



OIF OPTICAL
INTERNETWORKING
FORUM

**Common Electrical I/O - Protocol (CEI-P)
Implementation Agreement**

IA # CEI-P-01.0

March, 2005

Implementation Agreement created and approved
by the Optical Internetworking Forum
www.oiforum.com



The OIF is an international non profit organization with over 130 member companies, including the world's leading carriers and vendors. Being an industry group uniting representatives of the data and optical worlds, OIF's purpose is to accelerate the deployment of interoperable, cost-effective and robust optical internetworks and their associated technologies. Optical internetworks are data networks composed of routers and data switches interconnected by optical networking elements.

With the goal of promoting worldwide compatibility of optical internetworking products, the OIF actively supports and extends the work of national and international standards bodies. Formal liaisons have been established with The ATM Forum, IEEE 802.3, IETF, ITU-T Study Group 13, ITU-T Study Group 15, MEF, NPF, ATIS-TMOC, ATIS-OPTXS, TMF, UXPi and the XFP MSA Group.

For additional information contact:
The Optical Internetworking Forum, 39355 California Street,
Suite 307, Fremont, CA 94538
510-608-5928 ☎ info@oiforum.com

www.oiforum.com

Working Group: Physical and Link Layer

TITLE: Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

SOURCE: Winston Mok

Technical Editor
PMC-Sierra
8555 Baxter Place
Burnaby, BC, Canada
V5A 4V7
Phone : +1 604 415 6121
Email : mok@pmc-sierra.com

Mike Lerer

PLL Chair
Xilinx
Box 636
Londonderry,
NH 03053 USA
Phone : +1 603 548 3704
Email : mlexerer@fpga.com

Karl Gass

PLL Vice-Chair
Sandia National Labs
PO Box 5800
MS 0874 Albuquerque
NM 87185, USA
Phone : +1 505 844 8849
Email : kgass@sandia.gov

Notice: This Technical Document has been created by the Optical Internetworking Forum (OIF). This document is offered to the OIF Membership solely as a basis for agreement and is not a binding proposal on the companies listed as resources above. The OIF reserves the rights to at any time to add, amend, or withdraw statements contained herein. Nothing in this document is in any way binding on the OIF or any of its members.

The user's attention is called to the possibility that implementation of the OIF implementation agreement contained herein may require the use of inventions covered by the patent rights held by third parties. By publication of this OIF implementation agreement, the OIF makes no representation or warranty whatsoever, whether expressed or implied, that implementation of the specification will not infringe any third party rights, nor does the OIF make any representation or warranty whatsoever, whether expressed or implied, with respect to any claim that has been or may be asserted by any third party, the validity of any patent rights related to any such claim, or the extent to which a license to use any such rights may or may not be available or the terms hereof.

© 2005 Optical Internetworking Forum

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction other than the following, (1) the above copyright notice and this paragraph must be included on all such copies and derivative works, and (2) this document itself may not be modified in any way, such as by removing the copyright notice or references to the OIF, except as needed for the purpose of developing OIF Implementation Agreements.

By downloading, copying, or using this document in any manner, the user consents to the terms and conditions of this notice. Unless the terms and conditions of this notice are breached by the user, the limited permissions granted above are perpetual and will not be revoked by the OIF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE OIF DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY, TITLE OR FITNESS FOR A PARTICULAR PURPOSE.

1 Table of Contents

0	Cover Sheet	1
1	Table of Contents	4
2	List of Figures	5
3	List of Tables	6
4	Document Revision History	7
5	Introduction	9
6	Requirements	11
7	Objectives	12
8	Common Definitions	12
8.1	Introduction	12
8.2	Frame Format Definition	12
8.3	Framing Layer Definition	23
8.4	Aggregation Layer Definition	30
8.5	Adaptation Layer Definition	39
9	Lane Aggregation Application - no transcoding	39
9.1	Introduction	39
9.2	Adaptation Layer Definition	40
9.3	Aggregation Layer Definition	41
9.4	Framing Layer Definition	43
10	Future OIF Interface Application	45
10.1	Introduction	45
10.2	Adaptation Layer Definition	45
10.3	Aggregation Layer Definition	45
10.4	Framing Layer Definition	45
11	Summary	45
12	Glossary	45
13	References	45
13.1	Normative references	45
13.2	Informative references	45
14	Annex A: Default Values of Variables	45
15	Appendix B: Framing Process in a CEI Protocol Receiver	46
16	Appendix C: Sample CEI Protocol Frame	47
17	Appendix D: List of companies belonging to OIF when document is approved	48

2 List of Figures

FIGURE 1	CEI PROTOCOL LANE AGGREGATION APPLICATION.....	9
FIGURE 2	FOUNDATION OF FUTURE OIF INTERFACE APPLICATIONS	10
FIGURE 3	CEI PROTOCOL LAYERS	11
FIGURE 4	CEI PROTOCOL FRAME FORMAT	13
FIGURE 5	FREE RUNNING SCRAMBLER.....	14
FIGURE 6	TRANSMITTER LINK STATE STATE MACHINE DIAGRAM – DUPLEX	17
FIGURE 7	TRANSMITTER LINK STATE STATE MACHINE DIAGRAM – SIMPLEX	18
FIGURE 8	RECEIVER LINK STATE STATE MACHINE DIAGRAM –DUPLEX.....	19
FIGURE 9	RECEIVER LINK STATE STATE MACHINE DIAGRAM - SIMPLEX.....	21
FIGURE 10	ALIGNMENT OF CEI TRAINING PATTERN TO CEI PROTOCOL FRAMES	22
FIGURE 11	FRAME GENERATION PROCESS – LOGICAL BLOCK DIAGRAM.....	23
FIGURE 12	SAMPLE FRAMING ALGORITHM	26
FIGURE 13	FRAMING STATE MACHINE DIAGRAM.....	27
FIGURE 14	LOGICAL AGGREGATION LAYER INTERFACES - MULTIPLEXING	31
FIGURE 15	LOGICAL AGGREGATION LAYER INTERFACES – INVERSE-MULTIPLEXING.....	32
FIGURE 16	GENERIC HDLC FRAME OF SUPERVISORY CHANNEL	34
FIGURE 17	PARAMETER READ/WRITE FRAME OF SUPERVISORY CHANNEL.....	35
FIGURE 18	CEI PROTOCOL LANE AGGREGATION APPLICATION.....	40
FIGURE 19	CLIENT ASSIGNMENT TO CEI PROTOCOL FRAME	41
FIGURE 20	CLIENT STATUS MESSAGE FORMAT.....	43
FIGURE 21	LOGICAL FRAME CONSTRUCTION IN LANE AGGREGATION APPLICATION.....	44
FIGURE 22	FRAME RECEIVER PROCESS – LOGICAL BLOCK DIAGRAM.....	47
FIGURE 23	SAMPLE CEI PROTOCOL FRAME	48

3 List of Tables

TABLE 1	OVERHEAD ASSIGNMENT IN A FRAME	14
TABLE 2	OVERHEAD LOCATION OF STATE[2:0].....	15
TABLE 3	TRANSMITTER LINK STATE REPORTING	16
TABLE 4	TRAINING PATTERN (FROM CEI ELECTRICAL IA)	22
TABLE 5	CEI PROTOCOL FRAME FORMAT REQUIREMENTS	24
TABLE 6	SAMPLE FRAMING ALGORITHM	26
TABLE 7	CEI PROTOCOL RECEIVER FRAME AND MULTIFRAME REQUIREMENTS	28
TABLE 8	CEI PROTOCOL RECEIVER SCRAMBLER SYNCHRONIZATION REQUIREMENTS....	29
TABLE 9	CEI PROTOCOL BI-DIRECTIONAL LINK REQUIREMENTS	29
TABLE 10	CEI PROTOCOL FORWARD ERROR CORRECTION REQUIREMENTS	30
TABLE 11	CEI PROTOCOL MULTIPLEXING REQUIREMENTS	33
TABLE 12	ASSIGNMENT OF TYPE FIELD IN SUPERVISORY CHANNEL PACKETS	34
TABLE 13	CEI PROTOCOL SUPERVISORY INTERFACE REQUIREMENTS	35
TABLE 14	FIELDS OF PARAMETER READ/WRITE PACKETS	35
TABLE 15	COMMAND FIELD OF PARAMETER READ/WRITE PACKETS	36
TABLE 16	TARGET ADDRESS (TADDR) FIELD OF PARAMETER READ/WRITE PACKETS	37
TABLE 17	PARAMETER IDENTIFIER (PID) FIELD OF PARAMETER READ/WRITE PACKETS	38
TABLE 18	CEI PROTOCOL CONNECTION IDENTIFIER REQUIREMENTS	38
TABLE 19	MODE CODE ASSIGNMENT	39
TABLE 20	CEI PROTOCOL ADAPTATION LAYER REQUIREMENTS	41
TABLE 21	CEI PROTOCOL AGGREGATION LAYER REQUIREMENTS	42
TABLE 22	CLIENT STATUS CODES	43
TABLE 23	CEI PROTOCOL CLIENT STATUS REQUIREMENTS.....	43
TABLE 24	CEI PROTOCOL FRAMING LAYER REQUIREMENTS.....	44

4 Document Revision History

Working Group: Physical and Link Layer

SOURCE: Winston Mok

Technical Editor
PMC-Sierra
8555 Baxter Place
Burnaby, BC, Canada
V5A 4V7
Phone : +1 604 415 6121
Email : mok@pmc-sierra.com

Mike Lerer

PLL Chair
Xilinx
Box 636
Londonderry,
NH 03053 USA
Phone : +1 603 548 3704
Email : mlexerer@fpga.com

Karl Gass

PLL Vice-Chair
Sandia National Labs
PO Box 5800
MS 0874 Albuquerque
NM 87185, USA
Phone : +1 505 844 8849
Email : kgass@sandia.gov

DATE: March, 2005

Issue No.	Issue Date	Details of Change
OIF2003.401.00	November 2003	Initial editor's version taken from OIF2003.264.05. Fixed typos.
OIF2003.401.01	January 2004	Fixed typos and presented to the OIF January 2004 Plenary in San Diego
OIF2003.401.02	February 2004	Incorporated changes authorized by Motions at the OIF January 2004 Plenary in San Diego. Fixed typos. Added Abstract and References. Clarified peer-to-peer relationship between CEI and CEI-P. Clarified HDLC frame relationship to ISO 3309. Clarified use of a subset of Fire Code parity bit for framing. Updated Transmit State Transition Diagram. Clarified that Forward Error Correction of Overhead bits is unnecessary. Added tutorial Appendix on receiver framing process. Consolidated Mx counters. Updated default values of M1 and M2. Added description of Training Pattern. Updated presentation location of STATE[2:0]. Updated MF encoding when multi-frame is absent. Clarified scrambler operation in TREQEST,

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

		TPRESENT and IDLE states. Updated SS[16:0] to be independent of scrambler enable/disable.
OIF2003.401.03	May 2004	Incorporated changes authorized by Motions at the OIF May 2004 Plenary in Budapest. Updated Transmitter state machine. Added Receiver state machine. Added packet format for read/write of link parameters over the Supervisory channel.
OIF2004.229.00	July 2004	Incorporated changed from breakout group. Changed frame format from having sub-frames A to D to a single unit that is modeled on sub-frame A. Overhead previously carried in overhead of sub-frames B and C (Connection ID and Mode) are moved to the Supervisory Channel. Number of Supervisory bits per frame reduced from 16 to 4. Updated framing algorithm description. Deleted option to leave payload (including T bits) unscrambled in normal operation.
OIF2004.229.01	July 2004	Fixed typos
OIF2004.229.02	July 2004	Adopted as the new baseline text for CEI-P Implementation Agreement, Draft 5.0 Updated cover page and revision history
OIF2004.229.03	October 2004	Updated with Straw Poll comments.
OIF2004.229.04	November 2004	Updated with results of straw ballot comment resolution session at the October 2005 Plenary at Alexandria, VA.
OIF2004.229.05	December 2004	Fixed typos.
OIF2004.229.06	January 2005	Clarified Multi-frame framing.
OIF2004.229.07	March 2005	Fixed typos. Version used to create IA after Principal Member ballot passed.

5 Introduction

The CEI Protocol Implementation Agreement defines protocols that take advantage of faster electrical interfaces developed by the CEI project. The CEI Electrical Implementation Agreement and the CEI Protocol Implementation Agreement are peer documents. Adherence to one does not force adherence to the other. For example, a 10G SONET framer may connection directly to an optical module using CEI electricals with SONET scrambled data. In this case, CEI Protocol would be absent. It is also possible to use CEI Protocol without CEI Electricals. An example would be to encapsulate TFI-5 frames with CEI Protocol to provide forward error correction capability. The target applications of CEI Protocol are Lane Aggregation and Physical layer management of future OIF interfaces.

Lane Aggregation refers to the multiplexing of existing lower speed links onto a reduced number of high-speed links. A pair of devices would interface to existing lower speed client sources, transfer the client information over CEI Protocol links, and re-generate the client information onto the client sinks (See Figure 1).

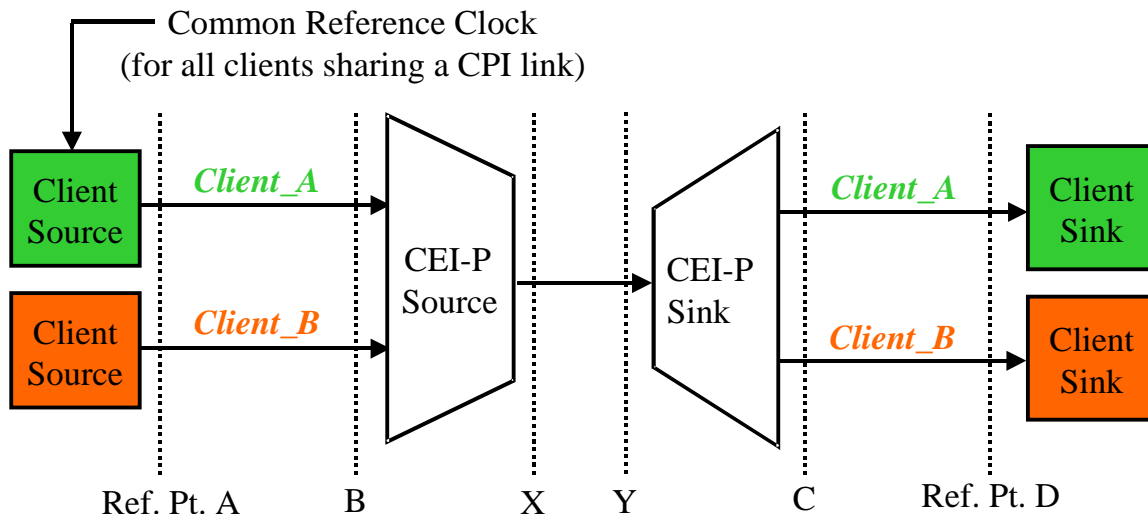


Figure 1 CEI Protocol Lane Aggregation Application

Existing OIF interfaces have a wide variety of foundation protocols. For example, de-skew in SFI-5 differs markedly from that in SPI-5. Error monitoring in TFI-5 has little in common with SFI-5, despite mainly targeted at SONET/SDH traffic. As links speed rise, and error performance becomes more stringent, forward error correction may become necessary. The CEI Protocol Implementation Agreement will provide a common set of framing, de-skew, error correction and detection protocols for future OIF interfaces (See Figure 2).

