

NETEVENTS

APAC PRESS & SERVICE PROVIDER VIP SUMMIT

Draft

Round Table Session V Metro Aggregation: a 100 Gigabit Opportunity

Chaired by: Manek Dubash, Editorial Director, NetEvents

Panellists:

Andrew Bond-Webster	Vice President Asia Pacific, Infinera
Junjie Li	Board Member, OIF, Director, Optical Communication Research Centre, China Telecom Beijing Research Institute
Gint Atkinson	Vice President, Network Strategy & Architecture, KVH Colt

Andrew Bond-Webster

[Skip] straight to the end and give you the answer, but I'll just walk you through a little bit as to what we are seeing. First of all a few minutes on who Infinera is. Our focus is DWDM optical transport networks. We are a global organisation, headquartered in California. We today have customers in around 73 countries, 144 customers, 62 of whom are using our 100G long-haul DWDM OTN integrated OTN platform called DTN-X. We launched that in 2011 and started delivering and shipping in 2012. As you can see our customer base goes across the whole spectrum. Tier 1's the cable MSOs in North America, the likes of Cox, Time Warner Cable, Charter both of whom have been in the news recently and also the wholesale and enterprise carrier providers.

One of the more interesting segments that we see developing is the internet content providers. This is a market that has begun to evolve and become very relevant and prominent for us over the last two or three years, to the point that we've actually established a business unit specifically focused on the internet content providers, by whom we mean the likes of Amazon, Microsoft, Google, Facebook. Don't draw the conclusion that they are all our customers, three of those four are one of whom we publicly announced which is Facebook. But what we are seeing is that they're driving a lot of the development. They are driving a lot of the capacity demands in the region

in Asia and globally. And also they've been very vocal to us in terms of the products and the solutions that they're looking to deploy within their networks.

One interesting fact I thought I would just share with you in terms of what is happening within the network and how content providers are driving traffic across the networks. This is an example from Facebook, which is where one kilobit enquiry or post or like on Facebook, the amount of traffic and how that multiplies through the network. It multiplies through, one kilobit multiplies through their network 930 times. That's just one like. So if you think about the amount. Every time you hit like on your laptop, on your phone or whatever the impact on the network and the traffic that it drives on the network for Facebook means they are continually, having to scale up and to grow their networks in a very rapid, rapid fashion. Another example I think there was a gentleman who jumped out of a plane 24 miles up and the impact on the traffic, the spike in traffic across the internet that that drove was huge.

So what we are seeing is both east to west traffic and also north to south traffic growing across the networks. And the content providers, the network providers whether you are a provider to the content provider or whether you're a content provider with your own network, the challenge you have is how do you scale up those networks to continue to deliver a quality of service experience to the user. So what we are seeing is that that's multiplying through not only at the core network, long-haul network but also within the Metro transport environment. And you need that bandwidth you need that capacity to be available at the time at which those spikes in traffic take place for example.

From a Metro perspective in Infinera, we characterise it in two different ways. One is the Metro aggregation market, and where that's the user to the data centre. And we see that economics primarily is the trigger for the growth and the evolution on those Metro aggregation markets. And they're driven primarily by the telcos and by the cable companies where you can see that the services as we see them they are looking for Mesh services, they are looking for OTN type-based services. And typically they're looking for n by 100G super channels.

Now Infinera has a piece of technology called the PIC, which is a Photonic Integrated Circuit. We have developed that capability and that chip to the point that we can now create -- that chip is now multi-directional, which enables us to deploy that and expand that into the Metro arena, which we were not able to do before.

The second part that we are seeing is what we've termed as the Metro cloud market, which is the data centre to data centre interconnect market. And that is where the likes of Facebook, the likes of Microsoft they need to move huge amounts of data point-to-point between their data centres. They need something that's very simple, very scalable and enables them to react very quickly to changes in demand and to the changes in data flows. And, as I say, we've termed that as the Metro cloud. Now for both markets what we believe and to satisfy both requirements scalable optics is absolutely fundamental to being able to address those markets.

We have, as I say, we have a piece of technology called the Photonic Integrated Circuit. Very simplistically that is where we can take 650 optical functions and put

them on a chip, the chip being the size of your fingernail. It's highly reliable. It's highly scalable. And now we've also introduced a degree of flexibility by being able to make that PIC multi-directional. So that in a Metro environment, a Metro aggregation environment for example, that enables you at a hub site to have that PIC which is today 500G on a line card and have that operate on a hub spoke type basis.

