CTF Bridges

Cut-Through Forwarding process functions and their supporting internal sublayers

Christophe Mangin

2024/3/15

MITSUBISHI ELECTRIC R&D CENTRE EUROPE

1 Some background

- Reception path
  - From the MAC client interface up to the EISS

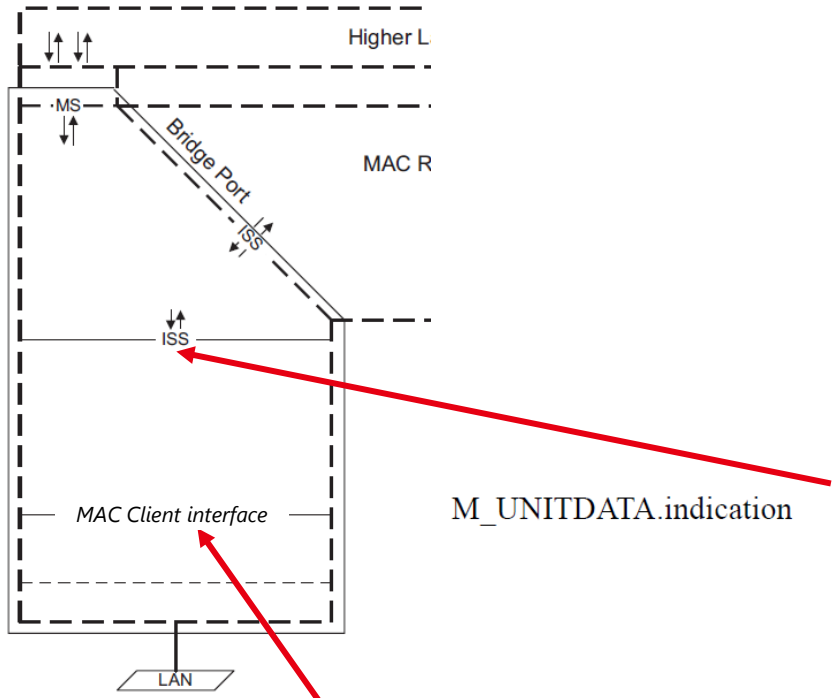

MA_UNITDATA.indication

(
destination_address,
source_address,
mac_service_data_unit,
priority
)

M_UNITDATA.indication

(
destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
frame_check_sequence,
service_access_point_identifier,
connection_identifier
)

EM_UNITDATA.indication

(
destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
vlan_identifier,
frame_check_sequence,
service_access_point_identifier,
connection_identifier,
flow_hash,
time_to_live
)

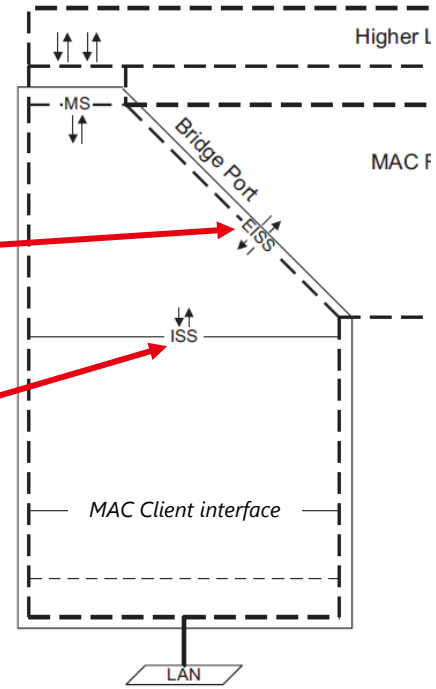

2 CTF-capable (E)ISS

MITSUBISHI ELECTRIC
Changes for the Better

2.1 CTF-ISS

CTF-capable Internal Sublayer Service

CTF-ISS services primitives

- CTF-M_UNITDATA.indication (

as per 802.1AC →

destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
frame_check_sequence,
service_access_point_identifier,
connection_identifier

status

)

- CTF-M_UNITDATA.request (

as per 802.1AC →

destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
frame_check_sequence,
service_access_point_identifier,
connection_identifier

status

)

