IEEE 802.3da SPMD:

LLDP for MPoE proposal

Peter Jones, Cisco Systems

February 2025

1 Overview

1.1 Goals: Reporting and controls for MPoE systems

- Leverage previous work on PoE TLVs, Clause 79.3.2
- Minimum needed.
    - MPD and MPSE status advertisement.
    - MPD request for power and MPSE allocate power.

1.2 Change log

- 1/6/2025
    - submitted for 802.3da D2.0 comment resolution
- 1/22/25
    - Removed Measurements/Telemetry section to separate document.
- 2/5/2025
    - Update definitions MPD voltage and low voltage reporting.

1.3 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 Clause 9.1.1. There are several conditions which prompt transmission:

- Periodic background transmission
  - 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 (**Error! Reference source not found.**) 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 and budgeting

MPSEs may advertise their power capacity, and their total allocated power in the MPSE Status TLV (4.2.2).

MPDs may request changes in their power allocation using the MPD normal power, and temporary power fields in the MPD Status TLV (4.2.4). 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 may maintain 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 temporary power fields in 4.2.3.

The MPSE power budgeting goals may 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.
- Requests for additional “temporary power” are processed considering node power priority and temporary power requests from other nodes on the mixing segment.

3.1.2 Withdrawing Power

If the MPSE knows it is going to stop providing power to the port, it may send a status TLV using the Withdrawing Power Notification and Withdrawing Power delay to let MPDs prepare to lose power.

3.2 MPDs

3.2.1 MPD power allocation

3.2.1.1 Static vs Dynamic power allocation

MPDs are permitted to consume power based on their “worst case” static allocation when they boot up. When designing/installing/modifying the mixing segment, the following elements need to be evaluated:

- the static allocations of all nodes
- the maximum power delivery of the mixing segment
- the capacity of the MPSE to provide power to the mixing segment.

MPDs may support additional functions for dynamic power budgeting/management using LLDP. It’s recommended that a mixing segment use either static power management, or dynamic power management on a given mixing segment.

3.2.1.2 Static vs Normal power allocation

An MPD may have a higher “Static” power than it needs to perform its normal function. The MPD may signal its “Normal” power to the MPSE to enable better power management.

3.2.1.3 Temporary power allocation

An MPD may request a “Temporary” power allocation for a given duration. This allocation may be larger than its “Normal” power to enable additional work to be performed, or lower than its “Normal” power, e.g., sleeping. Signaling this to the MPSE enables better power management on the mixing segment.

3.2.2 MPD Status using LLDP

The MPD Status TLV (**Error! Reference source not found.**) may be used to inform the MPSE about the MPD. It includes:

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

3.2.3 Changing Power Allocation using LLDP

An active MPSE receiving a power request may perform a budgeting/allocation function (outside the scope of the standard) to determine the power to allocate to the MPD.

An active MPSE may change power allocations based on local factors, e.g., losing a power supply.

MPD power allocation changes trigger transmission of the updated MPSE power allocation table for the mixing segment (**Error! Reference source not found.**).

3.2.4 Sleep and Shutdown using LLDP

If the MPD is about to sleep or shut down permanently, it may inform the MPSE using the “Temporary power” fields in the MPD Status TLV (4.2.3)

Knowing when the MPD will change its power requirements enables the MPSE to better manage power allocation.

3.3 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 Clause 79 Proposed Changes

4.1 Formatting notes

Current clause 79 TLV definitions include a figure showing the format of the LLDP TLVs, e.g.,

Figure 79–2 shows the format of this TLV.

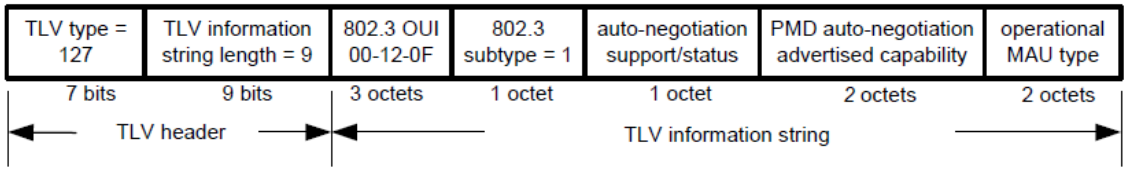

Figure 79–2—MAC/PHY configuration/status TLV format

I propose to use tables to show this information. They are easier to create/read/modify.

