

# CSD Compatibility Criterion for Cut-Through Forwarding

Roger Marks  
(EthAirNet Associates; Huawei)

[roger@ethair.net](mailto:roger@ethair.net)

+1 802 capable

2022-06-21

Copyright policy note: “Q: Do I need permission if I am using IEEE standards for comparison or discussion only? A: IEEE SA Working Groups are permitted to present (i.e., display during a Working Group meeting) and review information from an approved IEEE standard.”

• <https://standards.ieee.org/faqs/copyrights/working-group-and-activity-chairs/>

# Compatibility CSD

- *IEEE 802 criteria for standards development (CSD)*
- *14.2.2 Compatibility*
  - *Each proposed IEEE 802 LMSC standard should be in conformance with IEEE Std 802, IEEE 802.1AC, and IEEE 802.1Q. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with IEEE 802.1 Working Group prior to submitting a PAR to the IEEE 802 LMSC.*
    - *a) Will the proposed standard comply with IEEE Std 802, IEEE Std 802.1AC and IEEE Std 802.1Q?*
    - *b) If the answer to a) is no, supply the response from the IEEE 802.1 Working Group.*
- In some cases, may be useful to consider compatibility with other relevant 802 standards as well.
- Note: Throughout this document, italicized text is a direct quote.

# My Recommendations

- A consensus understanding the MAC Service is a precursor to understanding compatibility.
  - This may require interpretations of existing standards.
- In addition to dealing with CSD Compatibility, let's look for ways to improve the architectural clarity.
  - This can be a learning opportunity for all of us.
- No parties should claim exclusive rights to interpretation.
  - 802.3 and 802.1 participants can scrutinize 802.1 standards, ask questions, and examine responses.
    - If we don't understand, help us understand.
  - 802.1 and 802.3 participants can scrutinize 802.3 standard and ask questions, and examine responses.
    - If we don't understand, help us understand.
- Please don't take offense when questions are raised.
  - When it comes to architectural issues, there are some inconsistencies and ambiguities we need to resolve.

# MAC Service

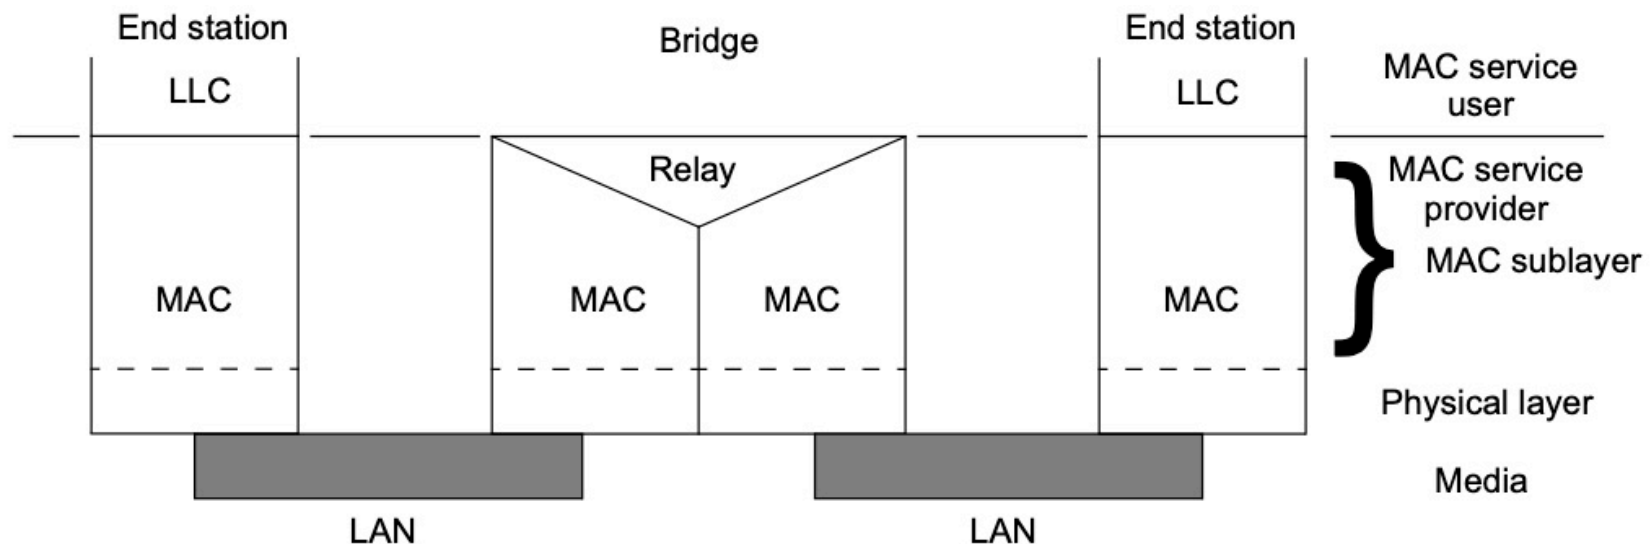
- Starting point for analysis: Seek a common understanding, among key standards, regarding the MAC Service:
  - How is the MAC Service described in various standards?
  - Clients of the MAC Service
  - The primitives and their names
  - The primitive parameters
  - When the primitives are generated

