EFM OAM Tutorial

Current as of IEEE P802.3ah/D3.2™

Kevin Daines

Editor, EFM OAM Sub-task force

Agenda

- Overview
- OAM Protocol Data Units (OAMPDUs)
- Events
  - Critical Link Events
  - Link Events
- Variable Retrieval
- Remote Loopback
  - Internal block diagram
  - Starting and exiting timing diagrams
- Organization Specific Extensions
- Discovery
- Active & Passive Modes

Overview: Parent Organizations

- IEEE 802 LMSC
  - *L*ocal Area Network/*M*etropolitan Area Network *S*tandards *C*ommittee
- IEEE 802.3 CSMA/CD
  - *C*arrier *S*ense *M*ultiple *A*ccess with *C*ollision *D*etect (CSMA/CD) Working Group
    - Commonly referred to as the Ethernet Working Group
- IEEE P802.3ah Ethernet in the First Mile Task Force (EFM)

Overview: OSI Layer Stack

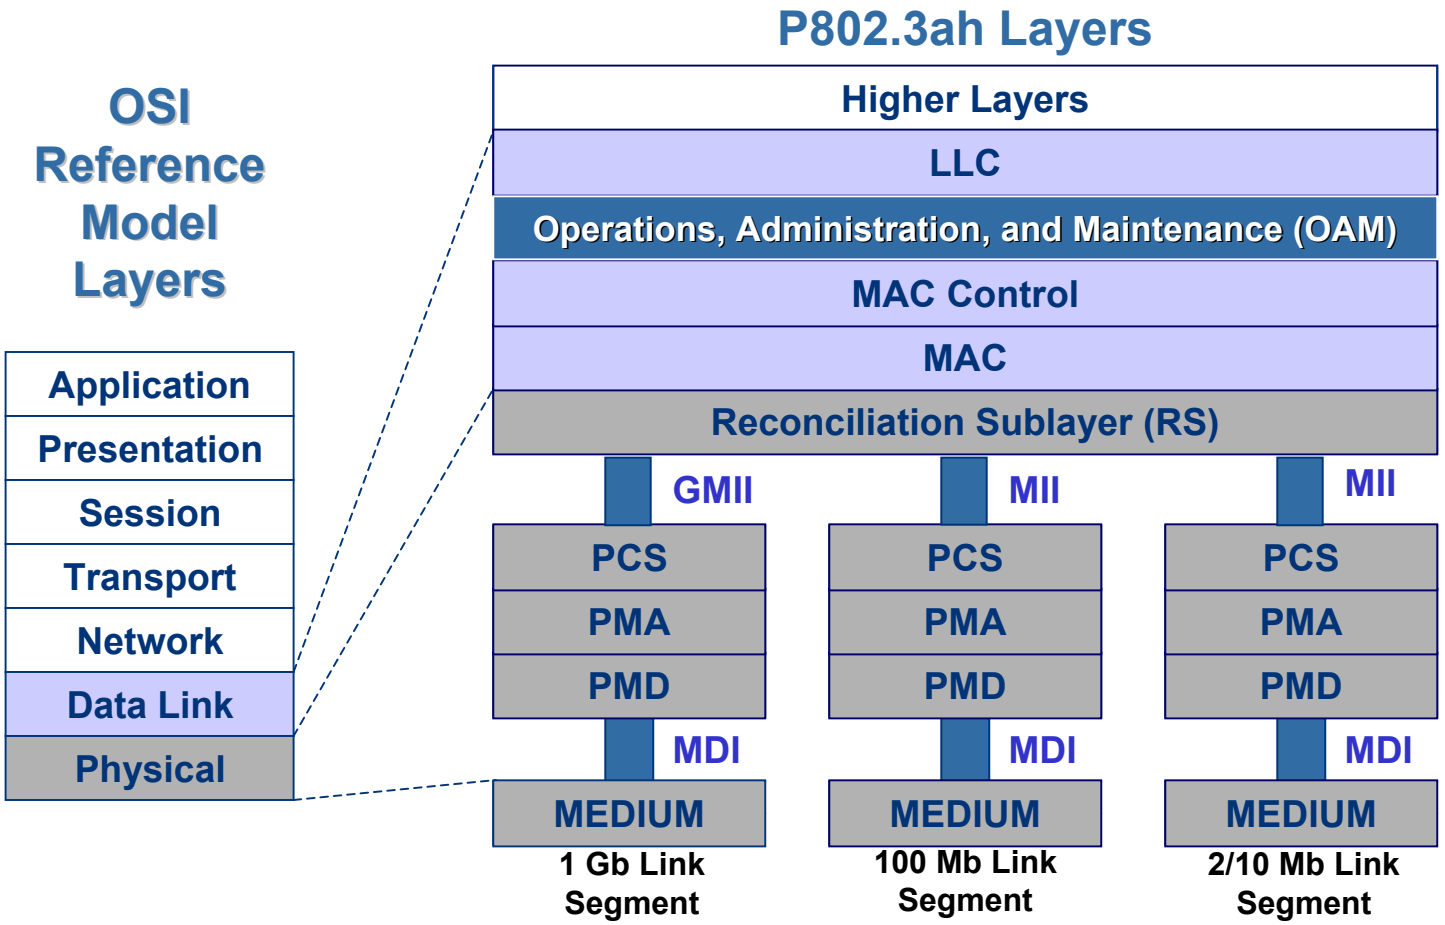

OAM = Operations, Administration, & Maintenance
MDI = Medium Dependent Interface
(G)MII = (Gigabit) Media Independent Interface
PCS = Physical Coding Sublayer
PMA = Physical Medium Attachment
PMD = Physical Medium Dependent

Overview: Objectives

- OAM provides mechanisms to:
  - Monitor link operation and health
  - Improve fault isolation
- Method: OAM data conveyed in basic (*untagged*) 802.3 Slow Protocol frames
  - Sent between two ends of a single link
    - *Note: called a "DTE" in 802.3 terminology*
  - Slow Protocols allows S/W implementation
- Fills major requirement to reduce EFM OpEx

Overview: Non-objectives

- Does not provide capabilities for:
  - Station management
  - Protection switching
  - Provisioning
    - *No SET functions*
  - Bandwidth allocation
  - Speed/duplex negotiation
  - End-to-end OAM communication
    - *802.3 scope restricted to single links*

Overview: Compatibility

- Optionality
  - OAM is optional; software and/or hardware implementations possible
  - May be implemented on one or more ports within a system
  - Individual OAM features are optional
- Supported media
  - All point-to-point (P2P) and emulated P2P links supported
- 802.3x MAC Flow Control (PAUSE)
  - Inhibits all traffic *including* OAMPDUs
- 802.3z 1000BASE-X Auto Negotiation
  - Support for unidirectional fault signaling is *mutually exclusive* with 802.3z Auto Neg
    - 802.3z Auto Neg *must be disabled* for fault signaling to be sent over 1000BASE-X unidirectional links

OAMPDU: Unidirectional

- 802.3ah/Clause 66 adds *optional* feature to allow optical links to operate unidirectionally
  - *Legacy links become inoperable when one direction fails*
  - Newer links can send Information OAMPDUs unidirectionally to signal link fault
    - 100BASE-X PCS, 1000BASE-X PCS & 10GbE RS supported

