CIR
ENLIGHTENED ANALYSIS

Next-Generation Ethernet: From 100 Gbps to 400 Gbps and Beyond

A New CIR Report Coming November 2013

PO Box 4353, Charlottesville, VA 22905
www.cir-inc.com
sales@cir-inc.com
Tel: 434-872-9008 Fax: 434-872-9014

Next-Generation Ethernet: From 100 Gbps to 400 Gbps and Beyond

Report Summary

With 40 GigE commonplace and 100 GigE no longer bleeding edge, optical networking is ready to embark on a quest for 400 GigE; perhaps an interim step on the way to Terabit Ethernet.

CIR has been tracking the commercial development of Ethernet since 1991, when the now-ubiquitous GigE technology was the avant-garde. We have subsequently analyzed every generation of Ethernet all the way to 100 Gbps. So we understand user needs, technology requirements and deployment patterns. And we are especially adept at identifying the commercial opportunities flowing from the latest developments in optical Ethernet.

Based on this more than two decades, CIR presents in this report our view on how 400 GigE will be commercialized, how the 400 GigE business case will be made, and where the opportunities are to be found for carriers, equipment makers and component makers.

This report will make compelling reading for optical networking firms at the equipment, subsystem and component levels of the Ethernet value chain. The report contains a quantitative forecast over the next decade for 400 GigE under several different scenarios, along with an assessment of strategies of the firms that CIR believes will be the key players in this space.

Table of Contents

Executive Summary
E.1 The business case for 400 GigE
E.2 400 GigE as enabling technology: End User Opportunities
E.2.1 Data centers and LANs
E.2.2 Carrier networks
E.3 Eight firms to watch in the 400 GigE space
E.3 When will equipment manufacturers deploy 400 GigE ports?
E.4 Optical component and subsystem opportunities from 400 GigE
E.5 Onwards to Terabit Ethernet

Chapter One: Introduction
1.1 Background to this report
1.2 Objective and scope of this report
1.3 Methodology of this report
1.4 Plan of this report
Chapter Two: Evolution of 400 GigE Technology

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

2.1 Lessons learned from the 100 Gbps struggle
2.1 Coordinating groups, standards and MSAs
2.1.1 IEEE 400 Gigabit Ethernet Study Group
2.1.2 CDFP hot-pluggable MSA
2.1.3 Optical Internetworking Forum (OIF)
2.1.4 Impact of 400 GigE on the development of other networking standards
2.2 Interfaces and transceivers
2.2.1 16 x 25 Gbps or 50 Gbps?
2.2.2 Optical interface issues
2.2.3 Power consumption
2.3 Connectors and cable
2.3.1 How far on copper?
2.3.2 Support on MMF and SMF
2.3.3 Parallel single-mode fiber
2.4 Possible signaling schemes
2.4.1 Possible use of novel modulation schemes
2.4.2 Use of forward error correction
2.5 Key points from this chapter

Chapter Three: Drivers and Retarding Factors for 400 GigE Demand
3.2 Positive impact of current data trends on the need for next-generation Ethernet
3.2.1 “Big data”
3.2.2 Cloud computing
3.2.3 Video streaming
3.2.4 Impact of content delivery networks (CDNs) and “content farms”
3.3 Network architecture factors driving the need for 400 GigE
3.4 Factors retarding the need for 400 GigE networks
3.5 Key points from this chapter

Chapter Four: Potential Markets for 400 GigE
4.1 400 GigE in the data center
4.1.1 Growth in data center size, processing power and storage capacity
4.1.2 400 GigE in future LAN switches
4.2 400 Gbps in the metro network
4.2.1 400 GigE for metro aggregation
4.3 400 GigE in the core network
4.3.1 400 GigE in WAN switches and routers
4.3.2 OTN, mesh networks and 400 GigE
4.4 Key points from this chapter

Chapter Five: Business Cases and Scenarios for 400 GigE
5.1 Economics of 400 GigE
5.1.1 Cost objectives for 400 GigE
5.1.2 Component costs
5.2 Time frames for development and deployment of 400 GigE

