



400G, OTN and Next-Generation Transport: A Market and Technology Forecast

January 2014

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

Table of Contents

Executive Summary

- E.1 Three Factors Driving 400G Backbones: Bandwidth Hunger, Applications and the 100G Paradigm
 - E.1.1 Mobile Broadband: A New Paradigm for Network Traffic
- E.2 Opportunities Created in the Transport Layer: The Need for Intelligence at Every Level
 - E.2.1 Role of Software Defined Networks: A Key Enabling Technology?
 - E.2.2 Danger of Market Overshoot
- E.3 Component Level Opportunities
 - E.3.1 The Growing Power of DSP and Waveform Engineering
 - E.3.2 Proprietary Network Processors in a 400G Environment
 - E.3.3 Integration and Opportunities for Optical Components Firms at 400G
- E.4 The Optical Networking Equipment Perspective on 400G
 - E.4.1 Four Factors That Will Shape the 400G Optical Platform Market
 - E.4.2 Five Firms that will Shape the 400G Transport Space
- E.5 Summary of Eight-Year Forecasts of 400G Transport Markets

Chapter One: Introduction

- 1.1 Background to this Report: Outstanding Questions About 400G Transport Deployment
 - 1.1.1 Trials and Addressable Markets for 400G Transport Aren't the Same Thing
 - 1.1.2 Service and Standards Environments Add More Uncertainties for 400G Transport
 - 1.1.3 Equipment and Components Firms: Together Again at 400G?
- 1.2 Objective and Scope of this Report
- 1.3 Methodology of this Report
- 1.4 Plan of this Report

Chapter Two: Analysis of Enabling Technologies and Standards for 400G Transport

- 2.1 Multiplexing, Switching and Routing Options in the 400G Transport Environment
 - 2.1.1 OTN Boxes, Routers and Optical Packet Switching in 400G Network
 - 2.1.2 Differing Visions of 400G: Sometimes 400G is 200G and Sometimes It is 500G
 - 2.1.3 WDM, Superchannels and Bandwidth Allocation
 - 2.1.4 A Role for Space Division at 400 Gbps?
- 2.2 New Directions for Modulation in the 400G Environment
 - 2.2.1 DPSK, QPSK, BPSK and PM-QSP
 - 2.2.2 DP-16 QAM and Beyond
- 2.3 Raman Technology in 400G Backbones
- 2.4 400 Gbps Coherent Technology: the Once and Future
- 2.5 Standards
 - 2.5.1 ITU-T, OTN at 400 Gbps
 - 2.5.2 Optical Internetworking Forum (OIF)
 - 2.5.3 IETF

2.5.4 Relevance of the Emerging 400 Gigabit Ethernet Standard for Carriers

Chapter Three: Equipment Vendor Strategies for 400G

- 3.1 Alcatel-Lucent (France/United States)
 - 3.1.1 Photonic Service Switch (PSS)
 - 3.1.2 Photonic Service Engine (PSE)
 - 3.1.3 FP3 400G Chip
 - 3.1.4 Wavelength Tracker
 - 3.1.5 Nextgen (Australia)
 - 3.1.6 Orange-France Telecom/RENATER (France)
 - 3.1.7 SaskTel (Canada)
 - 3.1.8 Shaw Communications (Canada)
 - 3.1.9 Telefónica España (Spain)
 - 3.1.10 Zain (Saudi Arabia)
- 3.2 Ciena (United States)
 - 3.2.1 WaveLogic 3
 - 3.2.2 Ciena 6500
 - 3.2.3 BT (United Kingdom)
 - 3.2.4 Sprint (United States)
 - 3.2.5 Comcast (United States)
 - 3.2.6 Verizon (United States)
- 3.3 Cisco (United States)
 - 3.3.1 NCS Platform
 - 3.3.2 nPower X1 Network Processor
- 3.4 Coriant (Germany)
 - 3.4.1 FlexiGrid ROADM nodes
 - 3.4.2 hiT 7100 and 7300
 - 3.4.3 Netia (Poland)
 - 3.4.4 Telekom Austria – A1 (Austria)
- 3.5 Cyan (United States)
 - 3.5.1 GlobalConnect (Denmark)
- 3.6 Ericsson (Sweden)
 - 3.6.1 MHL 3000 Platform
 - 3.6.2 SPO 1410 Platform
 - 3.6.3 Smart Service Routers 8000
 - 3.6.4 Telefónica España (Spain)
- 3.7 Fujitsu and NEC (Japan)
 - 3.7.1 NTT (Japan)
 - 3.7.2 Verizon (United States)
- 3.8 Huawei (China)
 - 3.8.1 NE5000E Router
 - 3.8.2 WDM Transport Solutions
 - 3.8.3 EXATEL (Poland)