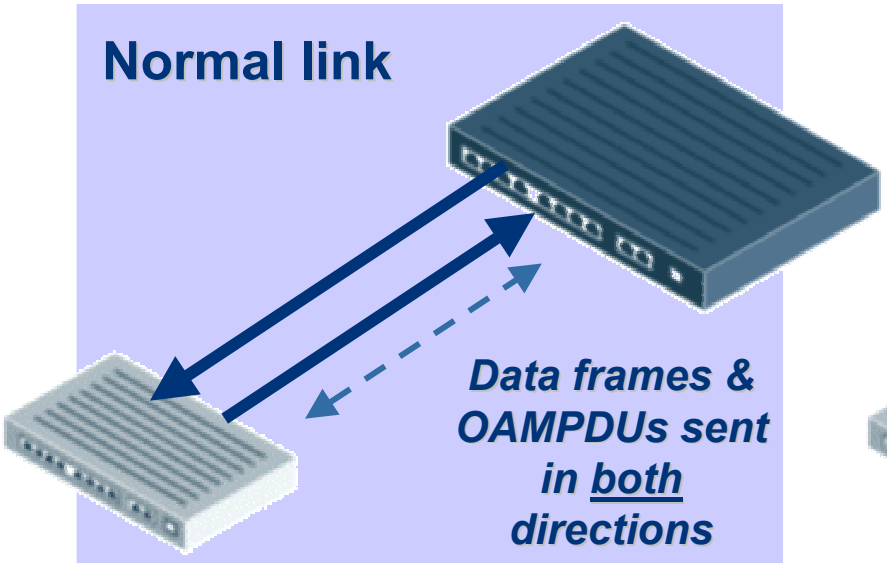

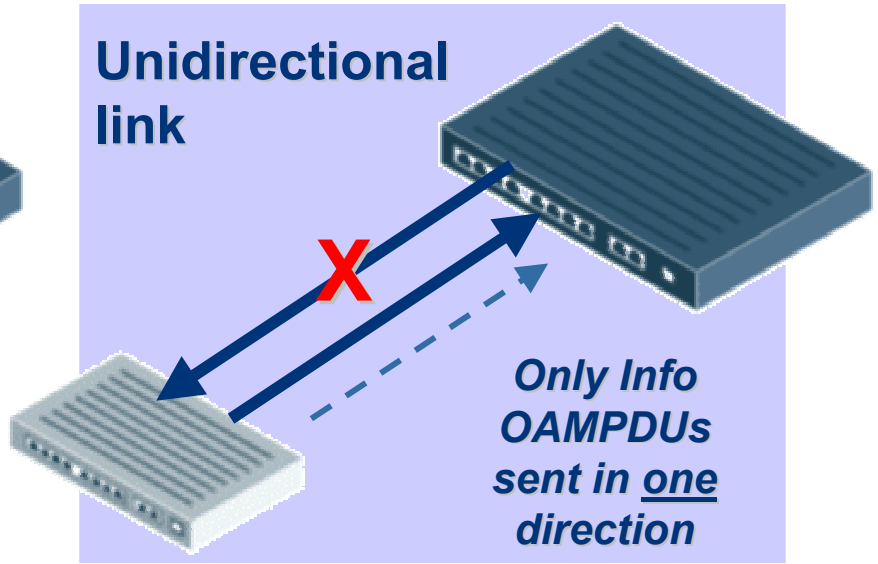

OAMPDU: Forwarding - NOT

- Only traverse a single link
  - Not forwarded by bridges
- Communication beyond a single link left to higher layers

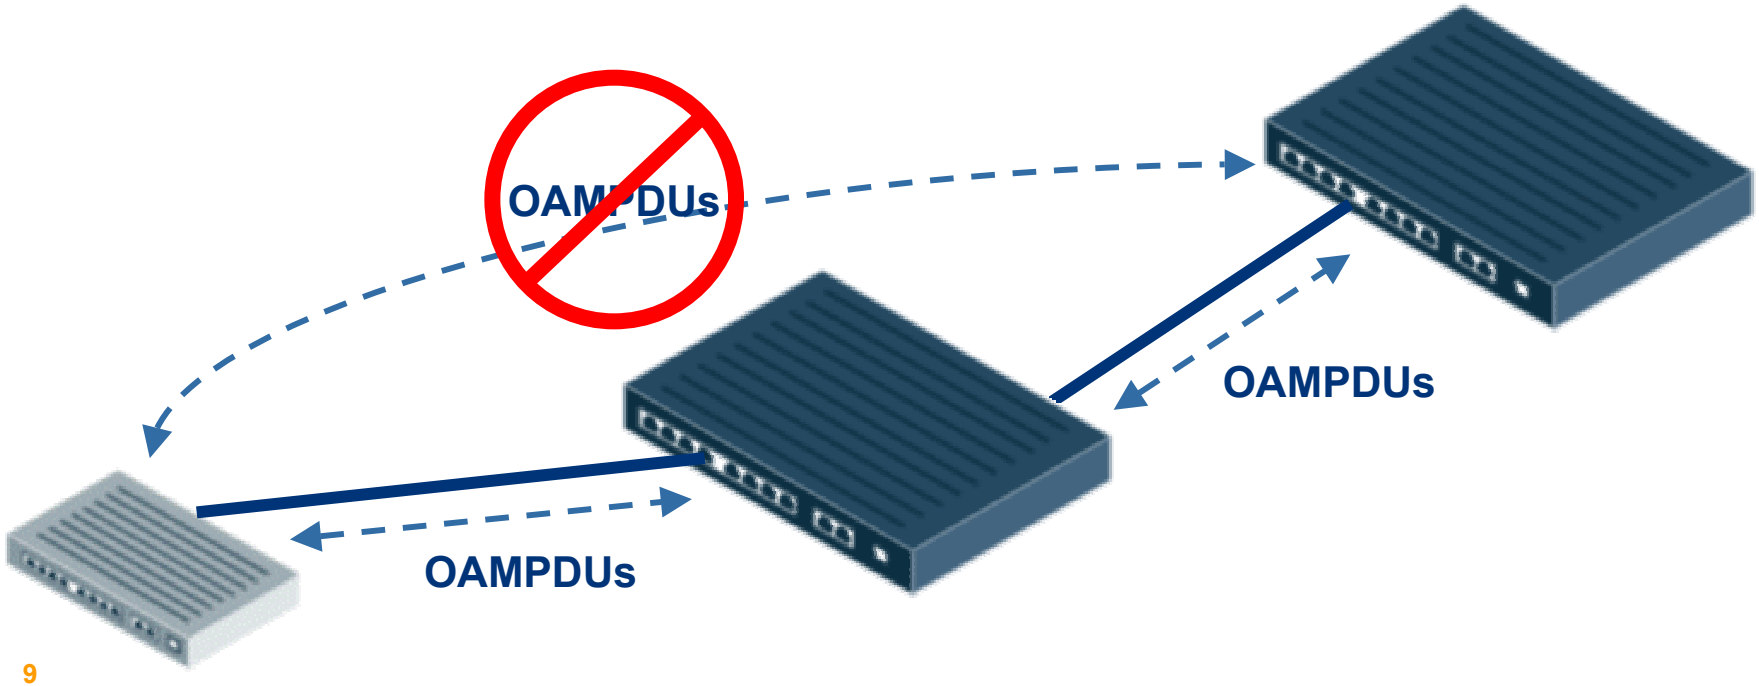

OAMPDU: Size/Rate

- Must be standard frame length
  - 64-1518 octets
  - Maximum PDU size determined during Discovery process
- *Must be untagged*

Octets	
6	01-80-c2-00-00-02 [*Slow Protocol*]
6	MAC Source Address
2	Type=88-09 [*Slow Protocols*]
1	Subtype = 0x03 [*OAM*]
2	Flags field
1	Code
42-1496	Data/Pad field
4	Frame Check Sequence
64-1518	

