LLDP in a topology with transparent links

Paul Congdon

Norm Finn

LLDP usage beyond discovery

- The exchange of certain LLDP TLVs are useful for diagnostic operations and network management decisions
  - For Example: Duplex miss-match can be detected from the 802.3 MAC/PHY Configuration/Status TLV
- The diagnostic operations likely assume the LLDP peers are connected to the same physical medium
- Devices that transparently forward LLDP frames break this assumption.
- We are standardizing devices that transparently forward LLDP frames
  - TPMRs
  - Provider bridges

Scenario

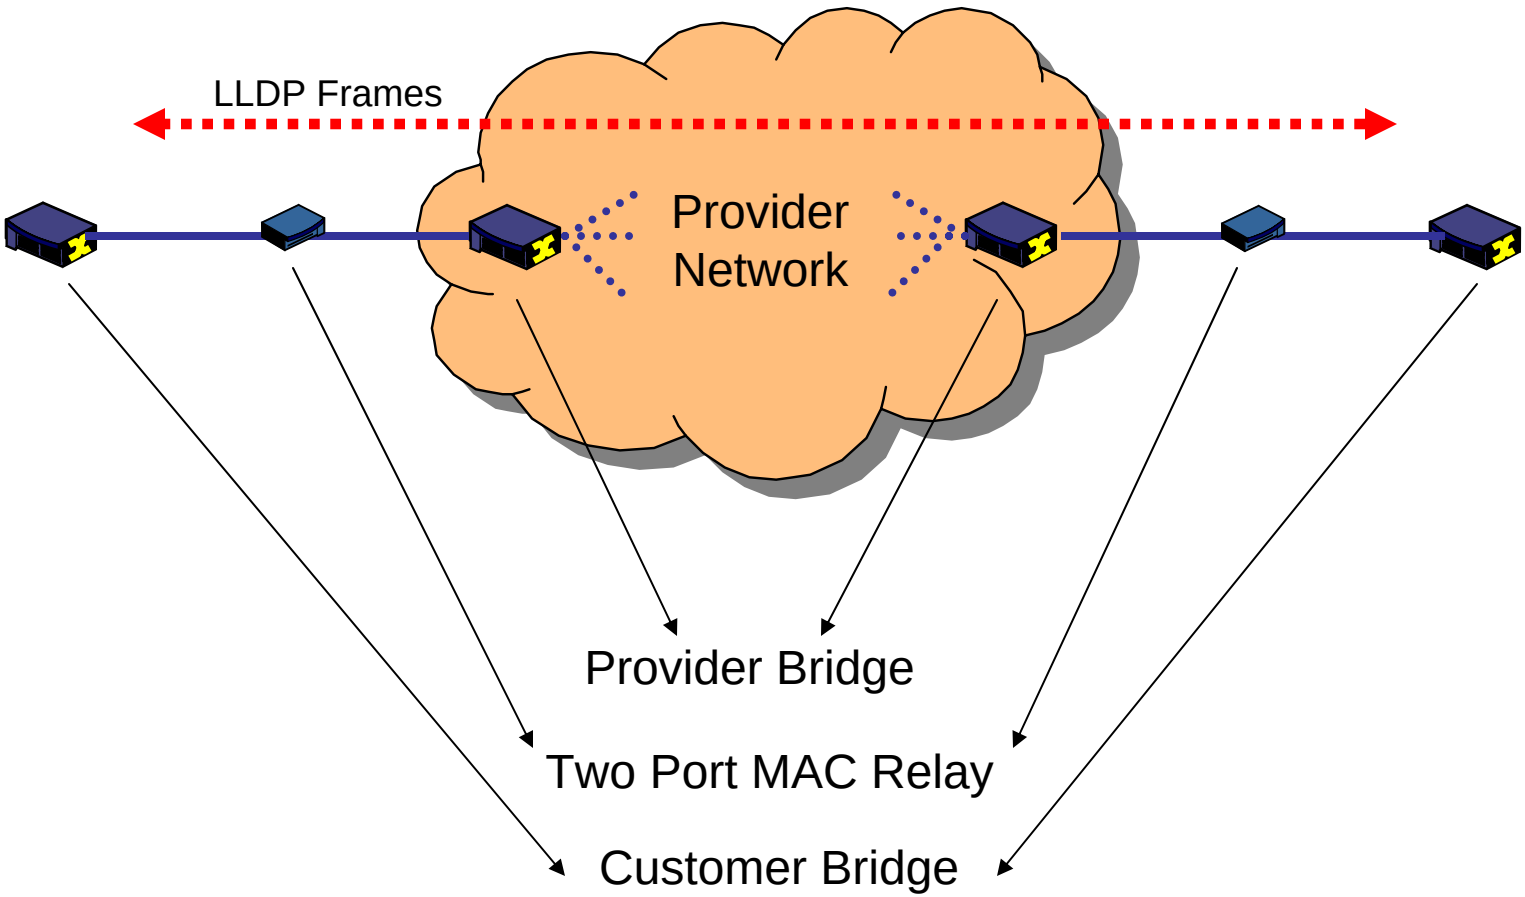
LLDP Frames
Provider
Network
Provider Bridge
Two Port MAC Relay
Customer Bridge