Receipt of a MAC frame from a LAN

- When *determined*, *status* is
  - derived from the frame reception status signaled by the underlying LAN during the reception of the frame or,
  - set by one of the functions constituting the MAC relay function
- *status* conveys the following information:
  - the absence of error in the received frame or,
  - the presence of an error in the received frame, and its reason:
    - invalid FCS
    - invalid frame length
    - non integral number of octets
    - error condition detected by one (or more ?) of the MAC relay functions
    - …

Transmission of a MAC frame to a LAN

Status is a copy of the *status* parameter provided by the CTF-M_UNITDATA.indication (pending or determined)

CTF-ISS services primitives

- Status parameter properties
  - Similar to the service_access_point_identifier and connection_identifier parameters
    - not a parameter of the peer-to-peer service
      - not conveyed to the communicating peer system
    - local to the system within which a given service request or service indication occurs
    - its value is not conveyed in any external protocol, including management protocols
  - 2 states
    - “determined” : status value = error, no_error
      - If status = error, the value indicates the error reason
    - “pending” : status value is not determined (yet)

- Principle
  - In a receive port, the parameters of the CTF-M_UNITDATA.indication primitive are **progressively obtained** from the underlying LAN or **progressively determined** based on the parameters obtained from the underlying LAN
- Temporal behavior
  - the CTF-M_UNITDATA.indication primitive is generated as soon as the set of parameters progressively provided by the underlying LAN comprises **all the parameters**, or parts thereof, required by the execution of **all the functions** constituting the MAC relaying entity the primitive is presented to.
  - the earliest time a CTF-M_UNITDATA.indication primitive can be generated is when the destination_address parameter is complete and the MAC relay entity only comprises a frame filtering function (VLAN-unaware bridge).
  - the latest time a CTF-M_UNITDATA.indication primitive can be generated is when all its parameters are complete and determined, i.e. all obtained from the underlying LAN.
    - At that time, the status parameter is **determined**, it is pending otherwise.
  - at the time the CTF-M_UNITDATA.indication primitive occurs, some of its parameters can be incomplete or pending.
    - their content or value is progressively made available to the MAC relay entity's functions
- Support of the CTF-ISS with signaled priority (IEEE 802.1Q 6.20)
  - the CTF-M_UNITDATA.indication primitive is generated when the initial octets of the mac_service_data_unit (VLAN-tag) are available

MITSUBISHI ELECTRIC
Changes for the Better

2.2 CTF-EISS

CTF-capable Internal Sublayer Service

CTF-EISS services primitives

- CTF-EM_UNITDATA.indication

as per 802.1Q →

(
destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
vlan_identifier,
frame_check_sequence,
service_access_point_identifier,
connection_identifier,
flow_hash,
time_to_live
status
)

- CTF-EM_UNITDATA.request

as per 802.1Q →

(
destination_address,
source_address,
mac_service_data_unit,
priority,
drop_eligible,
vlan_identifier,
frame_check_sequence,
service_access_point_identifier,
connection_identifier,
flow_hash,
time_to_live
status
)

The *status* parameter is handled the same way as at the CTF-ISS

- Principle
  - In a receive port, the parameters of the CTF-EM_UNITDATA.indication primitive are **progressively obtained** from the CTF-ISS or **progressively determined** based on the parameters obtained from the CTF-ISS
- Temporal behavior
  - the CTF-EM_UNITDATA.indication primitive is generated as soon as the set of parameters progressively provided by the CTF-ISS comprises all the parameters, or parts thereof, required by the execution of **all the functions** constituting the MAC relaying entity the primitive is presented to.
  - the earliest time a CTF-EM_UNITDATA.indication primitive can be generated is when the vlan_identifier parameter is complete and the MAC relay entity only comprises a frame filtering function.
    - vlan_identifier complete = initial octets of the mac_service_data_unit parameter containing a VLAN tag
  - the latest time a CTF-EM_UNITDATA.indication primitive can be generated is when all its parameters are complete and determined, i.e. all obtained from the CTF-ISS.
    - At that time, the status parameter is determined, it is pending otherwise.
  - at the time the CTF-EM_UNITDATA.indication primitive occurs, some of its parameters can be incomplete or pending.
    - their content or value is progressively made available to the MAC relay entity's functions

3 CTF Bridge port transmit and receive

Bridge port connectivity (802.1Q)