- Maximum of (10) OAMPDUs per second
  - Max rate defined in Annex 43B as modified by EFM
  - May be sent multiple times to increase likelihood of reception by remote device (e.g., in the case of high bit errors)

OAMPDU: Flags field

- Length: 2 octets
- Legend:

Critical Link Event bit

State information bit

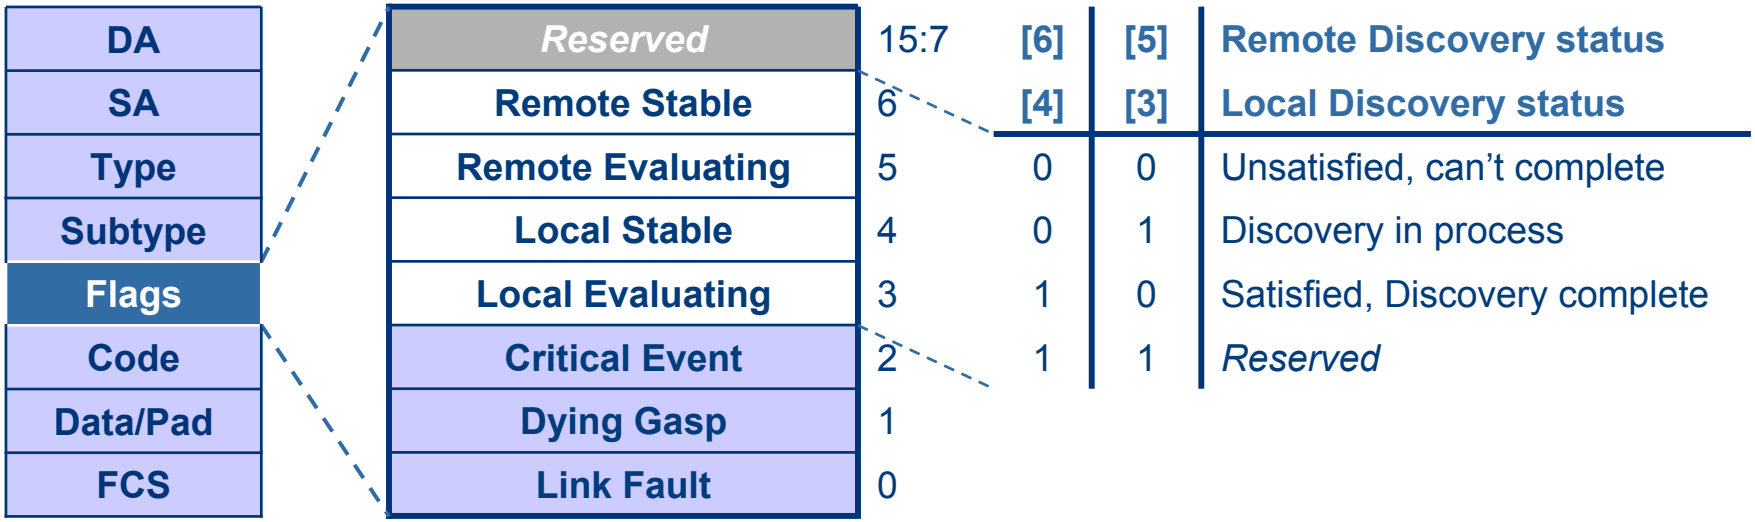

OAM Critical Link Events

- Link Fault
  - Signal remote device that receive path is broken
  - Sent once per second in Information OAMPDU
- Dying Gasp
  - Signal remote device that unrecoverable local fault (e.g., power failure) has occurred
  - May be sent immediately/continuously
- Critical Event
  - An unspecified critical event has occurred
  - May be sent immediately/continuously

OAMPDU: Codes

Code	OAMPDU	Length
0x00	Information	*varies*
0x01	Event Notification	*varies*
0x02	Variable Request	*varies*
0x03	Variable Response	*varies*
0x04	Loopback Control	64 octets
0x05-0xFD	*Reserved*	
0xFE	Organization Specific	varies
0xFF	*Reserved*	

- Unknown/unsupported OAMPDUs sent to OAM client
  - *Different than 802.3x behavior, which filtered unsupported opcodes*

OAMPDU: Information

- Code: 0x00
- Data field: Information TLVs
- Length: varies

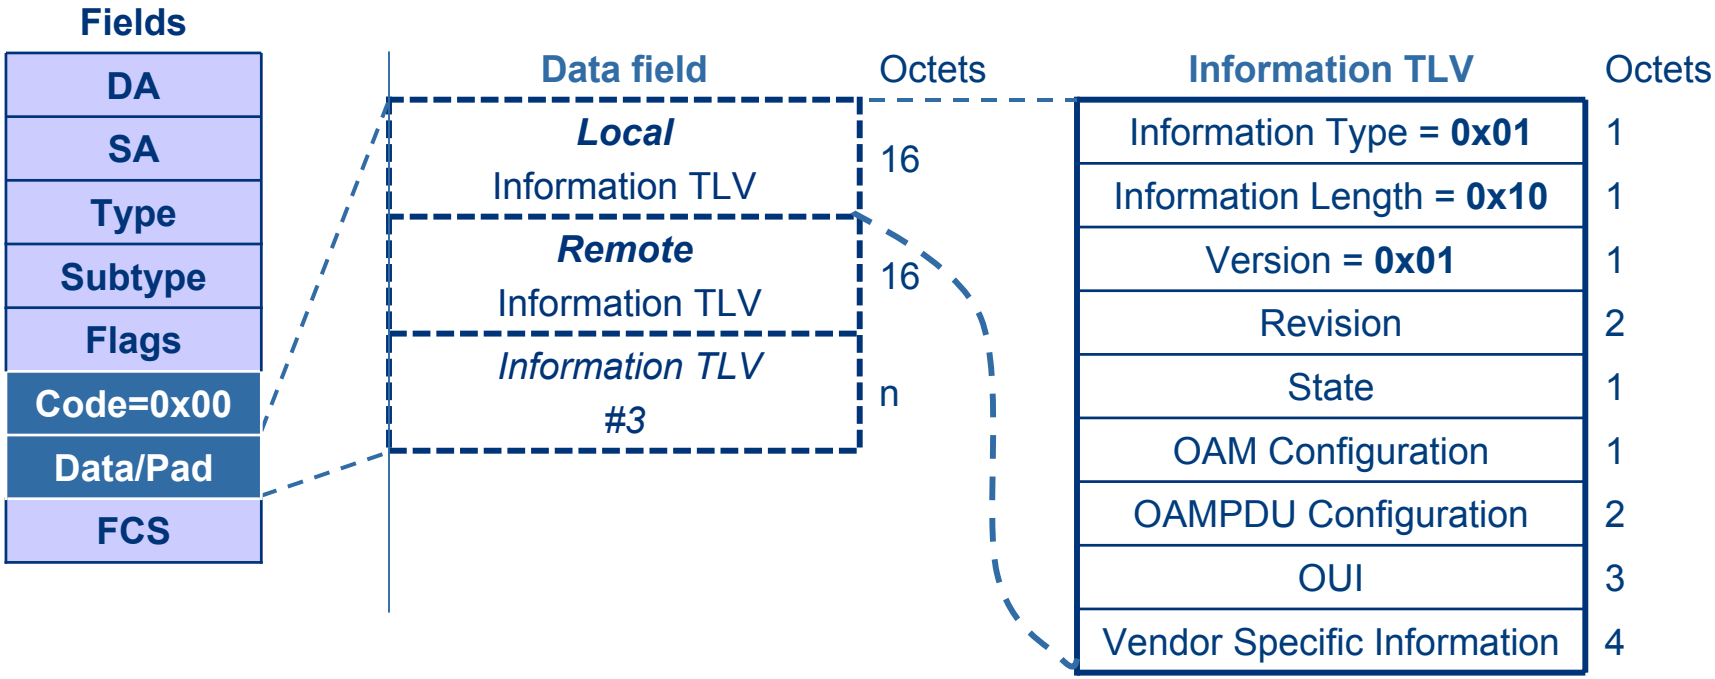

Information TLVs

Information Type	Information TLV Name
0x00	*End of TLV marker*
0x01	**Local Information**
0x02	**Remote Information**
0x03-0xFD	*Reserved*
0xFE	**Organization Specific Information**
0xFF	*Reserved*

