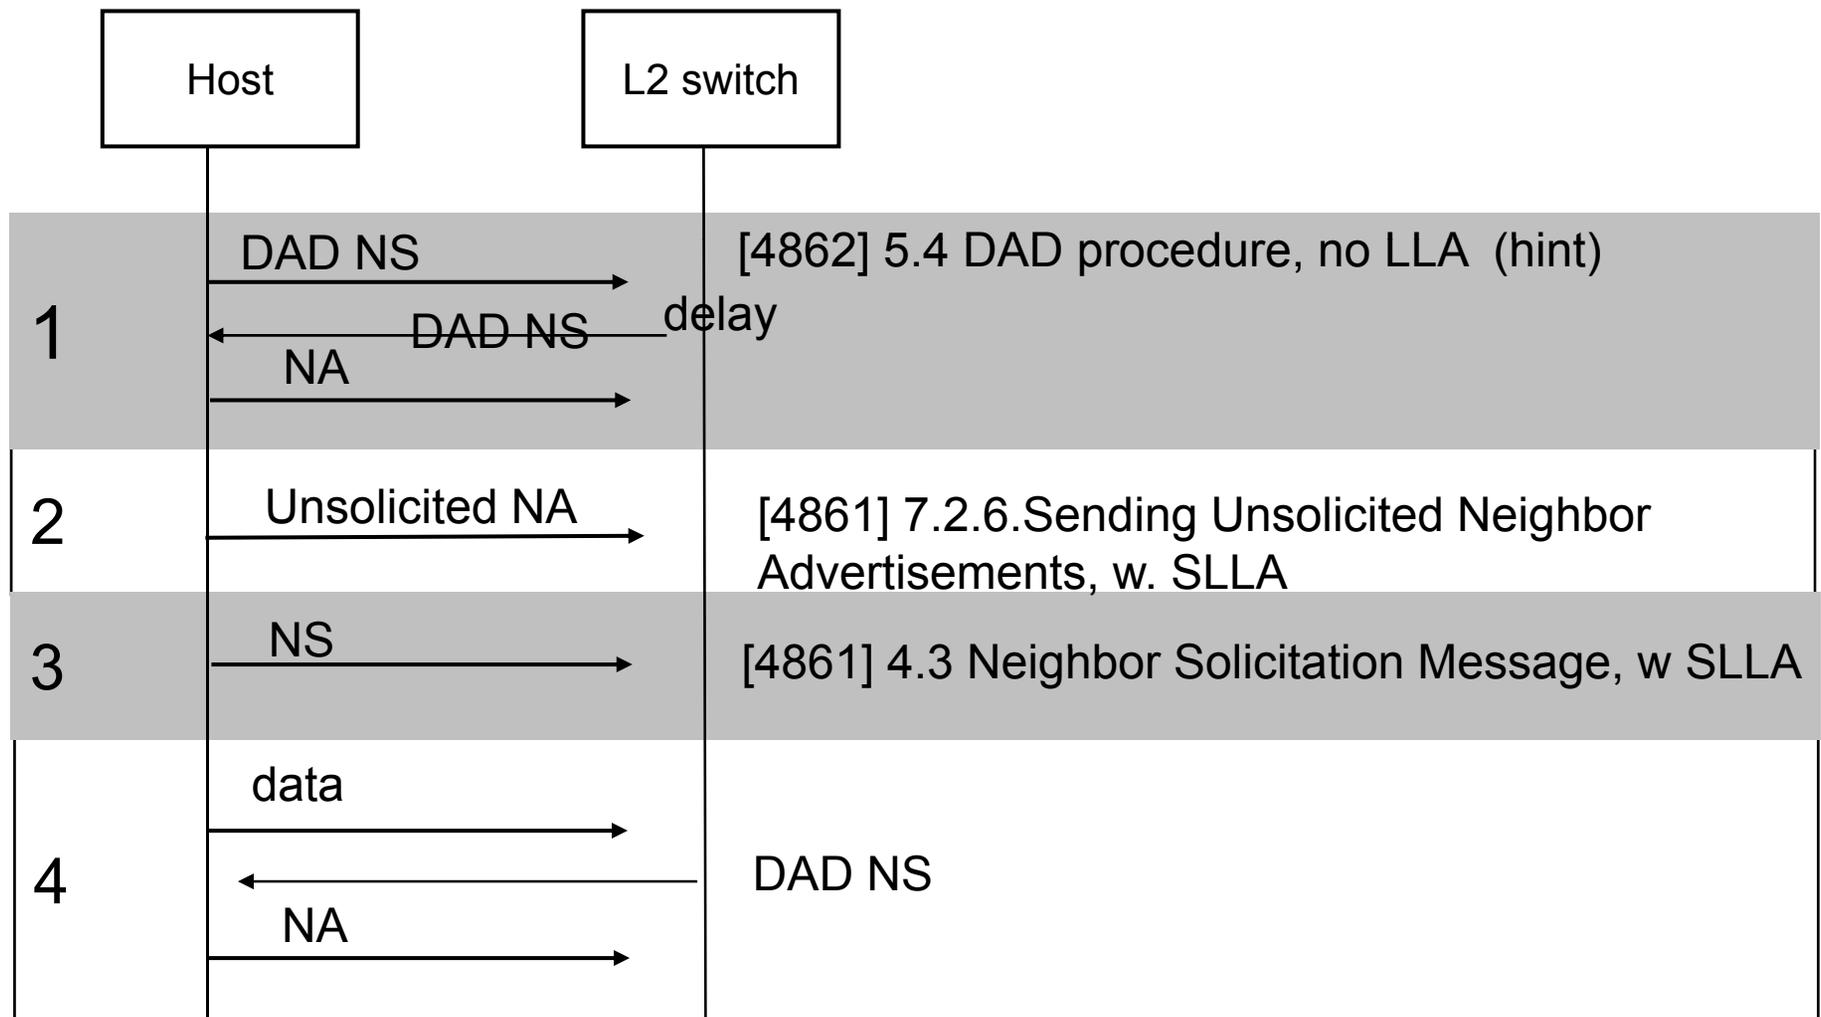
- Bindings are created “by L3 and above” information, not L2
- For SLACC, inspect all [NDP] messages
- Relies heavily on configuration
- Collision algorithm: “Prefer” some bindings over others (not LCFS, not FCFS)
- Maintain state accuracy for better movement handling
- CGA friendly
- Movement scenarios are a function of all the above

