

This document is a **DRAFT version of recollected responses from the IEEE 802 Plenary Tutorial on Cut-Through Forwarding (CTF) on July 7, 2021**

DCN 1-21-0042-01-ICne

Associated Slide	Comment #	Comment	Response (Revised)
29	1	What is the reason to assume daisy-chain instead of star topology in this application? both latency and power delivery would be much better with star	Daisy chain networks/network segments are more cost efficient due to physical constraints. Star topologies would imply significant higher cost for cabling. This observation likewise holds for industrial automation use-cases (see slide 19).
13	2	this example is not a structure (one window per stream) that will be used. We should examples that are more commonly used. And we should discuss in such cases about other ways to do this.	The use of the enhancements for scheduled traffic (IEEE 802.1Qbv), as shown on slide 13, eliminates all interfering traffic and is a valid use of IEEE 802.1Qbv, at least under the already standardized S&F behavior. Likely this would also be a valid use with CTF. It has been clarified after the tutorial that the comment solely relates to industrial automation applications, whereas the slide shows just one example how all interference can be avoided. The commenter is encouraged to discuss use of the enhancements for scheduled traffic in industrial automation applications in the IEC/IEEE 60802 joint project.
37	3	37 (or the entire presentation): I haven't heard any mention of the difficulties of CTF with heterogeneous port speeds. Cutting through from a slow to a fast port doesn't work well and from a fast to a slow port can have smaller returns. For industrial and AV applications I can believe homogeneous port speeds are common, but for datacenter networks, it seems harder to have uniform rates over the whole path.	The general observation on the difficulties on link speed transitions is correct. * The suggestion on slide 53 is to fallback to S&F in case of slow-to-fast link speed transitions. * In certain DCN areas (slide 38), InfiniBand has a significant higher market share than Ethernet. CTF is present in InfiniBand products (slide 39). Making CTF part of the tools in Ethernet competitive in certain areas of data center networks, as explained in the data center networks use case section of the tutorial.

Associated Slide	Comment #	Comment	Response (Revised)
42	4	Is CTF a best effort solution? Can it be accepted that CTF "falls back" to S&F when the gross port is busy sending another (larger) packet. Would it be fine if CTF would be "slightly disturbed" by smaller interfering packets? What would you go for?	<p>There is no general answer to this question. The answer depends on the use cases the associated latency requirements of the applications:</p> <ul style="list-style-type: none"> * There are use cases where no or at least very low interference must be guaranteed. An example how all interference can be avoided is shown on slide 13. The requirement of none or at least very low interference is found in some industrial automation applications and ProAV. * In other markets such as DCN, occasional higher interference can be acceptable.
18	5	If 100 Mb/s is the factory standard AND Store and Forward is such a timing problem why don't you use 802.3 REPEATERS which are already standardized?	
22	6	Automation systems integrate audio (e.g. PTT) and video (e.g. quality insurance) into their networks. Good to see that the requirements seem to be similar!	N/A

Associated Slide	Comment #	Comment	Response (Revised)
43	7	<p>This statement is NOT true. "CTF: Reduces delays of teh critical frames themselves (regardless whether interference by other frames is present or not)." One could make this a true statement by re-stating "Reduces delays of the critical frames themselves when the egress port is in idle", or similar statement.</p>	<p>"CTF: Reduces delays of teh critical frames themselves (regardless whether interference by other frames is present or not)." is a true statement. Under CTF, a frame is earlier available for subsequent transmission at the egress port. This delay reduction is independent of the delay such a frame experiences afterwards at the egress port by interference. Assuming the same amount of interference under S&F and CTF in the egress port, the delay reduction by earlier availability due CTF remains unchanged, and an overall delay reduction by CTF remains. The significance of the delay reduction by CTF varies, and decreases at higher levels of interferences. This is illustrated in the introduction section of the tutorial (slide 10 ff.), based on the equations and assumptions shown in the annex of the presentation.</p>
37	8	<p>(Reply to comment #4:) In ProAV we require deterministic predictable behaviour with known constant end-to-end latency. If that cannot be guaranteed then use cases/markets for a solution are vastly reduced.</p>	<p>N/A (see comment #4)</p>
39	9	<p>the IB numbers are shorter than the FEC latency of 25G and the Ethernet numbers also seem as if FEC is not used. Is it assumed that FEC is not used with CTF? Data center speeds of 50 Gb/s per lane and above require FEC.</p>	<p>The Ethernet number is from tolly report, the report says FEC is used. <<TBD>></p>

