Cut-Through Forwarding (CTF): Towards an IEEE 802.1 Standard

Johannes Specht

Introduction

Recap: Nendica weekly meeting on September 2, 2021 (1)

(https://mentor.ieee.org/802.1/dcn/21/1-21-0051-00-ICne-ctf-discussing-next-steps.pdf)

Based on the earlier work and associated discussions, it may be reasonable to distinguish between two topics, while working on both in 802 Nendica and with the WGs (IEEE WG 802.1 and IEEE WG 802.3)

- **Model**
  There is an existing model spanning across IEEE Std 802.3 and IEEE 802.1 (MAC Service Interface, Pascal Code, [E]ISS, timing, etc.).
- **Conformant implementations**
  Defined by the external visible behavior. This behavior can be derived from the elements of the model in IEEE 802.1 and IEEE 802.3 Standards.

A way to work on these topics

- Identify questions
- Work on technical answers in 802 Nendica
- Ask questions to the WGs/experts from both WGs

Recap: Nendica weekly meeting on September 2, 2021 (2)

(https://mentor.ieee.org/802.1/dcn/21/1-21-0051-00-ICne-ctf-discussing-next-steps.pdf)

Some initial questions

- **Q1**: Where are the limits of conformant implementations/what would violate the standardized external visible behavior?:
  - IEEE 802.1 conformant Bridge implementations appear to be no issue, at least if the behavior of CTF Bridges would be standardized in a new standalone IEEE 802.1 Standard (see earlier slides).
  - IEEE 802.3 conformant MAC implementations appear to be in the focus.
- **Q2**: Is there need/interest in IEEE WG 802.1 and IEEE WG 802.3 to introduce a common model with support for CTF?
- **Q3**: Could the CTF-Bridge behavior be specified in IEEE 802.1 …
  - … based on 802.3 conformant implementations, and …
  - without introducing ambiguity (compared to the already standardized Bridge behavior)?
- **Q4:** Assuming the answer to Q3 is TRUE, and the initial answer to Q2 is FALSE, is it possible to move forward towards standardizing CTF-Bridge behavior while permitting later addition of a common model if the answer to Q2 changes to TRUE?

Proposed Activities

1) IEEE 802 Nendica: Already ongoing

- Discuss administrative and technical aspects of potential lower layer modelling across IEEE 802.3 and IEEE 802.1 with support for CTF.
- Requirement (from **Q2**): There is sufficient need/interest in IEEE WG 802.1 and IEEE WG 802.3 to work on such a modelling.

2) IEEE WG 802.1: Proposed

- Motion to develop PAR&CSD for an IEEE 802.1 project to standardize CTF as standalone IEEE 802.1 standard (not amendments to 802.1 Standards). Proposed items in a PAR scope in the context of this slide set:
  - Include IEEE Std 802.3-2018 compatible ***real implementations*** in the supported MACs.
  - Incorporate/standardize IEEE 802.1 aspects of a model with support for CTF [see item 1) above], if such a model becomes available during such a project.

3) IEEE 802 Nendica/IEEE WG 802.1: Now … before the proposed motion in item 2)

- Discuss individual contributions that may support the aforementioned motion, and may be considered during development of the PAR&CSD [item 2)] by IEEE WG 802.1.

This slide set

Note: The location [IEEE 802 Nendica/IEEE WG 802.1] for item 3) should be discussed due to the ongoing structural changes.

Support for real implementations of IEEE Std 802.3-2018

Q1: Where are the limits of conformant implementations/what would violate the standardized external visible behavior?:

- IEEE 802.1 conformant Bridge implementations appear to be no issue, at least if the behavior of CTF Bridges would be standardized in a new standalone IEEE 802.1 Standard (see earlier slides).
- IEEE 802.3 conformant MAC implementations appear to be in the focus.

Q3: Could the CTF-Bridge behavior be specified in IEEE 802.1 …

… based on 802.3 conformant implementations, and …

without introducing ambiguity (compared to the already standardized Bridge behavior)?