# Specifications addressing MAC Service

- Sources considered here:
  - IEEE Std 802-2014
    - *IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture*
  - IEEE Std 802.1Q-2018
    - *IEEE Standard for Local and Metropolitan Area Networks — Bridges and Bridged Networks*
  - IEEE Std 802.1AC-2016
    - *IEEE Standard for Local and metropolitan area networks — Media Access Control (MAC) Service Definition*
  - IEEE Std 802.3-2018
    - *IEEE Standard for Ethernet*
  - ISO/IEC 8802-2
    - *Logical Link Control*
    - referred to herein as IEEE Std 802.2-1998 (withdrawn)

# MAC Service per IEEE Std 802

- *The MAC sublayer provides one or more MAC service access points (MSAPs) as interfaces to the LLC sublayer in an end station.*
- *Figure 7 shows the position of the bridging functions within the MAC sublayer; note particularly that relaying and filtering are considered to belong entirely within the MAC sublayer.*



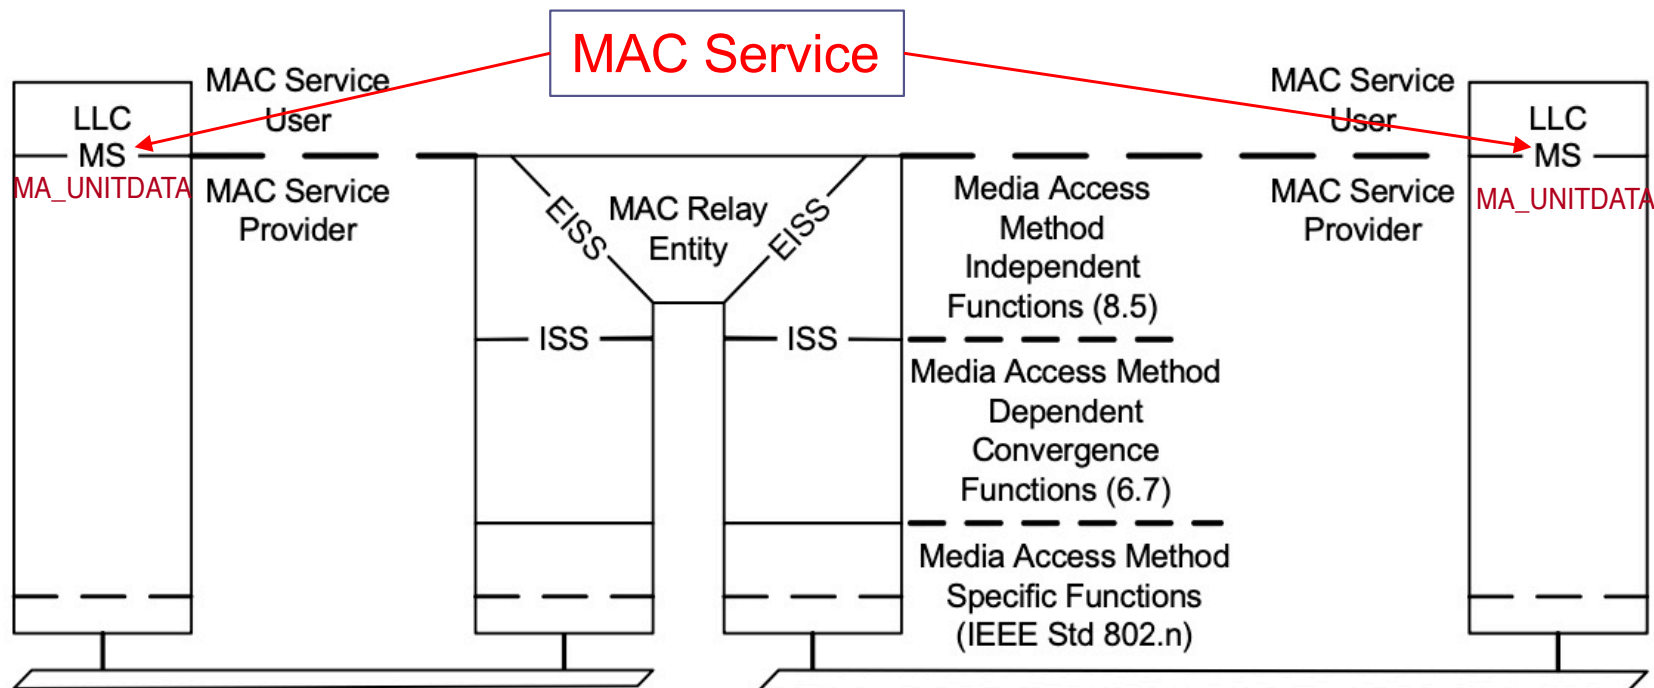
**Figure 7—Internal organization of the MAC sublayer with bridging**

# MAC Service per IEEE Std 802.2

- *services required of the MAC sublayer... to allow the local LLC sublayer entity to exchange LLC data units with peer LLC sublayer entities*
  - *MA-UNITDATA request:*
    - *requests the transfer of an MSDU from a local LLC sublayer entity to a single peer LLC sub-layer entity [multiple for multicast]*
    - *parameters: SA, DA, data...*
  - *MA-UNITDATA indication:*
    - *defines the transfer of a MSDU from the MAC sublayer entity to the LLC sublayer entity... In the absence of errors, the contents of the data parameter are logically complete and unchanged relative to the data parameter in the associated MA-UNITDATA request primitive*
    - ***When generated:*** *The MA-UNITDATA indication primitive is passed from the MAC sublayer entity to the LLC sublayer entity or entities to indicate the arrival of a frame at the local MAC sublayer entity. Frames are reported only if at the MAC sublayer they are validly formatted, received without error, and their destination address designates the local MAC sublayer entity.*
- **Note:** 802.2 seems to use “MA-UNITDATA” and “MA\_UNITDATA” interchangeably; most of the other standards do also.

