

CI 73 AN Baseline Proposal

Kent Lusted, Intel Corporation

Jeff Slavick, Broadcom

Adee Ran, Cisco

Arthur Marris, Cadence

Supporters

- Phil Sun, Credo
- Henry Wong, Alphawave Semi
- Dave Cassan, Alphawave Semi
- Priyank Shukla, Synopsys
- P.-R. Li, MediaTek
- Zvi Rechtman, NVidia

Clause 73 Auto-Negotiation Process

- Provides mechanism for links to negotiate the fastest common rate that both sides advertise.
- Useful for plug-and-play linking up
- Backplane and copper cable assembly support only

Problem Statement

- The 3dj Task Force needs to support at least 4 new Ethernet PHY types:
 - “200G-R1”
 - “400G-R2”
 - “800G-R4”
 - “1600G-R8”
- Maybe new FEC modes, as required (Exact number is TBD)
- The IEEE 802.3 Auto-Negotiation link codeword Base page for Backplane and Copper Cable Assembly (Clause 73) does not have enough free bits

AN73 Base Page – After 3df

- 2 bits are left unassigned in the base page
- Need to advertise at least 4 more PHYs (that we know of right now)

Table 73–4—Technology Ability Field encoding

Bit	Technology
...	
A18	400GBASE-KR4 or 400GBASE-CR4
A19	800GBASE-KR8 or 800GBASE-CR8
A19A20 through A21	Reserved

AN73 Message Base Page, see Figure 73-6															
D0	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15
S0	S1	S2	S3	S4	E0	E1	E2	E3	E4	E5	C0	C1	RF	ACK	NP
D16	D17	D18	D19	D20	D21	D22	D23	D24	D25	D26	D27	D28	D29	D30	D31
T0	T1	T2	T3	T4	A0	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10
D32	D33	D34	D35	D36	D37	D38	D39	D40	D41	D42	D43	D44	D45	D46	D47
A11	A12	A13	A14	A15	A16	A17	A18	A19	A20	A21	F4	F2	F3	F0	F1
[A20:A21] remain open after 3df															

Overview

- In Annex 73A, define a new Clause 73 Next Page (NP) message code 2
 - The contents of the new Next Page message code will expand the Clause 73 Technology Ability Field (TAF) and FEC Capability field to include more bits without changing the base page format nor the resolution protocol.
- Use the two remaining reserved bits in the base page as follows:
 - One unused bit is assigned to indicate that additional bits of the CI 73.6.4 Technology Ability Field (TAF) will be transmitted in a NP message code 2.
 - The other unused bit is assigned to indicate that additional bits of the CI 73.6.5 FEC Capability field will be transmitted in a NP message code 2.

Solution Details - 1

- Define bits D41:D42 of the Link codeword Base page to be as follows:
- D41 = (Extended_TAF_1 capability)
 - 0 = device does not have additional “Technology Ability” fields
 - 1 = device has additional “Technology Ability” field bits defined in Next Page message code 2
- D42 = (Extended_FEC_1 capability)
 - 0 = device does not have additional “FEC Capabilities” fields
 - 1 = device has additional “FEC Capability” bits defined in Next Page message code 2
- If any of the extended capabilities bits is set to 1, then bit D15 (NP) shall also be set to 1, and a next page with message code 2 shall follow
 - If a next page with message code 2 is received when the extended capabilities bits in the base page were set to 0, or not immediately following the base page, it may be ignored

Solution Details - 2

- A new Message Next Page of type code 2 is defined as follows:



AN73 Message Next Page, see Figure 73-7

D0	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15
Message Code Field (11 bits)											T	ACK2	MP	ACK	NP
D16	D17	D18	D19	D20	D21	D22	D23	D24	D25	D26	D27	D28	D29	D30	D31
Extended_TAF_1[0:15]															
D32	D33	D34	D35	D36	D37	D38	D39	D40	D41	D42	D43	D44	D45	D46	D47
Extended_TAF_1[16:27]												Extended_FEC_1[3:0]			

- The “Technology Ability and FEC extension” message type shall consist of only a Message Next Page.
- The message code field is 2 (M10:M0 = 00 0000 0010)
 - Follows same order as IEEE Std. 802.3-2022 Clause 73.7.7.1: M0 = D0, M1 = D1 ... M10 = D10
- [D15:D11] retain their function per Cl 73.7.7.1.



