

## 40/100Gb Ethernet, times four or times ten? What is the best way forward for Ethernet

*Date: Wed, 03/10/2010 - 19:57*

Those who reckoned that 100Gig Ethernet was overkill in a world still migrating to 10Gigs, were taken aback by CIR's recent report that huge companies like Google and Amazon already have a "desperate need" for 100Gb connections and that early adoption of the costly new technology would not be held back by the current economic climate. So what about those who say "why bother with 40Gigs when 100Gb is on its way?" According to CIR, the adoption of 40GbE is predicted to start in high-end servers and high-performance computing applications, and up to 80% of the market will be SR4 until CR4 over-copper becomes more viable. But both will rapidly lose share once serial 40Gbit Ethernet becomes a reality in 2014



Michael Howard, principal analyst and Co-founder, Infonetics

Research

For the first time in the history of Ethernet, a Higher Speed Study Group in 2006 decided that two new rates were needed: 40 gigabit per second for server and computing applications and 100 gigabit per second for network aggregation applications. Has CIR proved them right? Or should we be facing economic reality for now, and simply concentrating on lowering the cost of 10Gb and Infiniband solutions?

Michael Howard, principal analyst and Co-founder, Infonetics Research, knows a whole lot more about these issues than the man writing these program notes, so let's listen carefully as he begins the debate with an update on what is happening and why, before introducing the panel of speakers at NetEvents EMEA Press Summit, Barcelona, to argue the case for 40Gb, 100Gb, both – or neither.

**Panellists: Phil Tilley, VP Marketing EMEA; Alcatel-Lucent; European Marketing Co-chair, MEF; Lars Friedrich, Vice President and General Manager Lab & Production Business Unit, JDSU CommTest Transport Solutions; Andreas Stern, Director, Business Development Director, Spirent Communications; Bo Fjelkner, Principle, Nordics Sales Engineering, Verizon Business**

## Michael Howard

This session is about 40/100Gb Ethernet. But a lot of the session is about 40Gb/100Gb without Ethernet on it because that's where a lot of the development has been done so far. Again, I don't have a 45-minute presentation. I have a 35-minute presentation, but I'll do it in three minutes.

Just to set the scene, the background, the 40Gb market is in motion already. Some of our participants here have had 40Gb products of some sort of another for many, many years. It's mostly in play in the optical space. It's not 40GbE. It's not 100GbE so far. So the main application for 40Gb is taking four 10Gb and putting them onto a 40Gb, and will be for a long time.

So there is no standard yet for 40GbE. That's going to be completed this year, probably. And as part of the same standard, 100Gb and 100GbE are being standardised at the same time. So a big point to make, and we may hear this from others on the panel here, but it's the data centre that's driving 40GbE, just about purely. So 100Gb technology demos are here. 100Gb Ethernet is the next big thing. But not 100GbE, but 100Gb, not only demos are here, but we have Verizon on the panel today, who will tell us about the only live commercial deployment of 100Gb in any carrier in the world. It's somewhere near here.

So in the industry, I don't want to go into too much detail, but there's a lot of focus on 100Gb as well. But it's going to take a few years to get to, we think, a real big market for 100Gb. So we hear these statements all the time, service providers, data centre owners say they'd buy 100Gb today if it were available today. Well I don't think that's quite true because the technology has to be standard. It has to be reliable for the carriers to deploy it. And most importantly, it has to be reasonably priced, which it is not yet. Some carriers will buy it, absolutely, and have already.

So it's going to be, we think, still a couple more years before the products are going to arrive that are going to be really available to the whole market, and a couple more before the prices come down close to what people thought is the usual price demarcation. To buy 100Gb, the price range has to be at least the same as 10 times 10Gb or less. But this is proving not to be true. 10GbE prices have come down so fast that this relationship cannot hold in the market.