Model v.s. Implementations (1)

It is important to distinguish, however, between the model and a real implementation. The ***model*** *is* ***optimized for simplicity and clarity*** *of presentation, while any realistic* ***implementation*** *shall place heavier* ***emphasis on*** *such constraints as* ***efficiency and suitability*** *to a particular implementation technology or computer architecture.* [4A.2.2 of IEEE Std 802.3-2018, "Overview of the Procedural Model"]

… ***it is the behavior of any MAC sublayer implementations that shall match the standard****, not their internal structure. The internal details of the procedural model are useful only to the extent that they help specify that behavior clearly and precisely.* [item b) in 4A.2.2.1 of IEEE Std 802.3-2018, "Ground rules for the procedural model"]

Model v.s. Implementations (2)

Model

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.* [item c) in 4A.2.2.1 of IEEE Std 802.3-2018, "Ground rules for the procedural model"]

Implementations

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.* [item c) in 4A.2.2.1 of IEEE Std 802.3-2018, "Ground rules for the procedural model"]

Observations & Considerations

- MAC *implementations* that handle frames serially appear conformant to IEEE Std 802.3-2018
- Serial handling fits well to the concept around "incomplete frames" proposed during the IEEE 802 Plenary Tutorial on CTF
  (see https://mentor.ieee.org/802.1/dcn/21/1-21-0037-00-ICne-ieee-802-tutorial-cut-through-forwarding-ctf-among-ethernet-networks.pdf , section "IEEE 802.1 Considerations")
- It is the behavior that matters. The behavior of a Bridge is visible via (a) management variables and **(b) by frame transmission, which can be a result of frame reception**.

Behavior
(on the relevant path from frame reception to frame transmission)

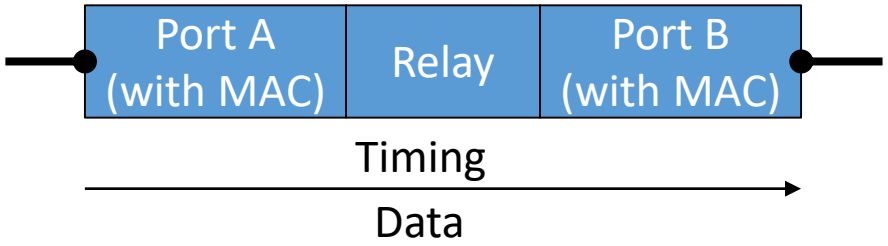

Two behavioral aspects

- **Data**
  - Reception
    - Which data is passed, potentially serially, from a Port A to the relay during reception?
    - Are contents of frames with invalid FCS available to the relay?
  - Transmission
    - Which data is to be passed, potentially serially, from the Relay to a Port B during transmission?
    - Can frames with invalid FCS be transmitted (e.g., avoid unintended “correction”)?
- **Timing**
  When are frames transmitted at a Port B as a result of frame reception at a Port A?

Data Flow: Receive Path

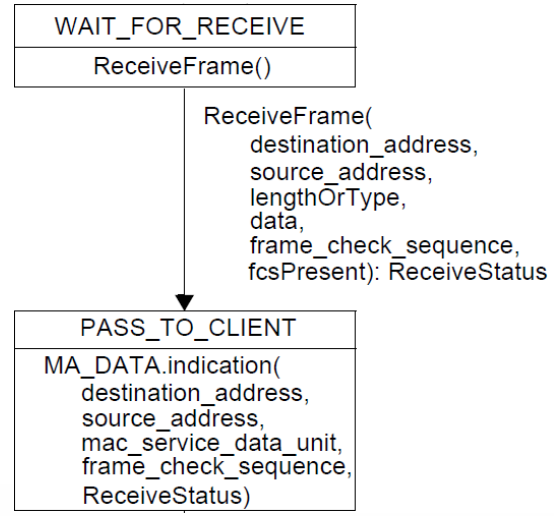

Source: Figure 4A-4 of IEEE Std 802.1-2018

From 4A.2.9 of IEEE Std 802.3-2018

procedure ReceiveLinkMgmt;
begin
  repeat
    StartReceive;
    while receiving *do* nothing; {Wait for frame to finish arriving}
    excessBits := frameSize *mod* 8;
    frameSize := frameSize – excessBits; {Truncate to octet boundary}
    receiveSucceeding := (frameSize ≥ minFrameSize) {Reject frames too small}
  until receiveSucceeding
end; {ReceiveLinkMgmt}

