



40 Gigabit Ethernet and 100 Gigabit Ethernet Technology Overview

November 2008

Authors:

John D'Ambrosia, Force10 Networks

David Law, 3COM

Mark Nowell, Cisco Systems

1. This work represents the opinions of the authors and does not necessarily represent the views of their affiliated organizations or companies.



Executive Summary

The IEEE P802.3ba 40 Gb/s and 100 Gb/s Ethernet Task Force has been working on the development of 40 Gigabit Ethernet and 100 Gigabit Ethernet. In October, 2008, the Task Force took a major step forward, as it generated Draft 1.0 of the amendment to the IEEE Std 802.3™- 2008 Ethernet standard. The Task Force has developed a single architecture capable of supporting both 40 Gigabit Ethernet and 100 Gigabit Ethernet, while producing physical layer specifications for communication across backplanes, copper cabling, multi-mode gigabit per second, and single-mode gigabit per second. This white paper provides an overview of the 40 Gigabit Ethernet and 100 Gigabit Ethernet project and the underlying technologies.

Introduction

For more than 30 years, Ethernet has evolved to meet the growing demands of packet-switched networks. It has become the unifying technology enabling communications via the Internet and other networks using Internet Protocol (IP). Due to its proven low cost, known reliability, and simplicity, the majority of today's internet traffic starts or ends on an Ethernet connection. This popularity has resulted in a complex eco-system between carrier networks, enterprise networks, and consumers creating a symbiotic relationship between its various parts.

In 2006, the IEEE 802.3 working group formed the Higher Speed Study Group (HSSG) and found that the Ethernet ecosystem needed something faster than 10 Gigabit Ethernet. The growth in bandwidth for network aggregation applications was found to be outpacing the capabilities of networks employing link aggregation with 10 Gigabit Ethernet. As the HSSG studied the issue, it was determined that computing and network aggregation applications were growing at different rates. For the first time in the history of Ethernet, a Higher Speed Study Group determined that two new rates were needed: 40 gigabit per second for server and computing applications and 100 gigabit per second for network aggregation applications.

The IEEE P802.3ba 40 Gb/s and 100 Gb/s Task Force was formed in January 2008 to develop a 40 Gigabit Ethernet and 100 Gigabit Ethernet draft standard. Encompassed in this effort will be the development of physical layer specifications for communication across backplanes, copper cabling, multi-mode fibre, and single-mode fibre. In October, 2008, the Task Force took a major step forward, as it generated Draft 1.0 of the amendment to the IEEE 802.3 specification.

This white paper provides an overview of the IEEE P802.3ba 40 Gb/s and 100 Gb/s Ethernet project and the underlying technologies.



The 40 Gigabit and 100 Gigabit Ethernet Objectives

The objectives of this project are listed below with a summary of the physical layer specifications provided in Table 1.

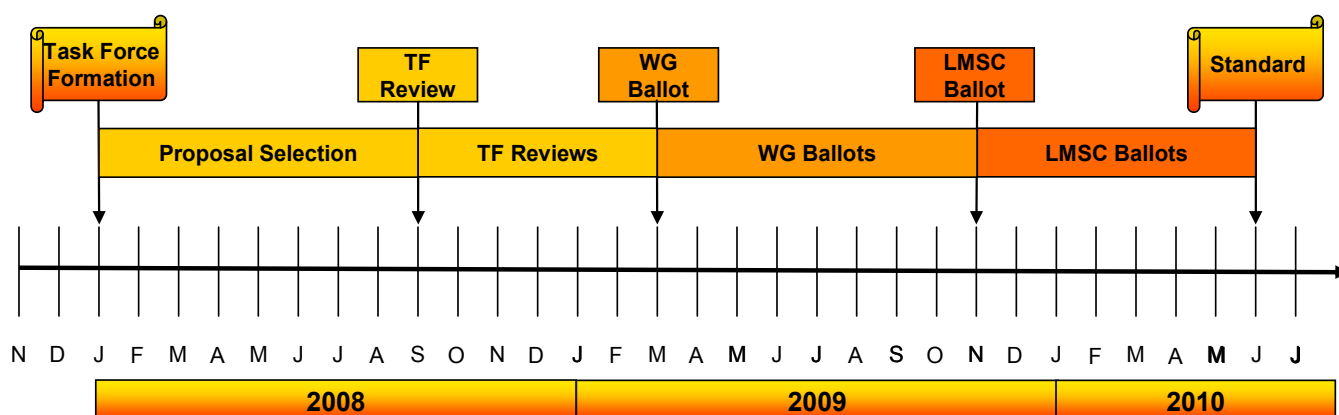
- Support full-duplex operation only
- Preserve the 802.3 / Ethernet frame format utilizing the 802.3 media access controller (MAC)
- Preserve minimum and maximum frame size of current 802.3 standard
- Support a bit error rate (BER) better than or equal to 10^{-12} at the MAC/ physical layer service interface
- Provide appropriate support for optical transport network (OTN)
- Support a MAC data rate of 40 gigabit per second
- Provide physical layer specifications which support 40 gigabit per second operation over:
 - ☐ at least 10km on single mode fiber (SMF)
 - ☐ at least 100m on OM3 multi-mode fiber (MMF)
 - ☐ at least 10m over a copper cable assembly
 - ☐ at least 1m over a backplane
- Support a MAC data rate of 100 gigabit per second
- Provide physical layer specifications which support 100 gigabit per second operation over:
 - ☐ at least 40km on SMF
 - ☐ at least 10km on SMF
 - ☐ at least 100m on OM3 MMF
 - ☐ at least 10m over a copper cable assembly