So some actuals, but our forecast of 40Gb and 100Gb -- now this is on WDM and Sonic Gear, not on routers, not in the data centre, but in carriers' backbone networks, essentially, that our to 2013, 40Gb ports -- and this is the number of ports, not dollars, not euros spent on it -- and then 100Gb ports. So that's the way we think it's going to play out.

Now here is a study we just did not too long ago of carriers around the world. We asked them about pricing. And this again is not for routers, not for data centre, but for WDM and Sonic Gear. So again, this is not 40Gb Ethernet or 100Gb Ethernet, but the optical interfaces.

So the question is, "When would you buy 100Gb?" So 17% said, "When the price is three times 40Gb." And let me say, today 40Gb prices are like six and eight times 10Gb. So 40Gb is nowhere near four times 10Gb. So carriers have to pay a premium for 40Gb or 100Gb. Of course there are a lot of practical reasons and then there is being first on the block or first in the world reasons to buy 100Gb. There aren't reasons to buy. And the total cost of ownership does play into it.

So at two and a-half times 40Gb, that's exactly the same price per gigabit -- if 100Gb was the same price per bit as 40Gb, then 17% more will buy it. And then two times 40Gb, you'd buy 100Gb, so in other words, you'd pay the price of 80Gb, that brings in 25%. Oh no, greater than that. But at two times 40Gb, you bring in a big piece of the market. So that would be still, today, at the same price as 40Gb, essentially, times 10Gb, it's going to be like 18 10Gbs, 16Gb, not 10 10Gbs. Again, because the 10Gb price has come down quickly.

So basically the give me two and a-half times the bandwidth at two and a-half times the price, that used to be the magic formula, doesn't work anymore.

I think we'll just stop there, except to point out one thing. And that is that if you look at the revenues brought in by 10Gb, 40Gb and 100Gb, last year again, on optical equipment, under 10Gb was quite a bit of the revenue. So 1Gb interfaces are still pretty prevalent, being purchased. The 10Gb itself is the majority of revenues spent. If we did ports, these ports would be very small. And then a tiny bit of 100Gb. But by 2013, we see that 10Gb is still big, but the amount spent on 40Gb and 100Gb is going to be big. In fact, we're predicting that in 2012 or 2013, but by 2013 anyway, more revenue is going to be spent on 100Gb than 40Gb. There's a lot of 40Gb in networks today.

So we met Phil Tilley already. We met Andreas. And I don't know if we've introduced [Bo Fielchner].

## Bo Fjelkner

My name is Bo Fjelkner. I work for Verizon. I'd just like to mention that Verizon has a long tradition

into Ethernet. We have a broad range of Ethernet offerings from the access, but also services like Layer 2 Ethernet on a global scale. [EPL] it's called. So Ethernet is very close to our heart. And it's been recognized by Metro Ethernet Forum and a number of awards, etc. So Ethernet, we like that.

#### **Michael Howard**

And Lars Friedrich, who I know wasn't introduced before because he arrived during the middle of lunch today. So welcome, Lars. He's the Vice President and General Manager of Lab & Production business unit at JDSU, of Deutschland.

#### **Lars Friedrich**

Yes, exactly, from [Inning] and [Stuttgart]. So JDSU, obviously making test and measurement gear, among other thing. And the business unit that I'm representing is making test gear for 100Gb, for 40Gb, and not necessarily field test gear, but lab and production. So that means we make gear that enables network equipment manufacturers and service providers to develop 100Gb and 40Gb solutions to troubleshoot and test them, and to produce them, essentially. So we tend to see the early curve of the technology development with our products.

#### **Michael Howard**

Great, thank you. Welcome. So I'd like to ask Bo, first off, tell us about your 100Gb link. How long is it? Where is it? Is there real traffic on it?