# Learning process

- Binding learnt by snooping all NDP traffic:
  - NS, NA, RS, RA [REDIR, CPS, CPA]
- L3 bound to NDP-LLA. SMAC used as a hint.
- Preference algorithm for collisions (First-come is “one” element of it)
- Entry state tracked by the switch
- End-devices optionally NUD’ed for maintaining binding table entries in most current state

# Inspect all NDP messages ...

Three (+1) chances to learn the binding

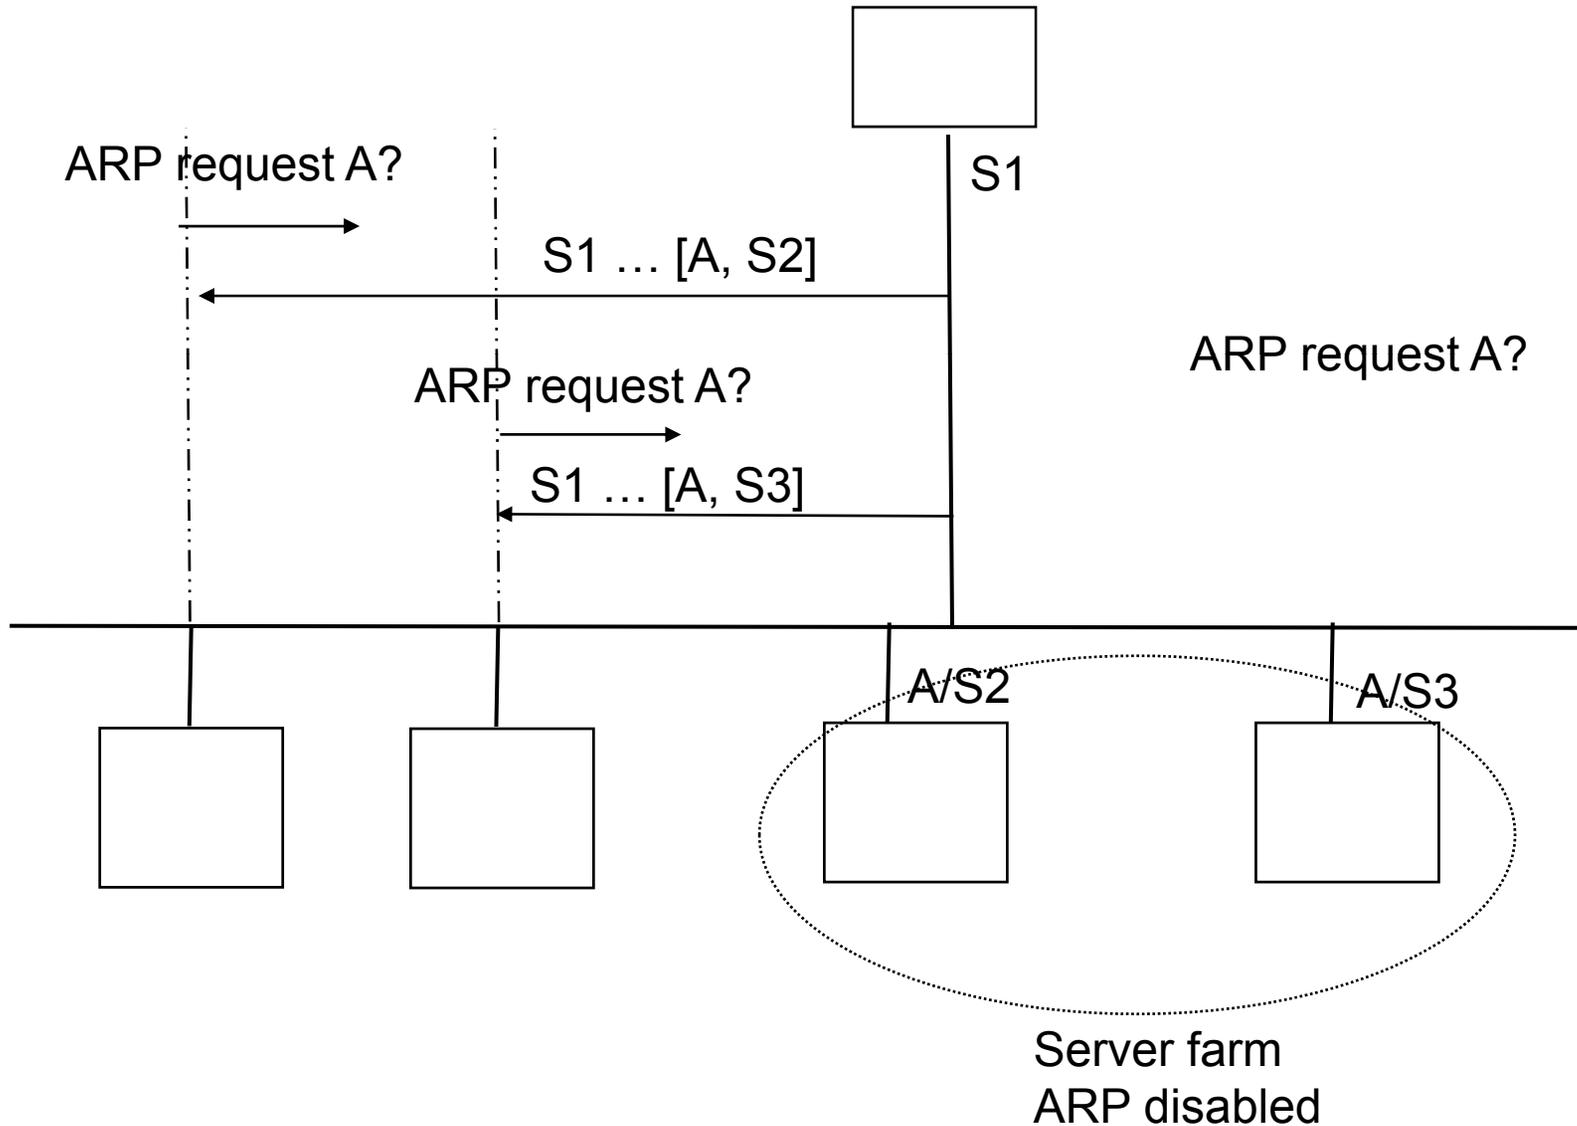


# Inspect all NDP messages ...

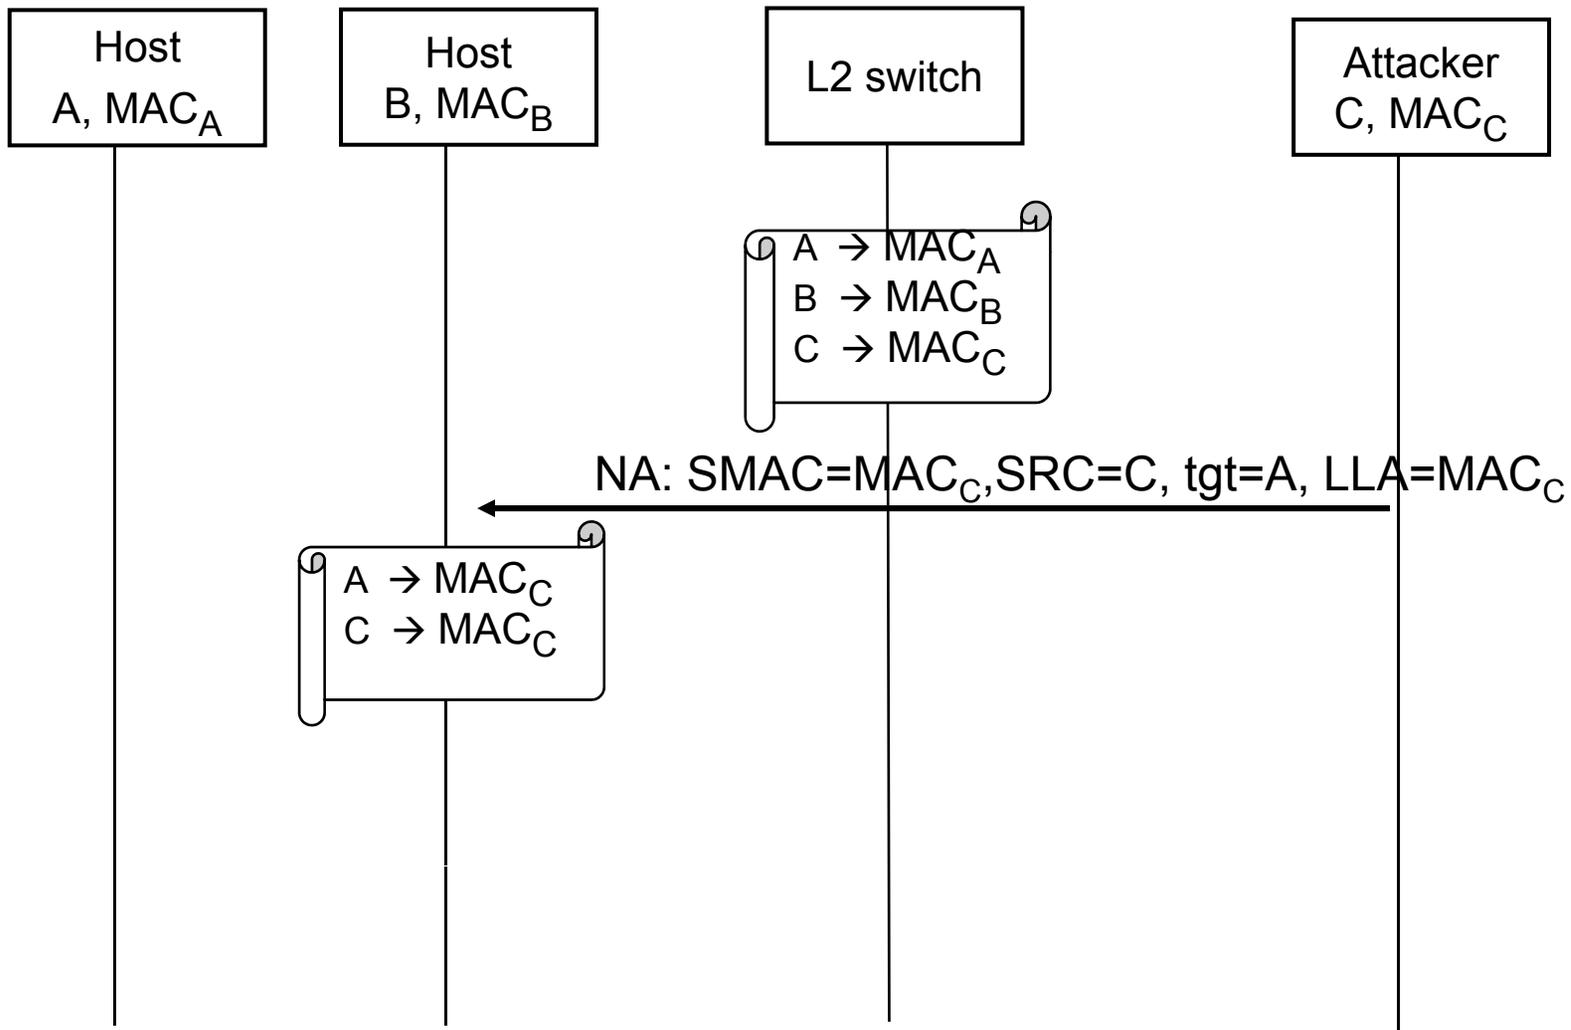
## L3 bound to NDP-LLA

- Currently, there is nothing that forces the SMAC and the LLA to be the same
- Some “border-line” scenarios can exploit this:
  - LLA server (NetPet)
  - Multiple SMAC, multiple LLA on same interface