- CTF VLAN bridge

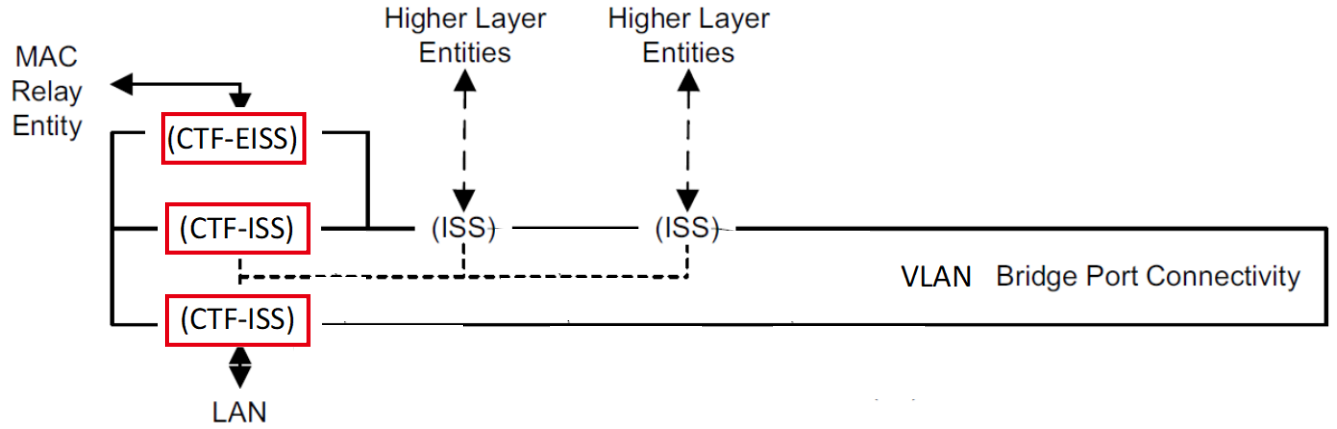

- CTF MAC bridge

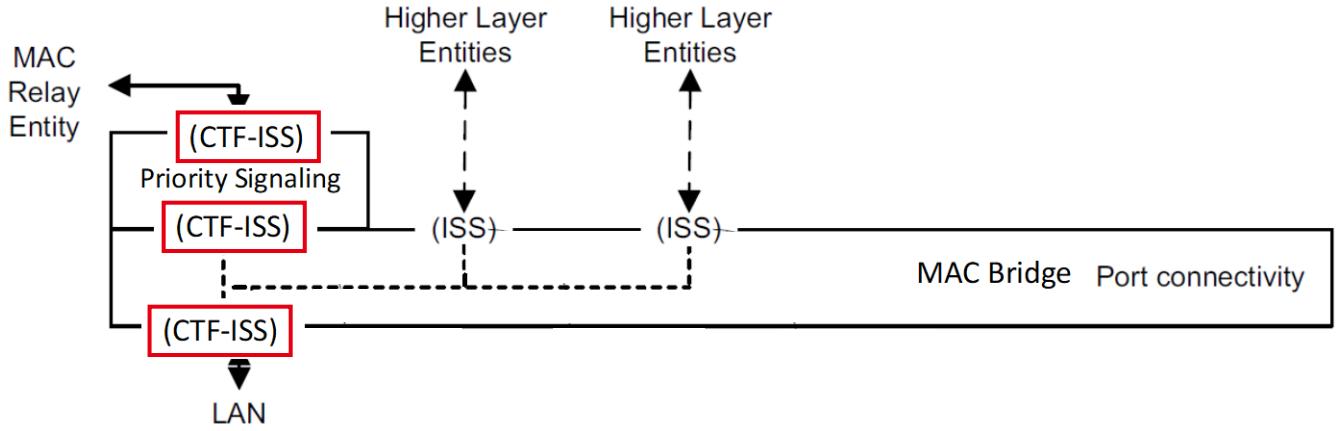

ISS to/from CTF-ISS primitives translation

- CTF-M_UNITDATA.indication to M_UNITDATA.indication
  - occurs when the status parameter is determined.
  - if status' value equals "no-error", the parameters of the CTF-M_UNITDATA.indication, except the status parameter, are used as parameters for the M_UNITDATA.indication primitive.
    - if not, the resulting M_UNITDATA.indication primitive is not generated, i.e. the frame is discarded.
- CTF-M_UNITDATA.request to M_UNITDATA.indication
  - occurs when the status parameter is determined.
  - if status' value equals "no-error", the parameters of the CTF-M_UNITDATA.request, except the status parameter, are used as parameters for the M_UNITDATA.indication primitive.
    - if not, the resulting M_UNITDATA.indication primitive is not generated, i.e. the frame is discarded.
- M_UNITDATA.request to CTF-M_UNITDATA.indication
  - occurs instantly.
  - the parameters of the CTF-M_UNITDATA.indication are copied from the M_UNITDATA.request primitive and the status parameter's value is set to "no-error".
- M_UNITDATA.request to CTF-M_UNITDATA.request
  - occurs instantly.
  - the parameters of the CTF-M_UNITDATA.request are copied from the M_UNITDATA.request primitive and the status parameter's value is set to "no-error".