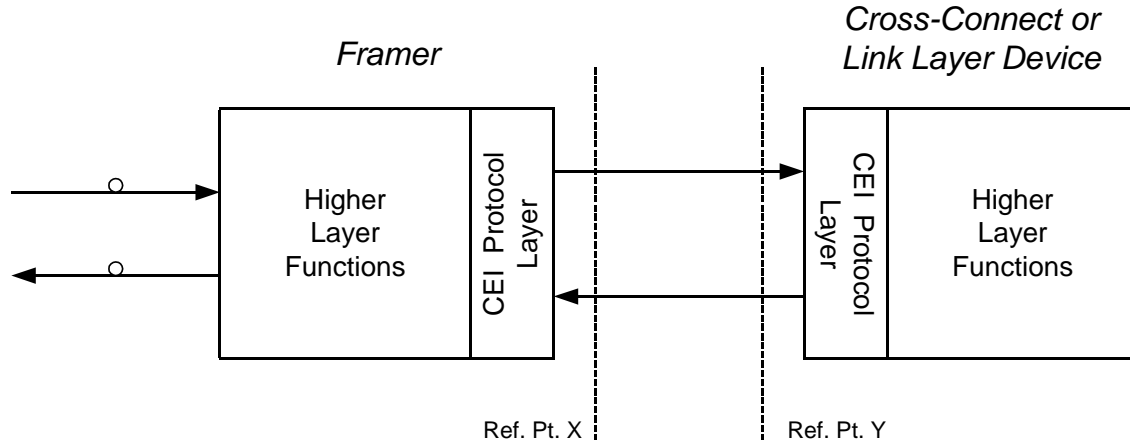


Figure 2 Foundation of Future OIF Interface Applications

The CEI Protocol is based on the layered concept. Three layers are defined; Adaptation, Aggregation and Framing (See Figure 3). In general, the Adaptation layer performs client signal alignment, performance monitoring and mapping functions to convert from client protocol formats to CEI Protocol formats. The Aggregation layer is responsible for multiplexing and de-multiplexing clients onto CEI Protocol links. It performs link de-skew, where necessary. Optionally, it performs connectivity monitoring and supports a supervisory interface between the source and sink devices of CEI Protocol link. Finally, the Framing layer performs payload alignment, error detection and optional error correction of CEI Protocol links. In this document, functionality common to all applications for each layer is described first. Then, each application is assigned its own clause where the functions of each the three layers are further defined. The flow of data from the Client to CEI Protocol links is defined as the Transmit direction. The flow from CEI Protocol links to the Clients is defined as the Receive direction.

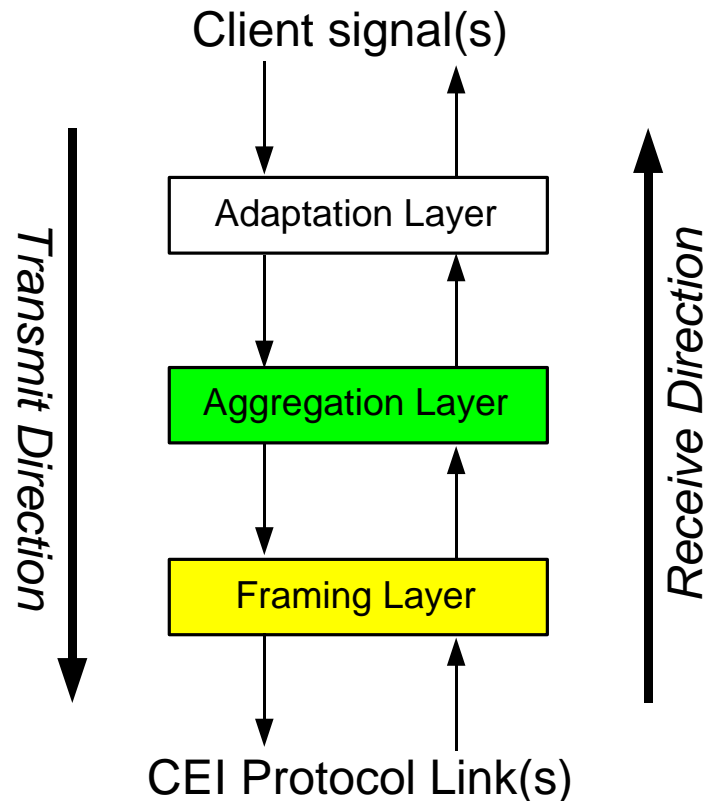


Figure 3 CEI Protocol Layers

6 Requirements

- The Implementation Agreement(s) shall allow single and multi-lane applications
- Shall conform to the data characteristics required by the CEI project (E.g., DC balance, Transition density, Max run-length)
- Support a wide range of client signals including SFI, SPI, TFI, 8b/10b Interfaces
- Protocol shall be the same within each CEI speed class (agnostic to SR or LR)
- Leverage layered protocol concepts developed for TFI-5
- In lane aggregation applications:
 - Ensure bit sequence integrity
 - Limit added client lane skew and wander
 - Provide mechanism to identify source client signals at de-mux point
 - Provide mechanism to de-skew CEI links

- In single lane applications :
 - Support SONET/SDH compliance at the optical carrier (OC) interface.

7 Objectives

- Support a wide range of client signals including SFI, SPI, TFI, 8b/10b interfaces and be agnostic to the client link protocol.
- Provides similar EMI spectrum for a wide variety of clients.
- Protocol shall be the same for each CEI speed class (agnostic to reach)
- Achieve low BER
- Provides layer pass-through for light-weight implementations
- In lane aggregation applications:
 - Limit added client lane skew and wander
 - Limit added latency

8 Common Definitions

8.1 Introduction

This clause defines the common functions of the CEI Protocol for the Adaptation, Aggregation and Framing layers that will be used by all applications. The layers may have options and parameters. The clauses pertaining to specific applications will select appropriate options and assign values to parameters.

8.2 Frame Format Definition

The CEI Protocol has a single frame format to address both low-latency and high-integrity applications. Parity check code from a Fire code generator polynomial is used for error checking and optionally, for forward error correction (FEC). The frame length of 1584 bits consists of 1560 bits of payload and 24 bits of overhead. Each frame is a Fire Code block.

8.2.1 Frame Format

Figure 4 shows a CEI Protocol frame. The length of the frame is 1584 bits. The frame has a regular structure. Each frame contains eight rows, and a Supervisory overhead bit is associated with each odd numbered row. At the end of each frame there is a 20-bit overhead field. Each row has three 65-bit payload fields (shown as T and 64 bits of payload in Figure 4). Transmission is from left to right within each row, and from top to bottom, between rows. The payload bits (1560 bits total) carry client payload originating from the Adaptation layer. The Adaptation layer of certain applications may insert stuff bits into the payload region for rate matching.

A number of existing protocols such as 1G Ethernet, 10G Ethernet (XAUI), InfiniBand, PCI-Express, Rapid I/O and Fibre Channel use 8b/10b encoding. It may be desirable for CEI Protocol links to be able to carry 8b/10b clients using a transcoding scheme that is similar to, but more efficient than, 64b/66b encoding. Twenty-four Transcoding overhead bits (T) are highlighted in Figure 4 to illustrate

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement example locations where the Adaptation layer may place transcoding tags for non-multiplexed payloads. For multiplexed payloads, the locations of transcoding tags will depend on the multiplexing structure. The ratio of one transcoding bit for every 64 payload bits matches the overhead requirements of GFP-T format in ITU-T G.7041. For applications that do not use transcoding, all 65 payload bits (including T) are available for carrying client payload.

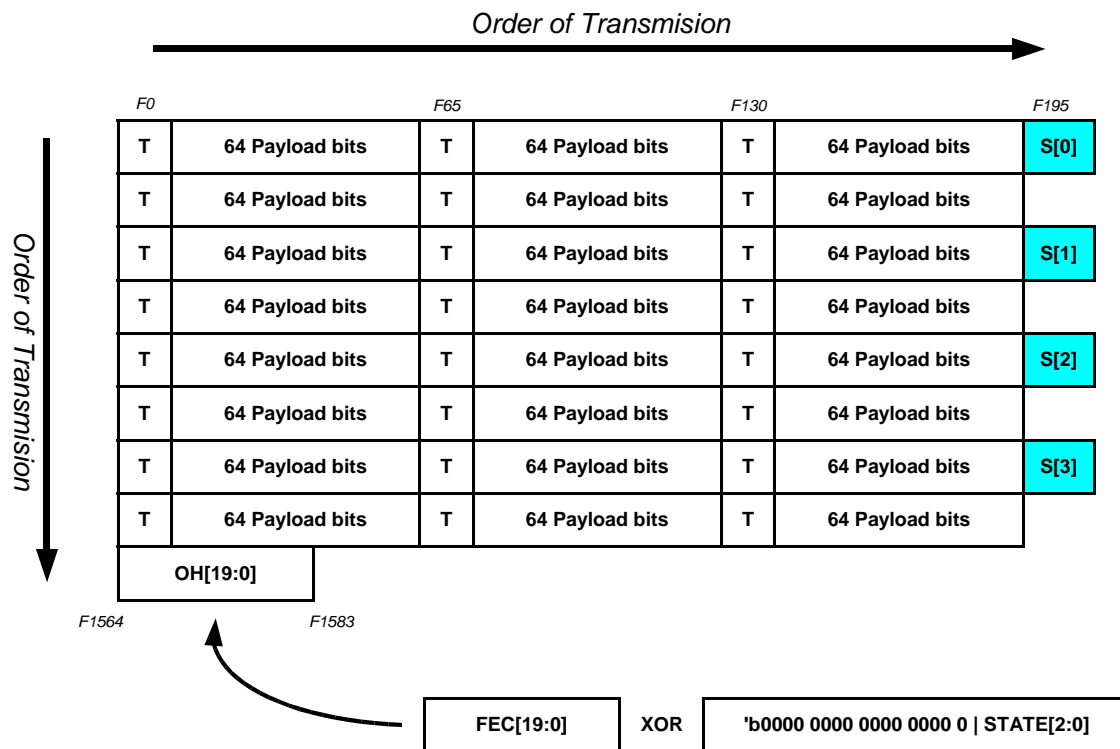


Figure 4 CEI Protocol Frame Format

The twenty overhead bits (OH[19:0]) located at the end of the frame (F1564 to F1583) hold the XOR of the Fire code parity check bits (FEC[19:0]) with the fixed zero field and the STATE[2:0] bits. Bit F1564 carries FEC[19], while F1583 carries the XOR of FEC[0] with STATE[0]. All bits in the frame are scrambled. The Overhead bits are assigned to the Framing layer in accordance to Table 1.

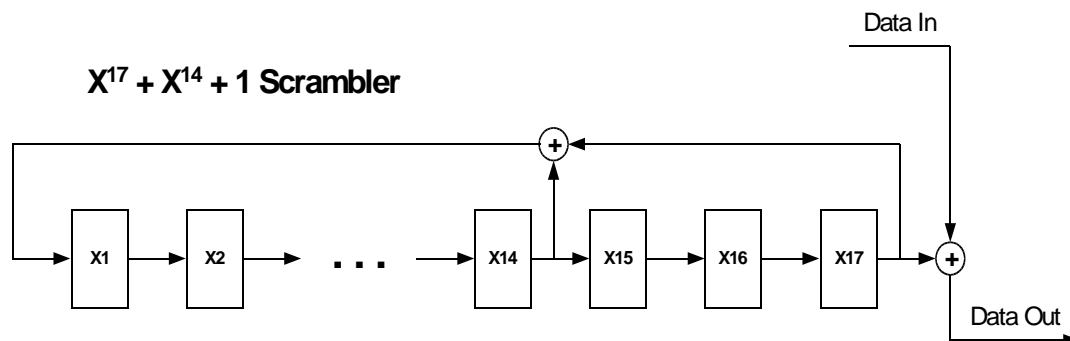
Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

Overhead	Location	Function
OH[19:3]	F1564 : F1580	Framing, Scrambler Synchronization, Error Detection, Forward Error Correction (Optional),
OH[2:0]	F1581 : F1583	Link State, Error Detection, Forward Error Correction (Optional)

Table 1 Overhead Assignment in a Frame

The optional Supervisory Interface is supported over the Supervisory overhead bits (S[0:3]) in Figure 4. The S[0:3] bits have not been assigned to any layer as the functions they serve tend to be management in nature and may cover both CEI Protocol link primitives and client primitives. When implemented, the Supervisory Interface is attached to the Aggregation layer. Examples of functions supportable by the Supervisory Interface include, but not limited to, inter-card data communications channel, per STS-1 status in TDM equipment, and flow-control flags/messages in packet-based equipment. The definition of the Supervisory Interface is application specific.

Frames are scrambled by a free running LFSR scrambler with a characteristic polynomial $g(X) = X^{17} + X^{14} + 1$ (See Figure 5). The scrambler output is applied to all Payload bits (including T bits), to all Supervisory Interface bits, and to Overhead bits the frame. Since the scrambler bits are XOR'ed with FEC[19:3] at the CEI Protocol Transmitter, the state of the scrambler can be recovered, at the CEI Protocol Receiver, by subtracting the calculated parity from the received scrambled parity at bits F1564:F1580.