From 4A.2.9 of IEEE Std 802.3-2018

function **ReceiveFrame** (
  ...
  function ReceiveDataDecap: ReceiveStatus; {Nested function; see body below}
...
function ReceiveDataDecap: ReceiveStatus;
…
  receiveSucceeding := LayerMgmtRecognizeAddress(destinationField);
  if* receiveSucceeding *then
  begin {Disassemble MAC frame}
    destinationParam := destinationField;
    sourceParam := sourceField;
    lengthOrTypeParam := lengthOrTypeField;
    dataParam := RemovePad(lengthOrTypeField, dataField);
    fcsParamValue := fcsField;
    fcsParamPresent := passReceiveFCSMode;
    …
    if exceedsMaxLength *then* status := frameTooLong
    else if fcsField = CRC32(incomingFrame) *then*
    if validLength *then* status := receiveOK *else* status := lengthError
    else if excessBits = 0 *then* status := frameCheckError
    else status := alignmentError;
  …
  ReceiveDataDecap := status
end; {ReceiveDataDecap}

→ Contents of valid and **several invalid frames including frames with invalid FCS,** are transferred to MAC clients (e.g., Relay) other than MAC control and LLC (see also 3.4 of IEEE Std 802.3-2018)

Data Flow: Transmit Path

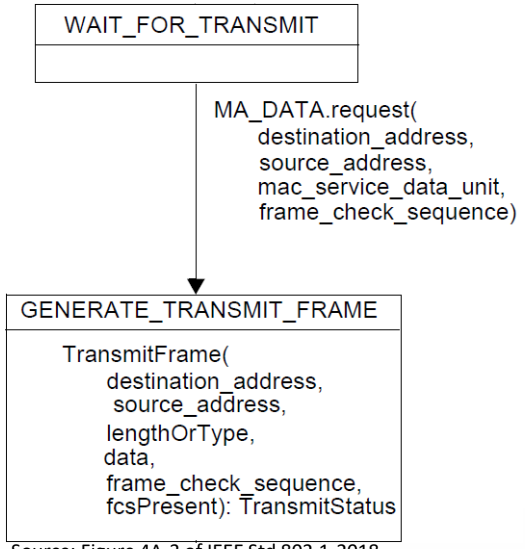

Source: Figure 4A-3 of IEEE Std 802.1-2018

From 4A.2.8 of IEEE Std 802.3-2018

```
process BitTransmitter;
begin {Inner loop}
...
  while transmitting do
  begin
    TransmitBit(outgoingFrame[currentTransmitBit]);
    currentTransmitBit := currentTransmitBit + 1;
    transmitting := (currentTransmitBit ≤ lastTransmitBit)
  end
...
```

From 4A.2.8 of IEEE Std 802.3-2018

```
function TransmitFrame (
  destinationParam: AddressValue;
  sourceParam: AddressValue;
  lengthOrTypeParam: LengthOrTypeValue;
  dataParam: DataValue;
  fcsParamValue: CRCValue;
  fcsParamPresent: Bit): TransmitStatus;
  procedure TransmitDataEncap; {Nested procedure; see body below}
  begin
    if transmitEnabled then
    begin
      TransmitDataEncap;
      TransmitFrame := TransmitLinkMgmt
…
```

From 4A.2.8 of IEEE Std 802.3-2018

```
procedure TransmitDataEncap;
begin
  with outgoingFrame do
  ...
  destinationField := destinationParam;
  sourceField := sourceParam;
  lengthOrTypeField := lengthOrTypeParam;
  if fcsParamPresent then
  begin
    dataField := dataParam; {No need to generate pad if the FCS is passed from MAC client}
    fcsField := fcsParamValue {Use the FCS passed from MAC client}
...
```

→ Contents of **frames, including frames with invalid FCS passed by the MAC client,** are transmitted

Timing from frame reception to frame transmission (1)

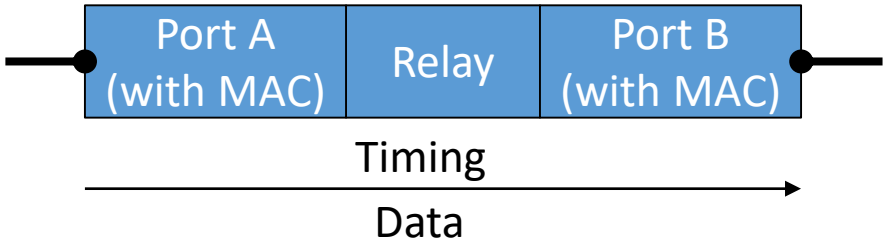

Existence of the timing

- Existence of a timing from frame reception at a Port A to an associated frame transmission at a Port B requires *something* (relay) between Port A and Port B.
- If absent, this behavioral aspect does not exist!