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

5.2.1 First generation: 16 x 25
5.2.2 Second generation: 8 x 50
5.2.3 Third generation: 4 x 100 Gbps
5.3 Possible ten-year scenarios and forecasts for revenue generation from 400 GigE
5.4 Timeframe for Terabit Ethernet
5.5 Key points from this chapter

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

Chapter One: Introduction

1.1 Background to this Report

With 40 GigE commonplace and 100 GigE no longer bleeding edge, the optical networking community is embarking on a quest for 400 GigE; almost certainly an interim step on the way to a Terabit Ethernet. There is already an IEEE 802.3 400 Gbps Ethernet Study Group and, at the time of this writing, it had met three times to discuss the many factors that will make 400 GigE a challenge to technologists. The current expectations are that formal standards for 400 GigE will emerge by 2017 or 2018, but no one is making promises in this regard.

As was the case with 100 GigE, the efforts towards 400 GigE are being driven by content providers, IXCs and others who say that they could actually use *Terabit* Ethernet. But Terabit Ethernet is something they are not about to get for many years. For some time to come these large users will have to make do with 100 GigE and then 400 GigE.

1.1.1 Big Bandwidth Users are Never Satisfied!

At present, CIR believes that there will not be any major surprises during the evolution of 400 GigE, unless (1) there are internal squabbles at 802.3 (as there was over 100 GigE with whether there should be a 40 GigE standard) or (2) major technological problems.

- With regard to (1), we think (or at least we hope!) 802.3 may have learned its lesson from the 40/100 GigE experience. *We can't be sure of this. IEEE 802.3 deals with leading edge technologies in its standards making and this inherently produces controversy.*

- As far as (2) is concerned, we note that while the electrical interfaces, optical interfaces, signaling, modulation, etc. for 400 GigE are all challenges, they will initially use technology that is off-the-shelf or close to being off-the-shelf. Indeed, one of the objectives of the IEEE's 400 GigE group is to ensure that precisely this will happen. *Thus it seems likely that the first 400 GigE deployments will use a 16 x 25 Gbps format initially making use of 25 Gbps lasers that are already commercialized.*

1.1.2 Demand Patterns for 400 GigE Will be Different than for 100 GigE

The forces behind the apparent need by large users are in some ways the same as the ones that drove the adoption of 40/100 GigE. Indeed, they are the same ones that have been driving the data rates offered by Ethernet since its inception in the late 1970s; that is to say more devices, more users, higher processor speeds and more demanding content.

This makes it sound like that there is "not much to see here" and that 400 GigE will take its place along with other Ethernets in a few years selling into the addressable market that 100 GigE currently holds. *However, CIR believes that demand patterns are now changing*

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

in a way that is expanding the addressable markets for next-generation Ethernet in a way that is qualitative not just quantitative.

Specifically, CIR believes that the trends indicate that medium-sized organizations and nc just the relatively few giant users—such as large governments, Facebook, Netflix, etc.,— are going to be candidates for high-speed Ethernets. To be clear, we are not talking about a radical shift here with small businesses using the higher Ethernet data rates. Nonetheless, there many medium size businesses around (the kind with say 1,000 employees or more), and few really large businesses (the ones who usually push for a next-generation Ethernet). If a relatively small proportion of the medium-sized business start to need the fastest Ethernets, the addressable market will expand significantly. Some evidence that this is likely to happen is highlighted in Exhibit 1-1.

The factors listed in Exhibit 1-1 are interrelated. What we are talking about here is: (1) rapid expansion in bandwidth demand, (2) almost exponential size growth in the size of the average data center, (3) Moore's Law continuing to push up the speed of processors, and (4) new kinds of content that cannot be easily handled under the old paradigms. *To some extent (1) through (4) are old news, but CIR contends that, for (1), (2) and (4) anyway, the expected changes are dramatic enough that they represent qualitative, not just quantitative change.*

Bandwidth demand expanding at double-digit rates: Within the IEEE 802 community, the commitment to move forward with the 400 GigE effort has largely been supported by reference to a study commissioned by IEEE 802.3 and referred to as the "Industry Connections Ethernet Bandwidth Assessment" (BWA).