In terms of the generations and the changes you can see the first generation was 10 x 10 giving you 100G on a PIC. That was then -- went to coherent 500G today, sliceable as we term it optics, sliceable PIC. We see the opportunity and we have already announced and tested a 1.2Tb PIC. So we see the ability to continue to grow the PIC technology to support the ever increasing demands in capacity way into the future.

In terms of the services and the services mix that we are seeing Layer 2 services growing at twice the rate of Layer 3, 300% growth in high-capacity Ethernet services. And we also see that users are deploying, or we expect, 84% of our customers will be deploying Packet Optical Transport solutions by 2017. And the growth rate in Asia of, or compound growth rate, Layer 2 services is over 10%.

So what are we seeing very simply for the need for Metro, Metro aggregation? Coherent super-channels I think are a must. The ability to scale those we believe the right answer to that is through the PIC, through integrated photonics. Clearly, within the Metro aggregation market you need a degree of granularity in order to be able to satisfy those markets. And that's where we believe the sliceable photonics comes in. You need to be able to support both point-to-point, data centre to data centre and also Mesh network technologies, particularly the Metro aggregation layer.

Packet, we see playing an ever important role and increasing role particularly in the high growth Ethernet services. And then ultimately that all needs to be programmable, I think there's been some discussion around SDN, the role of SDN, clearly that's going to play a role. And having all those capabilities under SDN control and programmability is where we see the future.

To answer the question do we see 100G happening in the Metro, very simply the answer is yes. We are seeing across the region in different markets different rates of evolution from 10 to 100. But it is happening. No question it is going to take place. It is happening now whether it's in Europe, the US, in Asia, we are seeing it in Japan, we are seeing it in Korea, we are seeing it in Australia. Talking to a gentleman [inaudible] earlier, China clearly moving to 100G, so it's happening, it's happening today. It's not is it going to happen. It's already happening.

Thank you.

Manek Dubash

Thanks Andrew. Thank you very much. Okay, I'd like to extend a special welcome to Junjie Li, from China Telecom who's not graced us with his presence before, I don't think. So Junjie, in fact perhaps call on you first, is the picture that Andrew paints in terms of the definition of the Metro, the speeds that are required, all those things is that familiar to you? Is that how you'd -- would you agree with that?

Junjie Li

First I will introduce some OIF efforts on the Metro [100G] then I will introduce some -- China Telecom's viewpoints about the Metro 100G. First of all I start with OIF. I think maybe many gentlemen know that several years OIF introduced the coherent DP-QPSK solution into the industry for the [pluggable] DWM and it become very successful.

And the next step for OIF efforts we are working on the, we call the Intermediate to reach 100G. Its, I think the target market for the managed networks. The reach target is for maybe 10 [kilometres] to maybe 500 or 600 kilometres. And OIF is working on this and I think this white paper will be released maybe later this year or early next year. So this is OIF's next step for 100G.

And during discussion [inaudible] OIF all the carriers and vendors believe that is [the pluggable] module will be the future for Metro 100G. So OIF is defining a new [IF] we call the CFP2-ACO this is a pluggable very small module for Metro 100G. And this project is planned to finish early next year, so this is the OIFs work and efforts on the Metro 100G I will introduce.

And from the viewpoint of China Telecom, we believe it's the time for introducing 100G into the Metro networks. In the past three years not only China Telecom I think the all three carriers in China we have deployed many, many 100G in our backbone. It was said maybe the China market is one-third of the global 100 market. I don't know if the figure is right, but actually we deployed many, many [pluggable] 100G.

And actually from last year some Metro networks in China Telecom started deployment of 100G. Why? I will give you some reasons. First is the cost, currently the, per bit cost for 100G transmission is lower than 10G.

Manek Dubash

Lower?

Junjie Li

Yes. And this conversion is just the device cost. We don't take the fibre into account. If you take the fibre into account I think the cost gap will be huge. So the second reason is that we want to increase the capacity per fibre. This is the second reason. And the third reason is that currently we deploy the coherent 100G in Metro networks. Coherent solutions has a very huge tolerance with [CD] and PMD so it's very easy to deploy and it's very easy to maintain. So this is the first time the maintenance department, the maintenance guys in carriers they always reject the new technologies, but this time they work on new technology because they find that the 100G is much easier to use, much easier to maintain. This is new.