	40 Gigabit Ethernet	100 Gigabit Ethernet
At least 1m backplane	✓	
At least 10m copper cable	✓	✓
At least 100m OM3 MMF	✓	✓
At least 10km SMF	✓	✓
At least 40km MMF		✓

Table 1 - Summary of Physical Layer Specifications for IEEE P802.3ba

Standards Timetable

Figure 1 shows the timeline that has been adopted by the IEEE P802.3ba 40 Gb/s and 100 Gb/s Task Force. Incubation of the efforts and initial discussions began inside the Ethernet Alliance in early 2006. The Task Force was formed in January 2008 and the standards effort proceeds well on track with ratification targeted for June of 2010.



The 40 Gigabit Ethernet and 100 Gigabit Ethernet Architecture

The IEEE P802.3ba amendment specifies a single architecture, shown in Figure 2, that accommodates 40 Gigabit Ethernet and 100 Gigabit Ethernet and all of the physical layer specifications under development. The MAC layer, which corresponds to Layer 2 of the OSI model, is connected to the media (optical or copper) by an Ethernet PHY device, which corresponds to Layer 1 of the OSI model. The PHY device consists of a physical medium dependent (PMD) sublayer, a physical medium attachment (PMA) sublayer, and a physical coding sublayer (PCS). The backplane and copper cabling PHYs also include an auto-negotiation (AN) sublayer and a forward error correction (FEC) sublayer.

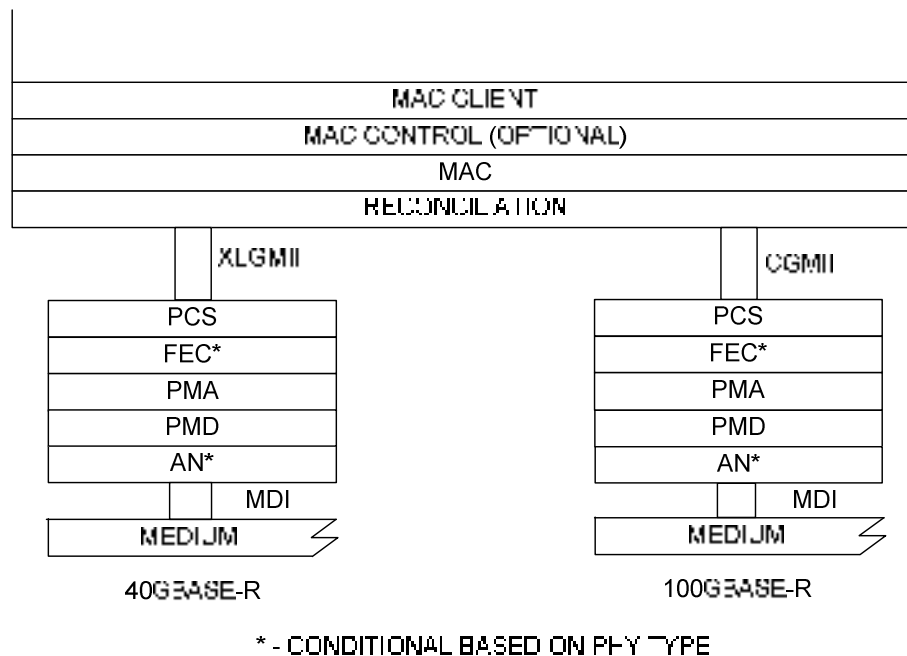


Figure 2 - IEEE P802.3ba Architecture

The Physical Coding Sublayer (PCS)

As shown in Figure 2, the PCS translates between the respective media independent interface (MII) for each rate and the PMA sublayer. The PCS is responsible for the encoding of data bits into code groups for transmission via the PMA and the subsequent decoding of these code groups from the PMA. The Task Force developed a low overhead scheme, referred to as “Multilane Distribution (MLD),” as the basis for the PCS for 40 Gigabit Ethernet and 100 Gigabit Ethernet.