#### **Bo Fjelkner**

I had to check that out before entering here. So the 12th of December, 2009, we announced in a press release that we are having commercial traffic between Paris and Frankfurt, [8 tonne driven], 93 kilometres on a single wavelength. And it's with the proper IP traffic over it. It actually goes parallel with other 10Gb links on the other wavelength on the same fibre. So it works very well. I've also found from the Internet that we don't test. We do have a [inaudible] [Alcatel-Lucent] here in end of 2007 between Tampa and Florida. So we work with a lot of suppliers. But we are very focused on 100Gb in the core.

#### **Michael Howard**

And again, just to point out, this is 100Gb wavelength. It's not 100GbE yet, no standard yet till later this year. So there is a deployment. Phil, would you like to talk about router deployments of 100Gb? There aren't any. Oh no, that's enough. But there might be some soon. Where do you think Alcatel-Lucent is in the router wars in 100Gb?

#### **Phil Tilley**

We sort of looked at it and went to the customer base and said, "Okay, for one of the first times in history, we're actually going to see two standards come out at the same time, 40Gb and 100Gb, launched and available as specified standards at the same time." Historically, you just generally go up in speed. But no, for having 40Gb and a 100Gb at the same time with standard interfaces creates a new challenge, a new interesting point. So we actually looked at it and said, "Okay, well we can develop both. So we'll do 40Gb and 100Gb." Of course, the natural thing, we talked to the customers, what do we think we're doing?

So we did, over a year ago, when we were at the point of deciding which card we develop first -- one has to prioritise developments, obviously -- went to the market and asked. And it was very clear from the customer base, looking at today's growth and expectations and bandwidth demand, looking at the video traffic on the network and all the other drivers, 40Gb is really just an interim stage. And actually, if we have to upgrade our optical gear on routers for 40Gb for a stage that's just purely in transition, it's not going to work out. We might as well take the steps. So a very clear message from all our customers is, "No, we want 100gb."

We're going to put new infrastructure in. It's going to cause a complete refresh of network infrastructures. We're going to look at the transmission side at the same time as we're looking at the routers. We'll put in and go for 100Gb. I was talking to one customer this week, saying they're buying today, per month, 150 10Gb Ethernet cards across their business on routers. Actually, no, that was across their whole business, routers, transmission. So you just see, they say they're buying so much 10Gb, they actually say, very shortly, "We're under pressure to get to 100Gb." So our customers say, "Go for 100Gb and we want it as soon as possible."

**Michael Howard**

So I asked Phil, do they know how much it's going to cost? And the answer?

**Phil Tilley**

They've got a feeling, an idea of how much it will be. Expectation, long-term, it will come down to 10 times 10. That's the expectation. Initially, yes, they're prepared to pay more. But the challenge and the driver to get it in there, to drive efficiencies and all sorts of other things, is just so large. It's just going to happen, from their perspective.

**Michael Howard**

Can you say one thing about what's the total cost of ownership equation for someone that says, "I'm willing to pay 20 times 10 to get 100Gb because in my network," what does it mean?

**Phil Tilley**

So in the network, obviously the equipment [back-planes] and all the equipment is being built and architected around 100Gb, so you'll have 100Gb or 200Gb per slot into your router equipment. If actually you can only get 40Gb out but you've got a capability to go 100Gb, all of a sudden you're burning an extra slot, which again is a huge cost. So that's where there is justification of 10 times 10 versus 20 times 10.

**Michael Howard**

And you're [burning to save] wavelength. The wavelength is either at 40 or --

**Phil Tilley**

Or at 100. So you're not running your network as efficiently. And of course, then there is power consumption, as we just talked about, the green IT. 40Gb isn't half the power consumption of 100Gb. It's more. There's a power thing as well.

**Michael Howard**

So how does this relate to your business?

**Andreas Stern**

We have looked into this as well 40Gb and 100Gb, or 100Gb. And we see in the market a huge demand for 40Gb as well. So looking into the test equipment market, we see three times the demand for 40Gb than we see for 100Gb. But it's clearly like Phil mentioned before. On the routing side, it is clearly 100Gb isn't it? 40Gb plays a huge role in the cloud computing data centre scenario.