This study projects the growth of Ethernet traffic in a number of critical applications sectors, between 2010 and 2020. The significance of 2010 is that this is when the primary 40/100 GigE standards were published.

BWA claims that demand for Ethernet bandwidth will grow in double digits in several large, critical sectors of the networking space. These include financial, cable, IP, core networks, server I/O and science.

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

EXHIBIT 1-1: DEMAND-SIDE FACTORS EXPANDING THE MARKET FOR HIGH-SPEED ETHERNET

Factor	Details	Impact on Potential 400 GigE Revenue Generation	Risks and Limitations
Bandwidth demand expanding at double-digit rates	According to an IEEE sector, key Ethernet user groups are expanding at rates from 32 percent to 95 percent CAGR	These data rates are high enough that they could quickly push existing 40/100 GigE users – and even some current 10 GigE users–to buy into 400 GigE once it is available.	There is a history of optimistic bandwidth forecasters of getting their projections wrong
Small data centers aren't so small anymore	Based on IBM data, in 1999 a large data center was considered to be one that was over 5,000 square feet. Today, the average data center can be as large as 200,000 square feet	The implications are that by the time the 400 GigE standards appear there will be a significant number of "medium-sized" data centers that need 400 GigE for some applications. In addition, medium-sized data centers are still likely to attract data center management talent that would prefer an off-the-shelf solution	The shift to cloud computing has both negative and positive impact for 400 GigE. Cloud providers need very large data centers. But cloud users may be able to reduce the size of their data centers or eliminate them completely
Effective processor speeds increasing	Moore's Law coupled with multicore processors and 3D chips are pushing up effective processors speeds. To this must be added the growing predominance of parallelism in large computers, which now have thousands of core processors	This has several potential impacts on the 400 GigE market. The most obvious is the need for high-speed interconnects. However, the high processors speeds at the core suggest that high-speed aggregation will be needed elsewhere in the data center or deeper in the network. In addition, as with the fact that "average" data centers are now a lot larger than they used to be, processor speed increases and distributed architecture means than the need for 400 GigE could be bigger than the equivalent demand for 100 GigE when this came to market	Unclear at the moment how far highly distributed architectures will penetrate into general business computers. More fundamentally, a question mark hangs over the future of processors with regard to speed. Moore's Law may run out of steam in a few nodes and it is not clear what the future of 3D chips or chips based on non-silicon materials will be
Content driving bandwidth demand and faster processors in dramatic ways	It has long been said that video would be the main driver for the deployment of very high-speed networks, but now this actually happening as the result of ubiquitous streaming video and videoconferencing. Other important content trends are the rise of cloud computing and "big data"	Video is now taken for granted in almost every kind of network and is often a big bandwidth hog. Cloud computing and big data are newer phenomena, but both could lead to the creation of large data centers with a need for 400 GigE in strategic deployments. "Big data" in particular, would most probably require highly parallel processing which may have some need for 400 GigE	Video is a bandwidth hog for sure. But, at this stage in the game, good-quality video can be run over relatively low bandwidth networks; 3G cellular, for example. Also it is not clear how much general business applications will be affected by "big data" or consequently how much it will have to adopt parallel processing. And as we have already noted above, cloud computing will have a mixed effect on data center deployment.

Source: Communications Industry Researchers, Inc.

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

As we discuss in some depth later in this report, the implications of the BWA study is that there are some segments of the network where expected traffic growth is so fast, that users may take a leap from using 10 GigE to 400 GigE technology by 2020.

Thus, from a marketing perspective, what we may be seeing in some sector is that the addressable market for 400 GigE would include not just current 100 GigE users—the obvious market for 400 GigE—but also potentially a major slice of today's 10 GigE users.