Known TLVs and their Sensitivity

TLV	Specification	Issue Sensitivity
End Of LLDPDU	802.1AB	N/A
Chassis ID	802.1AB	N/A
Port ID	802.1AB	N/A
Time To Live	802.1AB	N/A
Port Description	802.1AB	N/A
System Name	802.1AB	N/A
System Description	802.1AB	N/A
System Capabilities	802.1AB	N/A
Management Address	802.1AB	N/A
Port VLAN ID	802.1 (Annex F)	N/A
Port and Protocol VLAN ID	802.1 (Annex F)	N/A
VLAN Name	802.1 (Annex F)	N/A
Protocol Identity	802.1 (Annex F)	N/A
MAC/PHY Configuration/Status	802.3 (Annex G)	Problematic
Power Via MDI	802.3 (Annex G)	Problematic
Link Aggregation	802.3 (Annex G)	N/A
Maximum Frame Size	802.3 (Annex G)	??????????
LLDP-MED Capabilities	LLDP-MED	N/A
Network Policy	LLDP-MED	N/A
Location Identification	LLDP-MED	N/A
Extended Power via MDI	LLDP-MED	Problematic
Inventory (multiple)	LLDP-MED	N/A

Examples of what are we trying to discover…

1. What are my neighbor's auto-negotiation parameters?
2. What is my neighbor's power level information?
3. What is my neighbor's configured default VLAN?
4. Which VLAN has my neighbor, a bridge or router, configured to best carry voice traffic?
5. Who is my neighbor? (Customer Perspective)
6. Who is my neighbor? (Provider Perspective)

Objectives/Goals/Observations

Objectives

- Provide a reliable means to determine when information received in LLDP frames can assume to be associated with the same physical medium as that of the transmitter

Goals

- Work within 802.1 framework and support existing and new specifications and functions
- Minimize impact to existing implementations and current specifications wherever possible

Observations

- Sending devices should have to understand the topology to their expected destinations, but they know what type of devices should forward the information they send.
- Forwarding devices know their function and if they know when they ‘relay’ LLDP frames they can inform the receiver via intermediate marking

Possible Solutions to Consider

- Do not allow forwarding of LLDP frames (period).
- Define multiple addresses for LLDP frames (e.g. physical only (stopped by TPMR), provider level (stopped by all bridges), and customer level (transparent to provider bridge, but not to customer bridge))
- Define that forwarding devices ‘mark’ the LLDP frame when forwarded.
- Mandate that forwarding devices also send LLDP frames such that end-points can detect their presence and invalidate assumptions
- Others?

The ‘do not forward’ Solution

- Pros
  - Avoids the problem for ‘standard’ devices
- Cons
  - Restricts desired functionality
  - For end-to-end topology view it places full LLDP burden on forwarding devices
  - Many non-standard forwarding devices out there (e.g. IP-Phones)
  - We already specify forwarding in 802.1ad, would need to back that out

Use multiple addresses

- LLDP stations would use multiple addresses and send multiple PDUs (e.g. physical only (stopped by TPMR), provider level (stopped by all bridges), and customer level (transparent to provider bridge, but not to customer bridge)).
- Pros
  - Attributes that make assumptions about topology are put in appropriate PDU. Only the sender knows!
- Cons
  - Multiple instances of the protocol now running
  - Physical link constrained protocol must run on uncontrolled port only since forwarding devices are not participating in MACSec (maybe true for all cases anyway!)
  - Uses up more address space; are we sure this is the only set of forwarding constraints?

The ‘Mark’ing Solution

- Devices that are to forward LLDP frames somehow ‘mark’ that the frame has been relayed
- Possible ways to ‘mark’ the LLDP frame
  - Modify the destination multicast address
  - Modify the Ethertype
  - Insert or modify a TLV

NOTE: None of these work with MACSec is in operation!

Mandate that forwarding devices participate in LLDP

- All transparent forwarding devices will also participate in LLDP (at least as transmitters only). Receiving devices will act upon assumptions accordingly.
- Pros
  - Not much new to define
- Cons
  - Many non-standard devices out there
  - Links LLDP received information from multiple parties and independent PDUs (currently a no-no in the spec)
  - LLDP is not a reliable protocol, assumption linkage weak
  - Some LLDP PDUs may be on controlled port others on uncontrolled port.

Summary/Conclusions

- Seems as though marking solution is DOA because of MACSec and desire to keep TPMRs simple.
- Running LLDP on the uncontrolled port for physically constrained PDUs seems necessary.
- It would be beneficial for all transparent forwarding devices to participate in LLDP (at least as transmitters only)
- Multiple address solution may be extendable to other similar problems (LACP, 802.1X)
  - Assuming multiple address solution: Do we tell the world to change to a different address all of the existing applications that use the physical TLVs, or do we change to a different address for the new applications that might want to worry about whether a provider bridge is present or not?

Back-up Slides

Consideration: Multicast address or Ethertype modification

- TPMRs and Provider Bridges swap the destination MAC multicast address before encapsulating, tagging and forwarding the frame
- Pros
  - Doesn’t change the contents of the PDU
  - Once changed, downstream bridges may be able to forward using the bridge relay
- Cons
  - Burns address and ethertype space
  - Requires forwarding definitions for new addresses
  - Requires LLDP receiver to listen on new addresses

Consideration: Inserting and Modifying a new ‘hop-count’ TLV

- TPMRs and Provider Bridges either insert or modify a new TLV that indicates the LLDP frame was forwarded
- Pros
  - Supports an extensible way to document the number of invisible forwarding devices on the path
- Cons
  - YATLV = yet another TLV. PDU space already a concern
  - Required packet growth and/or modification inside of the PDU
  - Could require current sending station modifications