**Michael Howard**

And that is 40Gb Ethernet.

**Andreas Stern**

I talk only about 40Gb Ethernet. So this is the only [inaudible] we address. Different to GDSU, we are looking only into the Ethernet side of the market also.

**Lars Friedrich**

So the part obviously that we are looking at is the transport part, right? So it's the long distance kind of stuff that we test. And clearly, 40Gb has been there for -- we talked about it over lunch -- four or five years. But not as 40Gb Ethernet, as either SDH SONET or SDH SONET wrapped in OTN in [inaudible] seven or nine. And obviously that's there and that's a very healthy market right now. It continues to grow, right? Lots of developments on the line side and the technology developments. On the Ethernet side, now specifically for transport, not including the data centre part because that typically would not be part of our customer set. Clearly the demand to focus on is 100Gb. There is also interest in 40GbE. And the standards are sufficiently standardised. Even though they are not ratified yet, there are no major changes. Stuff is out there. Equipment is out there that is standard compliant with 802.3ba, essentially. So there, on the transport side, the focus clearly first is on 100Gb. But almost everything is switchable so that it's clearly a combined 100GbE/40GbE solution. But I think what is interesting on the transport side is that we're seeing here a little bit of a paradigm shift from the classical SONET circuit switch [inaudible] over to Ethernet now. And we see quite a lot

of the cost for the transport gear, but also in the data centre I imagine. It's in the interface. And with the new CFP-based interfaces, we've really seen a radical technology step from the classical SONET serial interface, one fibre, one wavelength, and basically the data stream modulated on there, just plain serially, to the CFPs.

And I think the general hope is that there will be quite a bit of cost savings from going to this new parallel CFP approach. It's too early to tell that. But clearly I think there is a hope for cost savings. So that means cost savings, if you compare the data rate that you have today, 40Gb in the SONET world to 40Gb Ethernet, you would see cost savings there. And then also relative to [what's in] times 2.5, if you look at cost per bit. Or do you buy 2.5 at 100Gb?

**Andreas Stern**

The standard is not too far away. The signature will be done in June this year, so it's really close.

**Michael Howard**

So like you said earlier, the standards are close enough that you can start developing products for it.

**Andreas Stern**

Actually, we've finished our product last year. And the biggest challenge was to get hands on the CFPs, to get hands on the optics. We could give this test equipment to our customer, but we have to loan the optics somewhere.

**Michael Howard**

So say a little more about that. So all of today's modern optical equipment has pluggable optics. That is, there's all the processing on the board, the logic to make 100Gb or take 10 10Gbs and put it into 100Gb or 4 10Gbs and put it into 40Gb. But the actual device that takes that electrical signal and puts it on opticals -- today the device that's going to be used for 100Gb is called a CFP. I don't know what that stands for. But basically that module is not available today. So June or something?

**Andreas Stern**

It is, in small quantities.

**Michael Howard**

As you were saying at lunch, it costs the same as a car. You can buy an automobile or a CFP. They're the same price.

**Andreas Stern**

As you said before, you had one in your hand.

**Michael Howard**

But I did. Actually I had Cisco's, Alcatel-Lucent's and one other vendor's 100Gb cards. And at Alcatel-Lucent, I had the 100Gb CFP. And I teased. I said, "Okay, I'm putting this in my pocket and walking away. But that's about the size of it."

**Phil Tilley**

So some big bodyguards walked out and stopped you. From our perspective, the race is now on where we see all the service providers across Europe wanting to actually try all this stuff. And actually, we're lining up a number of trials across Europe, starting in June when the CFP is starting to become readily available. And then it's a rush. And every service provider wants to actually go public and say, "We've upgraded our network to 100Gb. We're running 100Gb," because again, I think it's seen by our customers as significant investment in infrastructure to drive some operational efficiencies and show that they're actually really able to support high-bandwidth service efficiently. And it's going to be an interesting challenge, from our perspective, managing the number of trials that are going to go on, on this new stuff.