And the volumes and traffic [Mr. Gi] had discussed about data centre interconnection, and it's the same in carriers. In our networks data centre becomes a new centre of check flow. In some of our Metro networks the check flow in and out the data centres,

the percentage of these checks is 70% and it increases very, very quickly. So we understand our next solution DWDM networks should be focused on that centre.

And another [tangible] problem is that because in Metro networks we don't need the very high performance to [change] with several thousand kilometres with [inaudible] solution we can reduce the requirement for performance, but we also want to reduce the cost and the efforts to reduce the costs for Metro 100G is very welcomed. Currently, maybe the pluggable CFP2-[AP2] solutions and maybe lower cost, lower power coherent 100G solutions and our [inaudible] solutions we do some research with our vendors.

And the last issue I would like to introduce is the debate between the coherent and the non-coherent, because some vendors in the market they want to introduce non-coherent 100G into the market. They claim that it's much cheaper than coherent solutions. Currently we don't do some -- maybe test, but our opinion is that the coherent 100G maybe it's much better for, we call it the Metro Core and Metro Provision, because in set networks we needed to [do some], we would as management we needed to introduce [inaudible].

So we want 100G[inaudible] into one [inaudible] it's easier to manage. But for some other scenarios we call the [access FDM]. For example the data centre interconnection [inaudible] or we call the remote [RU], remote radio stations between [BBPO] to the RU is certain scenarios the [inaudible] only provide a point to point connection and the volume is very huge. So in that case the cost is more [inaudible] in that case maybe the non-coherent solution will have its market if they can reduce the cost.

So I think this is my [inaudible]. Thank you.

Manek Dubash

Thank you. Gint, thoughts, comments so far?

Gint Atkinson

Well KVH in Japan two and half years ago we were the first service provider to launch 100G service. So we made a commitment to offer a 100G service a long time ago. We had a lot of the over-the-top players were our customers for data centre and networking services and they increasingly needed more and more 100G bandwidth between data centres.

We then extended that capability all the way to the landing station, so that if you're an over the top provider you can get onto the KVH network at one of several landing stations in Japan and have complete diversity all the way into Tokyo, Osaka and through and between.

So our strategy has been intensely packet-optical driven for big providers with big needs. They can buy wavelength-based services and then late on migrate to a packet-based service. And you'll see we are building out our Singapore Metro. Our plans are to put minimally eight [LAM] with WSS out to every single access node. So we'll have packet at the very edge or optical or OTN and then inside our Metro network we

are going to have rings and partial Mesh's. What [LAM] they can all be upgraded from 100G to 400G to what's next 1.5G.

And this is the big thing that we are going for is building these extremely ultra-high capacity, ultra-low latency Metro network cores from the landing stations all the way on out to the very, very hot buildings that require a lot of capacity and data centres. And we are doing that in Singapore and Hong Kong, Japan is already there.

Manek Dubash

Okay. Questions at this stage? No. So how do you -- have we now reached the price point then on a per port or however you care to calculate it that 100G is now kind of mainstream?

Gint Atkinson

The truth is it really relies around putting a lot of pressure on our vendors and really getting them to deliver on a lot more than just the hardware. So I can highlight in Colt compared to KVH, we at KVH have a very limited range of vendors. We have some Sienna, Sienna does interesting stuff. Cyan in the Metro has played really well because we can get 98 [landers] in a tiny little bit of space.

And we are waiting for Infinera who is very soon going to be delivering that same type of density, but we are going to get all of the wonderful stuff we love about Infinera. We have Infinera on our core, so elastic, optical capability, being able to very easily operate the optical network, slice it up and virtualise it. Infinera has had this for a long time.

We've loved the DTN-X platform and we are waiting for it to get squeezed into a very small form factor which is, my understanding, coming very quickly. So I think what you'll see is our really exciting packet optical Metro which today is for the most part all Cyan will be able to grow with two vendors. And the two vendors that are going to be able to lead that are Infinera and Cyan.