- 3.8.4 Jazztel (Spain)
- 3.8.5 KPN International (Netherlands)
- 3.8.6 Mobily (Saudi Arabia)
- 3.8.7 True (Thailand)
- 3.8.8 Telefonica Chile (Chile)
- 3.8.9 400G Photonics Integration Program
- 3.9 Infinera (United States)
 - 3.9.1 DTN-X
 - 3.9.2 DANTE (Europe)
 - 3.9.3 TeliaSonera (Sweden/Finland)
 - 3.9.4 Zayo (International)
- 3.10 Juniper (United States)
- 3.11 TE SubCom (United States)
- 3.12 Xtera (United States)
- 3.13 ZTE (China)
 - 3.13.1 Recent 400G Trial
 - 3.13.1 Deutsche Telekom (Germany)

Chapter Four: Eight-Year Forecasts

- 4.1 Forecasting Methodology
 - 4.1.1 Assumptions About Network Requirements and Sources of Information
 - 4.1.2 Getting to Market in a Hurry: 400G Backbones without Standards
 - 4.1.3 Assumptions about Technology Evolution
 - 4.2 Eight-Year Forecast of 400G Transport Equipment: Core Analysis
 - 4.3 Eight-Year Forecast of 400G Equipment: Breakout by Type of Network
 - 4.4 Eight-Year Forecast of Optical Components for 400G Deployment
 - 4.5 Possible 400 Gbps Service Offerings
- Acronyms and Abbreviations Used In this Report
- About the Author

List of Exhibits

- Exhibit E-1: Eight-Year Forecast of 400G Transport Equipment: Core Analysis
- Exhibit 2-1: Goals of the 400 Gbps Ethernet Study Group
- Exhibit 4-1: Eight-Year Forecast of 400G Transport Equipment: Core Analysis
- Exhibit 4-2 Eight-Year Forecast of 400G Equipment: Breakout by Type of Network
- Exhibit 4-3: Eight-Year Forecast of Optical Components for 400G Deployment (\$ Millions)

Chapter One: Introduction

1.1 Background to this Report: Outstanding Questions About 400G Transport Deployment

Carriers worldwide are beginning to outgrow 100G transport networks and are taking the first tentative steps towards 400G backbones. CIR believes these ultra-fast networks will create major new revenue opportunities for optical components firms, silicon chipmakers and equipment companies alike. This new business will occur as 400G transport moves out of today's trial phase and 400G production networks are deployed. This shift is unlikely to happen in a big way for a few more years, but the foundations for success in the 400G space are being established now.

The challenges of 400G transport are made a little less daunting by the fact that 400G is as "second time around." That is to say, 400G is—most probably, anyway—the next phase in the 100G *core* era. Until 100G transport backbones began to appear a few years back, we had seen well over a decade in which transport networks were built around multiple 10G pipes. The technology that needed to be invented for a shift to a transport network infrastructure operating with 100G pipes can be reused to some extent at 400G.