The MLD scheme in the PCS has been designed to support all PHY types for both 40 Gigabit Ethernet and 100 Gigabit Ethernet. It is flexible and scalable, and will be able to support all PHY types currently under development in the IEEE P802.3ba project. Furthermore, the PCS will support future PHY types that may be developed that will be fueled by continuous advances in electrical and optical transmission. The PCS layer also performs the following functions:

- Provide frame delineation
- Transportation of control signals
- Ensure necessary clock transition density as needed by the physical optical and electrical technology
- Stripe and re-assemble the information across multiple lanes

The PCS leverages the 64B/66B coding scheme that was used in 10 Gigabit Ethernet. It provides a number of useful properties including low overhead and sufficient code space to support necessary codewords, which are also consistent with 10 Gigabit Ethernet.

The MLD scheme implemented in the PCS is fundamentally based on a striping of the 66-bit blocks across multiple lanes. The mapping of the lanes to the physical electrical and optical channels that will be used in any implementation is complicated by the fact that the two sets of interfaces are not necessarily coupled. Technology development in either chip interfaces or optical interface is not always tied together, and it was necessary to develop the concept of PCS lanes to allow the decoupling of the evolution of the optical interface widths from the evolution of the electrical interface widths.

The transmit PCS, therefore performs the initial 64B/66B encoding and scrambling on the aggregate channel (40 gigabit per second or 100 gigabit per second) before distributing 66-bit block in a round robin basis across the PCS lanes, as illustrated in Figure 3.

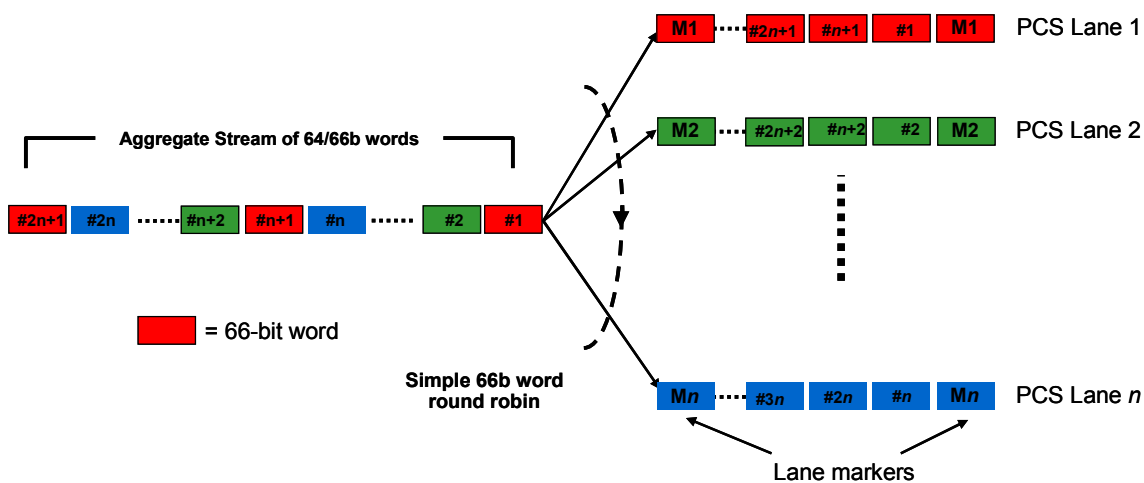


Figure 3 – PCS Lane Distribution Concept

The number of PCS lanes needed is the least common multiple of the expected widths of optical and electrical interfaces. For 100 Gigabit Ethernet 20 PCS lanes has been chosen. The number of electrical or optical interface widths supportable in this architecture is equivalent to the number of factors of the total PCS lanes. Therefore, 20 PCS lanes support interface widths of 1, 2, 4, 5, 10 and 20 channels or wavelengths. For 40 Gigabit Ethernet 4 PCS lanes support interface widths of 1, 2, and 4 channels or wavelengths.

Once the PCS lanes are created they can then be multiplexed into any of the supportable interface widths. Each PCS lane has a unique lane marker, which is periodically inserted. All multiplexing is done at the bit-level. The round-robin bit-level multiplexing can result in multiple PCS lanes being multiplexed into the same physical channel. The unique property of the PCS lanes is that no matter how they are multiplexed together, all bits from the same PCS lane follow the same physical path, regardless of the width of the physical interface. This enables the receiver to be able to correctly re-assemble the aggregate channel by first demultiplexing the bits to re-assemble the PCS lane and then re-align the PCS lanes to compensate for any skew. The unique lane marker also enables the deskew operation in the receiver. Bandwidth for these lane markers is created by periodically deleting interpacket gaps (IPG). These alignment blocks are also shown in Figure 3.



The receiver PCS receives all these multiple PCS lanes, realigns them using the embedded lane markers and then re-order the lanes into their original order to reconstruct the aggregate signal.

Two key advantages of the MLD methodology are that all the encoding, scrambling and deskew functions can all be implemented in a CMOS device, which is expected to reside on the host device, and minimal processing of the data bits other than bit muxing happens in the high speed electronics embedded with an optical module. This will simplify the functionality and ultimately lower the costs of these high-speed optical interfaces.