4 CTF Bridge forwarding process

- Same organization as in 802.1DU D0.2

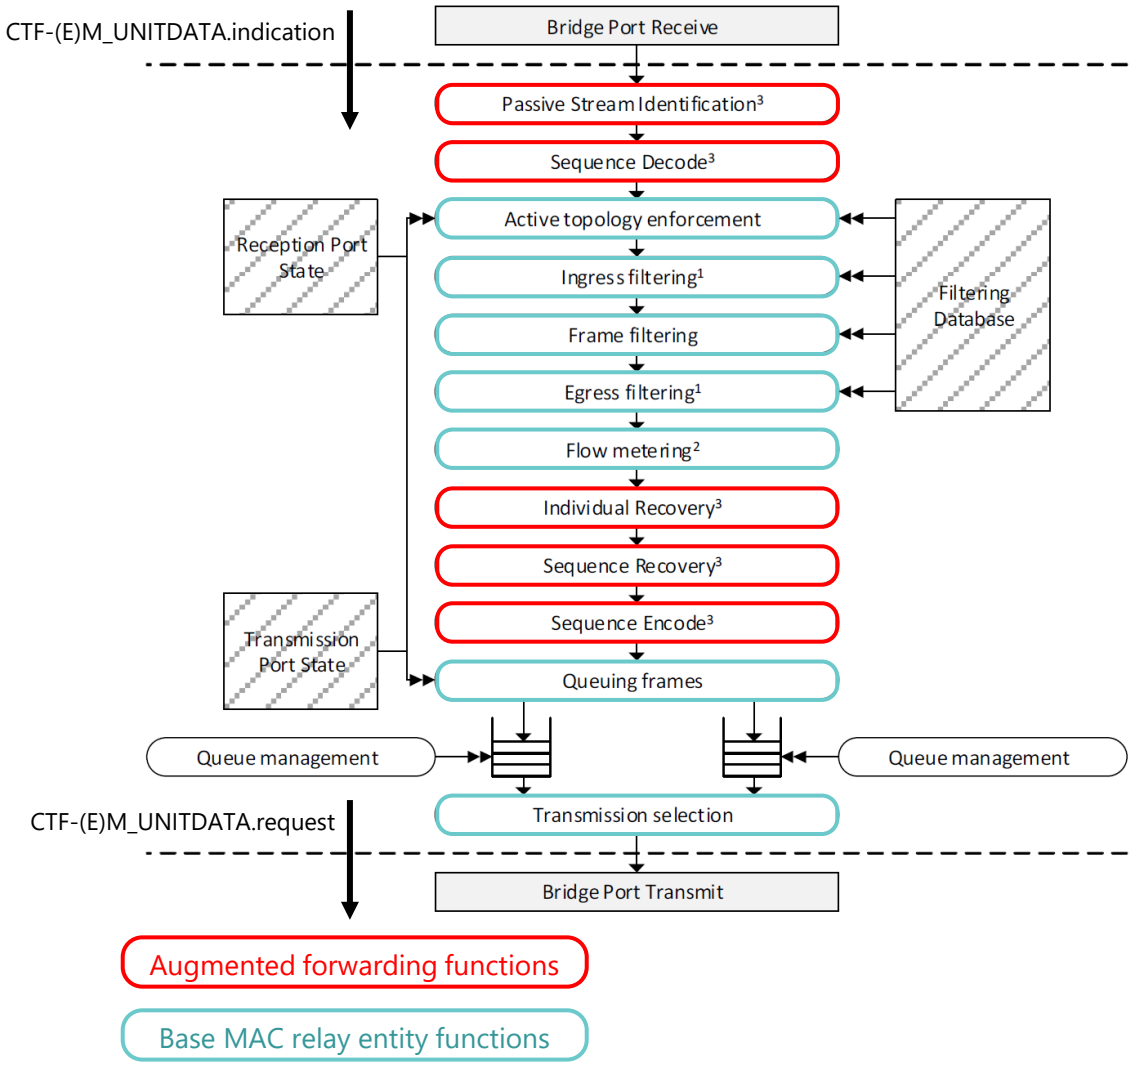

- Principle
  - The Forwarding Process is modeled as receiving a frame as the parameters of a CTF-(E)M_UNITDATA.indication and transmitting through supplying the parameters of a CTF-(E)M_UNITDATA.request.
  - As in an IEEE 802.1Q bridge, each frame submitted to the MAC Relay Entity is forwarded by executing the functions of the Forwarding Process.
  - Each function is described in terms of the action taken for a given frame being received on a given Port
  - A frame can be forwarded for transmission on some Ports and discarded without being transmitted at the other Ports.
  - In a Cut-Through bridge, the execution of these functions can occur before the frame is totally received, **each function using as input the set of parameters it requires for its execution**.
  - **The CTF-ISS or the CTF-EISS generates the indication primitives when the parameters required for the execution of** *all the functions of the MAC Relay Entity* **are available. The parameters not required by the execution of the functions may be incomplete or have a pending value.**
  - The **frame length** is an implicit parameter (not a primitive parameter), that is required by some functions for their execution. It is continuously evaluated by the function with the progressive provision of the frame parameters by the indication primitive.

Ingress augmented forwarding functions (802.1CB)

- Passive Stream identification
  - Performed at the CTF-ISS or CTF-EISS
  - Generates the stream_handle subparameter on frames passed up the stack based on a set of the frame's parameters.
  - Temporal behavior :