**Figure 5 Free Running Scrambler**

The state of a CEI Protocol link, is communicated between the Transmitter and Receiver pair by the STATE[2:0] bits. For robust communications, STATE[2:0] is

transmitted in every frame. (Note that STATE[2:0] is XOR'ed with FEC[2:0].)

Table 2 shows the overhead locations where each bit in STATE[2:0] is transmitted.

State	Overhead	Location
STATE[2:0]	OH[2:0]	F1581:F1583

Table 2 Overhead location of STATE[2:0]

The encoding of the State bits is shown in Table 3 below. When the associated receiver in the local device is in need of training, the transmitter will enter the Training Request state (TRequest). The transition to the TRequest state has precedence over all other transitions. The STATE[2:0] bits are set to 'b111 to indicate training is requested. The payload field (including T bits) will carry the training pattern defined by the CEI Implementation Agreement to train the remote receiver. The transmitter will stay in this state until the local receiver has completed training. At this point, the transmitter will transition to the TPresent state if the remote receiver continues to require training. The STATE[2:0] bits are set to 'b110 to indicate that the local receiver has completed training. Training pattern continues to be sent in order to train the remote receiver. When the remote receiver reports training completion (via STATE[2:0] != 'b111), the transmitter enters the Idle state. The transmitter will transition directly from the TRequest state to the Idle state, if the remote receiver has already completed training when the local receiver completes its training. In the Idle state, the Framing layer is capable of transferring data. However data may not be available to/from higher layers. The payload bits are set to all zeros and payload scrambling is enabled. The link is fully operational in the PScramble state.

The State bit STATE[0] (assigned to OH[0]) may be used to convey a multiframe alignment signal (MF) in the Idle and PScrambled states. The length of the multiframe is application specific and is static. The first frame in the multiframe is indicated by setting STATE[0] to logical one, and the remaining frames are indicated by setting STATE[0] to logical zero.

The TX Hold signal in Figure 6 is a software controlled signal. It is optional and may be used by software to synchronize the start of normal payload activity when multiple CEI-P links are aggregated to carry a single high-bandwidth client. In a simplex application, the remote receiver never requests training as it has no path to signal a need for training. The transmitter will dwell in the TPresent state for D₁ frames to allow time for the remote receiver to complete training.

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

State Name	STATE[2:0]] Encoding	Description
TRequest	'b111	Training pattern present. Training Requested. Payload scrambler disabled.
TPresent	'b110	Training pattern present. Local receiver trained. Payload scrambler disabled.
Idle	'b011	Local and Remote receivers trained. Payload scrambler enabled. First frame of a multiframe. MFAS marker frame.
Idle	'b010	Local and Remote receivers trained. Payload scrambler enabled. Second and subsequent frames of a multiframe.
PScramble	'b001	Link fully operational. Payload scrambler enabled. First frame of a multiframe. MFAS marker framer.
PScramble	'b000	Link fully operational. Payload scrambler enabled. Second and subsequent frames of a multiframe.
Reserved	Others	Reserved for future definition.

Table 3 Transmitter Link State Reporting

The transition diagram of the state reporting mechanism at a CEI-P transmitter for duplex applications is shown in Figure 6.

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

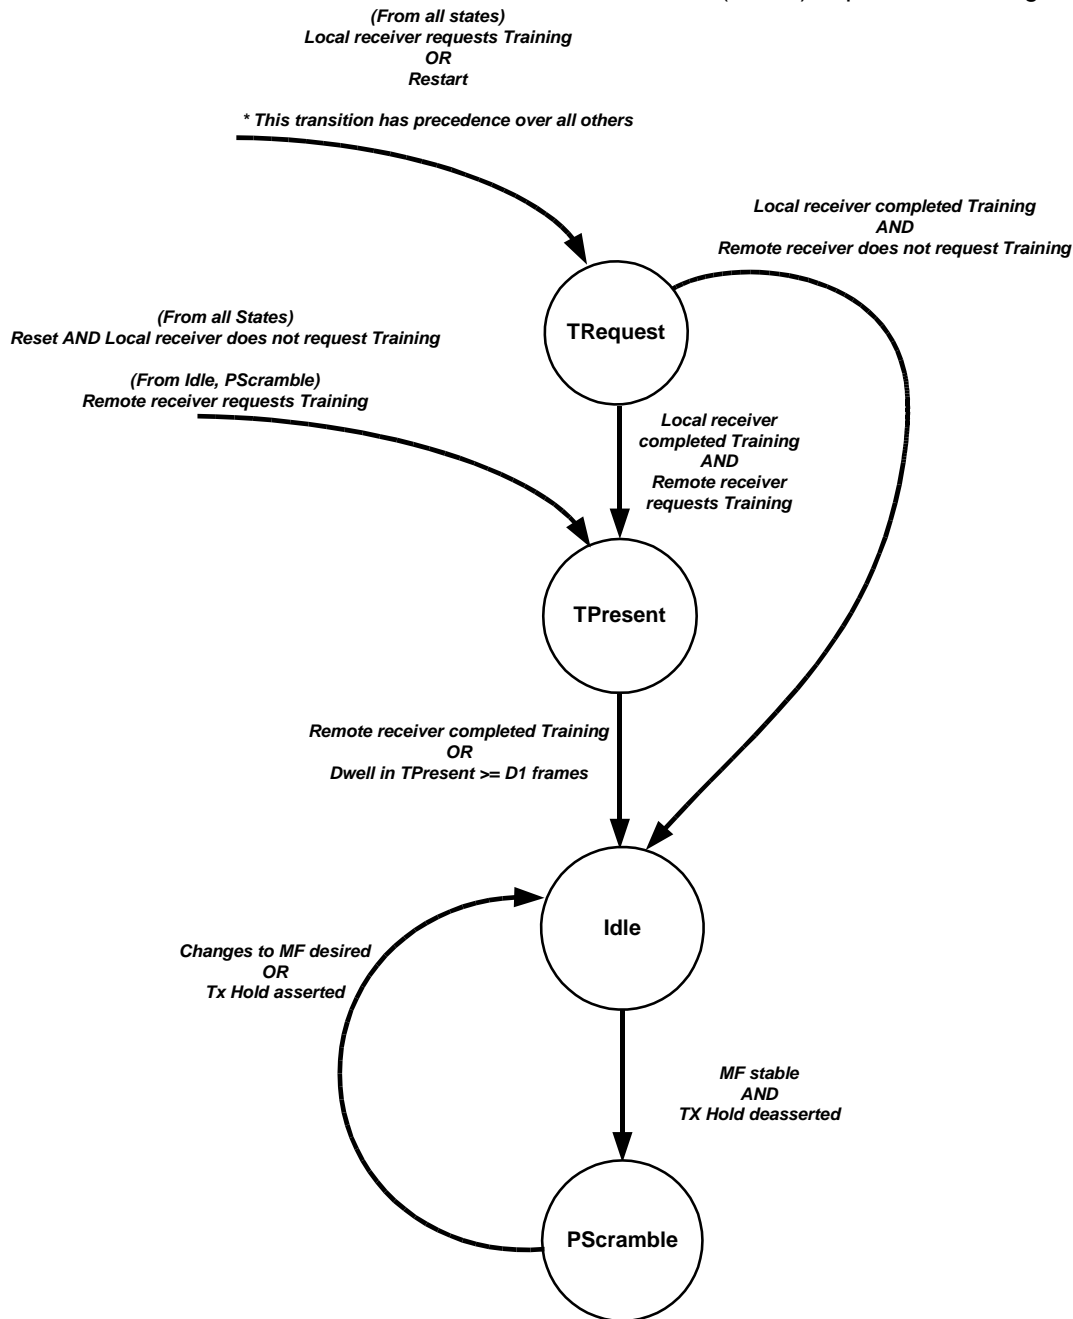


Figure 6 Transmitter Link State State Machine Diagram – Duplex

The transition diagram of the state reporting mechanism at a CEI-P transmitter for simplex applications is shown in Figure 7. In a simplex application, the CEI-P transmitter in the local device and the CEI-P receiver in the remote device are connected by a link. No return path from the remote device to the local device is associated with this CEI-P transmitter / receiver pair. Consequently, STATE[2:0] communications is unidirectional. The CEI-P Receiver is not able to request Training Pattern nor indicate training completion. Thus the TRequest state is not

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement present in Figure 7. An external means is needed to trigger Training Pattern insertion in the CEI-P transmitter.

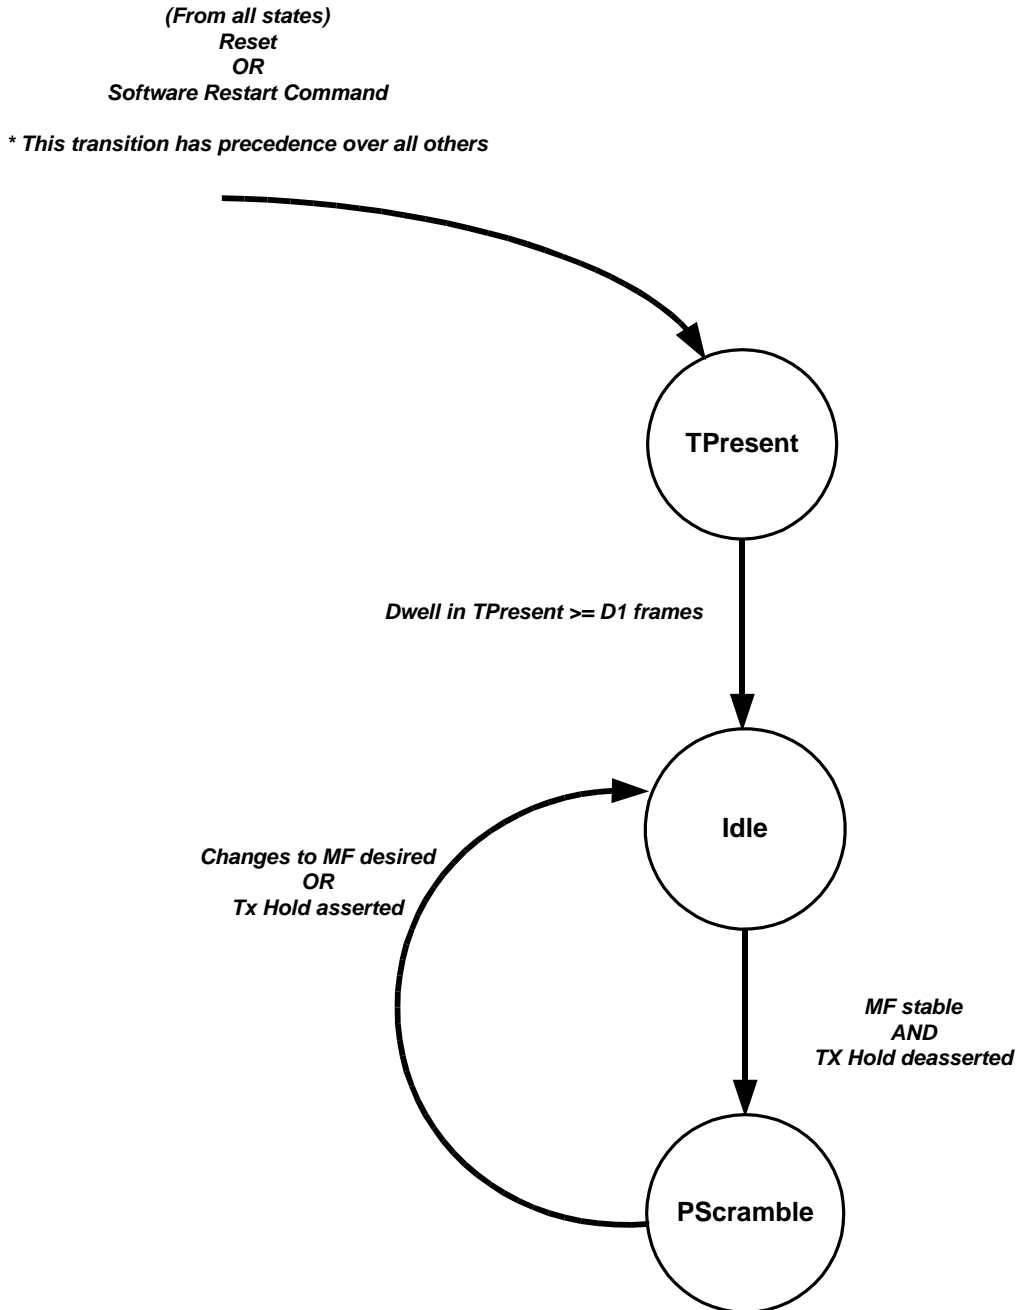


Figure 7 Transmitter Link State State Machine Diagram – Simplex

The Duplex Receiver link state machine (Figure 8) can be used in simplex or duplex CEI-P applications. It tracks the state of the remote transmitter (STATE[2:0]) of a CEI-P link. STATE[2:0] values are accepted when they are

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

received with a constant value for R1 consecutive presentations. The receiver link state machine enters the Seek state upon reset and when it is not able to track the state of the remote transmitter. The receiver transitions from the any state to the TRequest or the TPresent state when the accepted state from the transmitter is in those respective states. The receiver waits in the Idle state until the optional software controlled signal RX Hold is negated. It will transition to the Idle state when the RX Hold signal is asserted. As with the TX Hold counterpart, RX Hold is an optional software controlled signal. It is useful in synchronizing multiple CEI-P links that are aggregated to service a single high bandwidth client.