Manek Dubash

Okay, I'll come to you Andrew in a sec, Junjie do you want to comment on the cost question?

Junjie Li

Actually, I have introduced something but I believe the cost of 100G will be [inaudible] very quickly because the main cost of transmission device comes from the optical. For the chips if you increase the volume your cost reduce very, very quickly. And thanks to the coherent solutions all the complexities are put into the [inaudible] to the chips. So what I see is that currently the cost of 100G is -- probably the cost of 100G is lower than [100G] this is true in China market. And for Metro networks because we believe the volume of Metro 100G will be much larger than backbone, so the cost will be lower than backbone. This is our requirement. And I believe it will be reached soon.

Manek Dubash

Andrew, your stuff needs to be cheaper and smaller.

Andrew Bond-Webster

Yes I got that message loud and clear.

Gint Atkinson

I think smaller, just do it smaller.

Andrew Bond-Webster

I think just in general terms you look at 100 -- the cost of 100G platforms from when we first introduced the DTN-X in 2012 and started shipping, you look at the price points, that has come down consistently since that point in time, driven as Gint very eloquently said, they put pressure on us and there is competition in the market, which from a service provider perspective is great news.

If I look at the data centre market, the data centre to data centre interconnect market, the pressure on them is both on the CapEx side which comes directly back to us. So that's the challenge is how do we produce these high capacity boxes and capabilities. And today we have a platform called the Cloud Xpress which is a 2RU box which gives 500G point-to-point DWDM in a 2RU box. The CapEx component is only one piece of the equation, the other side of the equation is obviously the OpEx side, so space and power also plays a very significant role. And I think that's also where we've been pressured.

Our content provider customers have said we love the reliability of the PIC. We love the scalability of the PIC. We like the economics of the PIC. But your DTN-X platform is just too big, so we need to find -- you need to simplify it for us, which is what drove us to come up with the Cloud Xpress platform. Just talking about that as a solution its 2RU. It can scale up to 21 terabytes in a rack, so it's highly scalable from that perspective. From a power perspective, very simply it's also highly efficient. It takes less power than a hairdryer. So we have to recognise the pressure that our customers are under, respond to that pressure from a pricing perspective and give them solutions that meet those requirements.

At a Metro aggregation level, Gint has kindly made reference to the fact we need to take our core technology and make that more economically efficient if you will for the Metro space and the Metro aggregation. So we are taking, as I said in the presentation, our photonics, integrated photonics and making it sliceable so you can build these kinds of aggregation point, hub, spoke type aggregation networks with OTN switching as well into smaller form factors, so that they're more readily deployed and applicable within the Metro environment.

Manek Dubash

Inevitably with any technology, certainly with networking technology, eventually obviously mainly we are talking about within core networks right now and as you say

also DC to DC interconnects, that technology will expand outwards as it becomes cheaper. What's the next step for 100G do you think in terms of where you might deploy it outside of those obviously high capacity core network areas?

Andrew Bond-Webster

From our perspective we just see the need for capacity, the need for -- network traffic continues to increase day by day by day, whether it's -- pretty much whoever we talk to average you're looking at 30%, 40% growth in network traffic. So you need to -- we need to continue to be able to scale our platforms to meet those needs, make them economically efficient CapEx, OpEx perspective and make them expand volumes to the Metro.

I think the other side, and Gint you alluded to it as well, is from the user -- so from the line side we have to increase the capacity. What we are also seeing is that the user side, the interfaces that the users are going, and Gint you referred to 100, 10, 40, 100 what's next, 200, 400 who knows, but I think you're going to see that the user side also require increased capacity and import size.

Gint Atkinson

Yes there's a lot going on there that's going to affect the whole supply chain. We are still -- the bulk of our orders for 100G are still requiring a bundle of 10G hand-offs. So a lot of these content providers have not upgraded their routers to 100G yet, so they want it carved up. So that's something Infinera does well and most vendors know how to take optical and then carve it up into a handful of 10G Ethernet interfaces.

But again this is part of the space and power challenge. So we are expecting that Infinera is going to have all of the options available for us here in a very short period of time based on their acquisitions and what we've seen with new products. But you also have to surf the right wave. And this is where there's only a few vendors, the leading vendors that can take us in the discussion with, okay and how is all this going to evolve. Clearly our customers with 10G ports are going to upgrade to 100G Ethernet. Then they're going to go to what's next, but this an interesting problem the customer interface what does it look like?