➢ It is appropriate to standardize the timing from frame reception to frame transmission in an IEEE 802.1 Standard.

The entire timing path (PHY device input → MAC → Relay → MAC → PHY device output) appears to be in scope of IEEE WG 802.1:

- 12.29.2 of IEEE Std 802.1Q-2018
- 12.32.1 of IEEE Std 802.1Qcc-2018
- 12.31.3.4 of IEEE Std 802.1Qcr-2020
- …

Timing from frame reception to frame transmission (2)

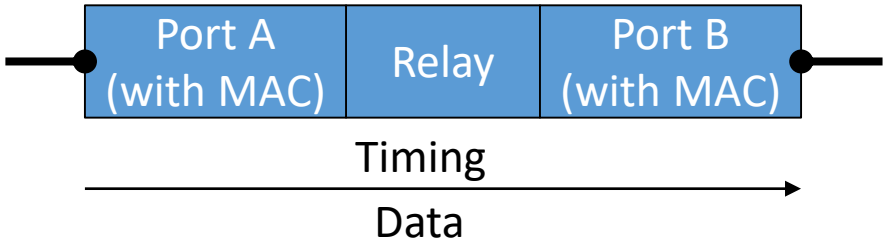

Real implementations of IEEE Std 802.3-2018

- The causal relationship and the (earliest) invocation times of the MAC service interface service primitives of the standardized model appears to be in conflict with the timing needed for CTF …
- … however:
  - Recap: Real implementations can handle frames serially
  - **Assumption: Real implementations of IEEE Std 802.3-2018 *are conformant* even if such implementations start such serial transfers earlier than the associated (earliest) service primitive invocation time.**
- Even if the assumption is wrong, it seems that CTF can be standardized in an IEEE 802.1 Standard nonetheless:
  - Recap: The timing from PHY device input to PHY device output appears to be in scope of an IEEE WG 802.1 standard.
  - If necessary at all for clarity of specification, another way to characterize timing is in relationship to the earliest invocation times of the standardized model. Figuratively, there may be some “minus”.

IEEE 802.1 Standalone Standard v.s. Amendments

IEEE 802 Plenary Tutorial: Recap. (1)