The average data center is getting larger: The latest generation of Ethernet has usually found its largest market in data centers where it is needed for aggregation or occasionally for interconnect. It is therefore where the first vendors offering 400 GigE will be looking for early revenues and it matters a lot in this context how big the average data center is.

In fact, data centers are getting larger; a lot larger. According to IBM, in 1999 a "large" data center was 5,000 square feet. Ten years later the area of a large data center was 500,000 square feet. IBM's own Range Technology Data Center near Beijing is 624,000 square feet! A "medium-sized" data center at this point in time might mean a data center as large as 200,000 to 300,000 square feet.

If the analysis above is correct there must already be a significant number of medium-sized data centers that need 100-GigE connections for aggregation or (perhaps) storage connectivity of various kinds. This is perhaps different from the situation a decade ago, when the then average data center had no need of the most advanced Ethernet technology.

Based on this analysis, it seems reasonable to suppose that by the time 400 GigE technology is ready for market, there will be a noticeable part of the market that will come from medium-sized data centers.

CIR also believes that as the average data center becomes that much larger there will be more demand for an off-the-shelf solution such as could be provided by 400 GigE, as opposed to some kind of cobbled together customized solution such as a LAG. The point here is that average data centers are run by average data managers who are unable or would not wish to deploy a customized connectivity solution.

Number of CPUs and racks set to increase dramatically: The speeds of individual processors continue to rise in accordance with Moore's Law, and the trend can probably be kept going with the use of multicore chips and perhaps also 3D chips.

In addition, the number of CPUs and number of racks is growing dramatically. In the 1980s, a "supercomputer" was made up of maybe ten core processors in one rack. Today, such a computer would be made up of maybe 10,000 CPUs in 100 racks. According to IBM again, in the 2018 to 2020 period, the company plans Exascale computing systems with tens of millions of processing cores. There are different opinions about when Exascale

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

systems will be with us. A consensus view is probably sometime in the 2020s, but few people would take a large bet on which part of that decade.

This has a number of implications for 400 GigE demand in the future. While, 400 GigE ha not been designed with interconnect applications primarily in mind, the increase in processor speeds combined with more parallelism could expand the needs for very high-speed interconnection by orders of magnitude.

Some of this interconnection would presumably be supported eventually at the 400 GigE rate. Because there is famously an interconnect bottleneck emerging as the result of ultra-fast processing, 400 GigE will clearly be of growing importance in the context of major data centers. But in such data centers, 400 GigE might have a broad range of aggregation uses.

Once again, CIR also believes there are some hints here that the demand for 400 GigE will be pushed beyond very specialist users. Today, the organizations that are planning to deploy data centers with highly parallel architectures use them for applications such as modeling of biological cell simulations, for example, or certain kinds of fluid dynamics.

This is not the kind of application that will ever reach the mainstream, for reasons that hardly need explaining. CIR believes, however, that content trends (see below) will increasingly make large-scale parallel computing much more common in general business computing.

Content changes driving it all: The technical changes listed above will be the proximate cause of 400 GigE's future market growth. However, it is important to note that none of these changes will occur unless driven by other important changes in the content environment. CIR believes that this is exactly what will occur.

At this point it is usual, in this kind of analysis, to point to the role that video has had in driving bandwidth. This impact is certainly undeniable; streaming video and videoconferencing are taking up a lot of bandwidth these days. But after many years in the making, the video "revolution" is no longer really anything new.

Perhaps the most important factor that CIR sees making a difference at the content level in the future is the arrival on the scene of "big data." The key point about big data for the purposes of the areas that we are concerned with here is that when very large data sets are processed, without being broken up into smaller sets, they can yield insights that would not be possible if the data sets were broken up, as they usually are. For this to happen, large-scale parallel processing—the type that would need 400 GigE—must be in place. And since the impact of big data will be felt in most larger businesses, big data is positive for 400 GigE.

Another important factor in understanding the potential market for 400 GigE is the network architectures that are appearing to support new developments in content delivery. There are some open issues here. If content delivery is very decentralized, there may be many

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

content delivery centers that do not need exceptional bandwidth. If content delivery is centralized, content delivery centers could become prime targets for 400 GigE.