  - ND-proxy
- DAD NS is used as a hint, to “lock” a new entry in the binding table. But LLA is the value bound to the IP@.
- Data packet could be used as a well as a hint
- Checking SMAC==LLA is a policy matter
- In the absence of LLA within T0, switch sends DAD NS

# Issue#1 with binding L3 to SMAC

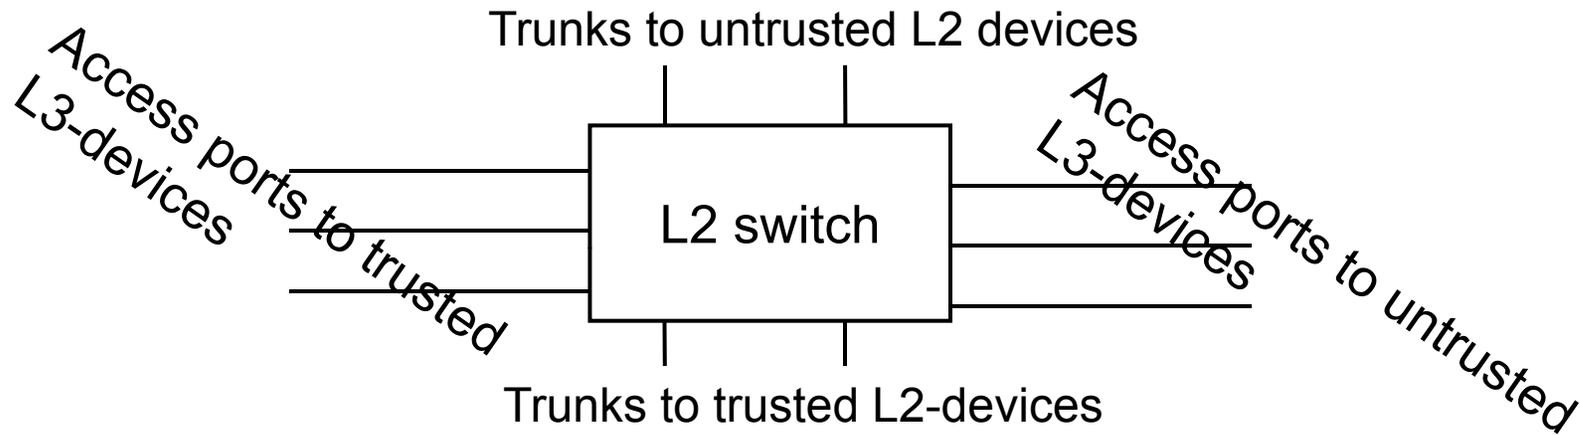


# Issue#2 with binding L3 to SMAC



# Switch port taxonomy

- L2 switch ports classified into:



- This classification is driven mostly by configuration
- Bindings evaluated based on:
  - where they are learn from
  - the credentials that they carry
- A lot of policies can be driven by this classification:
  - Binding entry replacement or update
  - movements
  - SAVI policy (drop, log, punt, rate limits, etc.)

