IETF 77

Coexistence of Address Assignment Methods or

HOW TO DEAL WITH BINDING COLLISIONS in an HETEROGENEOUS ENVIRONEMENT?

What is a binding collision?

- Entry [*IP Address, vlan*, anchor] exists in the binding table
- Collision happens when a candidate entry with same key [*IP Address, vlan]* and anchor' ≠ anchor is « discovered »

→How to choose one over the other?

FCFS? Discovery method? Best credentials? …?

What is an heteroneous environment?

- Different discovery methods (NDP, DHCP, data, Static, etc.)
- Different credentials carried by messages used by the various methods
- Different origins for messages used by various methods

→ In real world, no one-fits-all discovery method, credentials, origins.

Variety of methods for discovering bindings

- DHCP-snooping
- NDP snooping
- Data snooping
- Statically created
- « Local » to the switch (L2/L3)
- …

→Collisions within one method is usually well-understood/defined (FCFS, LCFS, etc.)

→Collisions between two methods is TBD

Variety of credentials carried by messages (and relatives) used for the discovery

- No credentials
- Consistent SMAC & Layer link-layer address
- Cryptographically proven
- Certificate proven
- EAP proven

Variety of origins for messages used for the discovery

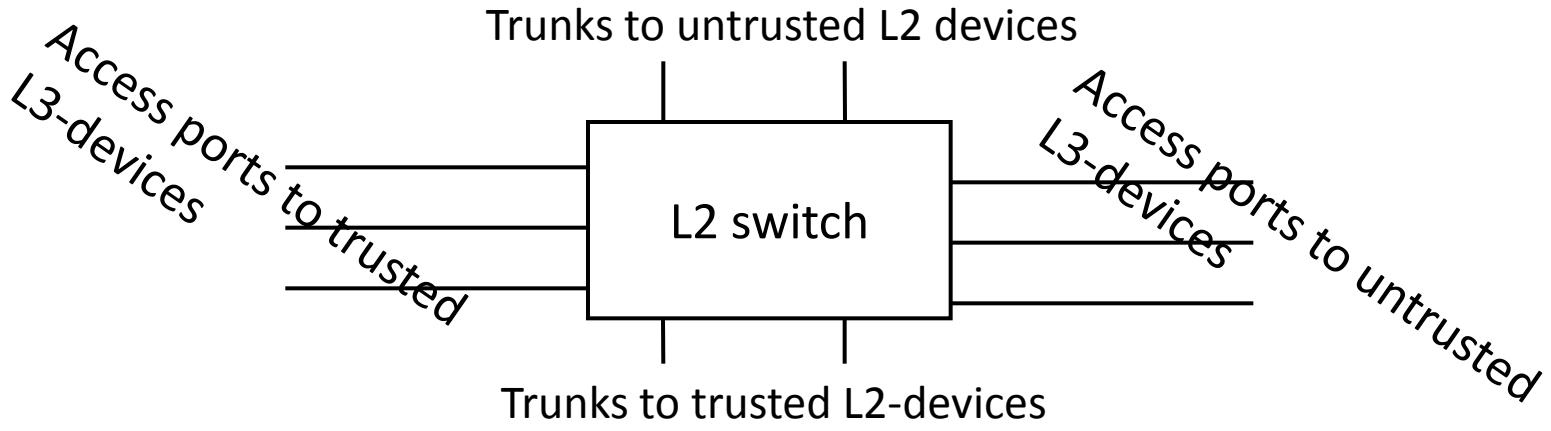

How to compile all variables?
How to compare different sets?

→DHCP-discovered vs NDP with CGA?

→Static entry vs DHCP-discovered

→NDP on trusted access vs DHCP on untrusted access

→ …

Preference level

A. We define preference “factors” , preference value and preference level:
   - A “factor” is associated with
     - a property of the port from which the entry was discovered
     - a property of the discovery method
     - or a property of the binding itself
   - Each factor is given a number $0 \leq f \leq n$: the bigger, the more prevalent
   - We compute the preference value of a factor as 2^{f}
   - We compute Preflevel = ∑preference_values associated with a binding

Factors

From least to most prevalent, proposed factor values /preference values are:

-	/ 0.	NDP-SNOOPING:	The entry was learnt by snooping NDP traffic (DAD, etc.)
0	/ 1.	LLA_MAC_MATCH:	LLA (found at L3) and MAC (found at L2) are identical
1	/ 2.	TRUNK_PORT:	The entry was learnt from a trunk port (connected to another switch)
2	/ 4.	ACCESS_PORT:	The entry was learnt from an access port (connected to a host)
3	/ 8.	TRUSTED_PORT:	The entry was learnt from a trusted port
4	/ 10.	TRUSTED_TRUNK:	The entry was learnt from a trusted trunk
5	/ 20.	DHCP_SNOOPING:	The entry is assigned by DHCP
6	/ 40.	CGA_AUTHENTICATED:	The entry is CGA authenticated
7	/ 80.	EAP_AUTHENTICATED:	The entry is EAP authenticated
8	/100.	CERT_AUTHENTICATED:	The entry is authenticated with a certificate
10	/200.	STATIC:	this is a operator configured entry (static or local)

Example

```
Binding Table has 3 entries, 3 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, DHC – DHCP

Preflevel flags (prlvl):

0001:MAC and LLA match     0002:Orig trunk            0004:Orig access
0008:Orig trusted access   0010:Orig trusted trunk    0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated    0100:EAP authenticated
0200:Operator assigned


      IPv6 address                          Link-Layer Adr   Interface   vlan   prlvl

ND    FE80::3C99:78CB:3EDC:47F7             AABB.CC01.F500   Et0/0       100    0045
ND    FE80::A8BB:CCFF:FE01:F600             AABB.CC01.F600   Et1/0       100    0005
ND    FE80::A8BB:CCFF:FE01:F700             AABB.CC01.F700   Et2/0       100    0005
ND    FE80::A8BB:CCFF:FE01:F800             AABB.CC01.F800   Et3/0       100    0003

ND    2001:DB8::3008:BC73:6873:F128         AABB.CC01.F500   Et0/0       100    0045
DHC   2001:DB8::F981:4906:29FB:78B5         AABB.CC01.F600   Et1/0       100    0024
S     2001:DB8::1                           AABB.CC01.F700   Et2/0       100    0200
ND    2001:DB8::BC10:1361:4712:AC5E         AABB.CC01.F800   Et3/0       100    0003
L     2001:DB8::2                           AABB.CC01.F100   SVI100      100    0200
```

Preference algorithm

B. Define the rules (applied in this order). Updating an entry attribute is:

1. Allowed, if no entry exist
2. Denied if existing entry is more prefered (with higher preflevel)
3. Allowed if existing entry is less prefered (with smaller preflevel)
4. Allowed, if received candidate on a trusted port
5. Denied if existing entry respond to pool (DAD NS)
6. Allowed otherwise

What's next?

- Current document is draft-levy-abegnoli-savi-plbt-02.txt
- One implementation …
- -01 reviewed/commented by 2 or 3 people
- What to do with this work?

  - Merge with « a » framework WG document?
  - Make it part of one of the existing WG?
  - Make it a separate WG document?
  - ?