

NETEVENTS

EMEA PRESS AND SP SUMMIT

Draft

Special Guest Speaker Presentation: Datacentre Interconnection and the Need for Speed

Speakers:

Nicolas Fischbach	Director of Strategy, Architecture & Innovation, Colt
Mike Capuano	VP Corporate Marketing, Infinera

Manek Dubash – Editorial Director, NetEvents

Good morning, NetEvents. Morning. Did you have a good time last night? Yes. Good.

Okay, so to kick off without further ado, I'm delighted to welcome our two keynote speakers, Nicolas Fischbach from Colt and Mick Capuano from Infinera. And we're going to be talking about the need for speed, datacentre interconnect. Come on down.

Mike Capuano

So I guess we're going to square off.

Nicolas Fischbach

Yes, that's what we do in [picture] dynamics. It's me with a tie so don't expect more. I think we can fight, but it's very interesting because we have a long-term relation with your guys and it's more friendship than fighting so it's going to just be like it is.

What I want to talk about a little bit in the context of datacentre interconnect and what Infinera is launching, announcing, is how we at Colt see the datacentre evolve. We've been around for 20 years, as I mentioned the other day. And we own actually 20 datacentres today spread across Europe, which you'll see on the map. And it's more for the guys at the back, but if I point over here, but all of these [three elements] you can see here are datacentres that we have which are spread over Europe. Some cities have one. But other cities, you look at London, they have three datacentres. They're all pretty large. Some of them actually filled up.

And everything you see in green on this picture, with the exception of the links going west to the US, those are all [pulled] by Infinera. So we've been deploying Infinera DTN equipment in 2009, if I'm not mistaken. And this year we started to roll out the DTN-X, which is the next-generation platform. And we started to do it in Germany over here. So the sub rings – what we call sub rings here, which is the fibre that sits in the ground that we own, it's being powered and given up by Infinera DTN-X to deliver 100gig services at pace and not 10gig any more like it was in the past.

And you see that, we go through a lot of evolutions there. At the bottom here I mentioned some of the things we participate in. So in MEF, you heard mentioned yesterday, the ONF, the Open Networking Foundation as well as the [NRV] we're on. So really trying to drive change. When we can pave the way, and sometimes we do it in partnership with those guys. Sometimes it's also [the room], so do not hate me.

So what we see, okay, that's the widescreen format, but basically to focus a little bit on datacentre connectivity both internally and externally. This is the evolution we've seen at Colt. But it applies pretty well to a number of datacentre operators. Do they own a network? Yes or no? Are they pure providers? Do they depend on others? Overall I think that's a picture you see happening across the world. It's not Europe-specific. You'll probably see something very similar in the US, probably more work to those in the US than in Europe in expanding. In Asia-Pac they are more driven by high bandwidth. But this kind of combination reflects really where we are today.

So back in the day, external connect with datacentre was purely internet. There was only client-to-server access on the internet. And everything else, like back-up storage and so on, was allocated high-speed services for some inside the datacentres. I'm talking about the stuff here on the left. What's the campus model? If you all remember, the Cisco campus model, three tiers, [co-edge] access. It was good in the early days because all the [CCIAs] knew exactly how to operate it. The problem was you are [inaudible] for the [calls]. You have the limitations of [VDANS]. I think you've heard the story for a very long time. Most of the connectivity was 1gig and the physical part was 1gig core or sometimes 1gig fibre. That's where we were early 2000s to probably late 2000s, early 2010s.

And then came the model we are in now. I'm saying now mostly to give a reference because plus/minus depending on where you are and who you are. So people who've switched to IP VPN, they use a lot of Ethernet with the growth of carrier Ethernet. The internets going to [get away] because it still needs the consumer to be able to access. You see a lot of dedicated high-speed services. And inside this datacentre you see the first evolution of people moving away from this campus model to this fabric model, you know, [lease pie] model, using an SDN overlay on top of it to address connectivity models, 10gig ports. Most of the people are somewhere here today.

And then that's where we're going. Internet's not going away. Ethernet still pretty strong in the middle space. There's evolution where you need an optical datacentre interconnect because there's much more east/west traffic. When you have more than

one datacentre in a single reach and you need to interconnect them, not just for back-up and restore but also to address this now east/west traffic demand from customers.

Also not between new datacentres but also with carrier [hotels], with the customers' own datacentre that they want to mix probably private cloud or IP cloud models. So that's pretty important. And then inside the datacentre at the bottom, you see this trend of going to a free software-defined datacentre. We're not there yet. [Inaudible] might tell you we are. That's still in the works. There's a lot of transformation going on in storage back-up, compute and so on, connectivity-wise. You're moving to either 40gig or 100gig depending on if you're going to make the leap or not.

