Editor's Report
60802 Draft 3.0

IEEE 802.1
October 2024

Jordon Woods, Self-affiliated

Overview

- Initial SA Ballot passed with 80% approval and an 80% response rate!
- 211 Total Comments (including 4 non-ballot comments).
  - General – 6
  - Editorial – 73
  - Technical – 132

Ballot Statistics

CATEGORY	All Respondents	
	TOTAL	%
Yes	68	80.95%
No	16	19.05%
Voting Yes or No	84	100.00%
Abstain	4	4.55%
Respondents	88	80.73%
Ballot Pool	109	
Non-voting	1	
No. of commenters	26	29.55%
No. of comments	211	
TR	98	46.45%
T	34	16.11%
ER	30	14.22%
E	43	20.38%
G & GR	6	2.84%

Responses

Name	Current Vote	Comments
Abdul Jabbar	Abstain	N
Adee Ran	Disapprove	Y
Alon Regev	Approve	N
Arumugam Paventhan	Approve	N
Atsushi Sato	Disapprove	Y
Bansi Patel	Approve	N
Bartien Sayogo	Approve	N
Benjamin Rolfe	Disapprove	Y
Boon Chong Ang	Approve	N
BUTCH ANTON	Approve	N
C Huntley	Disapprove	Y
Cam Posani	Approve	N
Christian Boiger	Approve	N
Christophe Mangin	Abstain	N
David McCall	Disapprove	Y
David Tepen	Approve	N
Dieter Proell	Approve	Y
Dorothy Stanley	Approve	Y
Ganesh Venkatesan	Approve	Y
Gary Stuebing	Approve	N
Gavin Lai	Approve	Y
Geoffrey Garner	Disapprove	Y
Glenn Parsons	Disapprove	Y
Greg Luri	Approve	N
Harry Bims	Approve	N
Hyeong Ho Lee	Approve	N
Janos Farkas	Approve	N
Janusz Zalewski	Approve	N
Jessy Rouyer	Approve	Y
Jhony Sembiring	Approve	N
Jingfei Lv	Approve	N
John Vergis	Approve	N
Jonathon Mclendon	Approve	N
Jordon Woods	Approve	Y
Lei Yang	Approve	N
Lokesh Kabra	Approve	N
Ludwig Winkel	Disapprove	Y
Maik Seewald	Approve	N
MARC EMMELMANN	Approve	N
Marcel Kiessling	Disapprove	Y
Marco Hernandez	Approve	N
Marek Hajduczenia	Approve	N
Mark Hantel	Approve	Y
Martin Mittelberger	Disapprove	Y

Name	Current Vote	Comments
Matthias Fritsche	Abstain	N
Max Turner	Approve	N
Maximilian Riegel	Approve	N
Michael Dood	Approve	N
Michael Montemurro	Approve	N
Oliver Holland	Approve	N
Oren Yuen	Approve	N
Paul Cardinal	Approve	N
Paul Nikolich	Approve	N
Paulo Goncalves	Approve	N
Pin Chang	Approve	N
Piotr Karocki	Approve	N
Pranav Jha	Approve	N
Radhakrishna Canchi	Approve	N
Raj Jain	Approve	N
Rajesh Murthy	Approve	N
Rich Boyer	Approve	N
Richard Bugg	Approve	N
Richard Mellitz	Approve	N
Richard Tse	Disapprove	Y
Richie Pearn	Approve	N
Rodney Cummings	Disapprove	Y
Scott Mansfield	Approve	Y
Scott Willy	Approve	N
Silvana Rodrigues	Disapprove	Y
Stefan Aust	Approve	N
Stephan Kehrer	Disapprove	Y
Stephen Haddock	Approve	N
Stuart Kerry	Approve	N
Sven Meier	Disapprove	Y
Sven Zeisberg	Approve	N
Tomoko Adachi	Disapprove	Y
Travis Breitkreutz	Approve	Y
Vern Brethour	Approve	N
Veselin Skendzic	Approve	N
Walter Struppler	Approve	N
Werner Hoelzl	Approve	N
William Armstrong	Approve	Y
Yanjie Gong	Approve	N
Yongbum Kim	Approve	N
Yongsen Ma	Abstain	N
Yu Yuan	Approve	N
Yukimasa Nagai	Approve	N
zhiman chen	Approve	N