    - The CTF-(E)M_UNITDATA.indication is presented to the set of Stream identification functions once the **frame parameters, or parts thereof, required by the set of Stream identification functions** are obtained from the underlying LAN.
    - The corresponding **stream_handle**-enhanced CTF-(E)M_UNITDATA.indication primitive is generated instantaneously
- Sequence decode function
  - When presented by the lower layers with a service indication (CTF-(E)M_UNITDATA.indication), the Redundancy tag Sequence decode function performs the operations as per IEEE Std 802.1CB 7.8 a) and b).
    - R-Tag search at the beginning of the mac_service_data_unit
    - 6-octet R-Tag extraction from the mac_service_data_unit
  - Temporal behavior :
    - The CTF-(E)M_UNITDATA.indication is presented to the Redundacy tag decode function once the **intial 6 octets of the mac_service_data_unit** parameter are obtained from the underlying LAN.
    - The regenerated CTF-(E)M_UNITDATA.indication primitive including the **sequence_number** subparameter and the possibly shortened mac_service_data_unit parameter is generated instantaneously.

Base CTF Bridge MAC relay entity (802.1Q)

- Active topology enforcement
  - Learning
    - Addition of dynamic entries in the FDB based on the source_address and possibly the vlan_identifier parameters
      - Same conditions as 802.1Q 8.4 and 8.6.1
    - Temporal behavior :
      - The CTF-(E)M_UNITDATA.indication is presented to the learning function once the **source_address (and the vlan_identifier)** parameter(s) are obtained from the underlying LAN.
      - **The update of the FDB is delayed until the status parameter is determined with a value of "no_error"**
        - i.e. upon reception of a complete valid frame
  - Initial set of potential transmission ports
    - As per 802.1Q 8.6.1

Base CTF Bridge MAC relay entity (802.1Q)