Add after Table 79–1—IEEE 802.3 Organizationally Specific TLVs

Table 79–2—Common TLV elements for 802.3

Field	Field size (bits)	Value	Notes
TLV Type	7	127	Organizationally Specific TLV
TLV information string length	9	TLV dependent	length of the information string in octets.
OUI	24	00-12-0F	IEEE 802.3
subtype	8	TLV dependent	See Table 79–1—IEEE 802.3 Organizationally Specific TLVs

As an example, to replace Figure 79–2—MAC/PHY configuration/status TLV format I would use the following table:

Table 79–3— MAC/PHY configuration/status TLV elements

Field	Field size (bits)	Value	Notes
Auto-negotiation support/status	8		See 79.3.1.1
PMD auto-negotiation advertised capability field	16		See 79.3.1.2
Operational MAU type	16		See 79.3.1.3

4.2 New MPoE TLVs

4.2.1 New subtypes

Change Table 79–1—IEEE 802.3 Organizationally Specific TLVs by adding the following rows and updating the “Reserved” subtype range to match:

IEEE 802.3 subtype	**TLV name**	Subclause reference
9	MPoE MPSE Status TLV	78.3.10
10	MPoE MPD Status TLV	78.3.11
11	MPoE MPSE Power Allocated TLV	78.3.12

4.2.2 MPoE MPSE Status TLV

Add subclause 78.3.10 with the following text:

78.3.10 MPoE MPSE Status TLV

Clause 189 Multidrop Power over Ethernet (MPoE) defines two optional power entities: a Multidrop Power Sourcing Equipment (MPSE) and a Multidrop Powered Device (MPD). These entities allow devices to draw/supply power from/to mixing segment defined in 188.8 using a Multidrop Power Interface (MPI).

The MPoE MPSE Status TLV allows MPSEs to advertise capabilities and status to other MPoE devices on the mixing segment. Table 79.N shows the format of the TLV.

Table 79–N—MPSE Status TLV elements

Field	Field size (bits)	Subclause reference	Notes
Capabilities and status	16	78.3.10.1	
MPSE Max Power	16	78.3.10.2	
MPSE Allocated Power	16	78.3.10.3	
Withdrawing Power Delay	8	78.3.10.4	
Reserved	8	78.3.10.5	

78.3.10.1 MPSE Capabilities and Status.

Field width	Bit	Function	Units	Value/meaning
16	0	MPSE Active		1 = active 0 = inactive
	1	Active MPSE Type		1 = Type 1 (50V Max) 0 = Type 0 (30V Max)
	2	Type 0 (30V Max) Supported		1 = supported 0 = unsupported
	3	Type 1 (50V Max) Supported		1 = supported 0 = unsupported
	4	Withdrawing Power Notification		1 = active. 0 = inactive
	15:*nn*	Reserved		

78.3.10.2 MPSE Max Power

Field width	Bit	Function	Units	Value/meaning
16		MPSE maximum power	0.1 W	Maximum power the MPSE can supply to the mixing segment.

78.3.10.3 MPSE Allocated Power

Field width	Bit	Function	Units	Value/meaning
16		MPSE allocated power	0.1 W	Power the MPSE has allocated for the mixing segment.

78.3.10.4 Withdrawing Power Delay

Field width	Bit	Function	Units	Value/meaning
16		Withdrawing Power Delay	secs	Seconds until the MPSE will stop providing power to the mixing segment. Ignored unless the “Withdrawing Power Notification” flag is set.

An MPSE can use the “Withdrawing Power Delay” in conjunction with the “Withdrawing Power Notification” flag to notify the other MPoE devices on the mixing segment that it will stop providing power after the specified period.

4.2.3 MPoE MPD Status TLV

Add subclause 78.3.11 with the following text:

78.3.11 MPoE MPD Status TLV

Clause 189 Multidrop Power over Ethernet (MPoE) defines two optional power entities: a Multidrop Power Sourcing Equipment (MPSE) and a Multidrop Powered Device (MPD). These entities allow devices to draw/supply power from/to mixing segment defined in 188.8 using a Multidrop Power Interface (MPI).