And there's much more constraint also on the power and what you can see is technical infrastructure. How do you address this concentration of heat in the airflow system, the number of users? There are physical space constraints and so on. So that's very important. That's getting where Mike is going to explain what they do about it.

And finally that's just a summary picture on how we see all this stack up. I've put the pictures of two datacentres, using this [lease pie] model, which that's the model you have with the compute sitting there. What we have is an integrated layer two/layer three environment that provides connectivity. And what we'll address today is this layer one optical domain network. I've put these links in red or actually orange here. And this is coming. Depending on where you are, Asia-Pac, there's a big trend. I think the US works to those. In Europe, my opinion is that it's just starting. So it's probably the perfect timing for you guys to launch this.

And where we're taking this is we just want to do more than just having physical links. This one is you actually want to automate and orchestrate this. So it's not any more about just having the point-to-point high-value services that are configured once, not being touched any more that cannot evolve. You want to make sure that whatever you put up here can be driven by what sits down here. So the whole integration, is it basic automation using scripts? Is it integration application? Is it SDN model? It really depends on the environment. But the bandwidth flexing requirements and the capacity management, the QS experience up here is super important.

And you know what the people want to do in the end is really drive all of this to build this what we call software-defined datacentres. These are all the building blocks, in our view, that compose it. And what I've done here is just highlight one that people sometimes forget is this datacentre interconnect piece, which, Mike, you're going to address. Over to you.

Mike Capuano

Great. Thanks. Thanks, Nico. Good context. So the title is datacentre interconnection and the need for speed. I'm going to just start at a high level to show you what we mean by east/west traffic.

So this is an example of an internet content provider, Facebook. If any of you happen to use Facebook or your kids use Facebook, you go and you want to see your page, your homepage, you send a 1KB request, ACTP request up to the Facebook datacentre. So that's kind of the user to the cloud. And then once you get inside the

cloud and you look at the way datacentres are built, in this case again by an ICP, there are multiple buildings, either in the metro or in rural environments. And if you look at this example, that one ACTP request turns into 930 server-to-server transactions. And this sort of thing is happening across all kinds of applications, both at these internet content provider datacentres and datacentres like Colt. So when we say east/west traffic, that's what we mean. We mean basically server-to-server transactions, not server to user.

And so this distributed compute model, where you've got multiple buildings. I think you said you have three in London. Facebook, in this example, has multiple buildings in the Bay area. Because of power restrictions they can't connect too many buildings to one substation so they've got to distribute them a few kilometres apart. It creates a lot of this east/west traffic.

Now when you think about datacentres, so that piece is building the building, and that could be in the metro or it can be long haul. That's called datacentre interconnect. And then we've got the rest of the network, which is the core network, which is typically greater than 600 kilometres, city to city. And then we've got the metro network. And you can see that cloud overlays all this. So you've got user to cloud distribution network core and metro and you've got datacentre to datacentre core and metro. Those are the different markets.

In this case I wanted to talk about metro because we've already seen – Infinera typically plays in the long-haul core. We've already seen that move to 100gig, and that started really – 100gig started really ramping in the core around 2012, and it's basically the currency of choice at this point because those long fibres are so precious, 100gig is the key.

So if we look at the metro, what we see is user to datacentre, based on our conversations with customers and based on analyst data, we see that taking off around the end of '15 or 2016, that's 100gig. Because right now in the metro there's generally a lot of fibre and generally 10gig economics are better than 100gig economics. And when that flips we think 100gig will take off in the metro for user to datacentre.

Datacentre to datacentre, there's more bandwidth pressure and we see that market taking off right now. And so on my last slide I'll show you a quick snapshot of a product we're developing just for that application because we think that's happening right now.

This is an analysis done by a researcher firm called ACG, where they actually took a look at the entire optical market and they segmented out just datacentre interconnect. And what they estimate is that by 2019 the metro datacentre interconnect market will be about 3 billion. So that is growing pretty fast. And then the long-haul datacentre interconnect market will be about a billion. So 4 billion by 2019. Depending on which analyst you look at, that can be around 25% or 30% of the overall market. So it's a pretty substantial and very high-growing market, as you can see from the compound annual growth rates on the side.

So then I'm going to switch topics. So that's datacentre interconnect. And in particular again I focused on the metro there. Now I want to switch topics back to the distribution network or the wide-area network from the user to the datacentre. And this is the core and the metro network. So you've all heard a lot about network function virtualisation and SDN. We believe this is a trend that it's not going to happen overnight, but we believe it's here to stay and it's going to be transformational. So we're also introducing another product that plays in this world.