- **Sent as Information TLVs within Information PDU**
  - **Local & Remote used for Discovery Process**
  - ***Optional* Organization Specific Information used for extension purposes**

Local/Remote Information

		7	6	5	4	3	2	1	0
1	Information Type	8-bit Type							
1	Information Length	0x10							
1	OAM Version	0x01							
2	Revision	16-bit Revision							
1	State	*reserved*					Mux	Parser Action	
1	OAM Configuration	*reserved*			Vars	Events	LB	Unidir	Mode
2	OAMPDU Configuration	*reserved*					Max OAMPDU Size		
		Maximum OAMPDU Size							
7	Vendor Identifier	24-bit Organizationally Unique Identifier							
		32-bit Vendor Specific Information							

OAMPDU: Event Notification

- **Code: 0x01**
- **Data field: One or more Link Event TLV(s)**
- **Length: *Variable***

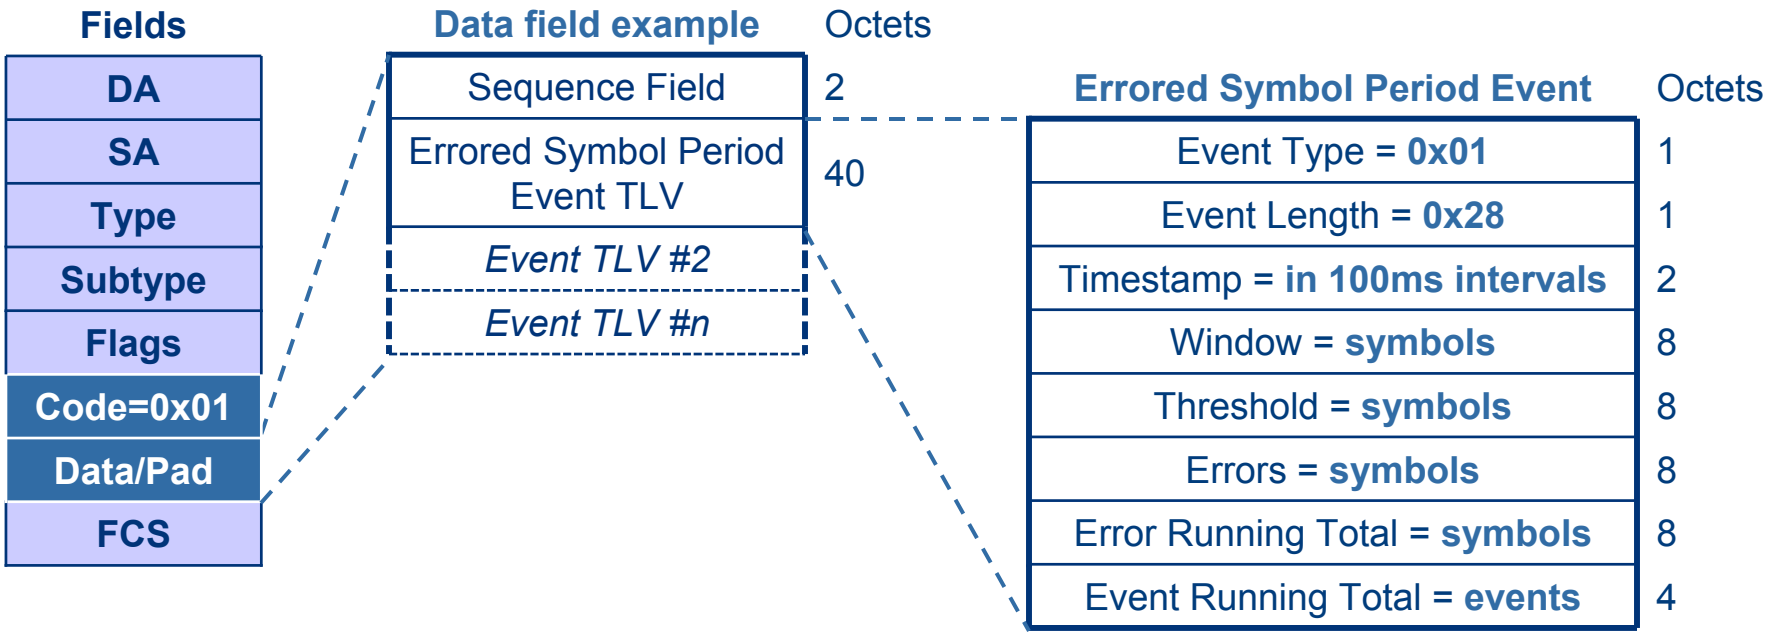

OAM Link Event TLVs

Event Type	Event TLV Name
0x00	*End of TLV marker*
0x01	**Errored Symbol Period Event**
0x02	**Errored Frame Event**
0x03	**Errored Frame Period Event**
0x04	**Errored Frame Seconds Summary Event**
0x05-0xFD	*Reserved*
0xFE	**Organization Specific Event TLV**
0xFF	*Reserved*

- **Sent as Link Event TLVs within Event Notification PDU**
  - **May be sent multiple times to increase likelihood of reception (e.g., in the case of high bit errors)**
  - **Includes time reference when generated**

Errored Symbol Period Event

- A window, measured in number of symbols, where number of errored symbols exceeded a threshold
- Type: 0x01
- Length: 0x28 (40 octets)
- Value:

Fields	Width	Description
Timestamp	16-bits	Time reference, in 100ms units, when generated
Window	64-bits	Lower bound: Symbols in 1 second Upper bound: Symbols in 60 seconds
Threshold	64-bits	Lower bound: 0 Upper bound: unspecified
Errors	64-bits	# of symbols errors in *Window*
Total Errors	64-bits	Total # of symbol errors since reset
Total Events	32-bits	Total # of events sent since reset

Errored Frame Event

- A window, measured in 100ms intervals, where number of errored frames exceeded a threshold
- Type: 0x02
- Length: 0x1A (26 octets)
- Value:

Fields	Width	Description
Timestamp	16-bits	Time reference, in 100ms units, when generated
Window	16-bits	Lower bound: 1 second Upper bound: 60 seconds
Threshold	32-bits	Lower bound: 0 Upper bound: unspecified
Errors	32-bits	# of frame errors in *Window*
Total Errors	64-bits	Total # of frame errors since reset
Total Events	32-bits	Total # of events sent since reset

Errored Frame Period Event

- A window, measured in received frames, where number of errored frames exceeded a threshold
- Type: 0x03
- Length: 0x1C (28 octets)
- Value:

Fields	Width	Description
Timestamp	16-bits	Time reference, in 100ms units, when generated
Window	32-bits	Lower bound: # of 64B frames in 1 second Upper bound: # of 64B frames in 60 seconds
Threshold	32-bits	Lower bound: 0 Upper bound: unspecified
Errors	32-bits	# of frame errors in *Window*
Total Errors	64-bits	Total # of frame errors since reset
Total Events	32-bits	Total # of events sent since reset

