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.
- 2/19/25
    - Updates during Management Ad Hoc Meeting 19 February 2025
    - Major items
        - Change Type 0/Type 1 reporting to bit map for future extensibility.
        - Added PICS
        - Harmonized with 802.3da D2.1
    - Submitted with D2.1 comments
- 2/21/25
    - Incorporate feedback
    - Align to 802.1AB-2016
    - Add “default” usage rule of only one 802.3 TLV per LLDPPDU.
    - Re-submitted with D2.1 comments

1.3 Open Items

- Relationship between data pair and power pair(s) as per “189.1.2 Relationship of MPoE to the IEEE 802.3 architecture”

1.4 Table of Contents

Contents

2 LLDP Basics

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

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 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

The MPSE Status TLV includes power capacity and total allocated power.

MPDs use the MPD Status TLV to request changes in their power allocation (i.e., normal and temporary power fields). 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.

An MPSE performing power management 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 temporary power fields in 4.3.3.

Examples of 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.
- 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 an MPSE knows it is going to stop providing power to the port, it can notify the MPDs the Withdrawing Power Notification and Withdrawing Power delay.

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 can 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 MPDs “Normal” power is less than or equal to it’s “Static” power The MPD can signal its “Normal” power to the MPSE to enable better power management.

3.2.1.3 Temporary power allocation

An MPD can request a “Temporary” power allocation for a given duration. This allocation could be higher or lower (including 0W) than it’s “Normal” power .

3.2.2 MPD Status using LLDP

The MPD Status TLV informs 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 can perform a budgeting/allocation function (outside the scope of the standard) to determine the power to allocate to the MPD.

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

Changes in MPD power allocation trigger transmission of the updated MPSE power allocation t .

3.2.4 Sleep and Shutdown using LLDP

If the MPD is about to sleep or shut down permanently, it can inform the MPSE using the “Temporary power” fields in the MPD Status TLV (4.3.3)

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

4 Clause 79 Proposed Text 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.

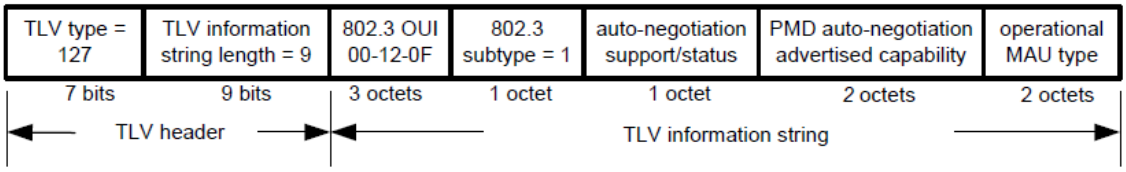

Figure 79–2—MAC/PHY configuration/status TLV format

Below I 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 Miscellaneous changes

4.2.1 Align to IEEE Std 802.1AB-2016

Replace the first paragraph of 79.1 with (changed text in ***bold italic***):

> The Link Layer Discovery Protocol (LLDP) specified in IEEE Std 802.1AB-***2016*** is a MAC Client protocol that allows stations attached to an IEEE 802 LAN to advertise to all other stations attached to the same IEEE 802 LAN: the major capabilities provided by the system incorporating that station, the management address or addresses of the entity or entities that provide management of those capabilities, and the identification of the station's point of attachment to the IEEE 802 LAN.

Replace the last paragraph of 79.1 with (changed text in ***bold italic***):

> Organizationally Specific TLVs can be defined by either the professional organizations or the individual vendors that are involved with the particular functionality being implemented within a system. The basic format and procedures for defining Organizationally Specific TLVs are provided in subclause ***8***.6 of IEEE Std 802.1AB-***2016***.

Replace text of 79.1.1.1 with (changed text in ***bold italic***):

> The Destination Address field of an IEEE 802.3 LLDP frame contains ***the "Nearest bridge"*** MAC address specified by 7.1 of IEEE Std 802.1AB-***2016***).

Replace text of 79.2 with (changed text in ***bold italic***):

> All IEEE 802.3 Organizationally Specific TLVs shall conform to the LLDPDU bit and octet ordering conventions of 8.1 of IEEE Std 802.1AB-***2016***.

4.2.2 Useage Rules

Add the following to the end 79.2 (changed text in ***bold italic***):

> ***Unless stated otherwise in the TLV definition, an LLDPPDU shall contain no more than one instance of each IEEE 802.3 Organizationally Specific TLV.***

4.3 New MPoE TLVs

4.3.1 New subtypes

Change Table 79–1—IEEE 802.3 Organizationally Specific TLVs by , adding new rows and adjusting reserved, as shown:

IEEE 802.3 subtype	**TLV name**	Subclause reference
…		
9	PLCA	79.3.9
10	MPoE MPSE Status TLV	79.3.10
11	MPoE MPD Status TLV	79.3.11
12	MPoE Power Allocated TLV	79.3.12
13-255	Reserved	

4.3.2 MPoE MPSE Status TLV

Insert subclause 79.3.10 with the following text:

79.3.10 MPoE MPSE Status TLV

The MPoE MPSE (see 189.4) Status TLV allows MPSEs to advertise capabilities and status to other MPoE devices on the mixing segment. Table 79-22 shows the format of the TLV.