**Michael Howard**

So just one point. I don't think you can say this officially, but the Verizon equipment, as they point out, is transport only. That is, the optical equipment only? It's 100Gb. It's Nortel gear. It's the only gear working at that speed today. And there are no router connections, right? There is no 100Gb router.

So Phil is saying they're going to have later this year the router connection. But inside the Verizon optical gear, they aggregate 10 10Gb streams and put it onto the 100Gb stream, inside the Nortel equipment.

**Bo Fjelkner**

Yes, to connect the routers.

**Michael Howard**

To connect the routers, each router by 10Gb, though.

**Bo Fjelkner**

Because it's commercial, like the traffic over this connection. What is good also to know, we've been very pushing on ultra long-haul. And the reason for this is that beneath this, we have an optical mesh. So our global solution is built on a physical layer with an extreme redundancy. So for instance, over Atlantics, we have seven paths. You can cut two cables. Customer traffic won't be impacted. We're extremely focused on IP. We can transport everything over IP. So we want the same reliability as we had in the old traditional network. So that's we're sort of testing [inaudible] 1,000-plus kilometres with a third supplier and so on.

**Phil Tilley**

And it's worth pointing out that actually, the challenge, if you've got 10 ports going from the routers, 10 times 10Gb, into this optical, what you actually do is you're obviously hashing, sort of spreading that IP, creating it or aggregating it across using a link aggregation protocol. And the trouble is, what you're not getting is 100Gb out there. You're more getting 60Gb, 70Gb because of the way you spread, and you aggregate and you wrap protocols around it. So the efficiency of 10 times 10 in that way, isn't as good as a straight 100Gb, which is why we want the straight 100Gb20.

**Bo Fjelkner**

We see the bandwidth increase. So with anything that we can optimise, it's good for us. When the day will fall, 100Gb, excellent. When we have the terabit, excellent.

**Michael Howard**

So there are a lot of routers connected at 40Gb today. So you probably have routers in Europe that are connected to your network at 40Gb. And those get also aggregated into the 100Gb. I think with that, we're about out of time. So we're going to take a question or two.

**Unknown Speaker**

I would like to know, if you need a lot of things in the fibre, in [inaudible]. Or is it the main thing? What kind of memory can you manage when you have a few 100Gb inside the routers? So what is the speed of memory to manage all of these [inaudible]? It must be crazy.

**Phil Tilley**

So from the router side, I'm not sure about the memory consumption. I don't think that particularly changes the number of advances. We've upgraded. We've got the routers capable of supporting the millions of routes required to be supported over the 100Gb.

In terms of the infrastructure, what we actually have in place now, the routes or platforms we have in place are all 100Gb capable. They've got the [back panes] and the speeds there, so it's really just the upgrades to the new cards that remain, from our side, on the wavelength, the fibre infrastructure. I'm not a fibre guy, so wavelengths.

**Michael Howard**

When I saw their line card that runs at 100Gb, they have two sets of chips, one for in, one for out, on the card. So the 100Gb in and 100Gb out are separate sets of chips. That's how they engineeringly designed it today.

**Unknown Speaker**

You talked about the primetime for 100Gb will be 2014. And last week the City of [inaudible] said that

they need a terabit, especially for the data centre, because we all talk about the core. But what about a data centre? And what is the processing power from the server that could take this traffic?

**Andreas Stern**

If you look into the [blade surface] available today, you have hundreds of cores available. You have hollow working [blade surface]. So by adding 100Gb streams, you come to your terabit. So it is really just using existing technologies or soon-existing technologies, standardised technologies, and then you will bring networks to the required terabit level. It will take years to step from 100Gb Ethernet to 1 Terabit Ethernet. Now this is a huge job.

**Michael Howard**

So the next level being looked at is 400Gb, correct?

**Andreas Stern**

Yes.

**Michael Howard**

I'd like to say thank you to the panel. We had a good session.