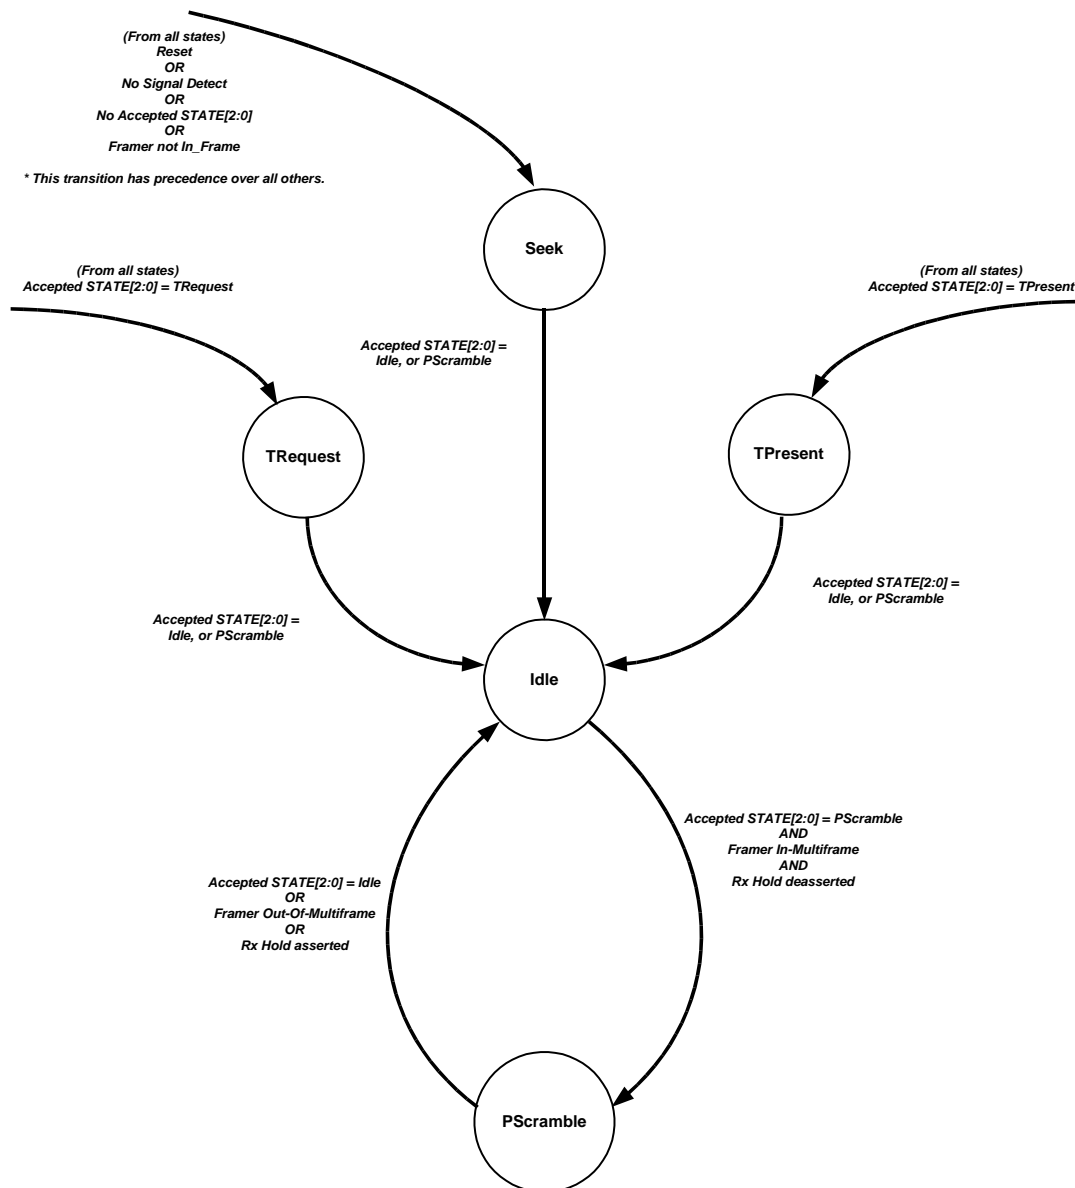


Figure 8 Receiver Link State State Machine Diagram –Duplex

The Simplex Receiver link state machine (Figure 9) is restricted to simplex CEI-P applications. It tracks the state of the remote transmitter (STATE[2:0]) of a CEI-P link. STATE[2:0] values are accepted when they are received with a constant value for R1 consecutive presentations. The receiver link state machine enters the Seek state upon reset and when it is not able to track the state of the remote transmitter. The receiver transitions from the any state to the TPresent state when the accepted state from the transmitter is in that state. The receiver waits in the Idle state until the optional software controlled signal RX Hold is negated. It will transition to the Idle state when the RX Hold signal is asserted. As with the TX Hold counterpart, RX Hold is an optional software controlled signal. It is useful in synchronizing multiple CEI-P links that are aggregated to service a single high bandwidth client.

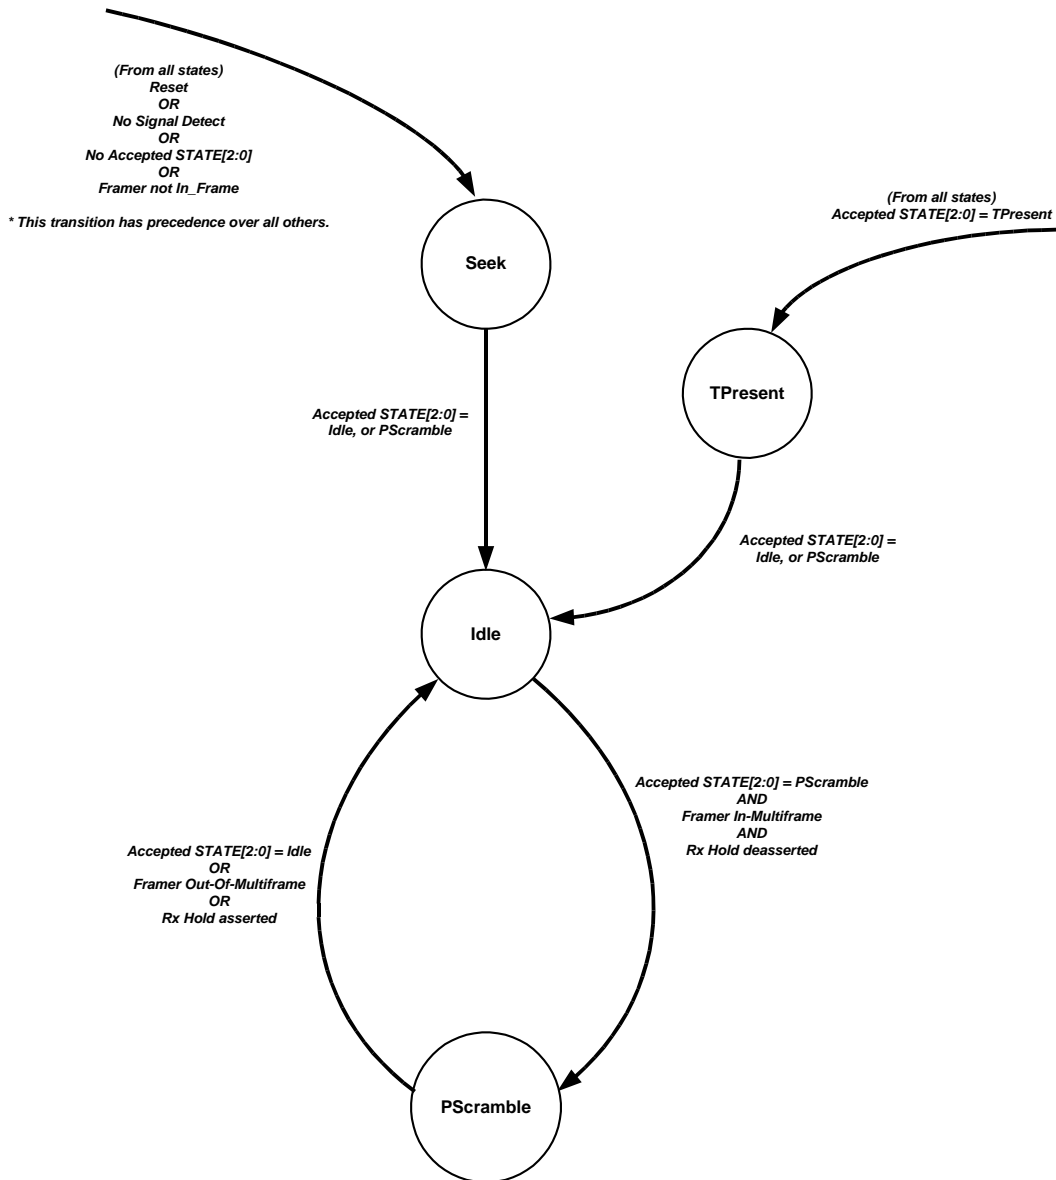


Figure 9 Receiver Link State State Machine Diagram - Simplex

The CEI Implementation Agreement has defined a Training Pattern for calibrating CEI electrical components. (Table 5-6 of the CEI IA is replicated here as Table 4 for reference only.) When sending the CEI Training Pattern, the sequence will be 391 bits long, consisting of a 192-bit fixed data field, followed by a 199-bit pseudo-random data field. The Supervisory Channel bits (S) are scrambled and form part of the pseudo-random data field. The first bit of the fixed data field is aligned to the first bit of the CEI Protocol frame. Each CEI Protocol frame carries four CEI Training Patterns. Table 4 Training Pattern (from CEI Electrical IA)

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement shows the alignment of the patterns. Implementations may support additional training patterns. The definition of these patterns is outside the scope of this IA.

Pattern (Hex)	Purpose
00 FF 00 FF 00 FF	48 bits - f/16 square wave
00 08 00	24 bits - positive impulse with 12 leading and trailing zeros
55 55 55 55 55	48 bits - f/2 square wave
FF FE FF	24 bits - negative impulse with 12 leading and trailing ones
00 FF 00 FF 00 FF	48 bits - f/16 square wave
At least 192 random or pseudo-random bits	Approximation of normal randomized data patterns

Table 4 Training Pattern (from CEI Electrical IA)

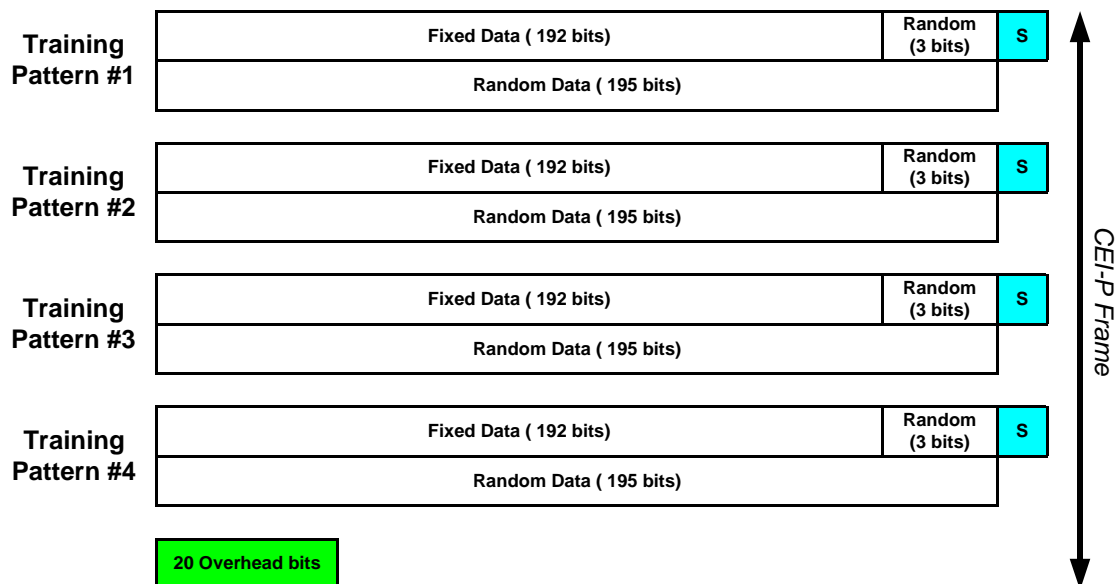


Figure 10 Alignment of CEI Training Pattern to CEI Protocol Frames

Figure 11 shows a logical block diagram of the frame generation process. Actual implementation may differ.

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

The generation of Forward Error Correction parity check code bits is mandatory and is appended to the end of the block in each frame. The FEC code is Fire code with a generator polynomial of $g(X) = (X^{13} + 1)(X^7 + X + 1)$, and is capable of correcting a single error burst of up to 7 bits, per frame. The parity check code is computed over all the Payload (including T bits) and Supervisory bits in the current frame, after scrambling. The Fire Code parity check code is XOR'ed with the associated STATE[2:0] bits, scrambled and inserted to the end of the frame.

The STATE[2:0] bits share some of the bit positions as the Fire Code parity bits. A single bit error in the payload will typically cause several parity bits, and the corresponding STATE[2:0] bits, to invert. Using the Fire code to correct errors in the STATE[2:0] bits, at a CEI Protocol Receiver, is therefore complex. Since the transmitted values of the STATE[2:0] bits are not permitted to change during normal operation, they should be protected from transmission errors using de-bouncing over multiple frames, rather than Forward Error Correction.

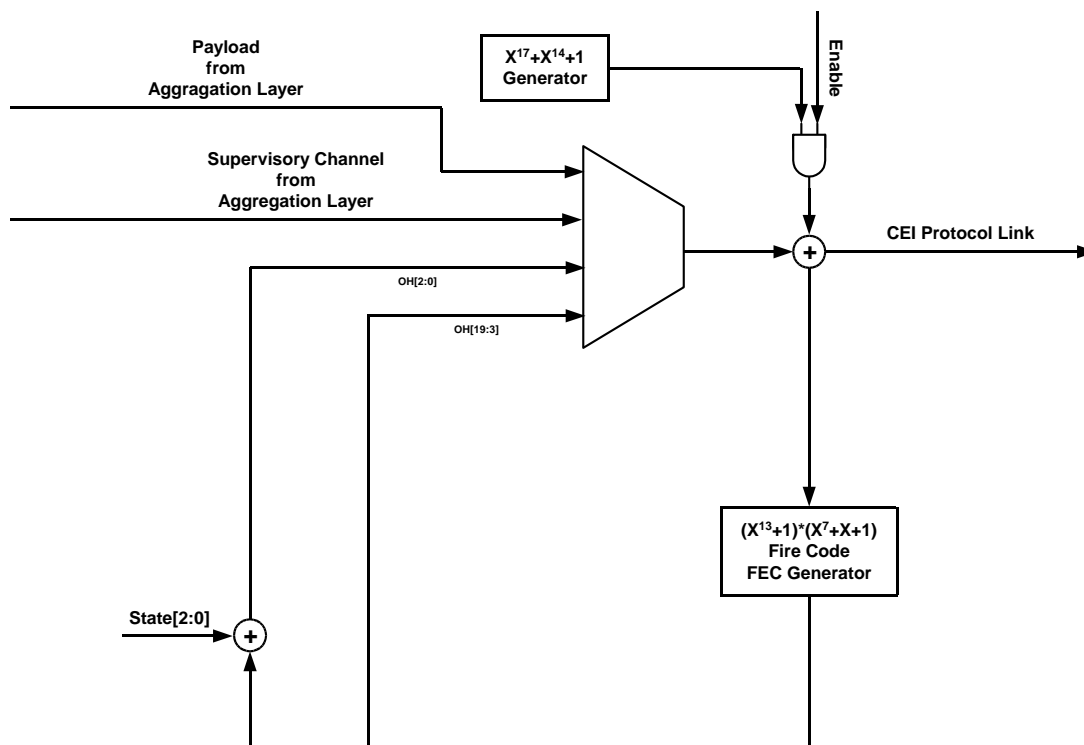


Figure 11 Frame Generation Process – Logical Block Diagram

8.3 Framing Layer Definition

The Framing layer is responsible for frame boundary delineation, error detection and maintaining rich transition density. Fire Code parity check code is defined to mark frame boundaries and to detect transmission errors. Transition density is

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement provided by a free-running LFSR scrambler. Optionally, the Framing layer at the receiver provides Forward Error Correction capability. The following requirements apply to the frame formats of the CEI Protocol frame.