So what we're seeing is we're seeing lots of service providers deploy cloud computer capabilities. They have all joined together and are working towards leveraging that to deliver network function virtualisation. And fundamentally, as network's grown, we've seen a lot of purpose-built appliances be developed and deployed. And that causes scaling problems. So the idea is if I can take these compute resources and run these functions on top of virtual machines in the cloud, I get a much more homogenous and saleable infrastructure and I can quickly move these functions around and deploy them quickly to scale.

So we see in the cloud we've got network function virtualisation and we have SDN control. And what we see at the lower layer of the network, because you can do all this stuff inside the datacentre but you still need to move from building to building. You still need to move from city to city. And to do that you have to create photons. You've got the transport data on a fibre. We see a convergence happening down at that lower layer. We call that intelligent transport. And that's fundamentally taking scalable optics, integrating switching and really just the right amount of packet to do that transport function. So this is not the same complexity as a full-blown router. But this is saying I'm going to take highly scalable optics, I'm going to take the right amount of packet to help them get traffic after it goes through the network function virtualisation cloud, from point A to point B down to the user eventually.

And so we see the higher-scale, most simple network being an NFV-powered SDN-controlled, SDN's talking down and controlling the intelligent transport network with scalable optics and just the right amount of packet. And we believe that long term that is the most efficient network that service providers can deploy. We do believe that it will reduce the amount of router ports out there in the network.

I'll just mention the two products that we introduced, and I talked about it yesterday and I'm talking to some of the other journalists today. On the left-hand side we have something called the cloud express. And this is a fairly different approach to [DWM] networking. Instead of a more telco-oriented refrigerator-sized rack system, it's a 2RU pizza box that delivers 1 terabit, 500gigabits of input and 500 gigabits of output in a rack-and-stack server-like experience for the datacentre.

And then on the right side it's the DTN-X that Nico mentioned. And for that product we're introducing packet capabilities, the right amount of packet to help take the initial step of delivering metro Ethernet forum services directly from the transport layer and preparing for a longer-term transformation to NFV, SDN across the scalable optics.

And that's it. Thanks. And I believe we're going to switch to Q&A now, right?

Interview & Q&A

Camille Mendler – Principal Analyst, Ovum

If anyone's got any questions right now, please put your hands up, otherwise I'm going to ask a few questions now.

Great. Well thank you for the overview. I just wanted to check just so we're all clear on this regarding some of the rationale for the datacentre interconnect. Now you were talking very much about the metro aspect. So how much of the interconnect market is really metro versus more longer-distance interconnect? I'm thinking, of course, has there been an impact of natural disasters in terms of people having – changing the distribution of the datacentres?

Mike Capuano

Yes. What we're seeing is that – so we've actually addressed the long-haul datacentre interconnect market typically with that larger platform, the DTN-X. And that's because 100gig happened first in the long haul because those fibres are very expensive and the economic trigger of 100gig being more economic and higher capacity in the long haul happened first.

What we see now is metro happening, which is more cost-sensitive. If I look at those ACG numbers, they're saying longer term about three-quarters is going to be metro and about one-quarter is going to be long haul. And what we're seeing is we're seeing multiple datacentre buildings in the metro, either all owned by the same provider, like Colt, or multiple datacentres plus co-los or exchanges that need to be quickly connected.

And if I talk, I've had examples from different carriers and internet content providers who say I've got some remote metro or I need to connect from my datacentre to, let's say, an Equinix cloud exchange. I don't have an optical expert out there. I just need them to be able to click, install the box. It's [sell the data]. It's just like a server. If it doesn't work, I pull it out and I put another one in. And so that's the model we're focused on.

Camille Mendler

Okay. That's cool. So within the metro, and thinking about those 20 datacentres that you've got, typically how many connections are we looking at, interconnections? So it's not a one to one, is it?

Nicolas Fischbach

No, no, you can at the start. To start it with [K]. So it's usually this combination of approximately based services where you actually attach those customers using such a platform in the metro for all the [agency] services, especially in the financial sector. So there it's pretty densely meshed, especially with the [stock market].

Camille Mendler

So the financial clients, just so we're clear for the audience, would pay you for those data – so it's interconnect between their datacentre and your datacentre.

Nicolas Fischbach

Yes. And even between them or us and the stock market, the traders. If you look at—

Camille Mendler

And you have a subsidiary, do you not, or a sister company that does ultra-low latency.

Nicolas Fischbach

They're actually integrated into Colt. So we have [got a marketplace] a couple of years ago that [inaudible]. But they're fully integrated into Colt now. So they used to be a separate operation. Now they're integrated because in the process we realised that they need to be part of the core of the business. You can't just keep them as a separate entity on their own.

Camille Mendler

And so for financial services, that ultra-low latency is what they actually – that's what differentiates them.