There's another problem which is most customers that order 100G wavelength service with an Ethernet hand-off, whether it's 10G or 100G Ethernet, very quickly they say by the way do you have a packet service that's almost as good. Now in the case of cloud service providers they're running on TCP-IP. And believe me they generally don't need a wavelength service but they are habitually ordering wavelength.

So this is something with, I expect Infinera like we are doing on our Cyan platform, in the future we can take a 100G lambda service and throw it into an Ethernet service and then reclaim that lambda into our backbone. This is the kind of stuff that Infinera does brilliantly on their platform today. So that's a difficult evolution. We see 400G wavelength, terabyte wavelength definitely quickly coming in the future. And the little thing is that customers want an Ethernet handoff, so where is Ethernet going, and that raises a lot of issues.

Manek Dubash

Junjie.

Junjie Li

Maybe my view is probably a little bit different, because for China Telecom we deploy 100G [for 100G]. It means we deploy 100G [equivalents] for 100G [reason] [inaudible] this is the situation. And we have seen very clearly [inaudible] 400G and the interface not only the backbone [inaudible] but also data centres, because big [packet] is always better for [inaudible] for servers. And I got the news that just last week the IEEE have decided, maybe decided the 400G Ethernet, the baseline for two kilometres and 10 kilometres - it's good news. And so we hope 400G Ethernet standard will keep to the planned schedule. Maybe they will delay maybe half a year but we still hope.

And for another issue, this is one issue for 400G, another issue is the licence for 400G. Licence for 400G we believe that from 100G to 400G licence we believe its evolution not a revolution but from 10G, 40G to 100G is leverage. And currently we have some tests with our vendors about maybe [TM16] something based on current 100G technology. They can provide 200 [equivalents] and [2] bandwidth for 100G. So we believe first kind of technologies will be adopted in 400G [inaudible].

And we are discussing about the Metro aggregation, you mentioned [how it works], yes I agree the big partners should adopt many, many maybe smaller check flows into --. So we are also very [interested] [inaudible] with the [POTA, the OTA] packed a device [integrated] as a [inaudible] capacity is the [inaudible] capacity and the [inaudible] together. This is the device we prefer in Metro.

And in Metro networks because the bandwidth the capacity requirement is very different in the different scenarios so we also need some, maybe different footprint, different capacities in different scenarios, so this is maybe a little bit complicated and different with [inaudible]. And you had discussed there earlier, introduced there too [500 box] and maybe smaller we will --. Thank you.

Manek Dubash

So what can Andrew do better then?

Gint Atkinson

More bandwidth, less space, less power and then lots of intelligence mapping, packet services to optical services, you're only one step away from that.

Andrew Bond-Webster

If I do that will you buy?

Gint Atkinson

Almost for sure, but.

Junjie Li

We are discussing about power this is also very important. And actually the past three years I believe the power per bit for 100G increased almost half. This has happened. So I believe it will happen.

Manek Dubash

Sorry you said power consumption is a big issue? Yes.

Andrew Bond-Webster

Yes we see that across the region with all our customers, everywhere it's how much power does it draw, because if they're in collocation centres for example they get how much they're being charged for power becomes a very significant part of the OpEx side. So it's always a matter, it's a balance between the CapEx which is clearly important but also the OpEx component is something from a total cost of ownership perspective plays a very significant role.

Facebook is one of our customers and we deployed the longest terrestrial route for Facebook in Europe. They put a data centre inside the Arctic Circle. Why, because they don't need cooling. Why don't they -- what's the benefit? It means it saves them power. So anything that's going to save them power anything that's going to have a value to them in OpEx is a direct value, goes straight to the bottom line, it's a very significant part of their cost.

Manek Dubash

So when I hit a Facebook 'like', how much Antarctic ice do I melt?

Andrew Bond-Webster

This is not a green conversation.

Gint Atkinson

So the opposite of that kind of on the Metro service provider side is when we go to a central office or a GC, we go in there and there's a limited amount of space and power. So even putting the cost aside, there's just the issue if the space isn't there and the power isn't there deal is over, we can't do a POP. So the value of, especially the [to you] solution, is really powerful because it means that we can get into a lot of these small, small POP locations that they just don't have this much space available there's only that much or to you or for you. So for a Metro player, Metro optical player like us it's huge. It's not just the cost, it's can we even get it in.