Table 79–22—MPSE Status TLV elements

Field	Field size (bits)	Notes
Capabilities and status	16	See Table 79-23
Supported Types	8	See Table 79-24
Active Type	8	See Table 79-25
Max Power	16	See Table 79-26
Allocated Power	16	See Table 79-27
Withdrawing Power Delay	8	See Table 79-28
Reserved	8	

Table 79–23— MPSE Capabilities and Status.

Field width	Bit	Function	Units	Value/meaning
16	0	Active		1 = active 0 = inactive
	1	Withdrawing Power Notification		1 = active. 0 = inactive
	15:2	Reserved		

Table 79–24— MPSE Supported Types

Field width	Bit	Function	Units	Value/meaning
8	0	Type 0 supported		1 = supported 0 = unsupported
	1	Type 1 supported		1 = supported 0 = unsupported
	7:2	Reserved		

Table 79–25— MPSE Active Type

Field width	Bit	Function	Units	Value/meaning
8	0	Type 0 active		1 = active 0 = inactive
	1	Type 1 active		1 = active 0 = inactive
	7:2	Reserved		

NOTE: Only one type can be active.

Table 79–26— MPSE Max Power

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

Table 79–27— MPSE Allocated Power

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

Table 79–28— 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 may use the following fields/notifications to notify other MPoE devices on the mixing segment that it will stop providing power after a specified period:

- Withdrawing Power Notification (Table 79–23)
- Withdrawing Power Delay (Table 79–28)

4.3.3 MPoE MPD Status TLV

Insert subclause 79.3.11 with the following text:

79.3.11 MPoE MPD Status TLV

The MPoE MPD (see 189.5) Status TLV allows MPDs to advertise to advertise capabilities, status and requests to other MPoE devices on the mixing segment. Table 79-29 shows the format of the TLV.

Table 79–29— MPD Status TLV elements

Field	Field size (bits)	Notes
Capabilities and status	16	See Table 79-29
Supported Types	8	See Table 79-30
Active Type	8	See Table 79-31
Static power	16	See Table 79-32
Normal power	16	See Table 79-33
Temporary power	16	See Table 79-34
Temporary power duration	16	See Table 79-35
Temporary power delay	8	See Table 79-36
Reserved	8	16-bit alignment
Instantaneous Voltage	16	See Table 79-37
Voltage Out of Range	16	See Table 79-38

Table 79–29— MPD Capabilities and Status

Field width	Bit	Function	Units	Value/meaning
16	1	Voltage monitoring		1 = supported 0 = unsupported
	2	Temporary power notification		1 = active. 0 = inactive
	3	Requested power priority flag		1 = Requested power priority valid 0 = Requested power priority invalid
	4:6	Requested power priority		0 = highest 7 = lowest
	15:7	Reserved		

Table 79–30— MPSE Supported Types

Field width	Bit	Function	Units	Value/meaning
8	0	Type 0		1 = supported 0 = unsupported
	1	Type 1		1 = supported 0 = unsupported
	7:2	Reserved		

Table 79–31— MPSE Active Type

Field width	Bit	Function	Units	Value/meaning
8	0	Type 0		1 = active 0 = inactive
	1	Type 1		1 = active 0 = inactive
	7:2	Reserved		

NOTE: Only one type can be active.

Table 79–32— MPD Static Power

Field width	Bit	Function	Units	Value/meaning
16		Static power	mW	The maximum power the MPD draws before MPoE power negotiation.

Table 79–33— MPD Normal Power

Field width	Bit	Function	Units	Value/meaning
16		Normal power	mW	Power the MPD needs to support its normal function in the range: 0W <= Normal Power <= Static power

Table 79–34— MPD Temporary Power Request

Field width	Bit	Function	Units	Value/meaning
16		Temporary power	mW	Power the MPD needs for a defined duration in the range: 0W <= Temporary power <= 100W Ignored if “Temporary power notification” is not set.

Table 79–35— MPD Temporary Power Request Duration

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

Table 79–36— MPD Temporary Power Request Delay

Field width	Bit	Function	Units	Value/meaning
8		Temporary power delay	secs	Delay before the MPD wants to change its power draw. Ignored if "Temporary power notification" is not set.