- Ingress filtering
  - As per 802.1Q 8.6.2
  - Temporal behavior
    - The CTF-EM_UNITDATA.indication is presented to the ingress filtering function once the **vlan_identifier** parameter is obtained from the underlying LAN.
- Frame filtering
  - Reduction of the set of potential transmission ports based on the destination_address and possibly the vlan_identifier parameters
  - Temporal behavior
    - The CTF-(E)M_UNITDATA.indication is presented to the frame filtering function once the **destination_address (and possibly the vlan_identifier)** parameter(s) are obtained from the underlying LAN.
    - if the query evaluation by the FDB results in “flooding”, the result of the frame filtering function is **delayed until the status parameter is determined**.
      - if the status parameter’s value equals “no_error” then the FDB query result is applied
      - otherwise, the frame is discarded
- Egress filtering
  - As per 802.1Q 8.6.4
  - Temporal behavior
    - The CTF-EM_UNITDATA.indication is presented to the egress filtering function once the **vlan_identifier** parameter is obtained from the underlying LAN.

Base CTF Bridge MAC relay entity (802.1Q)

- Flow classification and metering (1)
  - Stream filtering
    - As per 802.1Q 8.6.5.3
    - Temporal behavior:
      - The CTF-(E)M_UNITDATA.indication is presented to the Stream filtering function once the **stream_handle** is determined.
  - Maximum SDU size filtering
    - Based 802.1Q 8.6.5.3.1
    - Temporal behavior:
      - The function **continuously evaluates the SDU size** of the received frame based on its progressively received parameters. If the SDU size exceeds the value of the associated stream filter's Maximum SDU size parameter, the CTF-(E)M_UNITDATA.indication primitive's **status parameter is set to determined with an error value set to "NotPassingSDU"**. [Frame truncation ?]
  - Stream gating
    - Based 802.1Q 8.6.5.4
    - Temporal behavior:
      - When a CTF-(E)M_UNITDATA.indication is presented to a stream gate, the frame is relayed if the gate is in an open state. The frame is discarded otherwise.
      - If a stream gate closes prior to the determination of the status parameter of the CTF-(E)M_UNITDATA.indication (end of the frame under reception), the **status parameter is set to determined with an error value set to "SgGateClosed"**. [Frame truncation ?]
      - If the length of the frame reaches a value greater than IntervalOctetsLeft, the CTF-(E)M_UNITDATA.indication's **status parameter is set to determined with an error value set to "SgOctetsExceeded"**. [Frame truncation ?]

- Flow classification and metering (2)
  - Flow metering
    - Based 802.1Q 8.6.5.5
    - Temporal behavior:
      - The CTF-EM_UNITDATA.indication is presented to the flow mering function once the **stream_handle** parameter (hence the **drop_eligible** parameter) is obtained from the underlying LAN.
      - The frame is discarded if its **length**, at the time the CTF-EM_UNITDATA.indication is presented to the function, causes a discard decision to be taken by the flow metering algorithm.
      - The frame is marked as drop-eligible if its **length**, at the time the CTF-EM_UNITDATA.indication is presented to the function, causes a mark-as-drop-eligible decision to be taken by the flow metering algorithm. The **drop_eligible** parameter is updated accordingly
      - Otherwise, the flow metering function continuously evaluates the length of the received frame based on its progressively received parameters. If it reaches a value that triggers a discard or drop-eligible marking decision, the CTF-EM_UNITDATA.indication's **status parameter is set to determined with an error reason set to "FlowMeterDiscard" or "LateDropEligible"**. [Frame truncation ?, flow metering algo reset ?]

Egress augmented forwarding functions (802.1CB)

- Individual recovery
  - Performed at the ISS “coming up” from the internal LAN between the bridge forwarding function and an egress port
  - Temporal behavior :
    - The CTF-M_UNITDATA.indication is presented to the individual recovery function once the **status** parameter is determined (end of frame reception)
      - If the **status parameter’s value equals no_error**, the function is instantaneously invoked and the CTF-M_UNITDATA.indication regenerated accordingly.
      - If not, the frame is discarded.
- Sequence recovery
  - Performed at the ISS “coming up” from the internal LAN between the bridge forwarding function and an egress port.
  - Temporal behavior :
    - The CTF-M_UNITDATA.indication is presented to the individual recovery function once the **status** parameter is determined (end of frame reception)
      - If **the status parameter’s value equals no_error**, the function is instantaneously invoked and the CTF-M_UNITDATA.indication regenerated accordingly.
      - If not, the frame is discarded.
- Sequence encode function
  - When presented by the lower layers with a service indication (CTF-(E)M_UNITDATA.indication), the Redundancy tag Sequence decode function performs the operations as in IEEE Std 802.1CB 7.8 c) and d).
    - R-Tag information creation (incl. sequence number)
    - 6-octet R-Tag insertion in the **mac_service_data_unit**
  - Temporal behavior :
    - The CTF-(E)M_UNITDATA.indication carrying the **6-octet longer mac_service_data_unit**, is generated instantaneously.

