IEEE 802.3da SPMD: LLDP for MPoE proposal

Peter Jones, Cisco Systems
January 2025

1 Overview

1.1 Goals: Reporting and controls for MPoE systems

- Leverage previous work on PoE TLVs, Clause 79.3.2 and 79.3.8 (reuse or redefine)
- Minimum needed.
    - MPD and MPSE status advertisement.
    - MPD request for power and MPSE allocate power.
- Telemetry
    - Reuse 79.3.8 Power via MDI Measurements TLV

      OR

    - Define new Clause 30 objects to:
        - Report measurement capabilities
        - Trigger measurement actions
        - Report measurement results

1.2 Table of Contents

Contents

2 LLDP Basics

LLDP is specified in 802.1AB-2016 as amended by 802.1ABdh-2021.

Using LLDP for MPoE management will be optional for 802.3da, but without LLDP each node will be limited to a single “unit load”.

2.1 Destination Addressing

802.1AB Clause 7.1 allows for several different destination addresses:

- Group addresses
  - Nearest bridge
  - Nearest non-TPMR bridge
  - Nearest Customer Bridge
  - Any group MAC address
- Individual addresses
  - Any individual MAC address

2.1.1 802.3da Destination Addressing

802.3 10BASE-T1S/10BASE-T1M uses the “Nearest bridge” address. This is described as “Propagation constrained to a single physical link; stopped by all types of bridge.”. This ensures that these TLVs are confined to a single mixing segment.

2.2 Frame Transmission

Frame Transmission is specified in 802.1AB Clause 9.1.1. There are several conditions which prompt transmission:

- Periodic background transmission
  - 802.1AB Clause 9.1.1. The default value (LLDP MIB module lldpV2MessageTxInterval) is 30 seconds.
- New neighbor
  - This triggers 4 rapid transmissions to get the new neighbor up to date using the normal group address.
- Updated local information.
  - Transmission is triggered “immediately”, with a credit-based scheme to throttle transmissions if state is changing rapidly.

Note, for shared media LANs, the delay for the periodic background transmission includes allowance for a “jitter” component to avoid all nodes transmitting at the same time, see clause 9.2.2.

2.2.1 802.3da Frame Transmission

802.3da systems implementing LLDP will use the following transmission triggers:

- Periodic background transmission
- Triggered transmission
    - New neighbor(s)
    - Updated local information, e.g., MPSE power allocation map.
    - Triggered transmissions are delayed by 0.5 seconds to allow multiple updates to be combined into a single transmission.

3 LLDP operation for MPoE

3.1 MPSE

The MPSE Status TLV (4.2.1) includes:

- Capabilities, e.g., Type 0 (30V Max) Supported
- Status, e.g., MPSE Allocated Power
- Notifications, e.g., Power Down Notification

3.1.1 Power allocation

MPSEs advertise their power capacity, and their total allocated power in the MPSE Status TLV (4.2.1).

MPDs request power from an MPSE using the “MPD required power”, “MPD desired power” and “MPD desired power duration” fields in the MPD Status TLV (4.3.1). The MPSE collates all the requests and determines how the power to assign to each MPD. The budgeting/allocation function is outside the scope of the standard.

The MPSE maintains a table of the MPDs on the segment and their power requests and allocations, including MPDs not currently drawing power (e.g., asleep) based on based on the Power request withdrawal notification in 4.3.1.

The MPSE power budgeting goals include:

- The segment has sufficient power available (e.g., un-allocated) to allow new nodes to boot up and request additional power.
- The segment has power available for “sleeping” nodes when they wake up.
- “Desired power” is allocated to nodes based on their power priority and the expected schedule of nodes waking up.

3.1.2 Power Down Timer

If the MPSE knows it is going to stop providing power to the port, it sends out a status TLV with a non-zero duration to power down to let MPDs prepare to lose power.

3.2 MPD

The MPD Status TLV (4.2.1) includes:

- Capabilities, e.g., Type 0 (30V Max) supported
- Requests, e.g., MPD required power
- Notifications, e.g., Withdrawing Power notification
- Measurements, e.g., Voltage measurement

3.2.1 Requesting Power

When an MPD first powers up, it shall consume a maximum of one “unit load” (802.3da draft 2.0 clause 189.3). It may request additional power using the power request fields (“MPD required power”, “MPD desired power”, “MPD desired power duration”) in the MPS Status TLV (4.3.1)

An active MPSE receiving the power request performs a budgeting/allocation function (outside the scope of the standard) and determines the power to allocate to the MPD. If the MPD is allocated additional power, this change of local state triggers transmission of the updated MPSE power allocation table for the mixing segment (4.2.2).

At any time, the MPD can change its requested allocation by changing the power request fields and triggering an LLDP transmission. If this results in a change in the MPSE power allocation table, it triggers transmission of the updated table to the members of the mixing segment.

3.2.2 Sleep and Shutdown

If the MPD is about to sleep or shut down permanently, it can inform the MPSE using the “Delay until power draw ceases” and “Power request withdrawal” fields of the MPS Status TLV.

Knowing that MPD will stop consuming power ahead of time enables the MPSE to rapidly release the MPDs power allocation and potentially grant it to another MPD.

If the MPD is going to sleep for a defined period, the MPSE can take the sleep period into account in its power allocation function.