# MAC Service per IEEE Std 802.1Q (1/2)


- *Abstract: This standard specifies how the Media Access Control (MAC) Service is supported by Bridged Networks...*
  - *6. Support of the MAC Service*
  - *MAC Service provided in end stations attached to MAC Bridged Networks and Virtual Bridged Networks is the (unconfirmed) connectionless mode MAC Service defined in IEEE Std 802.1AC*
  - *describes the transfer of user data between source and destination end stations, via MA-UNITDATA request primitives and corresponding MA-UNITDATA indication primitives issued at MAC Service Access Points (MSAPs). Each MA-UNITDATA request and indication primitive has four parameters: Destination Address, Source Address, MAC Service data unit (MSDU), and Priority.*





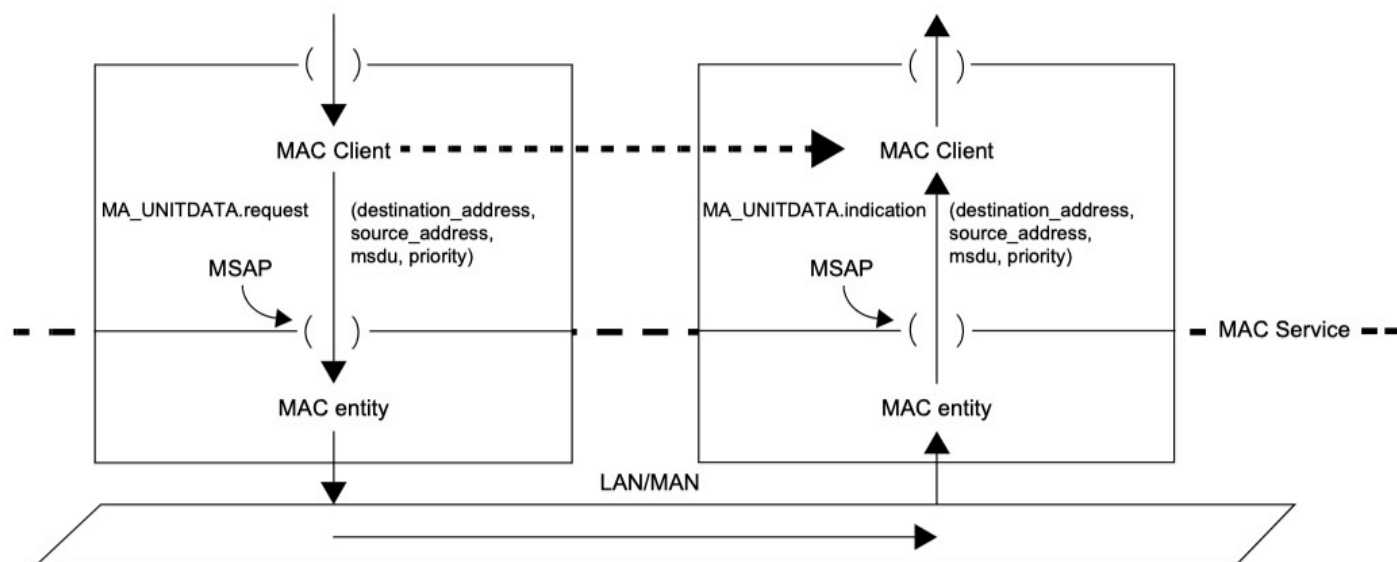
# MAC Service per IEEE Std 802.1Q (2/2)

- MAC services terminating at the bridge itself (special case), not further discussed here.