Egress queuing (802.1Q)

- Queuing frames
  - As per 802.1Q 8.6.6
    - Places the parameters of a CTF-(E)M_UNITDATA.request (to be generated) on an outbound queue
  - Temporal behavior :
    - The **parameters** provided by CTF-(E)M_UNITDATA.indication **available** at the time it is presented to the frame queuing stage are the parameters of the corresponding CTF-(E)M_UNITDATA.request.
    - If the speed of the link attached to the outbound queue is greater than the speed of the link the frame is received from, the frame is queued only once the **status** parameter is determined with a value of no-error.
- Queue management
  - As per 802.1Q 8.6.7
    - Removes frames from a queue (once transmitted) according to some rules
  - Temporal behavior :
    - A frame is removed from a queue if:
      - The CTF-(E)M_UNITDATA.request is generated with a determined **status** parameter of value **no_error** (fully received frame)
      - The **status** parameter is determined (**error or no_error**) after the CTF-(E)M_UNITDATA.request is generated.
      - The **status** parameter is determined with a value different of no_error before the CTF-(E)M_UNITDATA.request is generated.
- Transmission selection
  - As per 802.1Q 8.6.8
    - Determines frame transmissions ordering and times, dequeues frames and initiates transmission
  - CTF behavior :
    - Limited to the strict-priority algorithm (802.1Q 8.6.8.1)

Frames transmitting (derived from 802.1Q)

- Data transmit request
  - Once a frame is selected for transmission, the corresponding CTF-(E)M_UNITDATA.request is generated.
  - Temporal behavior :
    - If the frame is totally queued (all parameters of the CTF-(E)M_UNITDATA.indication are received) at the time it is selected for transmission, the parameters of the request primitive are populated with the parameters of the queued frame.
      - In that case, the status parameter has a value of no_error
      - The frame_check_sequence parameter is updated
    - If the frame is under reception, i.e. the status parameter is pending in the corresponding CTF-(E)M_UNITDATA.indication, the parameters of the corresponding CTF-(E)M_UNITDATA.request primitive are populated with the parameters already completed and those being progressively provided by the indication primitive, including the pending status parameter.
      - If the status parameter is determined with a value equal to “no_error”, the frame_check_sequence parameter is updated (valid value)
      - Otherwise the frame_check_sequence parameter is set to a specific value (other options possible ?)

Forwarding functions effect

Function	Input parameters	Status update	Late/early function result/decision	
Passive stream identification	destination_address, source_address, (vlan_identifier), [mac_service_data_unit]n	no	none	
Sequence decode function	[mac_service_data_unit]6	no	none	
Learning Active topology enforcement	source_address, (vlan_identifier)	no	FDB entry addition delayed, validated if status = no_error	
Ingress filtering	vlan_identifier	no	none	
Frame filtering	destination_address, (vlan_identifier)	no	none	
Egress filtering	vlan_identifier	no	none	
Stream filtering Flow classification and metering	stream_handle(1)	no	none	
Maximum SDU size filtering Flow classification and metering	*current frame size*	yes	set status to « error » => frame truncation ?	
Stream gating Flow classification and metering	*current frame size*	yes	set status to « error » => frame truncation ?	
Flow metering Flow classification and metering	*current frame size*, drop_eligible	yes	set status to « error », drop_eligible update algo backtracking ? frame truncation ?	
Individual recovery	status (= no-error)(2), [mac_service_data_unit]6	no	none	
Sequence recovery	status (= no_error)(2), [mac_service_data_unit]6	no	none	
Sequence encode function	source_address or [mac_service_data_unit]1	no	none	
Queuing frames	*whatever is available in the incoming indication primitive*	no	wait for status = no_error if egress port speed > ingress port speed	

[mac_service_data_unit]n : n first octets of the mac_service_data_unit parameter

(1) : stream_handle is available when the parameters required by the passive stream identifcation are available