The Physical Medium Attachment (PMA)

The PMA sublayer interconnects the PCS to the PMD sublayer, and contains the functions for transmission, reception, and (depending on the PHY) collision detection, clock recovery and skew alignment. Within this section, the description of the PMA sublayer will focus on the transmission, reception, and clock recovery aspects of the PMA function. The wide range of supportable interfaces and implementation options requires that, to fully explain the PMA function, it is necessary to explode the PMA function into some PMA sub-layers.

Figure 4 illustrates the general architecture for 100 Gigabit Ethernet, as well as examples of two other architectural implementations:

- 100GBASE-LR4, which is defined as 4 wavelength at 25 gigabit per second per wavelength on SMF
- 100GBASE-SR10, which is defined as 10 wavelengths across 10 parallel fiber paths at 10 gigabit per second on MMF.

These two implementations will be used to illustrate the flexibility needed by the PMA sublayer to support the multiple PMDs being developed for 40 Gigabit Ethernet and 100 Gigabit Ethernet.

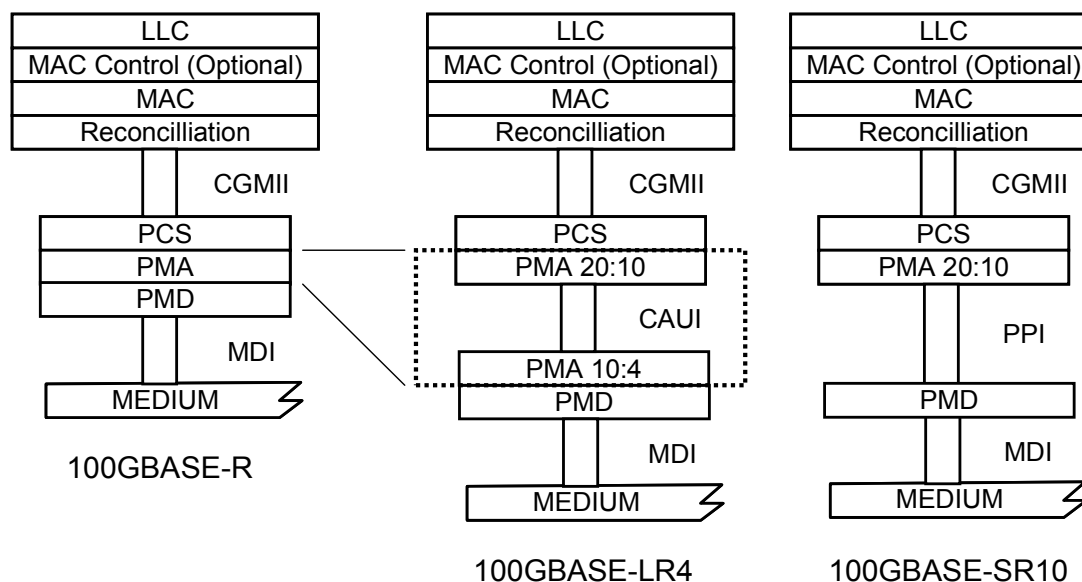


Figure 4 - Illustrations of 100GBASE-R Architectures



As described in the previous section, for 100 Gigabit Ethernet the PCS creates 20 PCS lanes. In the example implementation shown in Figure 4, the PMA functionality is split between two PMA devices that are interconnected via an electrical interface, known as the 100 gigabit per second attachment unit interface (CAUI), which is based on a 10 wide interface at 10 gigabit per second per lane. In this implementation the PMA sublayer at the top of the CAUI multiplexes the 20 PCS lanes into 10 physical lanes. The PMA sublayer at the bottom of the CAUI performs three functions. First, it retimes the incoming electrical signals. After the retiming the electrical lanes are then converted back to 20 PCS lanes, which are then multiplexed into the 4 lanes needed for the 100GBASE-LR PMD.

However the implementation of the 100GBASE-SR10 architecture is different. In this implementation a host chip is directly connected to an optical transceiver that is hooked up to 10 parallel fiber paths in each direction. The PMA sublayer resides in the same device as the PCS sublayer, and multiplexes the 20 PCS lanes into the ten electrical lanes of the Parallel Physical Interface (PPI), which is the non-retimed electrical interface that connects the PMA to the PMD.

In summary, the high level PMA functionality of multiplexing and clock recovery still exists but the actual implementation is dependent on the specific PMD being used.

40 Gigabit Ethernet and 100 Gigabit Ethernet Interfaces

The various chip interfaces in the IEEE P802.3ba amendment are illustrated in Figure 4. The IEEE P802.3ba amendment will specify some interfaces as logical, intra-chip, interfaces, as opposed to a physical, inter-chip, interfaces as they have been in the past. A logical interface specification only specifies the signals and their behavior. A physical interface specification also specifies the electrical and timing parameters of the signals.