# MAC Service per IEEE 802.1AC (1/2)

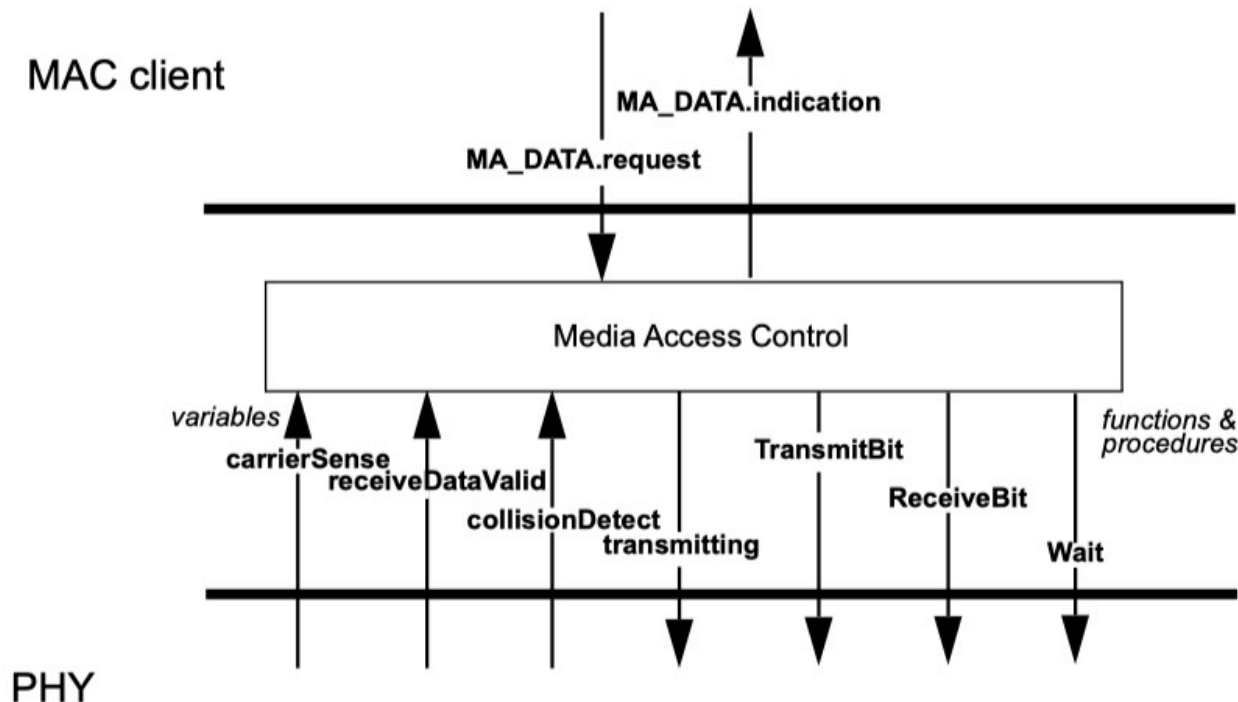
- *The MAC Service provides a connectionless-mode service for the transparent transfer of data between MAC Service users.*
  - *provides means for the MAC Service user to identify its MSAP and to specify the MSAP or MSAPs to which data is to be transferred.*



- *Two unit-data primitives are specified for the connectionless-mode data transmission service, an MA\_UNITDATA.indication and an MA\_UNITDATA.request, together with the parameters of those primitives. Each MA\_UNITDATA indication corresponds to the receipt of an error-free MAC frame from a LAN. A data request primitive is invoked to transmit a frame to an individual LAN.*
  - MA\_UNITDATA.request and MA\_UNITDATA.indication are not detailed, as they are in IEEE 802.2. For example, no “when generated” statement.
  - Without a “when generated” statement, it’s not apparent whether the MAC entity is responsible to filter frames by DA.