There is also the question of cloud computing, which CIR sees as promoting big dat centers (with a possible taste for 400 GigE) on the one hand, but at the same time allowing end users to get away with smaller data centers; a negative for 400 GigE. In addition, it is possible to get carried away with the importance of the trends listed above, a tendency that could cause manufacturers of equipment, subsystems and components to overshoot market demand.

1.1.3 Technology Opportunities in 400 GigE

These demand drivers are—or at least will be—producing new technology opportunities at the component and subsystems level. As is always the case, the first opportunities will accrue primarily to components and subsystems/modules makers.

At the present time, it seems that these opportunities will initially emerge in three different—but related areas—media, "formatting" and integration. This is not to say that there aren't other technology areas where 400 GigE needs to be developed and which may also lead to opportunities for component makers and others. These include: the implementation of a successful forward error coding (FEC) strategy, improved energy efficiency, low latency and compatibility with OTN.

Media: One of the biggest challenges for the 400 GigE standards effort will be the definition of reach. Put simply, it will be harder than ever to define a role for copper at anything other than for short interconnect ranges. CIR has always been impressed with how well copper has been able to survive as Ethernet data rates have grown. It would once have seemed inconceivable that 100 GigE could have been run over copper at all.

CIR believes that some kind of copper standard will be defined for 400 GigE, but does not expect such standards will be widely used. The twin-ax cable that would be used for such standards is expensive; no more twisted pair. Also, the use of copper in data centers is being thwarted by power consumption issues. We would expect that copper will be used over no more than 1-2 meters in a 400 GigE environment.

The opportunities for 400 GigE will be fiber/optical opportunities. The arrival of 400 GigE will spur demand for lasers operating at 25 Gbps and above. In addition, single-mode fiber will take on a new importance in the data center and may be required at 300 meters, as opposed to 500 meters in 100 GigE. This means that the SMF-in-the-data center market could swell considerably in that a lot more connectivity in the data center occurs at 300 meters than at 500 meters.

Formatting: If the past is anything to go by, the most contentious part of the standards setting process will have to do with formatting. Nothing is settled at the present time, but the most likely evolution has the initial 400 GigE implementation operating in a 16 x 25 Gbps format and then moving to and 8 x 50 Gbps standard and then on to 4 x 100 Gbps.

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

The 400 Gbps format is not likely to appear until 2018 or later. However, the timeframe for this will shape—and be shaped by—the emergence of active components that can serv- in the lower-channel 400 Gbps formats. So, for example, an early shift to 8 x 50 Gbp implies that 50-Gbps VCSELs would be available at a reasonable price.

The opportunity here occurs not only at the component level, but at the level of the module, since new MSAs must be defined. In addition, signaling and modulation software/coding decisions must be made that can also shift the balance of power in the optical networking industry, depending on the R&D directions that individual technology providers have taken.

Integration: Integration at the component level is another important consideration in assessing the opportunities stemming from 400 GigE. While no one is expecting the modules/interfaces for 400 GigE to be especially inexpensive, cost is obviously important in any LAN technology.

The necessary cost improvements can be achieved in a number of ways, but increased use of integration is certainly going to be one of them. This may turn out to be one of the most rewarding opportunities going forward in the sense that optical integration if used creatively might result in 400 GigE innovative products that can quickly take large market shares in the 400 GigE space as it emerges.

We also think that as 400 GigE emerges, it could prove to be a testing ground between the three current paradigms in optical integration, which are silicon photonics, GaAs and InP.

1.2 Objective and Scope of this Report

As we have already stated, the evolution of 400 GigE is at a very early stage. The IEEE group that is deciding on the appropriate standards for 400 Gbps met for the first time six months previous to these words being written. So the discussion of the commercial potential for 400 GigE equipment, subsystems and components *may* seem to be somewhat premature.

However, we do not believe that this is the case. CIR believes that there is little point in the IEEE pursuing 400 GigE unless at the end of the day there is a market for this kind of technology; and a market that goes well beyond the needs of a few ultra-large users.