The inclusion of logical interfaces supports system on a chip (SoC) implementations where various cores, implementing the different sublayers, are supplied by different vendors. The provision of an open interface specification through the IEEE P802.3ba amendment will help these cores to be integrated into a SoC in the same way that chips from different vendors can be integrated to build a system. While a physical interface specification is sufficient to specify a logical interface, there are cases where the interfaces are unlikely to ever be implemented as a physical interface so the provision of electrical and timing parameters are unnecessary.

There are three chip interfaces defined that have a common architecture for both speeds. The MII is a logical interface that connects the MAC to a PHY and the AUI is a physical interface that extends the connection between the PCS and the PMA. The naming of these interfaces follows the convention found in 10 Gigabit Ethernet, IEEE Std 802.3ae, where the 'X' in XAUI and XGMII represents the Roman numeral 10. Since the Roman numerals for 40 are 'XL' and the Roman numeral for 100 is 'C', the same convention yields XLAUI and XGMII for 40 gigabit per second and CAUI and CGMII for 100 gigabit per second. The final interface is the Parallel Physical Interface (PPI), discussed in further detail below, which is the physical interface for the connection between the PMA and the PMD for 40GBASE-SR4 and 100GBASE-SR10 PMDs.



40 Gigabit Media Independent Interface (XLGMII) and 100 Gigabit Media Independent Interface (CGMII)

The XLGMII, which supports the 40 gigabit per second data rate, and the CGMII, which supports the 100 gigabit per second data rate, are defined as logical interfaces between the MAC and the PCS which share a common interface specification, the only differentiation being the specified clock rate.

The interface provides 64 bit wide transmit and receive data paths. These 64 bit data paths are grouped into 8 lanes of 8 bits, with a control bit associated with each lane indicating if it is data or control information such as delimiters or idle being transferred during that clock cycle. There is a single clock associated with transmit and a single clock associated with the receive path. These clocks operate at one one-sixty-fourth of the data rate resulting in a 625 megahertz clock for 40 gigabit per second operation and a 1.5625 gigahertz clock for 100 gigabit per second operation.

Since this is a wide high speed interface it is not expected to be physically implemented as a physical, inter-chip, interface so is only specified as a logical, intra-chip, interface.

40 Gigabit Attachment Unit Interface (XLAUI) and 100 Gigabit Attachment Unit Interface (CAUI)

The XLAUI, which supports the 40 gigabit per second data rate, and CAUI, which supports the 100 gigabit per second data rate, are low pin count physical interfaces that enables partitioning between the MAC and sublayers associated with the PHY in a similar way to XAUI in 10 Gigabit Ethernet. They are self-clocked, multi-lane, serial links utilizing 64B/66B encoding. Each lane operates at an effective data rate of 10 gigabit per second, which when 64B/66B encoded results in a signaling rate of 10.3125 gigabaud per second.

The lanes utilize low-swing AC-coupled balanced differential signaling to support a distance of approximately 25 cm. In the case of XLAUI, there are four transmit and four receive lanes of 10 gigabit per second, resulting in a total of 8 pairs or 16 signals, in the case of CAUI there are ten transmit lanes and ten receive lanes of 10 gigabit per second resulting in a total of 20 pairs or 40 signals.

These interfaces serve primarily as chip to chip interfaces, for example to partition system design between the largely digital based system chip and more analogue based portions of the PHY chip which are often based on different technology. In addition, while there is no mechanical connector specified for XLAUI and CAUI in the IEEE P802.3ba amendment, these interfaces are also candidate interfaces for pluggable form factor specifications, enabling a single host system to support the various PHY types through pluggable modules. Due to this, these interfaces are being specified based on a channel of approximately 25 cm on FR4 printed circuit boards (PCB) strip line with one connector.

The pluggable form factor specifications themselves are beyond the scope of IEEE 802.3 and are developed by other industry organizations.

Parallel Physical Interface (PPI)

The PPI is a physical interface for short distances between the PMA and PMD sub-layers. It is common to both 40 Gigabit Ethernet and 100 Gigabit Ethernet, with the only differentiation being the number of lanes. The PPI is a self-clocked, multilane, serial links, utilizing 64B/66B encoding. Each lane operates at an effective data rate of 10 gigabit per second, which when 64B/66B encoded results in a signaling rate of 10.3125 gigabaud per second. In the case of the 40 Gigabit Ethernet, there are four transmit and four receive lanes of 10 gigabit per second, in the case of the 100 Gigabit Ethernet there are ten transmit lanes and ten receive lanes of 10 gigabit per second.

Physical Media Dependent (PMD)

Different physical layer specifications for computing and network aggregation applications are being developed. For computing applications, physical layer solutions will cover distances inside the data center for up to 100m for a full range of server form factors including blade, rack, and pedestal configurations. For network aggregation applications, the physical layer solutions include distances and media appropriate for data center networking, as well as service provider inter-connection for intra-office and inter-office applications. A summary of the physical layer specifications being developed for each MAC rate is shown in Table 2.