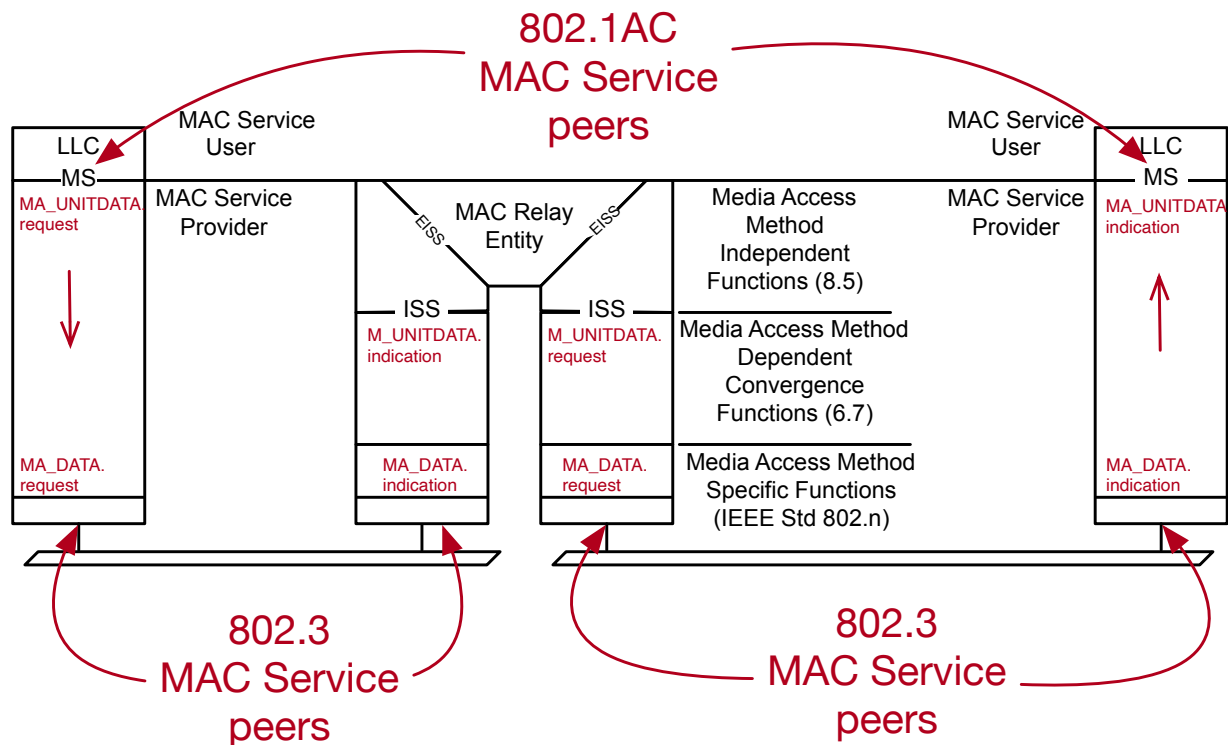
# MAC Service per IEEE 802.3 (1 / 3)

- *Media Access Control (MAC) service specification*
  - *The services provided by the MAC sublayer allow the local MAC client entity to exchange LLC data units with peer LLC sublayer entities.*
  - *MAC clients may include the Logical Link Control (LLC) sublayer, Bridge Relay Entity, or other users of ISO/IEC LAN International Standard MAC services*



# 802.1AC (2/2)

- (MAC Service primitives) can be used to transmit an independent, self-contained MSDU from one MSAP to another MSAP in a single service access
- [but, at a bridge ISS] Each MAC entity examines all frames received on the LAN to which it is attached. All error-free received user data frames give rise to M\_UNITDATA indication primitives.
- For the Ethernet Convergence Function:
  - When the convergence function receives an MA\_DATA.indication primitive from the underlying MAC Service, it generates a corresponding M\_UNITDATA.indication ...
  - When the convergence function receives an M\_UNITDATA.request primitive, it generates a corresponding MA\_DATA.request to the underlying MAC Service:...