Associated Slide	Comment #	Comment	Response (Revised)
57	10	there are also underrun situations faster transmitter than receiver	This is correct. The proposal on slide 53 suggests a fallback to S&F in this case to avoid such an underrun.
58	11	Your 802.3 layer diagram is not complete. It shows LLC but not TYPED MAC clients (e.g. VLAN).	
39	12	FEC is another component of static latency shared in common physical layers. Another reason to consider the importance of static latency	N/A (see comment #9)
62	13	Alon kept talking about "FRAMES" less than minimum size being discarded. This is incorrect. An EVENT of less than minimum size IS NOT A FRAME. It is considered to be an event fragment with no data validity.	Correct - term "frame" is use uncarefully at some places.
63	14	FEC is also a consideration - a FEC codeword contains several MAC frames and has to be stored in order to be decoded. PHYs which use FEC have basic latency muc larger than the frame size so CTF would not save a lot.	
62	15	Responding to comment #13:	N/A
62	16	i may presume the point is possibly about reducing the frame size	N/A (see comment #13)
67	17	Thank you. This is helpful information	Thank you for joining the tutorial!
62	18	as far i remember ; min frame size comes from ... csma/cd transmission on shared media ... which may not be really relevant today	N/A (see comment #13)

Associated Slide	Comment #	Comment	Response (Revised)
67	19	After more than 30 years of CTF deployment, are there known and important interop issues between CTF switch vendors?	<p>There are two interoperability aspects to consider (see also slide 15):</p> <ol style="list-style-type: none"> 1. Interoperability in terms of a unified management 2. Interoperability on the data plane <p>CTF is in products, but lacking unified configuration and status interface (aspect 1.). To a certain extent, this is related to aspect 2. As a simple example, consider erroneous error frame counting in a receiving bridge. Erroneous frames can be marked by CTF capable products (late error handling), but there are different ways this marking can be achieved (for example, a special FCS as indicated on slide 54). Dependent on how a bridge marks such frames, the next bridge can interpret different semantics in such markings and increase internal error counters differently.</p>

Associated Slide	Comment #	Comment	Response (Revised)
11	20	<p>starting with slide 11: I had assumed, as with preemption, this was most useful on lower link speed interfaces (2.5G the highest speed shown there). But I was surprised to see on slide 33 a claim of applicability to data center networks. Firstly, data center networks are much flatter (no daisy chaining), but on 100G, 400G links, packet transmission time is quite short. Sure, latency is important, but important factors on those links are FEC encode/decode, and I never recall packet transmission time being the limiting factor. Can you comment on the maximum link speed to which this should apply?</p>	
16	21	<p>Assuming that End Station may have the lower speed while the networks have the higher speed, the performance win is still big enough to request support of CTF in bridges in industrial automation</p>	N/A (see comment #19)
36	22	<p>DC/HPC networks use the fastest links available: 100G/400G today, going to 800G/1.6T next. At these speeds, by the time you figure out where to send the frame, it has already arrived and has been stored in you bridge, effectively making it S&F.</p>	<p>Forwarding table look-up is another aspect to be improved. Current forwarding involves L2 and L3 forwarding table, with inefficient look-up, such as LPM or EM HASH. https://www.ieee802.org/1/files/public/docs2021/cq-lv-semantic-address-with-P802.1CQ-0501-v01.pdf starts to discuss potential improvement.</p>

Associated Slide	Comment #	Comment	Response (Revised)
36	23	@comment #22 - You are pointing out the growing importance of static latency in networks and additional techniques also need to be considered to reduce the switch forwarding delays	N/A (see comment #22)
16	24	Today interoperability on the wire is achieved - configuration is vendor or organization specific. Possibility to mix of CTF and S&F components was always a requirement for automation systems	N/A (see comment #19)
22	25	Fully agree with comment #24	N/A (see comment #19)