This is encouraging, but any advantages that 400G transport may inherit from recent 100G efforts don't seem to us to help all that much when it comes to addressing the huge uncertainties that remain in 400G sales/marketing and product design strategies. These uncertainties are of three varieties.

The uncertainties associated with understanding how much the current enthusiasm over 400G transport trials can translate into actual deployment is one variety. Another is the relative absence of agreed standards for 400G public network transport. Yet another—and, as it turns out the most worrisome—is how 400G transport platforms can respond to a rapidly evolving service environment.

1.1.1 Trials and Addressable Markets for 400G Transport Aren't the Same Thing

It is easy to be impressed by the accelerating number of 400G transport trial announcements that get headline treatment in the trade press. What is especially notable about these trials is which service providers *aren't* involved. They are not just the usual suspects; the great ex-PTTs (such as France Telecom or BT) or AT&T or university research networks, although some of these are involved.

Instead, many of the 400G transport trials are at "Tier Two" carriers, often in smaller markets; Australia, Austria, Canada, Poland and the like.

CIR's long experience with such matters suggests that trials of this kind don't happen without key managers at a service provider being

enthusiastic supporters of a particular technology. So we think that these trials are proof that there are already plenty of 400G transport advocates at the service providers who genuinely want to see their networks upgraded to 400G in the next few years.

This is a good thing from the perspective of serious revenue generation from 400G transport in the near future. But to what degree can one extrapolate from the 400G transport trials to existence—and especially the timing—of actual addressable markets?

Page | 2

The motivations and financial arrangements behind trials are always murky. And the trial announcements are typically made with strong involvement of the 400G transport equipment providers; firms with something to sell. So other questions must be asked. How many carriers are really going to expensively upgrade their networks to 400G in the next, say, five years? How many may be able to make do with 100G or 200G?

And when the time for 400G deployments finally rolls around, will the initial substantial expenditures on them come mostly from long-haul and submarine routes; which would be the natural expectation for those of us who have watched the networks evolve over the past 20 years? Or will it be the metro networks that are the first to take up 400G transport?

1.1.2 The Service and Standards Environments Add More Uncertainties for 400G Transport

There is no surefire way to answer these questions—although we attempt to do precisely this in this report! In particular, the issue of which part of the network will first see 400G prove most strategically important is especially hard to answer, because the answer seems to rely on an applications layer environment that is—in CIR's opinion—entering a period dramatic change.

While established telecom hardware and service firms have all been through major transport network upgrades before and have an intuitive grasp of the risks and how to ameliorate them the new applications that 400G transport pipes will have to amalgamate are different in both their bandwidth and performance requirements, placing new demands on equipment and components suppliers and creating new opportunities, too.

The point here is that in the past, the traffic patterns that new transport platforms have had to cope with once deployed, were fairly well understood. This is simply no longer the case and for three main (interrelated) reasons.

The Internet-of-Things—adding unpredictability to the backbone: First, while the network has always had to cope with some degree of distributed intelligence, the level of distribution is about to shift upwards because of inter-related trends that are variously known as the Internet-of-Things (IoT), Machine-to-Machine (M2M), and by a few other names.

The essence of these concepts is that networked intelligence will be everywhere, especially in the form of sensors and actuators of

various types. The point here is that the number of nodes in the network will increase by a couple of orders of magnitude, and this cannot help but make a difference at the transport level.

No one can be entirely sure right now what exactly this difference will be. But it is clear that traffic on the backbone will be more unpredictable because each network node will in a sense have a "mind of its own."

Ubiquitous broadband and "big data"—adding bandwidth requirements to the backbone: Much of the traffic from IoT devices would definitely be characterized as narrowband—sensor output and control data is typically not all that bandwidth hungry. However, in aggregate, IoT traffic constitutes what is becoming known as "big data," which may indeed raise bandwidth issues at the backbone level.