# MAC Service per IEEE 802.3 (2/3)

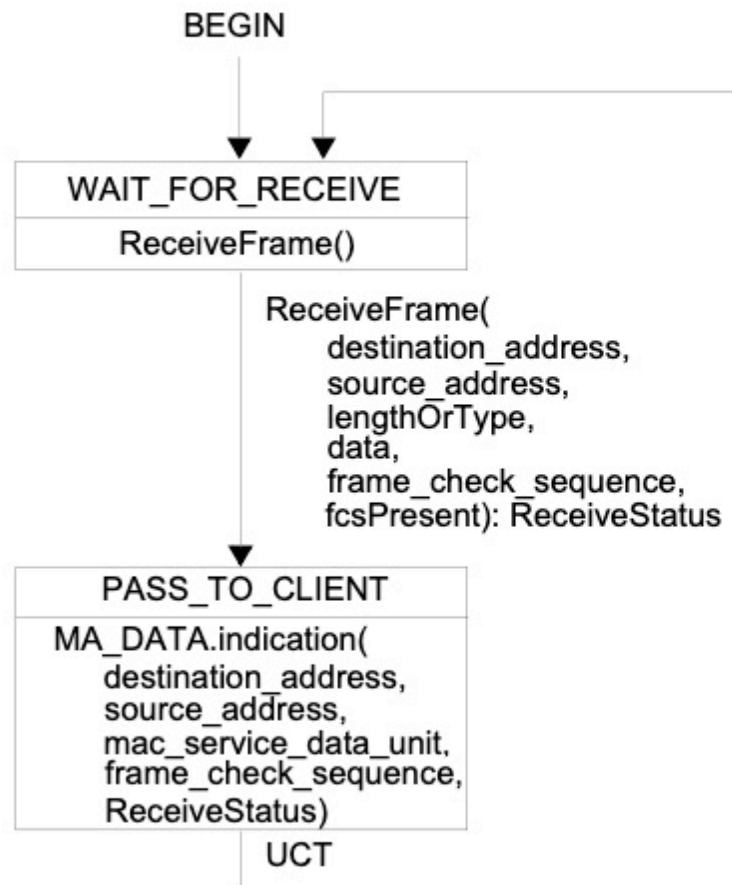
- *2.3 Detailed service specification [MAC Service]*
  - *MA\_DATA.indication ... defines the transfer of data from the MAC sublayer entity... to the MAC client entity or entities....*
  - *When generated: MA\_DATA.indication is passed from the MAC sublayer entity... to the MAC client entity or entities to indicate the arrival of a frame to the local MAC sublayer entity that is destined for the MAC client. Such frames are reported only if they are validly formed, received without error, and their destination address designates the local MAC entity.*
    - This description matches the functionality required at the end station.
    - But it seems no MA\_DATA.indication is passed to a bridge since the DA is not local.
- *4/4A Frame reception*
  - *ReceiveFrame ... calls the internal function ReceiveDataDecap to return the frame's fields to the MAC client if the frame's address indicates that it should do so. [see function LayerMgmtRecognizeAddress]*
  - *The MAC sublayer may also provide the capability of operating in the promiscuous receive mode. In this mode of operation, the MAC sublayer recognizes and accepts all valid frames, regardless of their Destination Address field values.*
    - but promiscuous receive mode seems to not provide the 802.3 MAC Service of Clause 2 and not generate MA\_DATA.indication, per Clause 2
      - **However, see next slide for the view from Clause 4/4A.**

# MAC Service per IEEE 802.3 (3/3)

*Note whenever there is any apparent ambiguity concerning the definition of some aspect of the MAC method, it is the Pascal procedural specification in 4A.2.7 through 4A.2.10 that should be consulted for the definitive statement. That excludes this figure.*

## 4A.3.2.2.4 MAC client receive interface state diagram

Figure 4A–4 specifies the behavior of the receive interface to the MAC client.



MAC passes received frames to client, regardless of frame errors and (in promiscuous mode) regardless of destination address.

Error details (in ReceiveStatus) are passed to Client, though “ReceiveStatus is not mapped to any MAC client parameter by the service interface defined in 2.3.2.” But, per 2.3.2.2, it “is used to pass status information to the MAC client entity.”

802.3 (2.3.2.5) says this parameter “is not mapped to any parameter and is ignored by MA\_UNITDATA.indication” (and the same for M\_UNITDATA.indication).

802.1AC maps MA\_DATA.indication to M\_UNITDATA.indication but does not refer to the ReceiveStatus field.