Non-ballot comments

- In addition, the editor has received a total of 6 comments from 2 different members of the Ballot pool who did not submit their ballots by the deadline
  - The editor proposes including these as non-ballot comments

Technical 49 5.7.1 1546

a) At least 3 flow meters make no sense in absence of a specified stream filter count, nor does omission of stream gates provide a functional flow classification and metering stage. Stream filters are just indirectly pulled in via reference to 8.6.5.3 of IEEE Std 802.1Q, but there is no clarity on the number of required stream filters. Stream gates are not mentioned at all - however, every stream filter requires a stream gate through which frames are passed to the flow meters. Further unclear to me is the resolution required to satisfy the example given in lines 1555 to 1556. Stream filters, gates and flow meters are instantiated per Bridge component. Now, is it intended to support sufficient stream filter for that example per Port or per Bridge component? Assuming it is per Port, the 802.1CB stream identification functions need further consideration. The stream identification function I identified so far for the distinction between unicast, multicast and broadcast is the mask and match stream identification. Identifying unicast can be trvially done by defining a mask and match stream identification and an associated stream filter (both are in a one-to-one relationship) by letting this mask and match identification check the DA's group address bit beeing cleared. Multicast and Broadcast then require other mask and match stream identification functions with that bit set. Distinction between multicast and requires many mask-and-match identification functions, and thus many stream filters. Broadcast is assumed to have all 48 DA bits set to 1, including the group address bit. Multicast is assumed to have at least one bit other than the group address bit set to 0. Expanding the combinations, this results in 50 mask and match stream identification functions/stream filters (F1 through F49) and 3 Flow Meters (M1 through M3) according to the following pattern:

| Traffic | DA (EUI48;ternary - "0","1" and "x" for "don't care") | Flow Meter | Stream Filter
00 | broadcast | 11111111-11111111-11111111-11111111-11111111-11111111 | M1 | F1
01 | broadcast | 0xxxxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx | M2 | F2
02 | broadcast | x0xxxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx | M2 | F3
03 | broadcast | xx0xxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx | M2 | F4
.
.
.
45 | broadcast | xxxxxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxx0xx | M2 | F46
46 | broadcast | xxxxxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxx0x | M2 | F47
47 | broadcast | xxxxxxx1-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxx0 | M2 | F48
48 | unicast | xxxxxxx0-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx | M3 | F49

b) The EIR, EBS and CM values in lines can only be a management recommendation. However, the way it is expressed right now expresses that flow meters setup these parameters, not the management.

c) The note in lines 3726 through 3730 appears inaccurate - one stream idenfication function for multicast seems insufficient, because IIRC there is no defined precedence in IEEE Std 802.1CB in case of multiple stream idenification functions matching to the same frame. The stream filters (IEEE Std 802.1Q) have a a defined precedence, but this does not help, because only one stream_handle (as determined by a "randomly" selected stream identification function out of many stream idenfication functions...) is passed to the stream filters, so the stream filter precedence ordering never becomes effective. Note that I assume that the example in lines 3726 through 3730 relies on precedence, such as having a single mask-and-match identification for broadcast (all DA bits set) and a second one for multicast (DA group address bit set, all other DA bits "don't care"), where the former takes precedence over the latter in case both match. But as explained, the lack of precedence would not result in the desired behavior.