# Preference value

## A. Defines preference “flags”

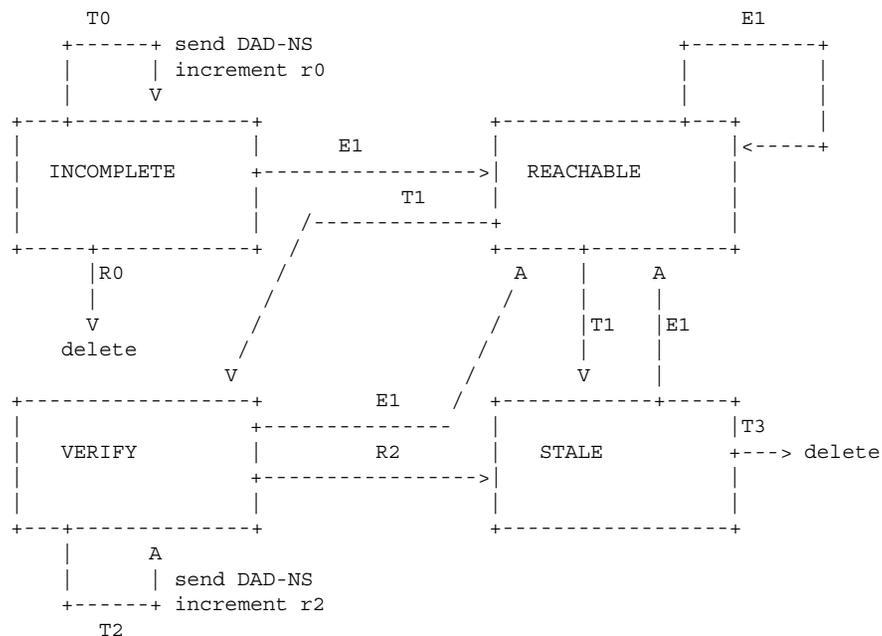
- A flag is either a property of the port from which the entry was learnt or a property of the binding itself
- $\text{Flag}_n = 2^n$
- $\text{Preflevel} = \sum \text{flags}$
- From less to most preferred, proposed “n” values are:
  1. TRUNK\_PORT: the entry was learnt from a trunk port (connected to another switch)
  2. ACCESS\_PORT: the entry was learnt from an access port (connected to a host)
  3. LLA\_MAC\_MATCH: LLA (found in NDP option) and MAC (found at layer2) are identical
  4. TRUSTED\_PORT: The entry was learnt from a trusted port
  5. TRUSTED\_TRUNK: The entry was learnt from a trusted trunk
  6. DHCP\_ASSIGNED: the entry is assigned by DHCP
  7. CGA\_AUTHENTICATED: The entry is CGA authenticated
  8. CERT\_AUTHENTICATED: the entry is authenticated with a certificate
  9. STATIC: this is a statically configured entry

# Preference algorithm

- B. Define the rules (applied in this order). Updating an entry attribute is:
1. Allowed when the plevel carried by the change is bigger than the plevel stored in the entry.
  2. Denied if the plevel carried by the change is smaller than the plevel stored in the entry
  3. Allowed if plevel  $\geq$  TRUSTED\_PORT
  4. Denied for entry in state INCOMPLETE if the change is not associated with the port where this entry was first learnt from.
  5. Allowed if the change is associated with a trusted port (and plevel are equal)
  6. Denied if the entry is in state REACHABLE or VERIFY (and plevel are equal)
  7. Allowed otherwise

# State machine

- The state of an entry is part of the collision algorithm
- For instance, entries in REACHABLE are “locked” to their attributes (movement disallowed unless  $\text{preflevel} \geq \text{TRUSTED\_PORT}$ )



Key: address/zoneid

Attributes: port, preflevel, owner, Ila

Others: state