Figure 4A–4—MAC client receive interface state diagram

# Example Issues

- IEEE Std 802.3 (§3.4) says:
  - *The contents of invalid MAC frames shall not be passed to the LLC or MAC Control sublayers. Invalid MAC frames may be ignored, discarded, or used in a private manner by MAC clients other than LLC or MAC control.*
    - IEEE 802.3 says “MAC clients may include the Logical Link Control (LLC) sublayer, Bridge Relay Entity, or other users of ISO/IEC LAN International Standard MAC services”
      - But the ban on passing invalid frames applies only to LLC and MAC Control.
      - So, are invalid MAC frames always passed (or can they be passed) to the bridge Convergence Function?
- It’s been said that the MAC does not know whether it’s sitting under an LLC entity or a bridge.
  - But the 802.3 MAC does know the difference:
    - If the client is LLC, the MAC filters by DA and generates MA\_DATA.indication.
    - If the client is the Convergence Function, it seems that promiscuous mode must be used.
- If the DA-filtering behavior depends on the nature of the client, then can other behaviors (e.g., error-filtering behavior) also depend on the nature of the Client?

# Extra Issue: Implementation Interface

- 802.3 §4A.2.9 :
  - *The ReceiveFrame operation is synchronous. The operation does not complete until a frame has been received.*
- 802.3 §4A.3.1 :
  - *The synchronous (one frame at a time) nature of the frame transmission and reception operations is a property of the architectural interface between the MAC client and MAC sublayers, and need not be reflected in the implementation interface between a station and its sublayer.*
- Question:
  - What is the distinction between the “architectural interface” and the “implementation interface,” and how do they relate?
  - Are there important aspects of the distinction between “architectural interface” and “implementation interface” that need to be recognized by adjacent sublayers; e.g. 802.1AC?



# Implementation Interface

- 802.3 Scope
  - *Speed specific Media Independent Interfaces (MIIs) provide an architectural and optional implementation interface to selected Physical Layer entities (PHY).*
  - This is the only other use of “implementation interface” in IEEE 802.3.
  - Since it’s in the Scope, it seems to be essential.
- 802.3 §4A.2.2.1:
  - *The handling of incoming and outgoing frames is rather stylized in the procedural model, in the sense that frames are handled as single entities by most of the MAC sublayer and are only serialized for presentation to the Physical Layer. In reality, many implementations will instead handle frames serially on a bit, octet or word basis. This approach has not been reflected in the procedural model, since this only complicates the description of the functions without changing them in any way.*
  - Is this use of “implementation” related to “implementation interface”?
- Is “implementation” related to “conformance”?
  - As in the term “Protocol implementation conformance statement (PICS)”?

# Questions

- Is the MAC Service as understood by 802.1AC the same MAC Service as understood by 802.3?
  - And what does 802.1AC mean by “underlying MAC service”?
- Should 802.1AC specify “when generated” for MA\_UNITDATA.indication?
- Is Ethernet Convergence Function of 802.1AC consistent with IEEE Std 802.3?
  - It is based on on MA\_DATA.indication, but is MA\_DATA.indication issued to the relay?
  - Should 802.1AC specify promiscuous receive mode at ISS?
- Does 802.3 specify how operation is differentiated when the client is the Convergence Function?
  - Does it specify that promiscuous receive mode is required?
  - Does it provide an alternative to the “MAC Service”?
  - Does it specify an alternative to MA\_DATA.indication?
  - Does 802.3 specify how the MAC identifies the nature of the client so it can determine when to use promiscuous receive mode?
- How is the “implementation interface” understood?
- Should the IEEE 802 architecture be amended to clarify the nature of the MAC Service and how the upper (LLC-to-LLC MAC) Service is related to the lower (LLC-to-bridge) service supported in a LAN link?
  - Should the terminology distinguish the end-to-end MAC from the MAC link?
- Is there a need for better understanding and harmonization of the architecture and interfaces among IEEE 802 standards?