Errored Frame Seconds Summary

- A window, in 100ms intervals, where number of errored frame seconds exceeded a threshold
- Type: 0x04
- Length: 0x16 (22 octets)
- Value:

Fields	Width	Description
Timestamp	16-bits	Time reference, in 100ms units, when generated
Window	16-bits	Lower bound: 10 seconds Upper bound: 900 seconds
Threshold	16-bits	Lower bound: 0 Upper bound: unspecified
Errors	16-bits	# of errored frame seconds in *Window*
Total Errors	64-bits	Total # of errors causing since reset
Total Events	32-bits	Total # of events sent since reset

Organization Specific Event

- Organizations may define events that are of variable length and are distinguished by the OUI
- Type: 0xFE
- Length: varies
- Value:

Fields	Width	Description
OUI	24-bits	Organizationally Unique Identifier
varies	*varies*	*varies*

OAMPDU: Variable Req/Resp

Variable Request

- Code: 0x02
- Data: Variable *Descriptors*
- Length: *Variable*

Variable Response

- Code: 0x03
- Data: Variable *Containers*
- Length: *Variable*

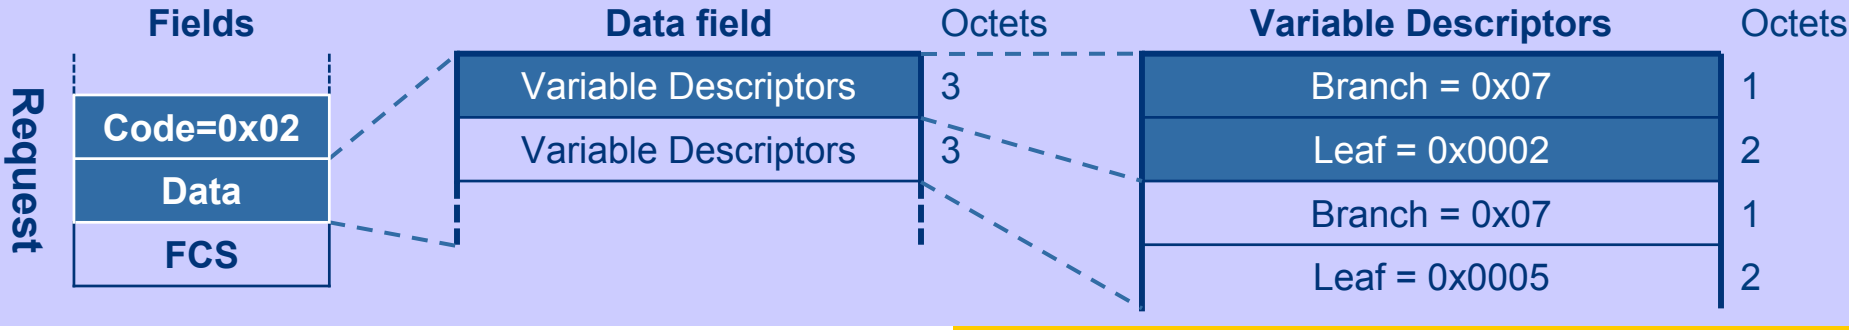

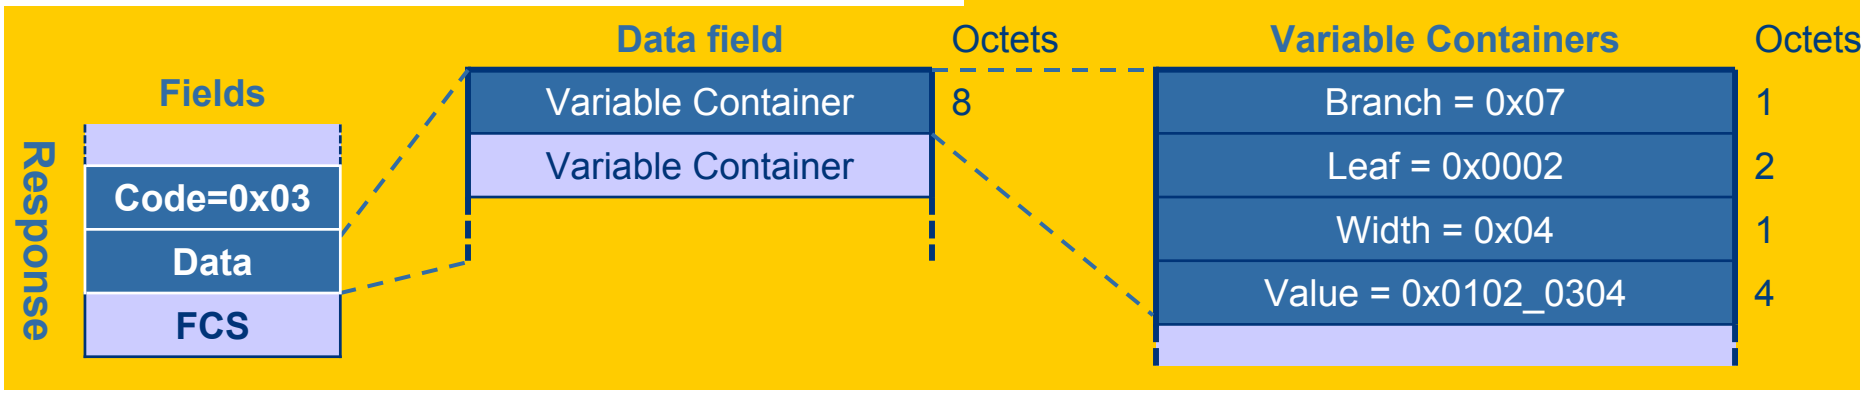

Variable Retrieval

- Transfer Ethernet counters and statistics via Variable Containers/Descriptors
- Variables are referenced using Annex 30A CMIP registration arcs
- Can be used to emulate L2 Ping
  - (i.e., Tx Variable Request, Rx Variable Response)

- Examples:

Variable	CMIP Registration Arcs	
	Branch	Leaf
aFramesTransmittedOK	0x07	0x0002
aFrameCheckSequenceErrors	0x07	0x0006
aOctetsReceivedOK	0x07	0x000E