All of this suggests a future in which backbones built between bandwidth-stressed points on the network find that bandwidth requirements among these points and even which points are stressed are constantly changing. We note in this context that this isn't entirely a new phenomenon, since broadband data traffic—increasingly represented by millions of video content creators—is already the order of the day.

New content providers take control away from the service providers: All of this means that the service providers are losing control of their networks. Even though the service providers never traditionally had much say over content creation, bandwidth-hungry applications in the past have been few or far between. And while the ability to create low-bandwidth voice traffic was always in the hands of millions of individuals, the service providers were in the position to tightly control how that traffic travelled across public networks.

The new world of IoT and highly distributed content providers threatens to disrupt future backbone facilities, and transport equipment firms—including those that provide 400G transport equipment—must respond to this, as must those companies that create enabling technology in the form of optical components and specialized chips. None of these firms are helped by the fact that at the present time the standards development work for 400G transport—at the OIF and ITU-T—is just getting underway.

1.1.3 Equipment and Components Firms: Together Again at 400G?

All this obviously creates considerable uncertainty—and hence business risk—for equipment makers and components firms that are becoming involved in the ramp up to 400G transport. They are in effect being asked to compete in a market where both the addressable market and the requirements of that market are extremely hard to calculate.

It is clearly hard for these firms to calculate the ROIs on their 400G transport network development projects in these circumstances. At the very least, the risk discounts that are injected into such calculations need to be considerable. In this difficult business

environment, a number of strategies suggest themselves.

400G platforms—the old versus the new: Any transport network equipment or component firm that chooses not to have anything to do with 400G is soon going to find itself in the ash heap of telecom history. Service providers aren't going to want to buy from equipment makers that do not have a good 400G story to tell.

At this stage in the game this story may not need to be complete. No telecom equipment firm yet has a platform that is specifically designed with 400G in mind. Most of the more advanced 400G transport platforms in the current generation of trials use technology primarily designed with 100G transport in mind.

By using earlier generations of telecom hardware to carry out 400G trials, equipment manufacturers can reduce their risk in this space significantly. Some of the trials are using platforms that began their life in the 10G era but have been significantly upgraded for these trials.

Providing for the new service environment—a source of competitive advantage: While 400G trials and upgrades are inherently about bandwidth, CIR believes that an important source of competitive advantage in this space for the equipment providers will be the ability of 400G transport platforms that are up to the task of supporting the very demanding service environment that we see emerging over the next decade.

As we discuss later in this report, each equipment vendor is going to have its own idea of how that can best be achieved, and we note that several—arguably most—leading telecom equipment vendors have not really shown their hands in this regard, although some have.

CIR believes that there is an interesting debate that is about to emerge and that will take place in various venues including standards groups about what is really required here. But at one level the answer appears to be as much as possible. This is to say that the 400G platforms that emerge over the next few years will have to provide the maximum amount of flexibility or—putting it another way—make the fewest assumptions about what traffic looks like.

In practice this means that 400G platforms must be able to support any service that the network can throw at it on any wavelength and in any direction. Given that services and applications will arrive from millions of sources over which the service provider has very little control, the traffic/wavelength management and performance management must be of a very high level and must be ubiquitous.

We have a lot more to say about this in the main body of this report. However, we note that what is being asked of 400G platforms is something very close to the agile network concept that was rather popular among telecom analysts in the 1980s, a time in which

the idea of a ubiquitously switched optical network was being widely discussed. It is interesting that Alcatel-Lucent has recently reintroduced that term in the context of 100-400G networking.

Optical components and silicon at 400G—the old made new again: Although no one really talks about all-optical switching as a key enabling technology for 400G, our sense is that a lot of other advanced optical technologies that have been theorized about for many years are now becoming essential at 100G and 400G transport. The most notable of these is probably coherent transmission, which is now used in almost all 100G transport networks and will probably always be used at 400G—coherent transmission is an idea that dates back to the 1980s.