Req. No.	Description
R1.1	The frame format of a CEI Protocol link shall be as shown in Figure 4. The first bit the frame is labeled F0 and the last bit F1583. F0 is aligned to the first payload bit (a.k.a., the first T bit), while F1583 is aligned to FEC[0], the least significant bit of the Fire code parity check code.
R1.2	Transmission of a CEI Protocol frame shall be sequential from bit F0 to bit F1583.
R1.3	In the Idle state, all payload bits (Including T bits) shall be set to zero (before scrambling) by a CEI-P transmitter, and shall be ignored by a CEI-P Receiver.
R1.4	A free-running stream cipher, with a polynomial of $g(X) = X^{17} + X^{14} + 1$, shall be used for scrambling. (See Figure 5.) The stream cipher state shall advance at every bit time, without regard to whether the bit is scrambled.
R1.5	The Overhead bits (OH[19:3]) shall be set to 'h00000 XOR FEC[19:3].
R1.6	When multiframe alignment is implemented, STATE[0] in the Idle and PScramble states shall be set high at the first frame of the multiframe and set low at all other frames.
R1.7	The FEC parity check code shall be computed using the Fire code polynomial $g(X) = (X^{13} + 1)(X^7 + X + 1)$. The parity check code shall cover the Payload bits (including T bits) and the Supervisory bits in the current Frame, after scrambling.
R1.8	The FEC parity check codes FEC[19:3] bits shall be inserted into frame locations F1564 : F1580, before scrambling. The FEC parity check codes FEC[2:0] bits shall be XOR'ed with the state code STATE[2:0] bits and inserted into frame locations F1581 to F1583, before scrambling.
R1.9	Payload scrambling shall be enabled in the PScramble and Idle states. The payload scrambler shall be disabled in the TRequest and TPresent states.

Table 5 CEI Protocol Frame Format Requirements

A CEI Protocol Receiver uses the Fire Code parity check code in the received data stream to delineate frame boundaries (See Table 6 Sample Framing Algorithm)

for sample framing algorithm). When the receiver is in the In-Frame state, Fire Code parity check code errors in M_1 consecutive frames will cause a transition to the Out-Of-Frame state. When the receiver is in the Out-Of-Frame state, it will sequentially search the received stream for candidate frame boundaries. A potential frame boundary is found if the associated 1584-bit block has correct Fire Code parity check code, after descrambling. The potential frame boundary is verified, for correct Fire Code parity check code over M_2 consecutive frames. At which point, the framer transitions to the In-Frame state. Implementations may use a only subset of the Fire Code parity check bits in the frame delineation process. The speed of frame delineation is proportional to the number of candidates (N_1) being searched simultaneously. It is outside the scope of this Implementation Agreement to specify N_1 .

An example of a framing algorithm is shown in Table 6 and below. It is provided here for illustrative purposes only. Alternative compliant implementations exist.

Step #	Description
1	Select an arbitrary bit as the candidate CEI-P frame boundary.
2	Compute a candidate Fire Code checksum from a candidate block of 1564 bits starting at the candidate frame boundary. Note : In the absences of transmission errors, if the candidate frame boundary matches the actual CEI-P frame boundary, the candidate Fire Code checksum would match the Fire Code checksum at the transmitter, before scrambling.
3	Generate a candidate scrambler state by XOR'ing the candidate Fire Code checksum with the next 17 bits received following the 1564-bit candidate block.
4	Load the de-scrambler with the candidate scrambler state and release it to advance normally.
5	Compute candidate Fire Code checksum #2 for candidate block #2. The first bit of block #2 is 1584 bits after block #1.
6	De-scrambled the next 17 received bits after candidate block #2 and compare with candidate Fire Code checksum #2. Note : In the absence of transmission errors, the 17 received bits, after de-scrambling should match the candidate checksum.
7	If the comparison in Step 6 yielded a match, a potential frame boundary is found. This frame is deemed the first frame in the search for consecutive matches. The Fire Code checksum is checked in subsequent frames until M_2 consecutive matches are

found. At this point the framer can declare In-Frame. The steps above naturally synchronize the de-scrambler in the receiver with the scrambler in the transmitter.

If the comparison in Step 6 yielded a mismatch, the candidate frame boundary is assume to be not aligned with the actual CEI-P frame boundary. An alternate candidate start of CEI-P frame bit is selected.

Table 6 Sample Framing Algorithm

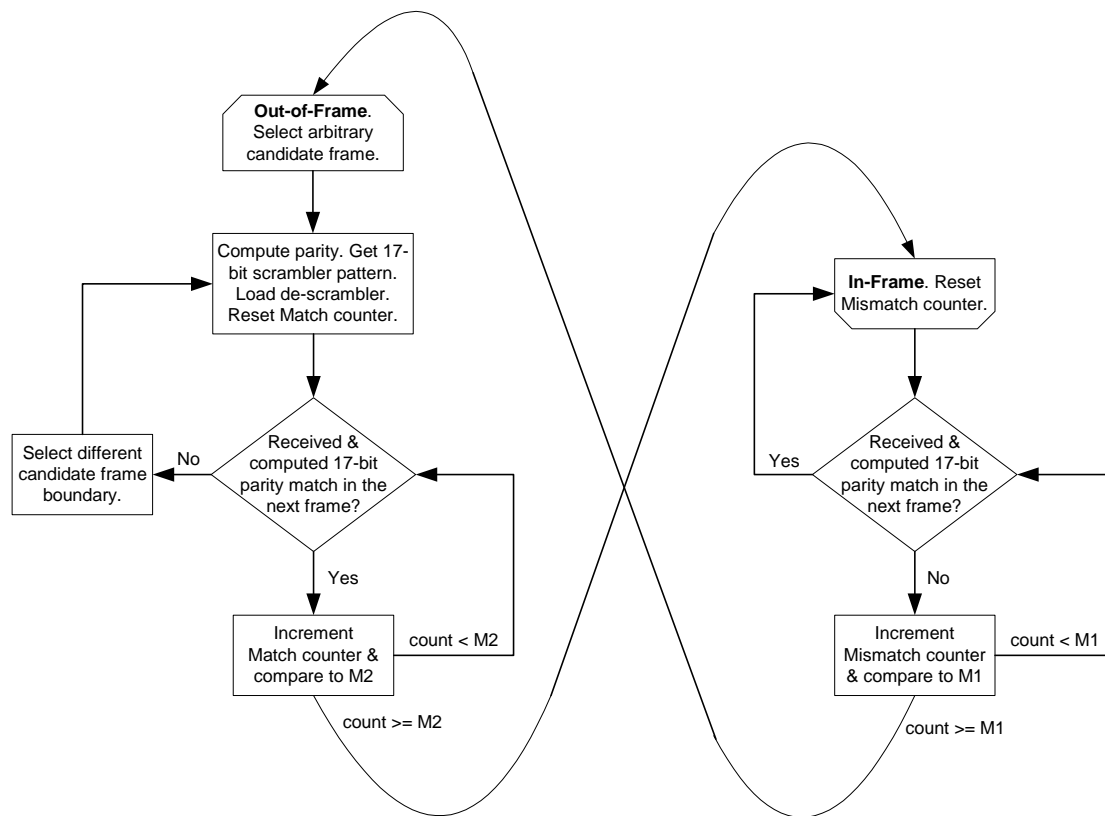


Figure 12 Sample Framing Algorithm




Figure 13 Framing State Machine Diagram

When the multiframe structure is enabled, the CEI Protocol Receiver uses the Multiframe Alignment Signal (MFAS) marker carried in STATE[2:0] of the received data stream to delineate multiframe boundaries. When the receiver is in the In-MultiFrame state, Multiframe Alignment Signal (MFAS) pattern errors in M_3 consecutive multiframe boundaries will cause a transition to the Out-Of-MultiFrame state. When the receiver is in the Out-Of-MultiFrame state, it will sequentially search the received STATE[2:0] for candidate multiframe boundaries. A potential multiframe boundary is found if a sequence of frames carried the MFAS marker in one frame, followed by $N-1$ frames without the MFAS marker; where N is the number of frames in a multiframe. The potential multiframe boundary is confirmed, and the receiver transitions to the In-MultiFrame state, when correct MFAS sequences (1 frame with MFAS, $N-1$ without) are detected for M_4 consecutive multiframe boundaries. A CEI Protocol Receiver is not required to automatically detect the length of the multiframe (N).

The following requirements apply to the frame and multiframe delineation process of a CEI Protocol Receiver.

Req. No.	Description
R2.1	A CEI Protocol Receiver shall follow Figure 12 Sample Framing Algorithm in framing to a CEI Protocol link.
R2.2	A CEI Protocol Receiver shall monitor Fire Code parity check

Req. No.	Description
	code. It shall transition from the In-Frame state to the Out-Of-Frame state when incorrect Fire Code parity check code is received in M_1 consecutive frames. It is acceptable to monitor a subset of the parity check bits.
R2.3	When a CEI Protocol Receiver is in the Out-Of-Frame state, it shall search for candidate frame boundaries. It shall transition to the In-Frame state when the current candidate frame boundary yields correct Fire Code parity check code for M_2 consecutive frames. It is acceptable to monitor a subset of the parity check bits.
R2.4	Consequential actions, such as data squelching, as a result of the CEI Protocol Receiver being in the Out-Of-Frame state, shall be defined by the each specific application.
CR2.5	For applications that define a multiframe, the MFAS marker (carried in STATE[2:0]) shall be set to 'b011 or 'b001, to indicate the first frame of the multiframe. the STATE[2:0] shall be set to values other than 'b011 or 'b001. For applications that do not define a multiframe, STATE[2:0] does not indicate MFAS markers.
CR2.6	For applications that define a multiframe, the CEI Protocol Receiver shall monitor STATE[2:0] for the MFAS marker. A CEI Protocol Receiver shall transition from the In-MultiFrame state to the Out-Of-MultiFrame state when the received STATE[2:0] does not follow the expected sequence of one frame with MFAS marker and N-1 frames without, for M_3 consecutive multiframes.
CR2.7	For applications that define a multiframe, the CEI Protocol Receiver shall transition from the Out-Of-MultiFrame state to the In-MultiFrame state when the received STATE[2:0] follows the expected sequence of one frame with MFAS marker and N-1 frames without, for M_4 consecutive multiframes.
R2.8	When the frame delineation process of a CEI Protocol Receiver is in the Out-Of-Frame state, the multiframe alignment process shall be locked in the Out-Of-MultiFrame state.
R2.9	Consequential actions, such as data squelching, as a result of the CEI Protocol Receiver being in the Out-Of-MultiFrame, shall be defined by the each specific application.

Table 7 CEI Protocol Receiver Frame and Multiframe Requirements

The Forward Error Correction parity check code bits (FEC[19:3]) are scrambled and placed in frame location F1564:F1580 by the CEI Protocol Transmitter. A

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

CEI Protocol Receiver can recover the scrambler sequence by computing an FEC parity check code locally and XOR'ing it with the received, scrambled OH[19:3].

The following requirements and options apply to the scrambler synchronization process of a CEI Protocol Receiver.

Req. No.	Description
R3.1	When a CEI Protocol Receiver is in the Out-Of-Frame state and has selected a new candidate CEI-P frame boundary, it shall compute a candidate state of the de-scrambler using bits F1564 to F1580 of the candidate frame. Comparison of Fire Code parity check code match/mismatch shall begin at the next frame.
R3.2	The candidate state of the scrambler shall be derived by subtracting the computed FEC parity of the candidate frame from the received value at OH[19:3] (bits F1564:F1580) of the candidate frame.
R3.3	Further consequential actions for the Out-Of-Frame state shall be defined by the each specific application.

Table 8 CEI Protocol Receiver Scrambler Synchronization Requirements

The training of a link is reported via the STATE[2:0] bits.

Req. No.	Description
R4.1	The state of a bi-directional CEI Protocol link shall be reported in the STATE[2:0] bits according to Table 3.
R4.2	State transitions shall follow Figure 6 and Figure 8.

Table 9 CEI Protocol Bi-directional Link Requirements

The Forward Error Correction process is a FEC block that is aligned to the CEI Protocol frame boundary. . Each FEC block is assigned 1560 bits of payload (including T), and four Supervisory overhead bit (S), and is protected by a 20-bit Fire Code parity check code. This code is capable of correcting a single burst of up to 7 bits in each frame. The generator polynomial is $g(X) = (X^{13} + 1)(X^7 + X + 1)$. The FEC parity check code is computed over the scrambled payload of each CEI Protocol frame. The parity check codes in each frame is then scrambled.