	40 Gigabit Ethernet	100 Gigabit Ethernet
At least 1m backplane	40GBASE-KR4	
At least 10m copper cable	40GBASE-CR4	100GBASE-CR10
At least 100m OM3 MMF	40GBASE-SR4	100GBASE-SR10
At least 10km SMF	40GBASE-LR4	100GBASE-LR4
At least 40km SMF		100GBASE-ER4

Table 2 - IEEE P802.3ba Physical Layer Specifications



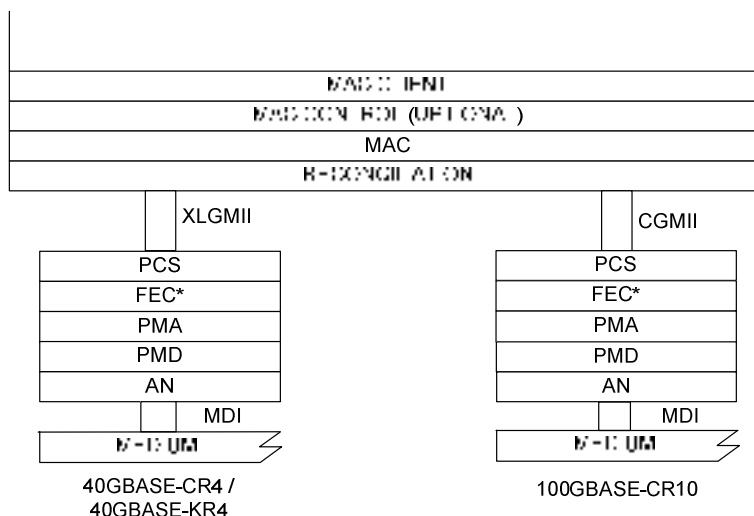
BASE-CR and 40BASE-KR4 Physical Layer Specifications

The 40GBASE-KR4 PMD supports backplane transmission while the 40GBASE-CR4 and 100GBASE-CR10 PMD support transmission across copper cable assemblies. All three of the PHYs leverage the Backplane Ethernet 10GBASE-KR architecture, developed channel requirements and PMD.

The architecture for the PHY types is shown in Figure 5. All three PHYs use the standard 40GBASE-R and 100GBASE-R PCS and PMA sublayers. The BASE-CR and 40GBASE-KR4 PHYs also include an auto-negotiation (AN) sublayer and an optional FEC sublayer.

The BASE-CR and 40GBASE-KR4 specifications also leverage the channel development efforts of the Backplane Ethernet project. The channel specifications for 10GBASE-KR were developed to ensure robust transmission at 10 gigabit per second. The 40 Gigabit Ethernet and 100 Gigabit Ethernet PHYs apply these channel characteristics to 4 lane and 10 lane solutions. The BASE-CR specifications will also leverage the cable assembly specifications developed in support of 10GBASE-CX4. For 40GBASE-CR4, two connectors have been selected: The QSFP connector which will support a module footprint that can support either copper-based or Gigabit per second optic based modules. The 10GBASE-CX4 connector has also been selected, which will enable an upgrade path for those applications that are already invested in 10GBASE-CX4.

The effective data rate per lane is 10 gigabit per second, which when 64B/66B encoded results in a signaling rate of 10.3125 gigabaud per second.. Thus, the 40GBASE-KR4 and 40GASE-CR4 PMDs support transmission of 40 Gigabit Ethernet over 4 differential pair in each direction over either a backplane or twin axial copper cabling medium, while the 100GBASE-CR10 PMD will support the transmission of 100 Gigabit Ethernet over 10 differential pair in each direction for at least 10m over a twin axial copper cable assembly.



* - Optional

Figure 5 – Backplane and Copper Cable Architecture



BASE-SR, BASE-LR, and BASE-ER Physical Layer Specifications

All of the optical PMDs being developed share the common architecture shown in Figure 6. While they share a common architecture, the PMA sublayer plays a key role in transmitting and receiving the number of PCS lanes from the PCS sublayer to the appropriate number of physical lanes that are required per the PMD sublayer and medium.