Page | 5

Another old optical technology that CIR expects to see a lot more of in a 400G environment is Raman amplification that will come into its own at 400G on long-haul routes. On the silicon side of the house, DSP chips also seem likely to increase their importance, both as way of implementing performance management and the new modulation schemes (e.g. dual-channel 16 QAM) that will accompany the deployment of 400G transport networks. DSP chips will also be used to reduce power consumption and adjust PMD.

This raises the interesting question of who will control component/chip technology going forward. At one time, optical component and (to a large extent) telecom silicon was manufactured by the leading optical equipment providers. Then most of this capability was divested or subcontracted to third parties.

But not surprisingly, the balance of considerations with regard to this issue has changed since then and CIR's sense of the market is that increasingly the networking equipment makers may want to control their own chips and optical components, or, at the very least, emphasize the role that these chips play in the uniqueness of their products.

This means more custom design and possibly more involvement in manufacture of the chips. There are already signs of this among the current providers of transport platforms for the 400G trials. At one time only Infinera was alone in making its core chip a key part of its product/marketing story for its long-haul platform. We note, however, that both Alcatel-Lucent and Ciena are also doing this.

Meanwhile, both Cisco and Huawei have acquired silicon photonics technology for in-house use, although this technology is presumably intended for a much broader range of products than just 400G platforms

1.2 Objective and Scope of this Report

The essay above on the state of the art of 400G transport poses many questions. The answers to these questions will shape where and how the money will be made in the next few years in the optical transport business. In this CIR report we provide answers to these questions with the objective of identifying

the key opportunities that will emerge in the 400G space over the next decade.

We also attempt to quantify these opportunities in a ten-year market forecast. However, we note that at this stage of the 400G transport market evolution—when the only real 400G transport is to be found in trials—any such forecasts should be regarded with some skepticism.

The objective of this report is therefore to provide a complete opportunity analysis and roadmap for 400G transport, including an assessment of innovations such as superchannels and the latest modulation schemes, and the role of coherent and Raman technology, as well as the emerging standards framework for 400G from the ITU and OIF. Based on this analysis, we identify the key capabilities that will be needed in next-generation core network platforms and how they are most likely to be provided.

The scope of the report is international, as it must be, given the fact that 400G transport trials are now taking place on four continents. We note, however, that the scope of this report is restricted to 400G transport over long-haul (including submarine) routes and metro networks. CIR has another report covering the future opportunities in 400 Gigabit Ethernet (400 GigE).

1.3 Methodology of this Report

The analysis and forecasting in this report are based on both primary and secondary data. The primary data comes largely from interviews of executives at firms that we believe will be critical to the evolution of 400G transport. Some of these firms were interviewed specifically with this report in mind. Others were interviewed as part of CIR's ongoing program of interviewing firms that are at the leading edge of optical networking technology development and deployments.

The secondary research that has gone into this report involves analysis of critical documents from standards groups as well as relevant press releases from firms involved in the 400G market. We have also reviewed all the appropriate articles in the trade and technical press.

For obvious reasons a considerable amount of attention was paid to the results of the 400 G transport trials, to the extent that these were made public. And in both our analysis and our forecasting, CIR's almost 35 years of coverage of telecom markets has been a potent factor. More details of our forecasting methodology are provided in Chapter Four of this report.

1.4 Plan of this Report

This report consists of four chapters and an Executive Summary. The latter is intended to summarize the key opportunities that CIR believes are emerging in the 400G transport space.

In Chapter Two we discuss the main enabling technologies for 400G transport and how these are likely to present opportunities for all the players in the 400G transport market. We also examine the strategic role of standards and how these standards are likely to evolve during the period under consideration in this report. Also covered in this chapter are some speculations on the future of 400G optical services.

Chapter Three analyzes the current state of the art in 400G transport, providing reviews—including our comments on all the leading 400G transport vendors and their current trials. Finally, in Chapter Four we provide eight-year forecasts of the 400G transport market.