Req. No.	Description
R5.1	A CEI Protocol frame shall consist of one FEC block covering the entire frame.
R5.2	The FEC parity check code shall be generated by the polynomial $g(X) = (X^{13} + 1)(X^7 + X + 1)$ over the scrambled payload (including T) bits, and the Supervisory overhead bits. At the start of each frame, the FEC parity check code generator shall be initialized to an all-zeros pattern.
R5.3	The 20-bit FEC parity check code of each FEC block shall be inserted into the FEC[19:0] Overhead bits, by the CEI Protocol Transmitter.
O5.4	When Forward Error Correction is enabled, a CEI Protocol Receiver should be able to correct a single burst of up to N (N = 7) errors within each FEC block.
CR5.5	When Forward Error Correction is implemented at a CEI Protocol Receiver, the function shall be configurable as being enabled or disabled.
O5.6	Implementation of Forward Error Correction at a CEI Protocol Receiver is optional.

Table 10 CEI Protocol Forward Error Correction Requirements

8.4 Aggregation Layer Definition

The Aggregation Layer is responsible for placing client data from the Adaptation Layer onto CEI Protocol links. Its function covers both multiplexing and inverse-multiplexing. Rate adaptation from the client signal rate to the CEI Protocol link rate is handled by the Adaptation layer. Thus, the interface between the Aggregation Layer and the Adaptation Layer can be timed by the CEI Protocol link(s). The Supervisory Interface is supported by the Aggregation layer.

In cases where clients have lower bandwidth than the CEI Protocol link, the Aggregation Layer can multiplex several clients onto a single link (See Figure 14). Logically, the Aggregation Layer process would have interfaces to multiple Adaptation Layer processes, one for each client, and a single interface to the Framing Layer process. The case where a client is a bus containing multiple physical links (E.g., 10GE over XAUI, SPI-4, SPI-5) is viewed as a having only one Adaptation layer serving that client. De-skew of signals within the client bus, if necessary, is the responsibility of the Adaptation layer. The assignment of client data bits to CEI Protocol frame bits is defined in the clauses of the specific application.

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

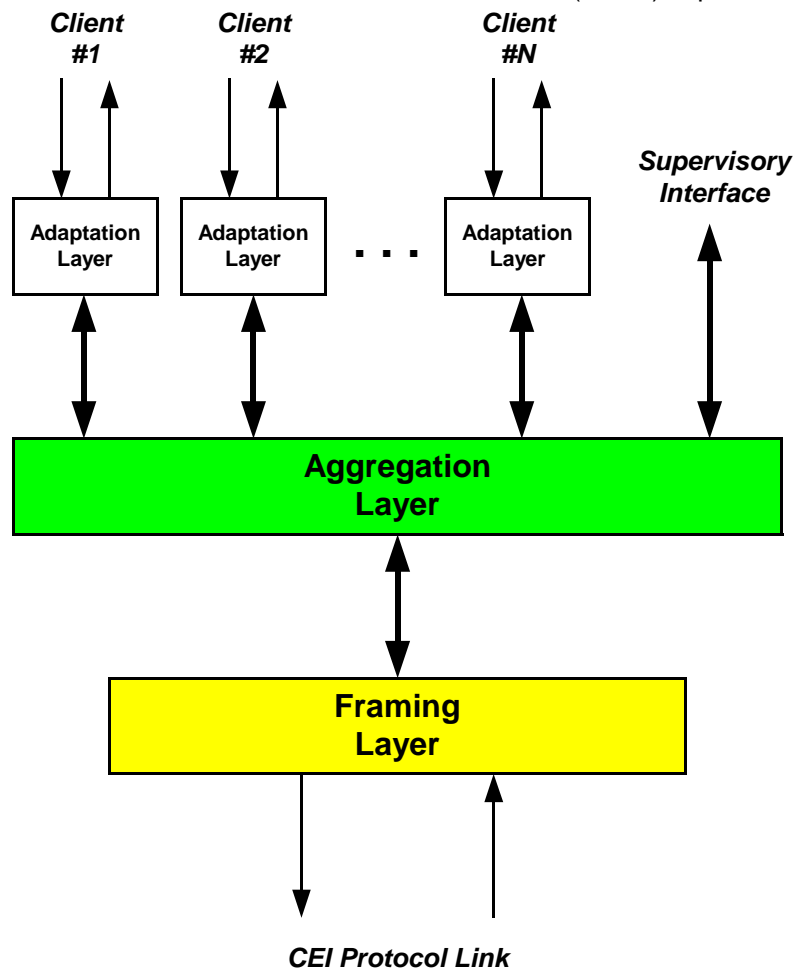


Figure 14 Logical Aggregation Layer Interfaces - Multiplexing

In cases where a client has higher bandwidth than a CEI Protocol link, the Aggregation Layer distributes the client data over multiple links (See Figure 15). Logically, the Aggregation Layer process would have a single interface to the Adaptation Layer process, and multiple interfaces to the Framing Layer processes, one for each link. The Aggregation Layer at a CEI Protocol Receiver is responsible for de-skewing the data from the multiple CEI Protocol links. The theoretical range that the frame format is capable of de-skewing is ± 792 bits. Individual applications will set the de-skew limits to appropriate values. If multiframe is implemented, the Aggregation layer in the CEI Protocol Transmitter will ensure that the multiframe boundaries in all N Framing layer processes are aligned. The assignment of client data bits to CEI Protocol links and to bits within the CEI Protocol frames is defined in the clause of the specific application.

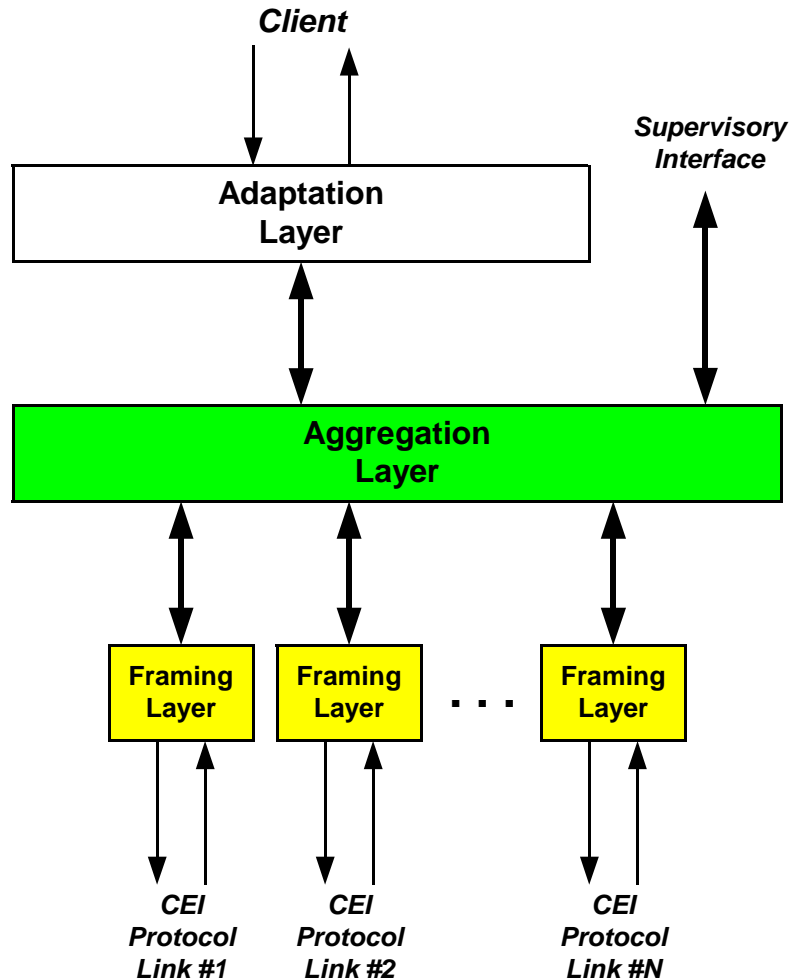


Figure 15 Logical Aggregation Layer Interfaces – Inverse-Multiplexing

Req. No.	Description
R6.1	For applications that multiplex multiple clients onto a single CEI Protocol link, the CEI Protocol Transmitter shall interleave the output of each Adaptation layer onto the Payload field of each CEI Protocol frame, in a round-robin fashion. The mapping and alignment of client data to the CEI Protocol frame (including grain size) shall be defined in the clauses of the specific application.

Req. No.	Description
R6.2	For applications that inverse-multiplex a client onto a set of CEI Protocol links, the CEI Protocol Transmitter shall distribute the client data onto the Payload field of the set of CEI Protocol frames, in a round-robin fashion. The mapping and alignment of client data to the CEI Protocol frames (including grain size), in the set of links, shall be defined in the clauses of the specific application.
R6.3	For applications that multiplex multiple clients onto a single CEI Protocol link, the CEI Protocol Transmitter/Receiver pair shall deliver data from each client source to the corresponding client sink.
R6.4	For applications that inverse-multiplex a client onto a set of CEI Protocol links, the CEI Protocol Receiver shall de-skew the set of CEI Protocol links. The range of skew to be tolerated shall be defined in the specific application.
R6.5	For applications that inverse-multiplex client onto a set of CEI Protocol links, the multiframe boundaries, if any, on the CEI Protocol links shall be aligned.

Table 11 CEI Protocol Multiplexing Requirements

Requirements R6.1 to R6.5 does not preclude applications where M clients share N CEI Protocol links. The M clients are expected to be frequency locked to each other. Such an application can be viewed logically as the serial instantiation of a Multiplexing Aggregation Layer followed by a Inverse Multiplexing Aggregation Layer. The M clients would first be concentrated onto a single Aggregation Layer Payload stream by a Multiplexing Aggregation Layer, and is then fed to a Inverse Multiplexing Aggregation Layer, where it is distributed to N Framing Layer functions.

8.4.1 Supervisory Channel

Optionally, the Aggregation layer provides a socket to support the Supervisory Interface. The use of this interface is application dependent. When implemented, data written to the Supervisory Interface are inserted into the Supervisory overhead bits (S[0:3]). The Supervisory overhead bits may be ganged into buses, or treated as individual signals. Data recovered from the Supervisory overhead bits are readable via the Supervisory Interface. The actual implementation of the Supervisory Interface and the mapping of Supervisory data onto S[0:3] is application specific. Examples uses of the Supervisory Interface include, data communications channel between cards, client status monitoring, and flow control messaging.

It is expected that a number of applications of the Supervisory channel can be supported by an HDLC data link. For those applications, a common HDLC frame format is shown in

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

Figure 16 Generic HDLC Frame of Supervisory Channel . Bit oriented HDLC in ISO-3309 [2] is used. Packets are delimited by flags ('b01111110) and protected by a CRC-16 parity check code. The Type field identifies the type and format of the message.

Flag (01111110)	TYPE (8bits)	Application Specific Message (Variable length)	CRC-16	Flag (01111110)
--------------------	-----------------	---	--------	--------------------

Figure 16 Generic HDLC Frame of Supervisory Channel

Type	Application	Description
00		Reserved
01	Lane Aggregation	Client Status
02	All	Parameter Read/Write
03-FC		Reserved
FD	Inter-chassis communications	System Vendor Specific
FE	Inter-card communications	Card Vendor Specific
FF	Inter-device communications	Device Vendor Specific

Table 12 Assignment of Type field in Supervisory Channel Packets

Req. No.	Description
O8.1	When implemented, the Aggregation layer should support a Supervisory Interface.
O8.2	When implemented, the Supervisory Interface in the Transmitter communicates with its peer in the Receiver via the Supervisory Overhead bits (S[0:3]). The organization of and the protocol carried over S[0:3] are application specific.
O.8.3	When an HDLC based Supervisory channel is implemented, the channel should use Supervisory Overhead bit S[0]. The organization of and the protocol carried over S[1:3] are application specific.
O8.3	When an HDLC based Supervisory channel is implemented, the CEI Protocol Transmitter should provide one or more software writable buffers to store transmit packets. Each buffer should be at least 64 bytes in size.
O8.4	When an HDLC based Supervisory channel is implemented, the CEI Protocol Transmitter should insert flags to delineate packets and insert zero stuff bits per reference [2]. The ability to compute and insert the CRC-16 parity check code automatically, in hardware, by the CEI Protocol Transmitter is not required.


Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

O8.5	When an HDLC based Supervisory channel is implemented, the CEI Protocol Receiver should provide one or more software readable buffers to store received packets. Each buffer should be at least 64 bytes in size.
O8.6	When an HDLC based Supervisory channel is implemented, the CEI Protocol Receiver should delineate packets with flags and zero stuff bits per reference [2]. The ability to compute and verify CRC parity check code automatically, in hardware, by the CEI Protocol Receiver is not required.
O8.7	The contents of the TYPE field should be one listed in Table 12.
O8.8	When the Connection Identification option is disabled, the CID[16:0] code should be held at all-zeros.
O8.9	When Mode reporting is disabled, the APM[7:0], AGM[3:0] and FRM[1:0] codes should be held at all zeros. Mode reporting is for further study and shall default to disabled in this revision of the IA.
O8.10	Implementation of the Supervisory Interface is optional.

Table 13 CEI Protocol Supervisory Interface Requirements

The Parameter Read/Write messages (Type=2) facilitate remote query and setting of parameters of devices at either end of a CEI Protocol link. The packet is fixed length and its format is shown in Figure 17 below. Figure 17 Parameter Read/Write Frame of Supervisory Channel

describes the individual fields within the frame.

**Figure 17 Parameter Read/Write Frame of Supervisory Channel**

Field	Description
Type	Set to 'h02 to identify packet as a Parameter Read/Write packet
Command	Action requests and responses
TAddr	Target entity of packet. A target is the set of CEI Protocol Transmitter, Receiver and the channel connecting them.
PID	Identity of parameter to be written or read
PVal	Value written to / read from specified parameter

Table 14 Fields of Parameter Read/Write Packets

The Command field defines the action taken by the Parameter Read/Write packet. Valid Command codes are listed in Table 15 below.

Command	Name	Description
00	NOP	
01	Read	Read specified parameter from target
02	ACK-Read	Response to Read. Value in PVal field.
03	Write	Write specified parameter with value in PVal field. Result of write reported in a ACK-Write packet.
04	Write-Incr	Increment specified parameter. Step size of increment defined by the sink device. Result of increment reported in a ACK-Write packet.
05	Write-Decr	Decrement specified parameter. Step size of increment defined by the sink device. Result of decrement reported in a ACK-Write packet.
06	ACK-Write	Reports result of Write, Write-Ince and Write-Decr commands
07	NACK-Error	Error detected in a Read, Write, Write-Ince and Write-Decr command.
08	NACK-NoParam	Specified parameter of Read, Write, Write-Ince and Write-Decr command is not supported by the targeted link.
09-FF	Reserved	

Table 15 Command field of Parameter Read/Write Packets

The Target Address (TAddr) field specifies the target entity of a Parameter Read/Write packet. The encodings of TAddr are listed in Table 16 below. The TAddr of ACK-Read, ACK-Write, NACK-Error, and NACK-NoParam response packets carry the TAddr value of the corresponding Read, Write, Write-Incr, and Write-Decr request packets.