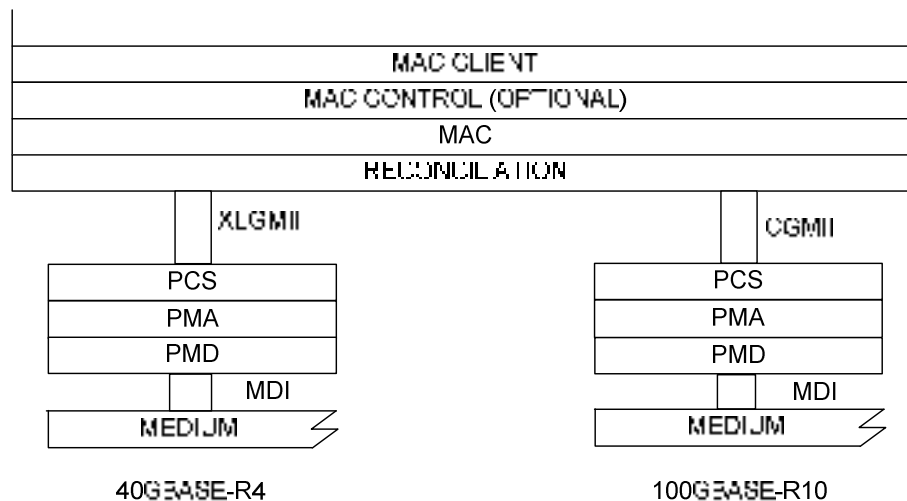


Figure 6 – 40GBASE-R and 100GBASE-SR Architecture

Below is a description of each of the different optical PMD's

- **40GBASE-SR4 and 100GBASE-SR10 PMD** - based on 850nm technology and supports transmission over at least 100m OM3 parallel gigabit per second. The effective data rate per lane is 10 gigabit per second, which when 64B/66B encoded results in a signaling rate of 10.3125 gigabaud per second... Therefore, the 40GBASE-SR4 supports transmission of 40 Gigabit Ethernet over a parallel gigabit per second medium consisting of 4 parallel OM3 fibers in each direction, while the 100GBASE-SR10 PMD will support the transmission of 100 Gigabit Ethernet over a parallel gigabit per second medium consisting of 10 parallel OM3 fibers in each direction.
- **40GBASE-LR4** - based on 1310nm, Coarse Wave Division Multiplexing (CWDM) technology and supports transmission over at least 10km over SMF. The grid is based on the ITU G.694.2 specification, and the wavelengths used are 1270, 1290, 1310, and 1330nm. The effective data rate per lambda is 10 gigabit per second, which when 64B/66B encoded results in a signaling rate of 10.3125 gigabaud per second. which will help provide maximum re-use of existing 10G PMD technology. Therefore, the 40GBASE-LR4 PMD supports transmission of 40 Gigabit Ethernet over 4 wavelengths on each SMF in each direction.



- 100GBASE-LR4 - based on 1310nm, Dense Wave Division Multiplexing (WDM) technology and supports transmission over at least 10km over single mode gigabit per second. The grid is based on the ITU G.694.1 specification, and the wavelengths used are 1295, 1300, 1305, and 1310nm. The effective data rate per lambda is 25 gigabit per second, which when 64B/66B encoded results in a signaling rate of 28.78125 gigabaud per second. Therefore, the 100GBASE-LR4 PMD supports transmission of 100 Gigabit Ethernet over 4 wavelengths on each SMF in each direction.

100GBASE-ER4 - based on 1310nm, WDM technology and supports transmission over at least 40km over single mode gigabit per second. The grid is based on the ITU G.694.1 specification, and the wavelengths used are 1295, 1300, 1305, and 1310nm. The effective data rate per lambda is 25 gigabit per second, which when 64B/66B encoded results in a signaling rate of 28.78125 gigabaud per second. Therefore, the 100GBASE-LR4 PMD supports transmission of 100 Gigabit Ethernet over 4 wavelengths on each SMF in each direction. To achieve the 40km reaches, it is anticipated that implementations will include semiconductor optical amplifier (SOA) technology.

Conclusion

Ethernet has become the unifying technology enabling communications via the Internet and other networks using IP. Its popularity has resulted in a complex ecosystem between carrier networks, data centers, enterprise networks, and consumers with a symbiotic relationship between the various parts.

While symbiotic in nature, the different applications in the Ethernet ecosystem are growing at different rates: server and computing applications are growing at a slower pace than network aggregation applications. This divergence in growth rates spurred the introduction of two higher rates for the next generation of Ethernet: 40 Gigabit Ethernet for server and computing applications and 100 Gigabit Ethernet for network aggregation applications. This will enable Ethernet with its proven low cost, known reliability, and simplicity, to continue to evolve and be the ubiquitous connection for traffic on the Internet. The targeted data for approval of the 40 Gigabit Ethernet and 100 Gigabit Ethernet amendment is June of 2010.

About the Ethernet Alliance

The Ethernet Alliance was formed by companies committed to the continued success and expansion of Ethernet technologies. By providing a cohesive, market-responsive, industry voice, the Ethernet Alliance helps accelerate industry adoption of existing and emerging IEEE 802 Ethernet standards. It serves as an industry resource for end users and focuses on establishing and demonstrating multi-vendor interoperability. As networks and content become further intertwined, the Ethernet Alliance works to foster collaboration between Ethernet and complimentary technologies to provide a totally seamless network environment. To learn more, please go to www.ethernetalliance.org.

Appendix A

The IEEE 802.3 Process

The standardization process in the IEEE 802.3 Ethernet working group is initiated with a call for interest in a proposal. If there is enough interest in the proposal, measured by a majority vote of IEEE 802.3 working group members, an IEEE 802.3 Study Group is formed. From this point on all votes related to technical matters require approval of at least 75% of those voting.