Dependent on whether or not the example in lines 1555 to 1556 is intended to define the number of required stream identification functions, implement one of the following options:

* Option 1 (stream identification function for the example required):
a) Change item l) to read "Support 49*N stream filters (IEEE Std 802.1Q-2022, 8.6.5.3), one stream gate (IEEE Std 802.1Q-2022, 8.6.5.4) and 3*N flow meters per Bridge component (IEEE Std 802.1Q-2022, 8.6.5.5), where N is the number of Ports of the Bridge component. It is recommended to setup the EIR, EBS and CM parameters of flow meters as follows:"

* Option 2 (stream identification function for the example NOT required):
b) Change item l) to read "Support 3*N stream filters (IEEE Std 802.1Q-2022, 8.6.5.3), one stream gate (IEEE Std 802.1Q-2022, 8.6.5.4) and 3*N flow meters per Bridge component (IEEE Std 802.1Q-2022, 8.6.5.5), where N is the number of Ports of the Bridge component. It is recommended to setup the EIR, EBS and CM parameters of flow meters as follows:"
c) Add the following sentence to the end of line 1556: "This example can be realized using 49 mask and match stream identification functions per Port and 49 associated stream filter instances (one-to-one relationship) that identify and map bnicast (1 stream identification function for identfying all frame with the group address bit cleared in the DA), broadcast (1 stream identification function for identifying all frames with all DA bits set) and multicast (47 stream identification function with the DA group address bit set and one other DA bit cleared for all 47 combinations).

Proposed changes applicable for both options:
d) If there is no defined precedence in case of multiple stream identification functions (as assumed), change "one mask and match stream identification for multicast traffic" in line 3727 to "47 mask and match stream identification for multicast traffic for unique between multicast and broadcast traffic".

e) Add the missing YANG nodes to 6.4.9.2.5.14 for stream gates. I did not look these up in detail, but I guess adding them should be straight forward.

Non-ballot comments

Category	Page	Subclause	Line #	Comment	Proposed Change	Must Be Satisfied
Technical	56	6.2.5	1842	The error generation limits are only specifying the timestamp accuracy of the timestamps in the Sync/follow_up frames. There are no error generation limits for the accuracy of the timestamps affecting the MeanPathDelay measurements.Any error in the MeanPathDelay will also affect the accuracy of the time delivery.	It is proposed to use the same metric for PTP performance as in other standards, specifically the ITU-T performance metric cTE and dTE. These specify the accuracy time delivery (using all the timestamps being used in the PTP protocol) instead of focusing only on the egress Sync timestamp.	NO
Technical	56	6.2.5	1842	Using a mean value as a spec limit is not testable. How long time must the measurement time be?? Seconds, minutes, days, years? If the error is +17ns for 5 minutes and -17ns the next 5 minutes it will meet the spec	It is proposed to use the same metric for PTP performance as in other standards, specifically the ITU-T performance metric cTE and dTE or even an MTIE mask. These specify the accuracy of the time delivery (using all the timestamps being used in the gPTP protocol) instead of focusing only on the egress Sync timestamp. Other metric can be used, but they must accurately specify the measurement period	YES
Technical	57	6.2.5	1850	Using a mean value as a spec limit is not testable. How long time must the measurement time be?? Seconds, minutes, days, years? If the error is +12ns for 5 minutes and -12ns the next 5 minutes it will meet the spec	It is proposed to use the same metric for PTP performance as in other standards, specifically the ITU-T performance metric cTE and dTE or even an MTIE mask. These specify the accuracy of the time delivery (using all the timestamps being used in the gPTP protocol) instead of focusing only on the egress Sync timestamp. Other metric can be used, but they must accurately specify the measurement period	YES
Technical	57	6.2.5	1850	Error generation for a PTP relay should be specified from input to the relay to output of the relay and not relative to the internal time of a connected Grandmaster. The allowed error from the the Grandmaster is higher than allowed by the PTP Relay.	It is proposed to use the same metric for PTP performance as in other standards, specifically the ITU-T performance metric cTE and dTE. These specify the accuracy time delivery (using all the timestamps being used in the gPTP protocol) instead of focusing only on the egress Sync timestamp. Performance should be specified with an ideal time input to the DUT (both Sync frames and Pdelay_req/resp and Pdelay_resp Follow_up.) and the resulting time error at the output port should be measured, making resulting time calculations based on the transmitted frames. The specification of the accuracy of RR, nRR and rateRatioDrift are OK except for the missing calculation period for the mean value.	YES

Thank you