TAddr	Description
0000	Specifies the link carrying the current Parameter Read/Write packet.
0001	Specifies the link carrying the response to the current Parameter Read/Write packet.
0002-0FFF	Reserved
1000-FFFF	Vendor specific

Table 16 Target Address (TAddr) field of Parameter Read/Write Packets

The Parameter Identifier (PID) field specifies the parameter to be accessed by a Parameter Read/Write packet. The encodings of TAddr are listed in Table 17 below. The PID of ACK-Read, ACK-Write, NACK-Error, and NACK-NoParam response packets carry the PID value of the corresponding Read, Write, Write-Incr, and Write-Decr request packets.

The PID code of “Announced Connection ID” is sent by the CEI Protocol Transmitter to identify the CID of the local link. It should only be sent as part of a Write packet and with a Target Address of 0x000 to specify the local link. The PID code of “Expected Connection ID” is used to set or query the expected CID of CEI Protocol Receivers at a remote device. PID codes “Announced MODE” and “Expected MODE” behaves similarly to their CID counterparts for announcing transmitter mode and setting receiver expectations.

PID	Description
0000 – 00FF	Receiver equalizer parameter number.
0100 – 01FF	Transmitter equalizer parameter number.
0200	Training Pattern
0201	Multiframe Length
0202	Announced Connection ID (CID)
0203	Announced MODE
0204	STATE
205	Expected Connection ID (CID)
206	Expected Mode
0207 – 0FFF	Reserved
1000 – FFF0	Vendor Specific
FFF1	CEI Clause (Read only)
FFF2	CEI Version (Read only)
FFF3	CEI-Protocol Version (Read only)
FFF4 – FFFE	Reserved

FFFF	JTAG IDCODE (Read only)
------	-------------------------

Table 17 Parameter Identifier (PID) field of Parameter Read/Write Packets

The Parameter Value (PVal) field is 64 bits wide. In cases where the parameter to be accessed is narrower, the value is left justified to PVal[63:63-(W-1)], where W is the width of the parameter. It is acceptable for the source to a Write packet to send a PVal with more bits of precision than the sink. In responses to such a Write command, the PVal field of the ACK-Write packet would echo the bits that are successfully captured and the complement of the bits that are not captured. For example, a source can compute equalizer tap settings to 7 bits of precision and sends the result on PVal[63:57]. The sink has only 6 bits of precision and captures PVal[63:58]. The ACK-Write response packet echoes PVal[63:58] and sets PVal[57:0] to the complement of what was in the Write command.

Optionally, the Aggregation Layer may insert a Connection Identification code (CID) into the Supervisory Channel of each CEI Protocol link. The CID may be monitored for correct connection between source and sink ports in systems with a large number of CEI Protocol links.

Req. No.	Description
O7.1	When implemented, the software configurable Connection Identification code should be reported over the Supervisory Channel at the CEI Protocol Transmitter.
O7.2	When implemented, the CEI Protocol Receiver should allow software access to the received Connection Identification code.
O7.3	Implementation of Connection Identification is optional.

Table 18 CEI Protocol Connection Identifier Requirements

The Connection ID bits are allocated to the Aggregation layer and are communicated using the facilities of the Supervisory Channel. CID is an optional feature. When implemented, Connection ID allows each CEI Protocol link within a system to be uniquely identified. The Connection Identifier is a software configurable field for labeling CEI Protocol links in a system. When CID is not enabled, CID bits default to all zeros. The Connection Identifier bits are expected to remain constant for the duration of a connection.

The Modes bits are allocated to the report the status of all three layers in CEI Protocol and are communicated using the facilities of the Supervisory Channel. Mode is an optional feature. When implemented, Mode bits provide overhead bits for each layer in a CEI Protocol link to report status information. The definition of the Mode bits is shown in Table 19. The Mode bits are expected to remain constant for the duration of a connection.

Bit Assignment	Description
MODE[16:9] = APM[7:0]	Application Layer Mode, Reserved
MODE[8:4] = AGM[4:0]	Aggregation Layer Mode, Reserved
MODE[3:0] = FRM[3:0]	Framing Layer Mode, Reserved

Table 19 Mode Code Assignment

8.5 Adaptation Layer Definition

Definition of Adaptation Layer is application specific.

9 Lane Aggregation Application - no transcoding

9.1 Introduction

This clause addresses the application of aggregating multiple lower rate clients onto a single CEI Protocol link without the use of transcoding. Minimal impact on the client signals is desired. Thus, additional skew and latency are minimized. The number of clients that can be multiplexed is limited to powers of two, up to 128. The clients must be frequency locked to each other and have bounded relative wander. Although the Lane Aggregation Application is itself protocol agnostic, devices may be protocol specific and still be compliant to this clause.

Figure 1 shows an example of multiplexing two clients onto a single CEI Protocol link (this figure is reproduced below as Figure 18). Client_A and Client_B must be frequency locked to each other and have bounded wander. It is not necessary to further define the requirements and parameters of the client interfaces as they are already specified in standards relevant to the clients. For example, if the clients are two XAUI links, their characteristics are specified in IEEE 802.3. The CEI Protocol link has a data rate that is twice the client rate, multiplied by 1.03125, to accommodate the CEI Protocol Overhead bits. This factor of 1.03125 is identical to the factor used in 64b/66b based protocols such as 10G Ethernet LAN PHY and OIF SFI-4, Phase 2.

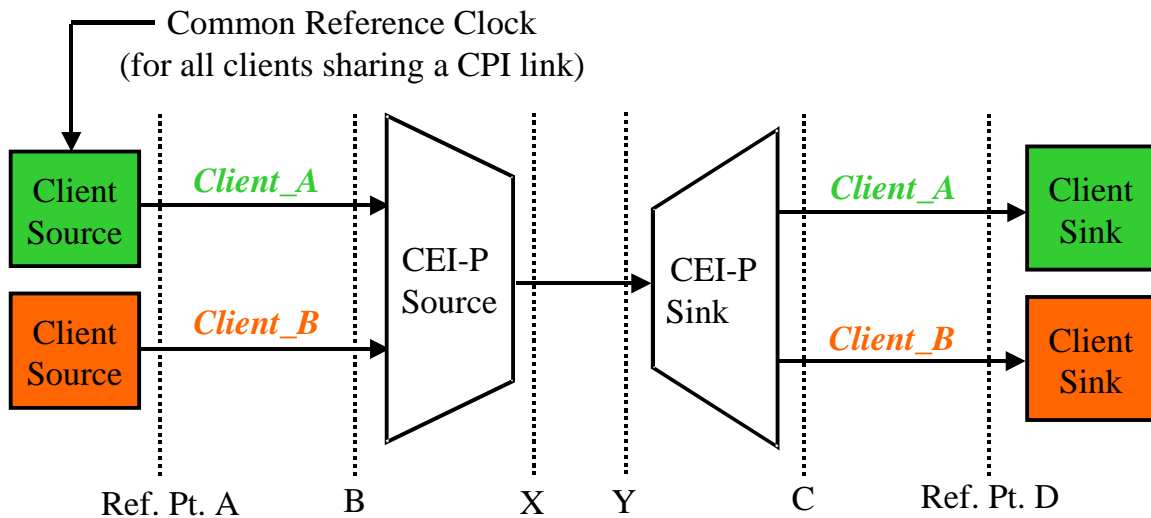


Figure 18 CEI Protocol Lane Aggregation Application

9.2 Adaptation Layer Definition

The Adaptation Layer process performs data recovery on the client signals and maps the bits directly into the Payload region of a CEI Protocol frame. Rate adaptation is unnecessary, as the CEI Protocol link is frequency locked to the clients.

In the transmit direction, performance monitoring of client signals may be performed, but it is not a requirement. Consequential actions due to client signal failures are defined in the relevant standard of the client. They will be different for each client protocol. It is neither necessary nor possible to define them here. In the receive direction, failures on the CEI Protocol link may trigger alarm indications on the client signal. The exact mechanism of the alarm indication is defined in the relevant standard of the client. They will be different for each client protocol. It is neither necessary nor possible to define them here.

The following requirements apply to the Adaptation Layer process.

Req. No.	Description
R9.1	The Adaptation Layer shall conform to the electrical requirements of the standards relevant to the client interface. However, it is acceptable for a device compliant to this Implementation Agreement to place a further requirement that the clients are frequency locked.
R9.2	At a CEI Protocol Transmitter, the Adaptation Layer process shall be able to tolerate relative wander of at least <i>TBD</i> UI between clients.
R9.3	In the transmit direction, consequential actions, due to failures

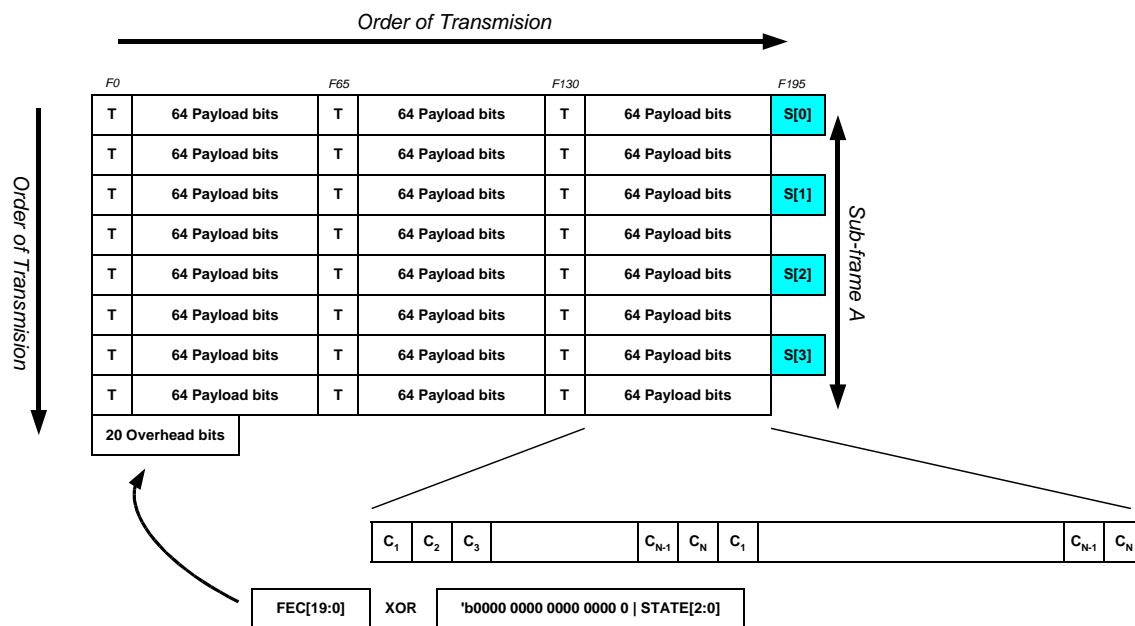
Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

	in the client signal, shall not be defined in this IA.
R9.4	In the receive direction, consequential actions, due to failures in the CEI Protocol link, shall not be defined in this IA.

Table 20 CEI Protocol Adaptation Layer Requirements

9.3 Aggregation Layer Definition

The Aggregation Layer function is responsible for multiplexing of client data onto a CEI Protocol link. Client data is bit-interleaved onto the Payload region, avoiding the T bits, in a round-robin fashion. Data from the first client (C_1) is always placed in the second bit of the CEI Protocol frame. Due to the limitation that the number of clients is a power of two, data from the last client (C_N) naturally lands in the last bit of the last Payload region in each frame. Figure 19 shows the assignment of data from clients C_1 to C_N onto a CEI Protocol frame. Processing of the Connection Identification code Overhead is defined in Clause 4.


Figure 19 Client Assignment to CEI Protocol Frame

Since the CEI Protocol links are independent in the Lane Aggregation applications, there is no requirement for the Aggregation Layer process to de-skew. In addition to those defined in Clause 6, the following requirements apply to the Aggregation Layer process.

Req. No.	Description
R10.1	At a CEI Protocol Transmitter, the Aggregation Layer process shall insert client data into the Payload region of a CEI Protocol frame in the order the data bits are received. Client data shall

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

	be aligned to the CEI Protocol frame and inserted in a bit-interleaved fashion. Data from the first client shall be in the second bit (F1) of the frame. The T bits shall be set to logical zeros.
R10.2	At a CEI Protocol Receiver, the Aggregation Layer process shall retrieve client data bits from the Payload region of a CEI Protocol frame and assign the bits to clients in the order the bits are received. Received data shall be distributed to clients in a bit de-interleaved fashion. Data from the first client shall be taken from the second bit (F1) of the frame. The T bits shall be ignored.

Table 21 CEI Protocol Aggregation Layer Requirements

The Lane Aggregation Source and Sink devices are optionally connected by a bi-directional data link using the Supervisory Interface. The data link uses the bit oriented HDLC protocol of ISO 3309 [2] and is expected to have very low bandwidth requirements. At the CEI Protocol Transmitter, system software constructs an HDLC packet and loads it into the buffer provided. The Transmitter inserts the delimiting flags and the necessary zero stuff bits and transmits the packet serially over Supervisory overhead bit S[0]. The remaining Supervisory overhead bits (S[3:1]) are reserved for future use and are set to logical zeros. The CEI Protocol Receiver monitors S[0] for packet arrivals. The received packet is recovered and the zero stuff bits are removed. The packet is then loaded into a buffer for reading by system software. It is acceptable for the CRC generation and verification to be performed by hardware or by software. As state in [2], the CRC generator polynomial is $X^{16} + X^{12} + X^5 + 1$ and is initialized to all-ones.

Currently, one message type is defined to report the status of the client signals. The format of the client status packet is shown in Figure 20. The TYPE field is set to 'h01 to indicate that the packet is a Lane Aggregation, Client Status message. The status of each client is encoded in a two-bit code (CS[1:0]) and are listed sequentially in the message body of the HDLC packet (See Figure 20). Client #1 is listed first, and Client #N is listed last. The encoding of the client status is shown in Table 22.