The primary goal of this report is therefore to assess whether there is really going to be a market for 400 GigE and to provide some guidance on just how big this market may ultimately turn out to be. We see this report as an independent check on where the commercial opportunities, business models and revenue streams will evolve in the new 400 GigE space that will emerge.

More specifically, the goals of this report are to identify (1) the technological options available for 400 Gbps, especially with regard to optical and electrical interfaces and

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

signaling/modulation schemes, (2) to better understand what segments of the networking market are likely to buy into 400 GigE and when.

Given the above, the scope of this report includes an analysis of future demand for 40 GigE from the various locations in the network where 400 GigE may be of some use. These are quite diverse and include interconnect, data center aggregation, content distribution centers, and even central offices.

The other aspect to the scope of this report is that our coverage attempts to provide an opportunity assessment at all levels of the value chain. However, given the current very early state of 400 GigE development there is an inevitable strong focus on components and modules.

From a geographical perspective this report is worldwide in scope. A companion volume covering optical interconnection at the rack- and board-to-board level is also available from CIR.

1.3 Methodology and Information Sources for this Report

CIR has been tracking the commercial development of Ethernet since 1991, when the now-ubiquitous GigE technology was the avant-garde. We have subsequently analyzed every generation of Ethernet all the way to 100 Gbps. So we understand user needs, technology requirements and deployment patterns. CIR published a similar report to this one on the development of 100 GigE at exactly the same stage of its evolution.

The information for this report has come from a number of different sources. CIR analysts regularly talk with influencers in the data communications and optical networking space and the views collected in this way are used to create realistic forecasting models and market assessments such as those in this report.

We have also conducted extensive secondary research for this report. This included an in-depth review of the papers presented at IEEE meetings related to 400 Gbps. In addition, we also surveyed relevant articles, blogs, financial reports, conference proceedings and corporate Web sites, etc., as well as previous CIR reports.

1.3.1 Forecasting Methodology

The details of our forecasting methodology are provided in the main body of this report. However, generally our approach is to examine the underlying addressable market for each of systems we consider and then consider reasonable deployment patterns. By bringing a pricing time series to bear on all of this, it is possible to create a market forecast.

In undertaking this analysis we have also attempted to place some size and growth projections on the 400 GigE market going forward, understanding that for quite a few more years the revenues will effectively be zero, and for a couple of years after that they will reflect sampling only. Nonetheless, CIR believes that an early consideration of revenue

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com

potential still makes sense, given that presumably the primary objective of 400 GigE is to make money for all of various players in the value chain.

CIR has been forecasting developments in the optical telecom and data communication business since 1985. As such, we have developed a sophisticated approach to forecasting that takes into consideration likely demand factors including realistic timetables for deployment.

The forecasting approach we bring to this report is similar to that we brought to CIR's first market analysis of 100 GigE.

1.4 Plan of this Report

This report consists of five chapters plus an Executive Summary. The Executive Summary is designed to profile the key opportunities and the timetable for their appearance. Our opportunity profiles are derived from the analysis in the rest of the report.

Chapter Two reviews the technology and standards evolution relevant to 400 Gbps. As noted, it reviews the state of the art in relevant component and module technology as well as the directions for 400 GigE that have been revealed so far in the few IEEE meetings that have been held.

In Chapter Three, we examine the factors that will shape the demand for 400 GigE in the future, delving deeper into the issues that we have already explored to some extent in this chapter. In Chapter Four, we provide an analysis of what the demand patterns might look like for the main segments of the 400 GigE market, if and when such a market appears. In carrying out this analysis we have tried to cover all segments of the network which will feel the impact of 400 GigE.

Finally in Chapter Five, we provide some guidance on the business cases that can reasonably be made for 400 GigE, along with some early forecasts of possible revenue generation from this technology.

PO Box 4353, Charlottesville, VA 22905
Tel: 434-872-9008 Fax: 434-872-9014
www.cir-inc.com sales@cir-inc.com