The MPoE MPD Status TLV allows MPDs to advertise to advertise capabilities, status and requests to other MPoE devices on the mixing segment. Table 79.N shows the format of the TLV.

Table 79–N— MPD Status TLV elements

Field	Field size (bits)	Subclause reference	Notes
Capabilities and status	16	78.3.11.1	
Static power	16	78.3.11.2	
Normal power	16	78.3.11.3	
Temporary power	16	78.3.11.4	
Temporary power duration	16	78.3.11.5	
Temporary power delay	8	78.3.11.6	
Reserved	8		16-bit alignment
Instantaneous Voltage	16	78.3.11.7	
Low Voltage	16	78.3.11.8	

78.3.11.1 MPD Capabilities and Status

Field width	Bit	Function	Units	Value/meaning
16	*0*	Type 0 (30V Max) MPD		1 = supported 0 = unsupported
	1	Type 1 (50V Max) MPD		1 = supported 0 = unsupported
	2	Active type		1 = Type 1 0 = Type 0
	3	Requested power priority flag		1 = Requested power priority field valid 0 = Requested power priority field invalid
	4	Temporary power notification		1 = active. 0 = inactive
	5:7	Requested power priority		0 = highest 7 = lowest
	8	Voltage monitoring		1 = supported 0 = unsupported
	15:*nn*	Reserved		

78.3.11.2 MPD Static Power Announcement

Field width	Bit	Function	Units	Value/meaning
16		Static power	0.1 W	Power the MPD draws before MPoE power negotiation.

78.3.11.3 MPD Normal Power Announcement

Field width	Bit	Function	Units	Value/meaning
16		Normal power	0.1 W	Power the MPD needs to support its normal function (may be less than static power).

78.3.11.4 MPD Temporary Power Request

Field width	Bit	Function	Units	Value/meaning
16		Temporary power	0.1 W	Power the MPD needs for a defined duration (may be less than normal power including 0 W for a sleeping device). Ignored if “Temporary power notification” is not set.

An MPD can use Temporary power/delay/duration MPSE change its allocation for a specified duration.

78.3.11.5 MPD Temporary Power Request Duration

Field width	Bit	Function	Units	Value/meaning
16		Temporary power duration	secs	Duration of the MPD’s intended power draw change. 0 means infinite. Ignored if “Temporary power notification” is not set.

78.3.11.6 MPD Temporary Power Request Delay

Field width	Bit	Function	Units	Value/meaning
16		Temporary power delay	secs	Delay before the MPD intends to change its power draw. Ignored if “Temporary power notification” is not set.

78.3.11.7 MPD Instantaneous Voltage

Field width	Bit	Function	Units	Value/meaning
16		Instantaneous voltage	1 mV	The instantaneous voltage observed by the MPD. Ignored if “Voltage monitoring” is not set.

78.3.11.8 MPD Low Voltage

Field width	Bit	Function	Units	Value/meaning
32		Low voltage event counter		The number of low voltage events observed by the MPD. A low voltage event is when the MPD state diagram (Figure 189–8) transitions from “PON_LOAD_ON” to “PNE_NO_POWER”.

4.2.4 MPSE Power Allocated TLV

Add subclause 78.3.12 with the following text:

78.3.12 MPoE MPSE Power Allocated TLV

The MPoE MPSE Power Allocated TLV allows MPSEs to advertise power allocation information to other MPoE devices on the mixing segment. Table 79.N and 79.N show the format of the TLV.

Table 79–N—MPSE Power Allocation TLV fixed elements

Field width	Bit	Function	Units	Value/meaning
8		Entry Count		The number of power allocation entries.
8		Reserved		16-bit alignment.

Table 79–N—MPSE Power Allocated TLV power allocation entries

Field width	Bit	Function	Units	Value/meaning
48		MPD MAC address		MPD MAC.
16		MPD granted power	0.1 W	Power the MPD is allocated.
16		MPD static power announcement	0.1 W	See 78.3.11.2.
16		MPD normal power announcement	0.1 W	See 78.3.11.3.
16		MPD temporary power request	0.1 W	See 78.3.11.4.
16		MPD temporary power duration request	seconds	See 78.3.11.5.
16		MPD temporary power delay request	seconds	See 78.3.11.6.

END OF DOCUMENT