An MPD may use the following fields/notifications to request a temporary change in allocated power:

- Temporary power notification (Table 79–29)
- Temporary power (Table 79-34)
- Temporary power duration (Table 79-35)
- Temporary power delay (Table 79-36)

Table 79–37— MPD Instantaneous Voltage

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

Table 79–38— MPD Voltage Out of Range events

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

NOTE: assumes PON_NO_POWER is renamed to “PON_OUT_OF_RANGE”.

4.3.4 MPoE Power Allocated TLV

Insert subclause 79.3.12 with the following text:

79.3.12 MPoE Power Allocated TLV

The MPoE Power Allocated TLV allows an MPSE to advertise power allocation information to other MPoE devices on the mixing segment. Table 79-39 and 79-40 show the format of the TLV.

Table 79–39— MPoE 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 79–40—MPSE Power Allocated TLV power allocation entry

Field width	Bit	Function	Units	Value/meaning
48		MPD MAC address		MPD MAC.
16		MPD granted power	mW	Power the MPD is allocated.
16		MPD static power	mW	See Table 79-32
16		MPD normal power	mW	See Table 79-33
16		MPD temporary power	mW	See Table 79-34
16		MPD temporary power duration	seconds	See Table 79-35
8		MPD temporary power delay	seconds	See Table 79-36
8		Reserved		16-bit alignment.

4.4 PICS changes

4.4.1 Major capabilities/options

Add the following rows to 79.5.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*MPSE	MPoE MPSE Status TLV	79.3.10		O	Yes[] No[]
*MPD	MPoE MPD Status TLV	79.3.11		O	Yes[] No[]
*MPA	MPoE Power Allocated TLV	79.3.12		O	Yes[] No[]

4.4.2 IEEE 802.3 Organizationally Specific TLV

In 79.5.4 IEEE 802.3 Organizationally Specific TLV, change TLV1, and insert TLV3 after TLV2 as show below (changed text in ***bold italic***):

Item	Feature	Subclause	Value/Comment	Status	Support
TLV1	Group MAC addresses	***79.1.1.1***	***Nearest bridge*** group MAC address listed in Table 7-1 of IEEE Std 802.1AB-***2016***	M	Yes[] No[]
TLV3	***TLV usage rules***	***79.2***	***No more than one instance of each TLV in an LLDPPDU unless specifically defined.***	***M***	***Yes[] No[]***

4.4.3 Per TLV PICS

Insert the following after 79.5.12 Power via MDI Measurements TLV

79.5.13 MPoE MPSE Status TLV

Item	Feature	Subclause	Value/Comment	Status	Support
MPSE1	Capabilities and Status	79.3.10	Bitmap as defined in Table 79-23	MPSE:M	Yes[] No[]
MPSE2	Supported Types	79.3.10	Supported MPSE types as defined in Table 79-24	MPSE:M	Yes[] No[]
MPSE3	Active Type	79.3.10	Active MPSE type as defined in Table 79-25	MPSE:M	Yes[] No[]
MPSE4	Max Power	79.3.10	MPSE maximum power in mW as defined in Table 79-26	MPSE:M	Yes[] No[]
MPSE5	Allocated Power	79.3.10	MPSE allocated power in mW as defined in Table 79-27	MPSE:M	Yes[] No[]
MPSE6	Withdrawing Power Delay	79.3.10	MPSE withdrawing power delay in seconds defined in Table 79-28	MPSE:O	Yes[] No[]

79.5.14 MPoE MPD Status TLV

Item	Feature	Subclause	Value/Comment	Status	Support
MPD1	Capabilities and Status	79.3.11	Bitmap as defined in Table 79-29	MPD:M	Yes[] No[]
MPD2	Supported Types	79.3.11	Supported MPD types as defined in Table 79-30	MPD:M	Yes[] No[]
MPD3	Active Type	79.3.11	Active MPD type as defined in Table 79-31	MPD:M	Yes[] No[]
MPD4	Static power	79.3.11	MPD static power in mW as defined in Table 79-32	MPD:M	Yes[] No[]
MPD5	Normal power	79.3.11	MPD normal power in mW as defined in Table 79-33	MPD:M	Yes[] No[]
MPD6	Temporary power	79.3.11	MPD temporary power request comprising of power level(mW), duration(secs) and delay (secs) as defined in Table 79-34, Table 79-35, Table 79-36	MPD:O	Yes[] No[]
MPD7	Instantaneous voltage	79.3.11	MPD instantaneous voltage as defined in Table 79-37	MPD:O	Yes[] No[]
MPD8	Voltage Out of Range events	79.3.11	MPD “Voltage Out of Range” events as defined in Table 79-38	MPD:M	Yes[] No[]

79.5.15 MPoE Power Allocated TLV

Item	Feature	Subclause	Value/Comment	Status	Support
MPA1	Entry Count	79.3.12	The number of power allocation entries in the TLV.	MPA:M	Yes[] No[]
MPA2	Entry	79.3.12	Table of per MPD power allocation entries.	MPA:M	Yes[] No[]

END OF DOCUMENT