Proposed Location in IEEE 802.1 Standards

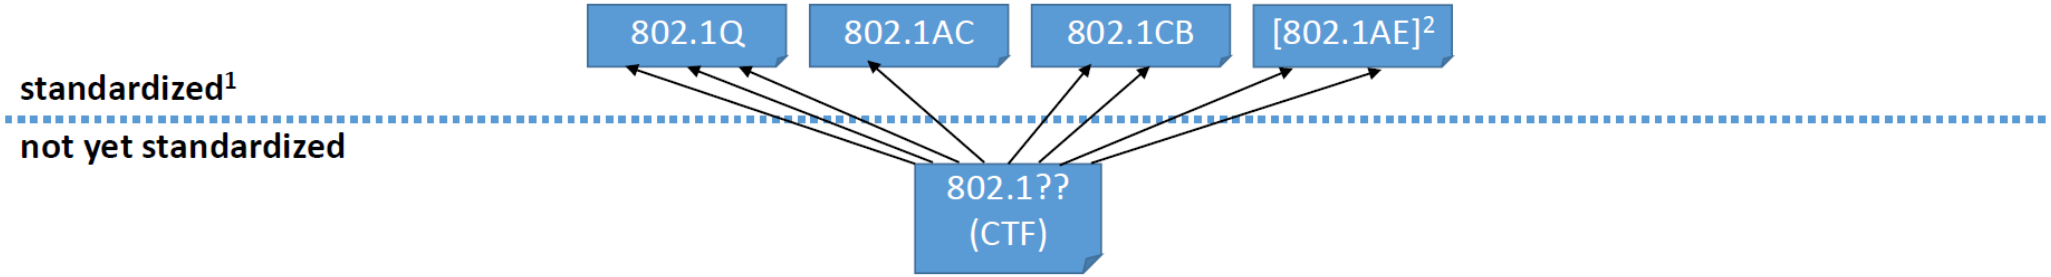

Separate stand alone IEEE 802.1 base standard for CTF

- **Single document**
  Avoids distribution of CTF across existing standards (compared to multiple amendment projects).
- **Exclusion, inclusion/re-use and adjustment of existing protocols**
  - Existing protocols not referred to are basically beyond specification.
  - If no adjustments for CTF are needed:
    Inclusion by reference (e.g., *"as specified in x.y.z of IEEE Std 802.1Xxx-20XX"*) can be sufficient.
  - If adjustments for CTF are needed:
    - Additional description of the differences can be sufficient.
    - Adjustments apply for CTF only; no side effects on existing protocols in absence of CTF support.

1) The latest published base standards, corrigenda, and a subset of the published amendments.
2) See later slide on security/privacy.

July 7, 2021 IEEE 802 Tutorial: Cut-Through Forwarding (CTF) among Ethernet networks 44

Relationship to the Criteria for Standards Development (CSD)

1.2.2 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 WG prior to submitting a PAR to the Sponsor.

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 WG.

The review and response is not required if the proposed standard is an amendment or revision to an existing standard for which it has been previously determined that compliance with the above IEEE 802 standards is not possible. In this case, the CSD statement shall state that this is the case.

A standalone 802.1 Standard for CTF would not be entirely in conformance with IEEE Std 802, IEEE 802.1AC, and IEEE 802.1Q. And this has advantages:

- Only a subset of protocols and protocol procedures in IEEE Std 802.1Q should be combined with CTF. The others would simply not be included (see previous slide).
- There is no risk for existing products conformant to IEEE Std 802, IEEE Std 802.1AC and IEEE Std 802.1Q to loose conformity by standardizing CTF.

IEEE 802 Plenary Tutorial: Recap. (2)

Proposed Content Categories and (some) Contents[1]