A study group has the task of generating and obtaining approval of the project documentation, this consists of a set of objectives that requires approval by the IEEE 802.3 Ethernet working group, a set of 5 Criteria that requires approval of the IEEE 802 LAN MAN Standards Committee (LMSC) Executive Committee (EC), the parent body of the IEEE 802.3 Working Group, and a Project Authorization Request (PAR) that requires approval of the IEEE Standards Association (IEEE-SA) Standards Board, a parent body of the IEEE 802 LMSC. See Figure 7

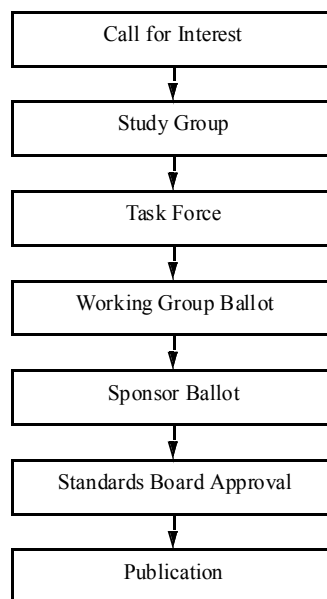


Figure 7 - IEEE 802.3 Standardization Process

The PAR, among other items, defines the scope and purpose of the project. Of the 5 Criteria the most important, and to somewhat interrelated, are

- Broad Market Potential: (broad sets of applications with multiple vendors and users)
- Technical feasibility (can it be built reliably with today's technology)
- Economic feasibility (know reasonable cost for performance).

Once the PAR is approved by the IEEE-SA Standards Board the work of the Study Group is complete and it is disbanded.



An IEEE 802.3 task force is then formed to produce the draft standard, the task force reviews presentations from participants and selects a set of baseline proposals to meet the project objectives. The baseline proposals are then used to produce the first draft. An iterative process of reviewing and refining the draft takes place until the task force believes the draft is complete with no open technical issues.

The draft then enters a two stage ballot process. The first stage is the working group ballot, the second stage is the sponsor ballot. During the working group ballot all IEEE 802.3 working group voters have the right to vote. The sponsor ballot is the 'public' review stage of the draft. Any IEEE-SA member can participate in a sponsor ballot. If someone is not a IEEE-SA member, they can still participate after the payment of a balloting fee although this fee is the same as membership.

In both ballots, the same process is used to ensure that consensus is achieved and that all objections have been considered. An initial ballot takes place where the entire draft is open for comment, approved voters may supply comments, disapproved voters have to provide comments explaining what needs to be changed for them to change their vote. As with all technical votes, the ballot will pass if there is at least 75% approval. A process of comment resolution then takes place where each comment is considered by the task force and if necessary changes are made to the draft. In the case of disapproved comments where the Task Force rejects the comment, a rebuttal to the comment must be written.

If there are changes to the draft, or there are rejected disapprove comments, a recirculation ballot will then take place. During this ballot, changes to the draft, as well as rejected disapproved comments and their rebuttals, are provided to the balloters. Assuming that the initial ballot passed, this is all that can be commented on, or voted disapprove on, during the recirculation ballot. If the initial ballot failed the entire draft remains open for comment.

This process of recirculation and comment resolution continues until the point where there is at least 75% approval, no further changes to the draft are required, and there are no new valid disapprove comments. Once this point is reached during working group ballot the draft moves forward to Sponsor ballot. Since at this point the draft has already gone through detailed technical review during working group ballot, there are often few significant technical changes during Sponsor ballot.

Once sponsor ballot is complete the draft is then forward the IEEE-SA Standards Board Review committee (RevCom) which reviews to make sure that the balloting process was correctly followed and based on that review makes a recommendation to the IEEE-SA Standards Board with respect to approval of the draft. It is the IEEE-SA Standards Board that finally approves the draft as an IEEE Standard.



Glossary

- AN - Auto-negotiation
- CAUI – 100 gigabit per second Attachment Unit Interface
- CGMII – 100 gigabit per second Media Independent Interface
- CWDM – Coarse Wave Division Multiplexing
- FEC - Forward Error Correction
- IEEE 802.3 Standard – the Ethernet Standard
- IEEE P802.3ba – the proposed amendment to the Ethernet Standard for 40Gb/s and 100 Gb/s Ethernet
- IP - Internet Protocol
- MAC – Media Access Control Layer
- MDI - Medium Dependent Interface
- MII - Media Independent Interface
- MLD - Multilane Distribution
- OTN - Optical Transport Network
- PCS - Physical Coding Sublayer
- PHY - Physical Layer Devices Sublayer
- PMA - Physical Medium Attachment Sublayer
- PMD - Physical Medium Dependent Sublayer
- PPI - Parallel Physical Interface
- RS - Reconciliation Sublayer
- WDM – Wave Division Multiplexing
- XLAUI - 40 gigabit per second Attachment Unit Interface
- XLGMII – 40 gigabit per second Media Independent Interface