OAMPDU: Loopback Control

- Code: 0x04
- Data field: Loopback Command (1 octet)
- Length: 64 octets

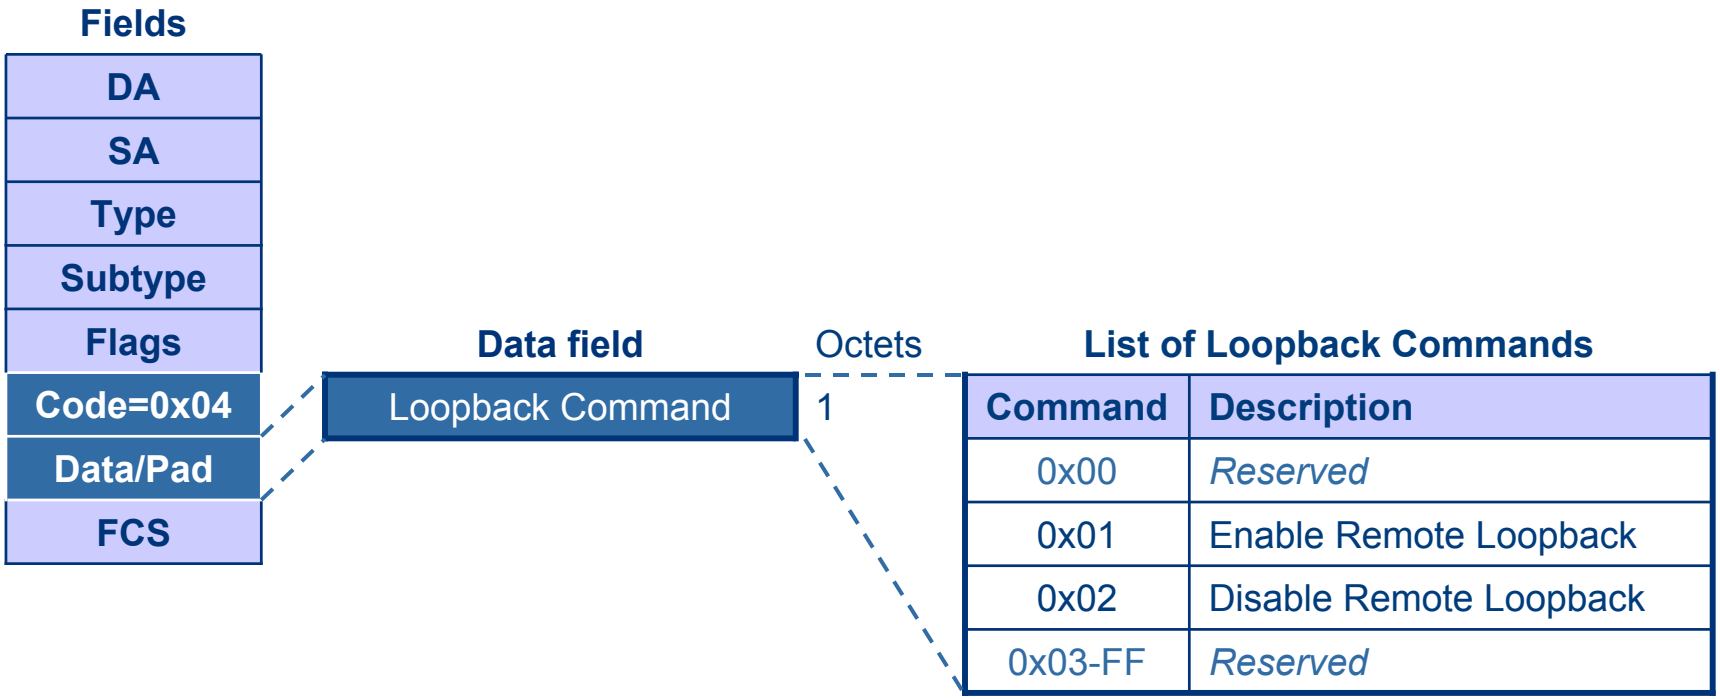

OAM Remote Loopback

- Local DTE sends arbitrary data frames
- Remote DTE returns data frames
- Frame BER equals bit BER to high probability when bit BER is better than 10^{-6}

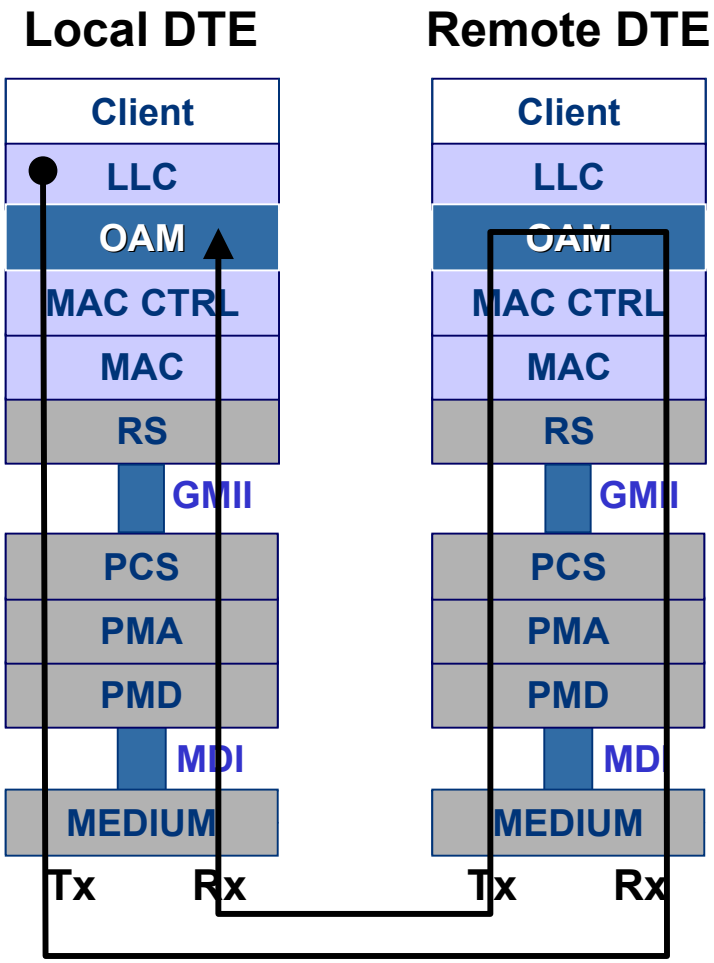

Can be implemented in H/W or S/W

OAM Sublayer Block Diagram

- **OAM client**
  - Configures OAM sublayer through Control
  - Processes received PDUs
  - Transmits PDUs
- **Control**
  - Provides interface with OAM client entity
  - Contains Discovery process
- **Parser**
  - Inspects received frames, sends PDUs to Control and based on configuration, sends:
    - Non-PDUs to upper layer

      or
    - Non-PDUs to Multiplexer
- **Multiplexer**
  - Multiplexes PDUs and non-PDUs