Figure 20 Client Status Message Format

CS[1:0]	Description
00	No Alarm
01	Alarm
10	Alarm - Reserved
11	Alarm - Reserved

Table 22 Client Status Codes

Req. No.	Description
O11.1	When implemented, the data link between the CEI Protocol Transmitter and Receiver should be carried over Supervisory Overhead bit S[0] and use bit oriented HDLC protocol. S[3:1] are reserved and should be set to all zeros by the Transmitter. The Receiver should ignore S[3:1]. When the data link is not implemented, S[0] should be set to zero.
O11.2	When Client Status message is implemented, the status of each client should be encoded in the CS[1:0] field according to Table 22.
O11.3	When Client Status message is implemented, the packet should conform to Figure 20. Client status codes (CS[1:0]) are listed sequentially beginning with Client #1 and ending with Client #N.
O11.4	Implementation of Client Status messaging is optional.

Table 23 CEI Protocol Client Status Requirements

9.4 Framing Layer Definition

The CEI Protocol links are frequency locked to the Client links. Thus, in the transmit direction the CEI Protocol link rate is equal to client link rate multiplied by the number of clients and further multiplied by a factor of 1.03125 to account for the Overhead bits. In the receive direction, the client link rate is equal to the CEI Protocol link rate divided by the number of clients and by the overhead factor of 1.03125. It is outside the scope of this Implementation Agreement to define how the client and CEI Protocol link rates are maintained in frequency lock. Examples of acceptable methods include referencing the client and the CEI Protocol to a common reference and using clock recovery techniques on one signal (client or CEI Protocol) to generate reference for the other (CEI Protocol or client).

The Lane Aggregation application may be required to support a multiframe in the Framing layer if required by the intended clients. The length of the multiframe is dependent on the client protocol. It is acceptable for a CEI Protocol Transmitter

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement and receiver to support only the multiframe lengths necessary for their intended clients. A CEI Protocol Receiver is not required to automatically detect multiframe length from the MFAS markers in STATE[2:0].

In addition to those defined in Clause 4, the following requirements apply to the Framing Layer process.

Req. No.	Description
R12.1	Client links shall meet the electrical characteristics as defined in the relevant standards.
R12.2	The CEI Protocol transmit link rate shall be given by : Client_Link_Rate * Number_of_Clients * 1.03125.
R12.3	The Client link rate at a CEI Protocol Receiver shall be given by : CEI_Protocol_Link_Rate / (Number_of_Clients * 1.03125)
CR12.4	A CEI Protocol Transmitter and receiver shall support multiframe in overhead bits STATE[2:0] if required by the client signal.

Table 24 CEI Protocol Framing Layer Requirements

The construction of a CEI Protocol frame is shown logically in Figure 21. Other implementations are possible.

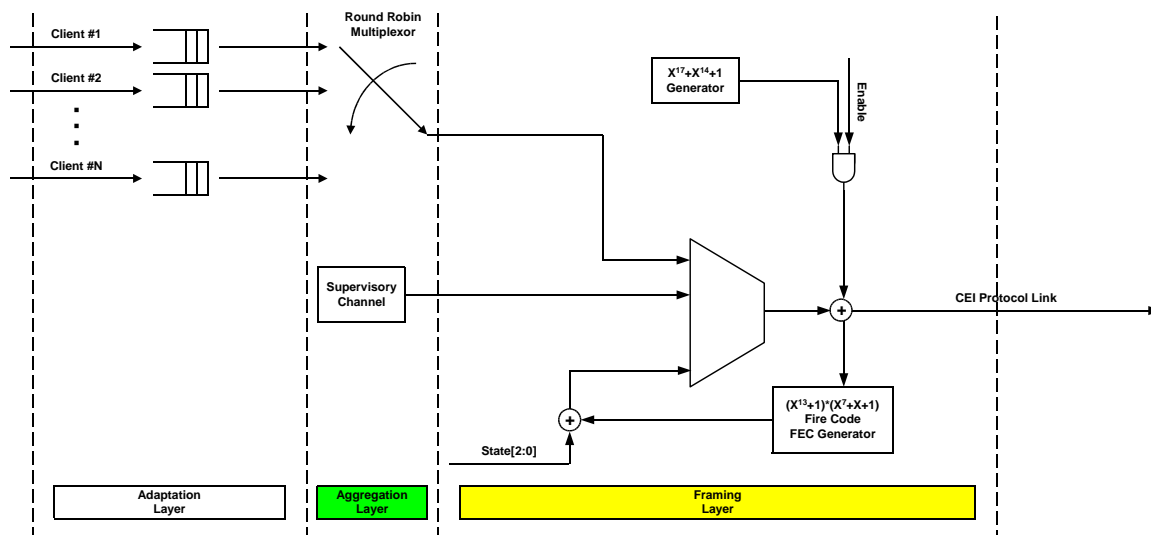


Figure 21 Logical Frame Construction in Lane Aggregation Application

10 Future OIF Interface Application

- 10.1 Introduction
- 10.2 Adaptation Layer Definition
- 10.3 Aggregation Layer Definition
- 10.4 Framing Layer Definition

11 Summary

12 Glossary

13 References

13.1 Normative references

- [1] S. Lin and D. Costello, *Error Control Coding: Fundamentals and Applications*, Prentice Hall, Englewood Cliffs, New Jersey, 1983
- [2] International Organization for Standards, ISO Standard 3309-1993, "Information Technology – Telecommunications and information exchange between systems – High-level data link control (HDLC) procedures – Frame structure", December 1993.
- [3] Optical Internetworking Forum, OIF2003.104.05, Common Electrical I/O (CEI) – Electrical and Jitter Interoperability Agreements for 6+ and 11+ Gbps I/O", February 2004.
- [4] Optical Internetworking Forum, OIF2003.253.01, Common Electrical I/O (CEI) – Electrical and Jitter Interoperability Agreements for 6+ and 11+ Gbps LR I/O", February 2004
- [5] International Telecommunication Union, ITU-T G.7041/Y.1303, "Generic framing procedure (GFP)", December 2001.

13.2 Informative references

14 Annex A: Default Values of Variables

This normative annex lists the default values of variables used in this Implementation Agreement..

Variable	Description	Value
N1	Frame boundary candidate searched	N/A

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

M1	Number of consecutive frames with Fire Code check parity errors to go Out-Of-Frame	15
M2	Number of consecutive frames with good Fire Code check parity to go In-Frame	4
M3	Bad MFAS patterns to go OOM	3
M4	Good MFAS to go IM	1
D1	Time-out for CEI Protocol Receiver Training	TBD
R1	Number of consecutive frame of consistent STATE[2:0] to accept received value	TBD

15 Appendix B: Framing Process in a CEI Protocol Receiver

This informative appendix describes the framing process in a CEI Protocol receiver.

The framing process (Eg., from power-up) involves locating the CEI Protocol frame boundaries and simultaneously synchronizing the stream cipher de-scrambler. A CEI Protocol receiver first searches for the frame boundaries by selecting a potential frame boundary and then computing the corresponding Fire Code Parity bits from the received data. The difference between the computed Fire Code Parity and the Overhead bits in the candidate frame is assumed to be due to scrambling. The scrambler state is initialized by the difference vector and allowed to free-run. The potential frame boundary becomes the candidate frame boundary if the computed Fire Code Parity bits matches the received bits, after de-scrambling, at the next frame. If the parity bits do not match, a new frame boundary is examined. Once a candidate framing boundary is found, the framer examines multiple Fire Codes Parities in a sequence of frames to build confidence in the candidate frame boundary.

Figure 22 below shows a logical model of a CEI Protocol receiver. Actual implementation may differ.

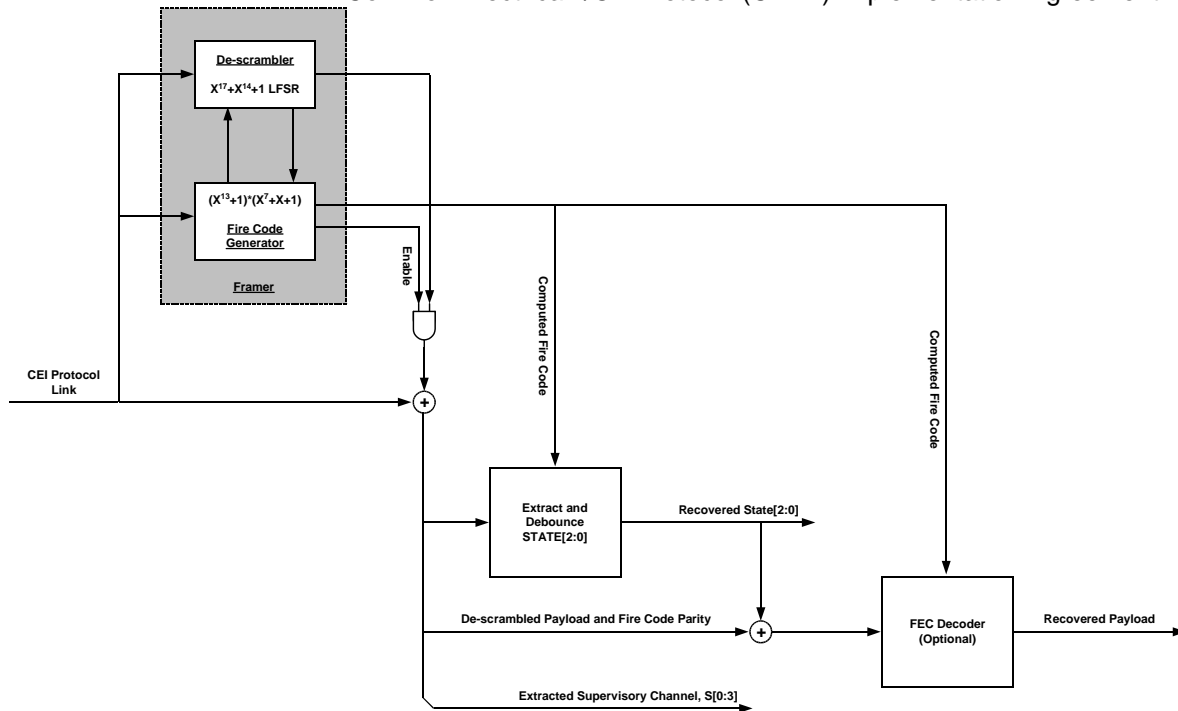


Figure 22 Frame Receiver Process – Logical Block Diagram

16 Appendix C: Sample CEI Protocol Frame

This informative appendix shows a sample CEI Protocol frame.

Figure 23 below shows a sample CEI Protocol Frame. In this frame, all Payload, Transcoding (T), Supervisory (S[0:3]) and State (STATE[2:0]) bits are set to zero prior to scrambling. Contents of the frame after scrambling, is shown below. At the first bit of the frame (F0), the scrambler is assumed to have an initial state of all-ones. The unscrambled Fire Code parity (FEC[19:0]) for this frame is 'h278B4' and the scrambler sequence used to scramble it is 'h430B2'. After scrambling the Fire Code parity is 'h64806'.

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

1	FFFF0003800FC038	1	C1FFE7006F8189C6	1	C7F1FC3F0EE3BCFF	0
0	B60C83344D4B1389	1	8FC7F8FC1F8E71FE	1	FE0907209F1233C1	
1	ADCF61DA3785C5D5	1	EB7A58544D7B1351	1	89DFE3907F89C1E3	1
1	9DF9B91A38E5FFE5	0	00CD03528CE375FC	0	DA1F057397779665	
1	254907809DC239E9	1	F6A621CB97CB94CB	1	32975316E5D0E5AB	1
1	91E98F52FAE296FB	0	2123941F9A71A4FE	0	830C4DB71007C81C	
1	88FAE796ED90C08B	0	C4D07129FC330ED5	1	78785DDD599BCAA2	1
1	3F696211394C7997	1	B5290E30BFDAA087	1	A4B940B9A2BA5BB4	
278B4 XOR 430B2 =						
64806						

Figure 23 Sample CEI Protocol Frame

17 Appendix D: List of companies belonging to OIF when document is approved

ADVA Optical Networking	Fujitsu	PMC Sierra
Aeluros	Furukawa America	Pontusys
Aevix Systems	Gennum Corporation	Princeton Optronics
Agere Systems	Harris Corporation	Quake Technologies
Agilent Technologies	Hi/fn	Quellan
Alcatel	Hitachi	RedC Optical Networks
Altera	IBM Corporation	Sandia National Laboratories
AMCC	IDT	Santur
Analog Devices	Infineon Technologies	SBC
Anritsu	Infinera	Scientific Atlanta
AT&T	Inphi	Siemens
Atrica Inc.	Intel	Silicon Laboratories
Avici Systems	Interoute	Silicon Logic Engineering
Azna	Iolon	ST Microelectronics
Big Bear Networks	JDS Uniphase	StrataLight Communications
Bookham Technology	KDDI R&D Laboratories	Sycamore Networks
Booz-Allen & Hamilton	Kodeos Communications	Syntune
Broadcom	KT Corporation	Tektronix
Cadence Design Systems	Lattice Semiconductor	Telcordia Technologies
Caspian Networks	LSI Logic	Telecom Italia Lab
China Telecom	Lucent	Tellabs

Common Electrical I/O - Protocol (CEI-P) Implementation Agreement

Ciena Corporation	Mahi Networks	Teradyne
Circadian Systems	Marconi Communications	Texas Instruments
Cisco Systems	MCI	T-Networks, Inc.
CoreOptics	MergeOptics GmbH	Toshiba Corporation
Cortina Systems	Mindspeed	TriQuint Semiconductor
Cypress Semiconductor	Mintera	T-Systems/ Deutsche Telekom
Data Connection	MITRE Corporation	Turin Networks
Department of Defense	Mitsubishi Electric Corporation	Tyco Electronics
Diablo Technologies	Molex	Verizon
Elisa Communications	Multiplex	Vitesse Semiconductor
FCI	Mysticom	W.L. Gore & Associates
FiBest Limited	Navtel Communications	Winchester Electronics
Flextronics	NEC	Xignal Technologies
Force 10 Networks	Nortel Networks	Xilinx
Foxconn	NTT Corporation	ZTE Corporation
France Telecom	OpNext	

Notes