4 LLDP TLVs for MPoE

4.1 Common Information Elements

- Type – indicates system power type (30V vs 50V)
- Power – units 0.1 W
- Voltage - units of 1 mV
- Current - units of 0.1 mA
- Energy - units of kJ
- Time – seconds or microseconds

4.2 MPSE

4.2.1 MPSE Status TLV

Field size	Bit	Function	Units	Value/meaning
16	0	MPSE Active		1 = active 0 = inactive
	1	Active MPSE Type		1 = Type 1 0 = Type 0
	2	Type 0 (30V Max) Supported		1 = supported 0 = unsupported
	3	Type 1 (50V Max) Supported		1 = supported 0 = unsupported
	4	Power Measurements (see 4.4.2)		1 = supported 0 = unsupported
	5	Withdrawing Power Notification		1 = active. 0 = inactive
	15:xx	Reserved		
16		MPSE Max Power	0.1 W	
16		MPSE Allocated Power	0.1 W	
8		Withdrawing Power Delay	seconds	Time delay before MPSE withdraws power
8		Reserved		

4.2.2 MPSE Power Allocated TLV

The “MPD grant entries” shall be sorted numerically by MAC address.

Fixed component

Field size	Bit	Function	Units	Value/meaning
8		Entry Count		Number of MPD grant entries
8		Reserved		

MPD grant entries.

Field size	Bit	Function	Units	Value/meaning
48		MPD MAC address		MPD MAC
8		MPD granted power	0.1 W	Power the MPD is permitted to draw.
8		MPD desired power	0.1 W	Power the MPD would like to draw.
8		MPD desired power duration	seconds	Duration the MPD asked for the desired power, 0 represents indefinite.
8		Reserved		

4.3 MPD

4.3.1 MPD Status TLV

Field size	Bit	Function	Units	Value/meaning
16	0	Type 0 (30V Max)		1 = supported 0 = unsupported
	1	Type 1 (50V Max)		1 = supported 0 = unsupported
	4:2	Power priority		0 = highest 7 = lowest
	5	Power request withdrawal notification		1 = active. 0 = inactive
	6	Voltage Measurement		1 = supported 0 = unsupported
	7	Power Measurements (see 4.4.2)		1 = supported 0 = unsupported
	15:xx	Reserved		
8		Required power	0.1 W	
8		Desired power	0.1 W	
8		Desired power duration	seconds	Duration for desired power, 0 means indefinite.

8		Power request withdrawal	seconds	Power request withdrawal duration, 0 indicates indefinite.
8		Delay until power draw ceases	seconds	Delay before MPD stops drawing power, 0 represents no scheduled transition
8		Reserved		
16		Voltage measurement	0.1V	

4.4 Measurements/Telemetry

We need to address the integration period for the measurements to ensure that we get "useful" and consistent data.

4.4.1 Power via MDI Measurements TLV

Editorial commentary

We could use this TLV as it is currently defined, and it's Clause 30 objects (aLldpXdot3LocMeasVoltageSupport to aLldpXdot3LocEnergyMeasurement, aLldpXdot3RemMeasVoltageSupport to LldpXdot3RemEnergyMeasurement), but that assumes the key consumer of this information is one of the nodes on the segment, or that PDs do not support a management interface. I think the key consumer is not local, and this should be Clause 30 management function. Using that technique, a management station can trigger the measurement on an interval of its choosing.

Apart from using voltage drop to estimate cable loss, what would an MPSE do with this information from the set of MPDs on the wire? To communicate voltage, drop to the MPSE, I've added "Voltage Measurement" attributes to the MPD Status TLV.

4.4.2 Clause 30 Measurement Proposal

I propose defining new Clause 30 objects for Power via MDI Measurement and not supporting the Power via MDI Measurements TLV for MPoE or the corresponding LLDP objects in Clause 30, 802.3.1 and 802.3.2. These objects should work equally well for Clause 145 PoE as well as MPoE.

4.4.2.1 Additions to MPSE and MPD managed object classes

Attribute	Type	Bit #	Function	79–21 Bit #	Units	Value/meaning
Measurement Capabilities	Bit String	0	Measurement Supported	N/A??		1 = supported 0 = unsupported
		1	Voltage support	159		1 = supported 0 = unsupported
		2	Current support	158		1 = supported 0 = unsupported
		3	Power support	157		1 = supported 0 = unsupported
		4	Energy support	156		1 = supported 0 = unsupported
		15:xx	Reserved			

Voltage accuracy	uint8			143:128	1 mV	
Current accuracy	uint8			127:112	0.1 mA	
Power accuracy	uint8			111:96	10 mW	
Energy accuracy	uint8			95:80	0.1 kJ	
Measurement Active	Enum					1 = active 0 = idle

4.4.2.2 Additions to MPSE and MPD action classes

Attribute	Type	Bit #	Function	79–21 Bit #	Units	Value/meaning
Perform Measurement	Enum			N/A		1 = Start 0 = idle

4.4.2.3 New “Measurement Results” class

Attribute	Type	Bit #	Function	79–21 Bit #	Units	Value/meaning
Valid	Bit String	0	Measurement valid	N/A		1 = valid 0 = invalid
		1	Voltage valid	147		1 = valid 0 = invalid
		2	Current valid	146		1 = valid 0 = invalid
		3	Power valid	145		1 = valid 0 = invalid
		4	Energy valid	144		1 = valid 0 = invalid
		5	Age valid			1 = valid 0 = invalid
		6	Duration valid			1 = valid 0 = invalid
		15:xx	Reserved			
Voltage	uint16			79:64	1 mV	
Current	uint16			63:48	0.1 mA	
Power	uint16			47:32	10 mW	
Energy	uint32			31:0	kJ	Energy consumed since last measurement
Age	uint16		Seconds since measurement was performed.		seconds	
Duration	uint16				Micro-seconds	Measurement period length

END OF DOCUMENT