Manek Dubash

Forgive my ignorance of how photonics works, but is there a, not being an optical engineer, does it work the same as Moore's Law as in you make it smaller you use less. What's the equation? How does it work Andrew?

Andrew Bond-Webster

I've got a colleague at the back who can probably answer that question better than I can. But from the point of view of being able to scale the PIC, as I say we started out 10 x 10 with 100G on a PIC. That then went to 100G coherent and long-haul. We've already taken it to 500G on a line card if you will on a PIC. So that, our focus is scaling that up that continues to maintain the space, the power benefit by being able to get 1.2 terabytes on a line card, which means with our existing chassis for example it's easy to expand that from the 5 terabyte today to 12 terabytes tomorrow.

Manek Dubash

1.2 terabytes with the same power consumption?

Andrew Bond-Webster

Probably.

Manek Dubash

It goes down [inaudible] per gig it goes down.

Andrew Bond-Webster

So it goes down per gig.

Manek Dubash

It goes down, can we quantify that?

Andrew Bond-Webster

No.

Gint Atkinson

Trust me it does. Being familiar with their product line it's -- I think its --.

Andrew Bond-Webster

Yes. On an aggregate basis obviously it goes up but on a per gig basis then it goes down.

Manek Dubash

Okay. Questions? Long-term future for 100G where's it going next?

Gint Atkinson

It's going into your TV. You've got 4K there and those built-in cameras the government spy agencies want to watch you watch your TV. There's just --

Manek Dubash

Don't laugh, it's happening in the UK.

Gint Atkinson

4K is there, there's just what happens when you push the Facebook button with 4K content. So I have 10G service at home and normal service is 1G for about SGD70 a month. You get 1G service and that includes standard set of IPTV channels. You can now get 10G service and that's just tripled the price.

Manek Dubash

Junjie?

Junjie Li

In China and in Japan we [also have] but we promote the 100M per user this is our target in the next year. And in large cities like Shanghai, Guangzhou, [Tianjin] we begin to provide 1G per home so I believe in the Metro area maybe. But I cannot imagine when we will provide 100G per home. It's amazing. But currently we are pushing 100M per home and 1G per home. And the FTTX or the FTTH market in China is also [fairly big] because we have pushed the carriers, we are pushed by the government and also pushed by the users they want more bandwidth.

For next step of 100G I think maybe -- currently we plan to adopt 100G to record the Metro obligation, maybe this [is a sight], maybe in Western you call the CO, central office and [inaudible] we are connected to many [inaudible] or [OT, or BBPO] or some enterprise connected to our building.

And another possible scenario for Metro success for 100G is the remote RU, because currently for 4G [OTE] we [inaudible] supply per [inaudible] is 2.5G for [FTD] and 10G for [TDD-LTE] and people are discussing about the 5G. I don't know what will 5G require but maybe 20, 30, 50 I don't know. But [at a guess] it will be much higher than 400G. So in that case maybe we need some 100G technology to do, and the volume versus [RU] not per [station] [inaudible] at least the 3RU, so maybe this will be in the future.

Manek Dubash

And now the poor hard-put-upon vendor, Andrew.

Andrew Bond-Webster

For us clearly you heard it from KVH, from China Telecom the demand for capacity, the need for them to be able to provide capacity to their users whether you're an individual or whether its enterprise our challenge is to continue to scale our platforms which enable them to deliver those high capacity services to their customers. That's where we've got to be.

Gint Atkinson

Music to our ears. I kind of think specifically the 100G user interface we are thinking that probably it's going to be servers with 100G interfaces in the data centre going up the rack and then into the optical network. That seems to make the most logical sense really where the 100G access interface is going to come.

Manek Dubash

But if you've got 100G to the server then you need to aggregate that in the rack and then you need more.

Gint Atkinson

Right and why not optical right there what's the packet optical architecture and then switching comes into it, service level switching. But I think the -- when we see a lot of 100G interfaces its most likely going to be on servers and data centres.

Manek Dubash

Okay, you heard it here first. Okay, thank you very much gentlemen.

[End]