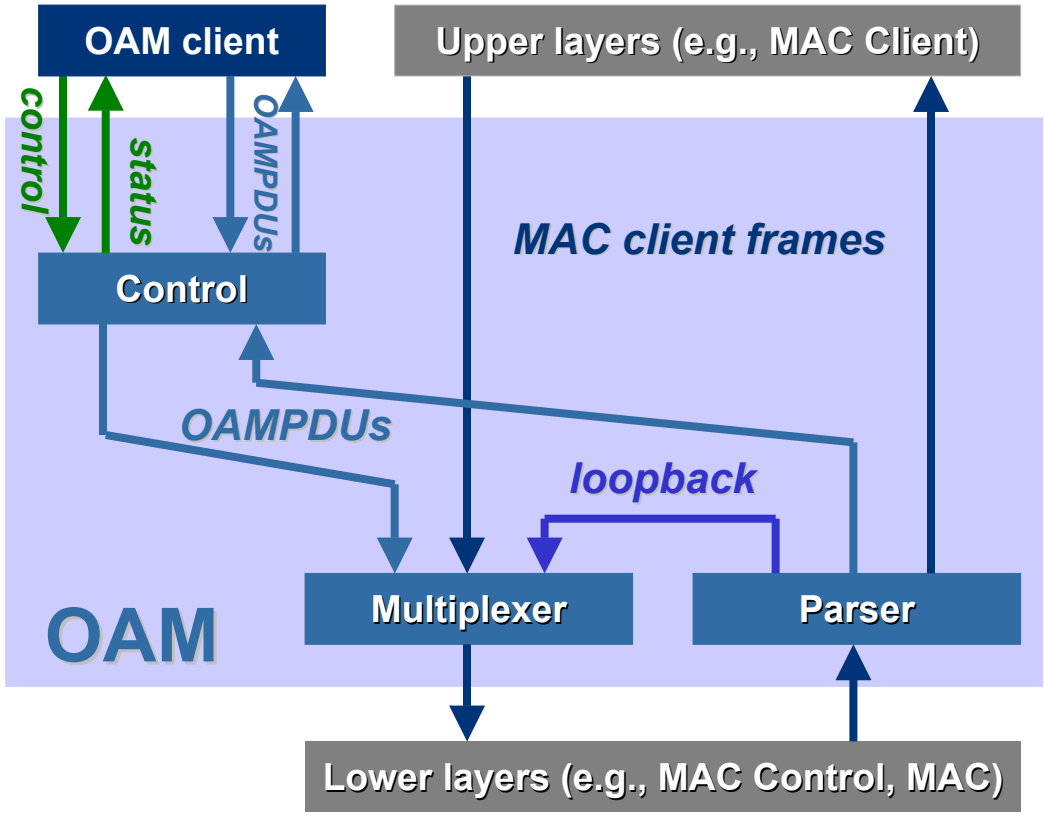

Starting Remote Loopback

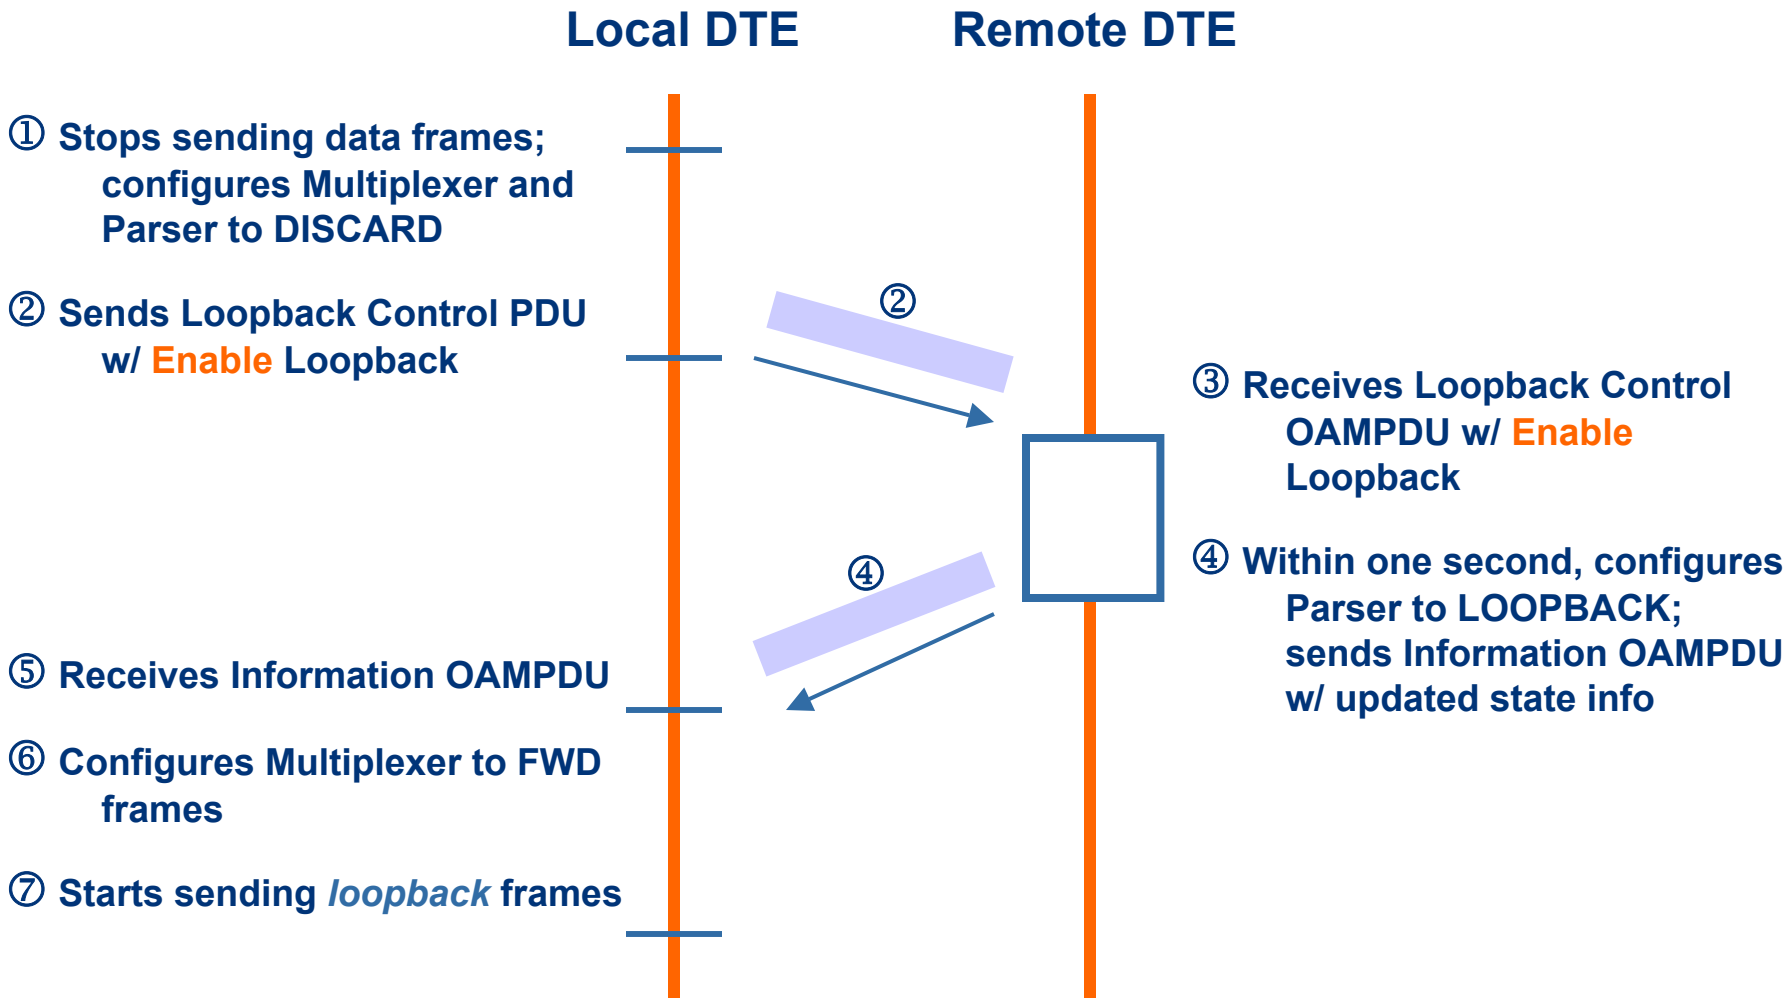

Source: Jee-Sook Eun, ETRI

Exiting Remote Loopback

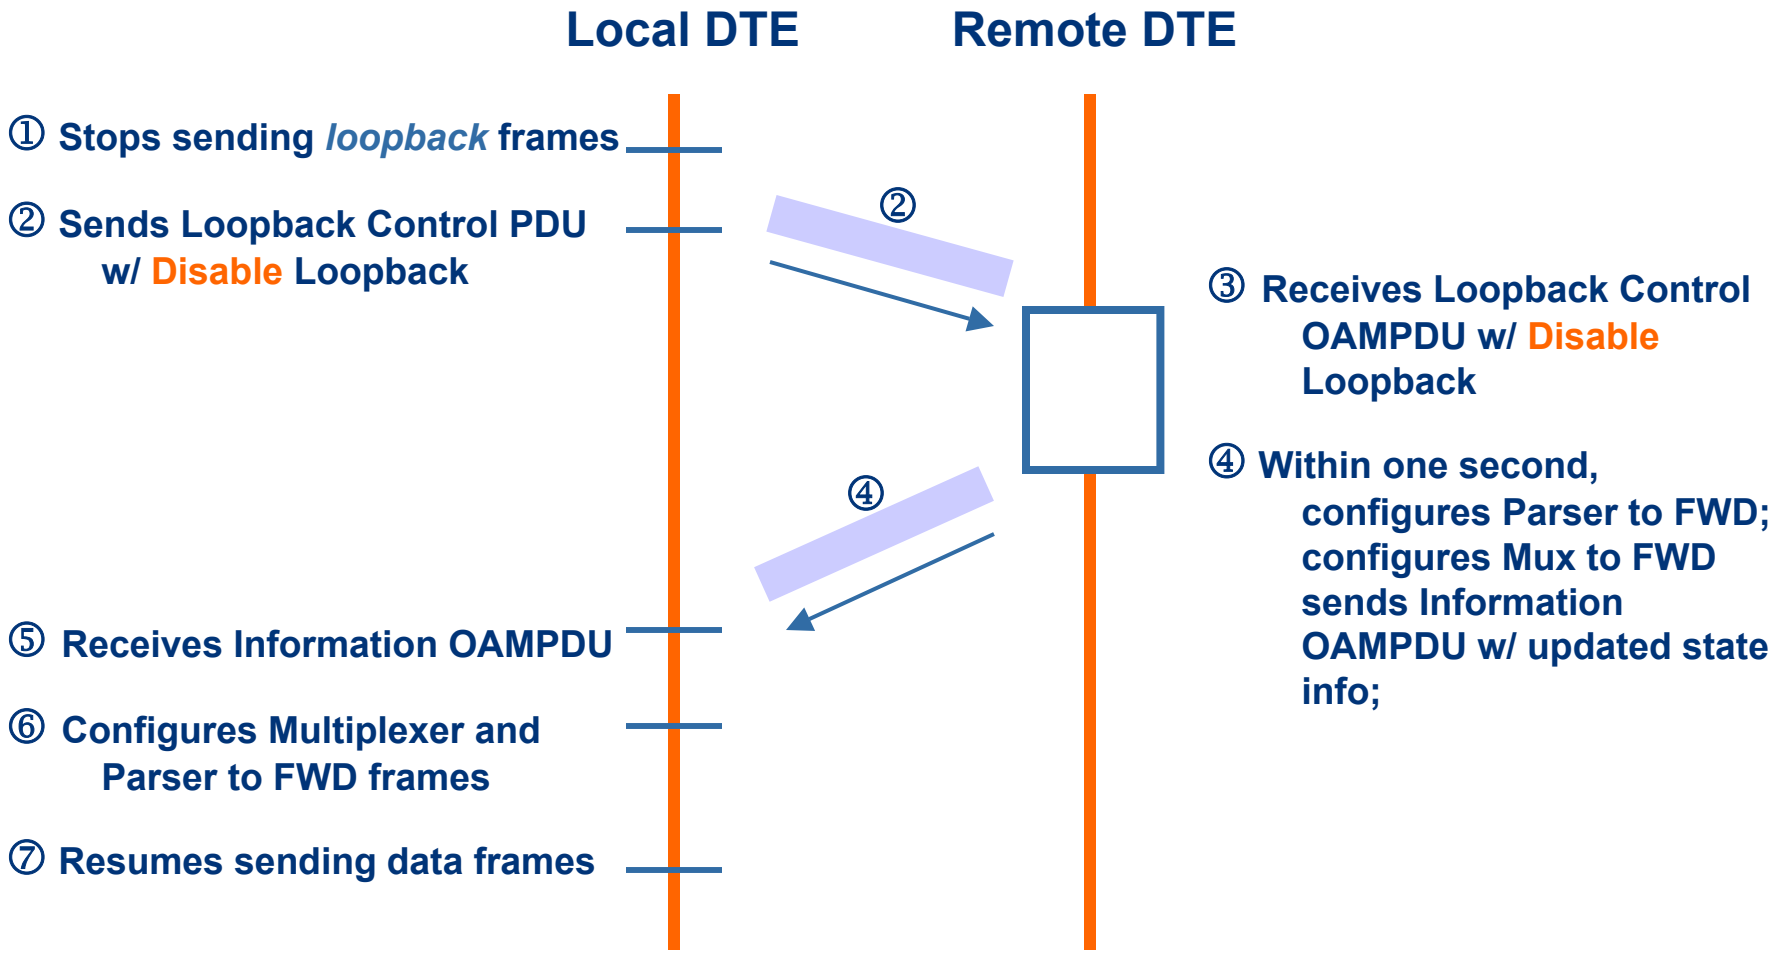

Source: Jee-Sook Eun, ETRI

OAMPDU: Organization Specific

- Code: 0xFE
- Distinguisher: IEEE 24-bit Organizationally Unique Identifier
- Data field: Organization Specific

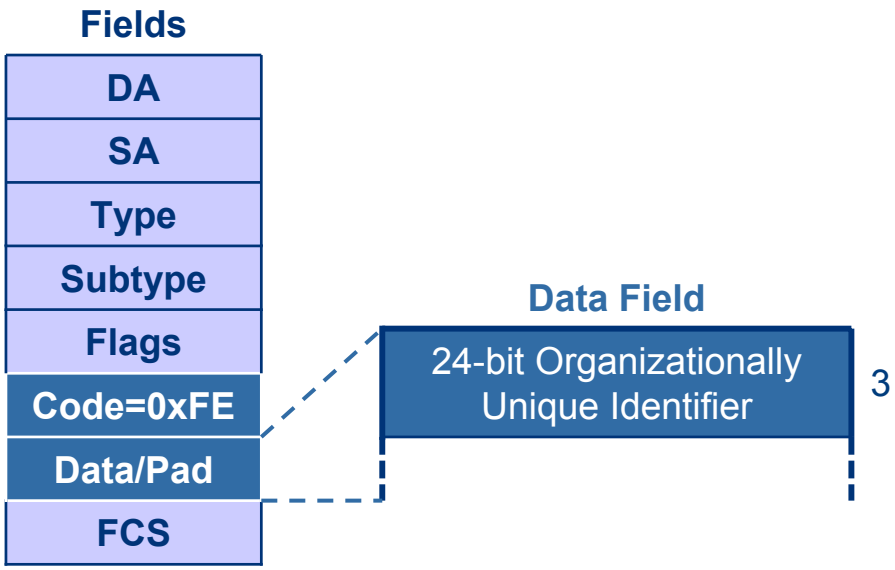

OAM Discovery

- Allows local DTE to detect OAM on remote DTE
- Once OAM support is detected, both ends of the link exchange state and configuration information
  - e.g. mode, PDU size, loopback support
- If both DTEs are satisfied with settings, OAM is enabled on link
- Loss of link and non-reception of PDUs for 5 seconds are causes of Discovery re-starting

OAM Active Mode

- A DTE in Active mode:
  - Initiates the OAM Discovery process
  - Sends Information PDUs
  - May send Event Notification PDUs
  - May send Variable Request/Response PDUs
  - May send Loopback Control PDUs
  - *Exceptions:*
    - Does not respond to Variable Request PDUs from DTEs in Passive mode
    - Does not react to Loopback Control PDUs from DTEs in Passive mode

OAM Passive Mode

- A DTE in Passive mode:
  - Waits for the remote device to initiate the Discovery process
  - Sends Information PDUs
  - May send Event Notification PDUs
  - May respond to Variable Request PDUs
  - May *react* to received Loopback Control PDUs
  - Is not permitted to send:
    - *Variable Request PDUs*
    - *Loopback Control PDUs*