(2) : status = no_error indicates that the frame is completely received (all primitive's parameters available and error-free)

Late : upon complete frame reception

Early : before complete frame reception

Wrap-up

- No extra CTF-specific (sub)layer
  - The CTF-(E)ISS provides for the CTF features
  - 802.1Q bridge model is kept as is
- CTF-(E)ISS primitives are enhanced (E)ISS primitives
  - 1 additional locally-significant parameter : status
- The forwarding functions don’t perform stalls
  - No frame “buffering” internal to the functions
  - Stalls only handled by the CTF-(E)ISS
- Forwarding functions capable of:
  - Signalling errors on the fly
  - Taking interim decision
    - Validated/invalidated a posteriori when their input parameters are fully determined
      - Typically at the end of a frame reception


ANY
QUESTIONS
?

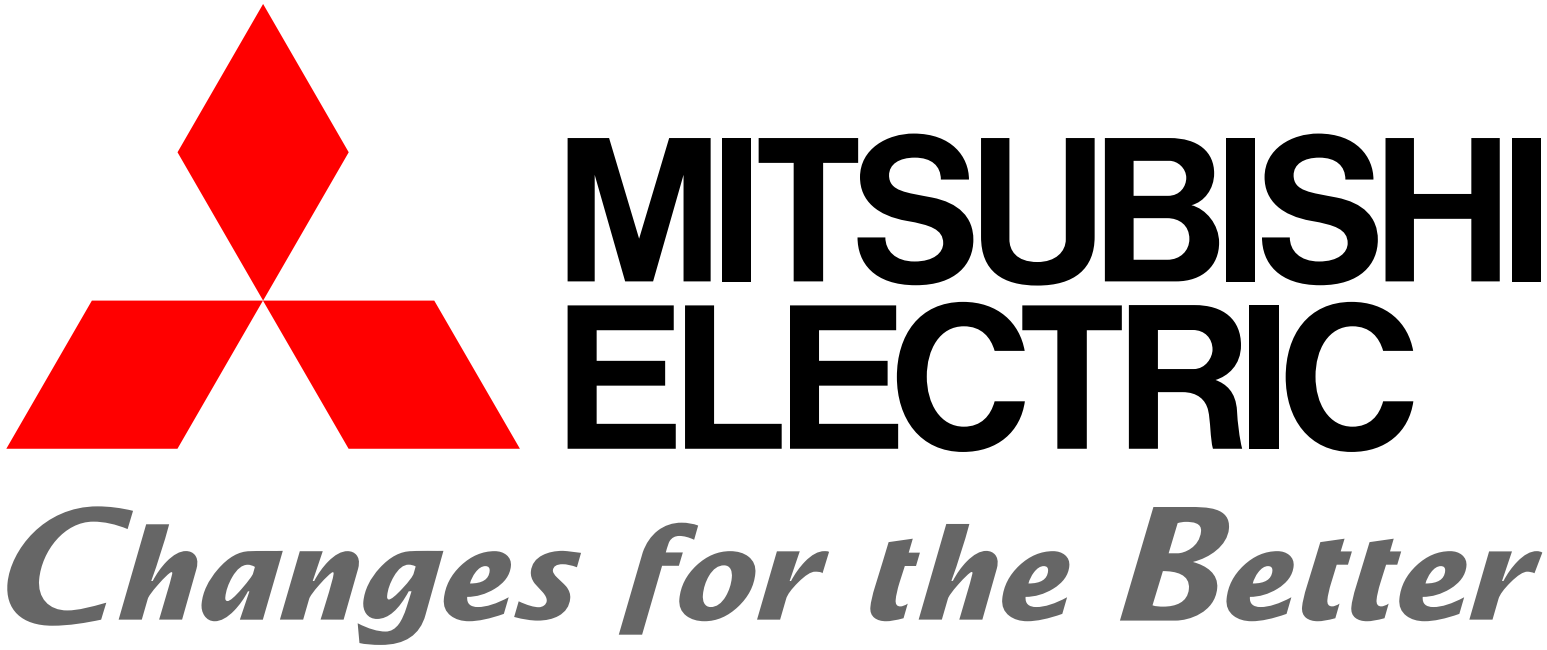
MITSUBISHI
ELECTRIC
Changes for the Better