Solution Details - 3

- The contents of the Unformatted Code Field U27:U0 (D43:D16) shall be as defined as the Extended_TAF_1[27:0] field with the following assignments
 - Bit ordering per slide 8
 - Bit 0 = 200GBASE-CR1 or 200GBASE-KR1
 - Bit 1 = 400GBASE-CR2 or 400GBASE-KR2
 - Bit 2 = 800GBASE-CR4 or 800GBASE-KR4
 - Bit 3 = 1.6TBASE-CR8 or 1.6TBASE-KR8
 - Bit [27:4] = reserved for future use



Solution Details - 4

- The contents of the Unformatted Code Field U31:U28 (D47:D44) shall be as defined as the Extended_FEC_1[3:0] field with the following assignments
 - Bit ordering per slide 8
 - Used for additional FEC capabilities, if needed (e.g. interleave FEC, low-latency FEC, high-correction capability FEC)
 - All bits reserved for future use at this time



Solution Details - 5

- Amend CI 73.6.4 to incorporate the extended Technology Ability Field
- Add a new Table 73-4a with the appropriate mapping of Extended_TAF_1 fields to Technologies
 - Add a note “If the Extended Technology Ability Field is not received or not sent, then its effective value is all zeros.”
- Amend CI 73.6.5 to include any newly defined FEC fields in the Extended_TAF_1 and the capabilities

Solution Details - 6

- In Figure 73-1, update to add 1.6TMII
- In Cl 73.5.1, update the second paragraph to add references to the new 1.6TBASE-CR8, 800GBASE-CR4, 400GBASE-CR2 and 200GBASE-CR1 clauses
- In Table 73-5, update Priority Resolution to insert “1.6TBASE-CR8 or 1.6TBASE-KR8” at the top (highest priority) and add the others to the table accordingly
 - 1.6TBASE-CR8 or 1.6TBASE-KR8
 - 800GBASE-CR4 or 800GBASE-KR4
 - 800GBASE-CR8 or KR8
 - 400GBASE-CR2 or 400GBASE-KR2
 - 400GBASE-CR4 or 400GBASE-KR4
 - 200GBASE-CR1 or 200GBASE-KR1
 - ...

Solution Details - 7

- In Cl 73.10.1, add new entries into the variable list:
 - “1.6TR8; represents the 1.6TBASE-CR8 or 1.6TBASE-KR8 PMD”
 - “800GR4; represents the 800GBASE-CR4 or 800GBASE-KR4 PMD”
 - “400GR2; represents the 400GBASE-CR2 or 400GBASE-KR2 PMD”
 - “200GR1; represents the 200GBASE-CR1 or 200GBASE-KR1 PMD”
- In the single_link_ready entry in Cl 73.10.1, add in the appropriate place:
 - “link_status_[1.6TR8] = OK”
 - “link_status_[800GR4] = OK”
 - “link_status_[400GR2] = OK”
 - “link_status_[200GR1] = OK”

Solution Details - 8

- In Table 73-7 Timer min/max value summary, add a new row for “link_inhibit_timer (when the link is 1.6TBASE-CR8, 1.6TBASE-KR8, 800GBASE-CR4, 800GBASE-KR4, 400GBASE-CR2, 400GBASE-KR2, 200GBASE-CR1, or 200GBASE-KR1)” with value TBD
- Modify Table 45-388 Backplane Ethernet, BASE-R copper status 2 register bit definitions to include an entry for “1.6TBASE-CR8, 1.6TBASE-KR8, 800GBASE-CR4, 800GBASE-KR4, 400GBASE-CR2, 400GBASE-KR2, 200GBASE-CR1, and 200GBASE-KR1”

Solution Details - 9

- In Cl 73.7.7, change the last sentence of the fifth paragraph into a new paragraph as follows:
 - “Any number of Next Pages may be sent in any order, with the exception that if a Next Page Message code2 is sent it shall be the first Next Page transmitted; however, it is recommended that the total number of Next Pages sent be kept small to minimize the link startup time.”
- In Cl 73.7.6, change the last sentence of the first paragraph to :
 - “The single PHY enabled to connect to the MDI by Auto-Negotiation shall be the technology corresponding to the bit in the Technology Ability Field and the Extended Technology Ability Field that is common to the local device and link partner that has the highest priority as defined in Table 73–5 (listed from highest priority to lowest priority).”
- Add a Clause 45 registers dedicated to Message Code 2 reception and storage
- Add or modify relevant PICS items

Summary

- A baseline proposal for CI 73 was proposed
- A straw poll (and possibly a motion) is being considered

Thanks!

Reference

Annex 73A

(normative)

Next page message code field definitions

This Annex defines the Next Page message code fields for devices using Clause 73 Auto-Negotiation. The message code field of a message page used in Next Page exchange shall be used to identify the meaning of a message. Table 73A–1 identifies the types of messages that may be sent. As new messages are developed, this table will be updated accordingly.

The Message code field uses an 11-bit binary encoding that allows 2048 messages to be defined. All message codes not specified are reserved for IEEE use or allocation.

Table 73A–1—Message code field values

Message code	M 10	M 9	M 8	M 7	M 6	M 5	M 4	M 3	M 2	M 1	M 0	Message code description
1	0	0	0	0	0	0	0	0	0	0	1	Null Message
5	0	0	0	0	0	0	0	0	1	0	1	Organizationally Unique Identifier Tagged Message
6	0	0	0	0	0	0	0	0	1	1	0	AN device Identifier Tag Code
10	0	0	0	0	0	0	0	1	0	1	0	EEE Technology Message Code. EEE capability is advertised using unformatted message code field in the Message Next Page (see 73A.4).