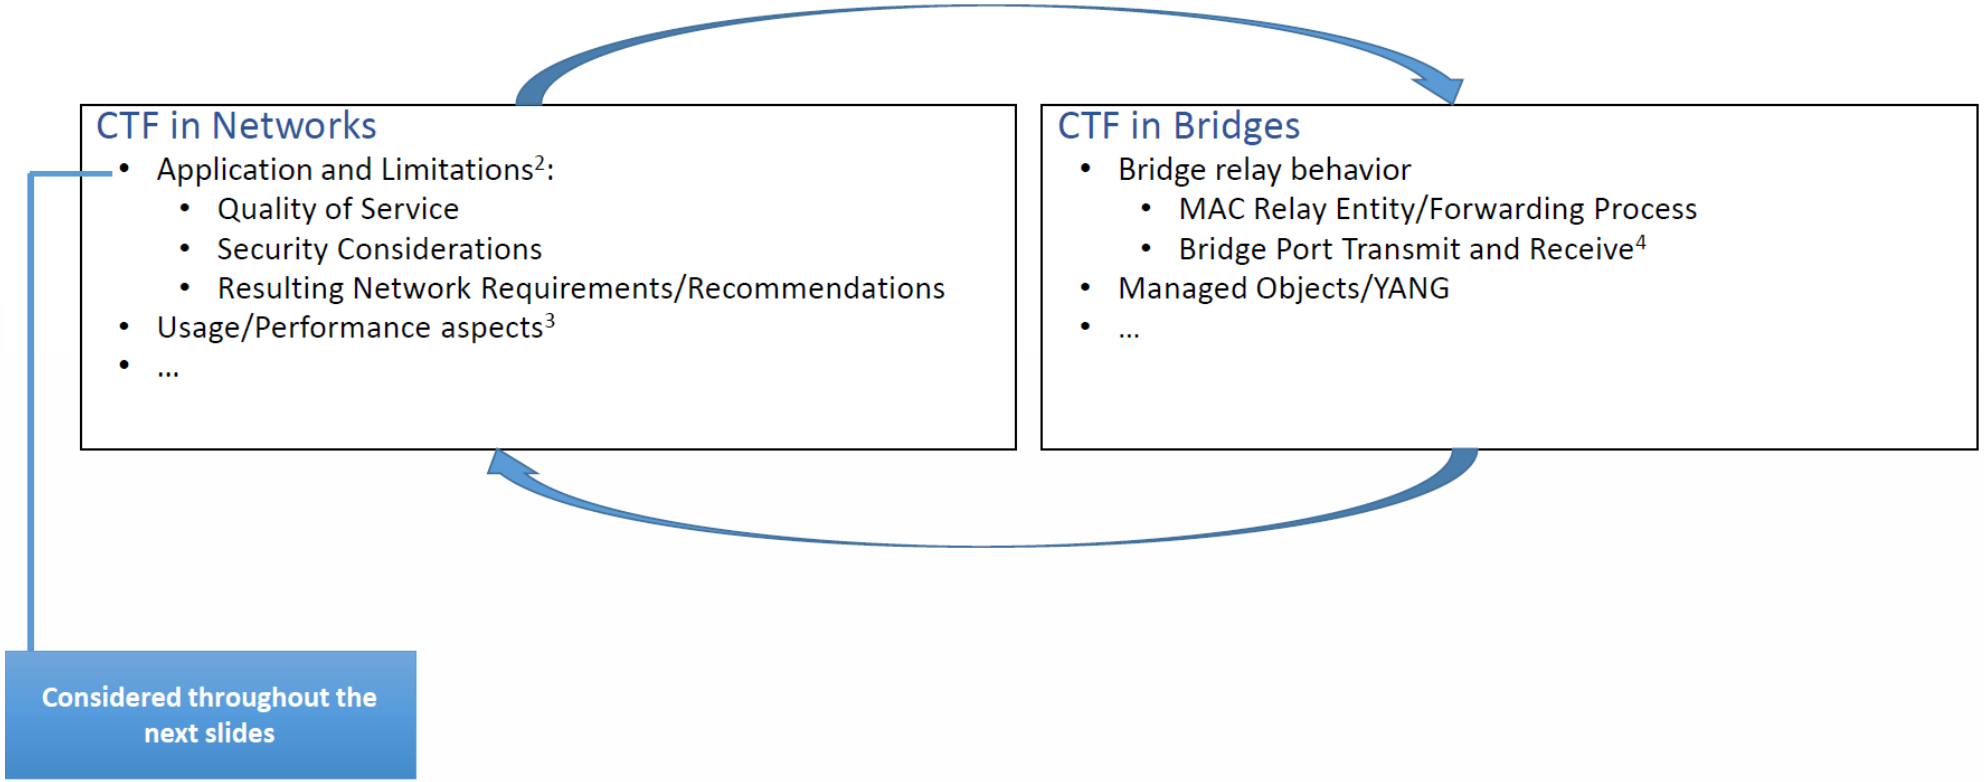

1) Illustration of technical contents, not a clause structure.
2) Issues introduced by CTF (cmp. https://ieee802.org/1/files/public/docs2017/new-tsn-thaler-cut-through-issues-0117-v01.pdf and ttps://www.ieee802.org/1/files/public/docs2019/new-seaman-cut-through-scissors-0119-v01.pdf) that can be addressed on a network level.
3) See the introduction of this slide set.
4) To the extent possible in IEEE 802.1.

Summary & Conclusions

Summary & Conclusions

Proposal for 802.1

- Motion to develop PAR&CSD for an IEEE 802.1 project to standardize CTF as standalone IEEE 802.1 standard (not amendments). Proposed items in a PAR scope in the context of this slide set:
  - Include IEEE Std 802.3-2018 compatible ***real implementations*** in the supported MACs.
  - Incorporate/standardize IEEE 802.1 aspects of a model with support for CTF, if such a model becomes available during such a project.

Presented

- ***Multiple reasons why an 802.1 standalone Standard for CTF would be preferable over multiple amendments.***
- ***Such a standard seems possible with or without a model with support for CTF across 802.1 and 802.3.***
- ***Work on such a model can continue in parallel; such a model requires sufficient need/interest in IEEE WG 802.1 and IEEE WG 802.3.***

Now … before such a motion

- Discuss individual contributions that may support the aforementioned motion, and may be considered during development of the PAR&CSD [item 2)] by IEEE WG 802.1.

Thank you for your Attention!

Questions, Opinions, Ideas?

Johannes Specht
Dipl.-Inform. (FH)

GERMANY
M +49 (0)170 718-4422
johannes.specht.standards@gmail.com