Nicolas Fischbach

Yes, correct. And at the same time they want to make sure they have an integration with the long haul. So what we discussed I think yesterday also is this integration between DTN-X for long haul and [Code Express] in the metro, because what we also realised as a service provider is we cannot keep running separate platforms, separate, segregated platforms because otherwise we're going to end up with the problem we have now, in the last 20 years, where you have the PDH, SDH, ATN, Ethernet, IP. We don't want to be—

Camille Mendler

[You do] mention frame relay, but that one.

Nicolas Fischbach

Yes, talking right now—

Camille Mendler

I'm sure there's [NX25] somewhere as well.

Nicolas Fischbach

Yes, yes, somewhere. What I'm trying to say is that integration price is key, not just because we don't want to select multiple technology platforms, but also from an

operations point of view, if you're not having seamless deliveries, seamless insurance, you cannot just keep stacking stuff. You have this picture where you show abstraction. Abstraction is nice, but there's only so much you can abstract and you still need the skills.

Camille Mendler

Okay. So you want to simplify. That's one of the benefits. Now let's go to back to the – we talked about financial services. Let's go back to that Facebook use case that you talked about. Was it not up in Sweden I think you were talking about?

Mike Capuano

Right. They have three buildings in Sweden, about 30,000 square metres each.

Camille Mendler

They're all merged together.

Mike Capuano

I was just using that as a–

Camille Mendler

As a hypothetical example if you were in that game.

Mike Capuano

Yes. They all have to be meshed together. And you can quickly do the math. If you have a few hundred thousand square foot–

Camille Mendler

And that's the download speed on your mobile or on your laptop, that improves that.

Mike Capuano

Yes. They get some benefit from being north of the Arctic circle which is they can do air cooling instead of liquid cooling. And their transactions will cross multiple buildings. They've got three buildings up there. That's the metro piece. And then they have to come down to Stockholm, that's the long-haul piece, so you can see your page.

Camille Mendler

But essentially, us as end user, we get a better experience as a result of them investing in datacentre interconnect.

Mike Capuano

That's right.

Camille Mendler

Right. And in terms of simplification, let's talk about simplification. You hinted that potentially simplification clearly in terms of parts of equipment that are going to be virtualised, that we talked about NFV briefly, and then there was a little discussion about routers. So can you talk a little bit about that, router ports and all of that good – what's the impact potentially in your view?

Nicolas Fischbach

Well the impact is the positive change that we see. So we can hit those guys pretty hard on package optical. Really this integration of the optical domain with the IP in the domain because in the core of the network there is no point in having the optical domain and then using routers as overlays when you can just do one core LSP switching or you have some basic–

Camille Mendler

Unless you hadn't depreciated that equipment, then you would have to–

Nicolas Fischbach

Yes, but it's depreciated, but bandwidth demand is growing. Okay, it's depreciated. But at some point, look, every three to five years, depending on your size, on your city, the tier model you have, you have to–

Camille Mendler

So simplify, simplify, throw out the boxes?

Nicolas Fischbach

Not all of it. Not all of it. So what we see at the edge, the intelligence has got to stay at the edge of the network. It doesn't move away. But the core boxes are just interconnected city at the [trend] point, why do you need a layer of three? Just with a layer one and use [contra plane], is it [BGP], is MPLS, is it some SDN? That depends. We haven't decided yet. And it depends on the history you have also. But that's what the change we see that needs to happen.

Camille Mendler

Okay. Now questions from the audience. Any questions from the audience? Yes, sir. At the back. It's Bob.

Bob Emmerson

I'm Bob Emmerson. I'm a freelance writer. At the end of the day, the people who pay for all this are the enterprises, the big enterprises. They get the financial [trust] at the moment. If you were going and talking to a CTO and you didn't want to use – drop down into too much detail, what would be your key messages to the CTO about

saying this is the goodies that are coming up in the next year or two and these are the benefits they're going to bring to your organisation?

Mike Capuano

I'm going to toss that over to Nico. So Infinera primarily focuses on service providers. So we ourselves go and empower service providers, who in turn would be having the conversation with the enterprise. We have a few large enterprises that have unique use cases that use our equipment.

Nicolas Fischbach

He gave me ten seconds to soak up the pressure. That's very good. Thank you. So I think CTOs or CIOs, we talk to I think the [I3] about – what they tell us, they want to concentrate the datacentres, because usually the larger enterprises we address, they have many, many datacentres for historical reasons. They want to simplify it. They need to reduce their cost base. They want to run [hybrid] clouds. So don't touch too much of their core BSS systems but still go with the flow and have a mobile app, some other end-user presence online, they want to start to give part of this to us, but in our datacentres. They want us to interconnect with AWS, with Azure because they want to experience also this part of the public cloud. So that's what we're hearing a lot.

Camille Mendler

So that's the hybrid cloud aspect.

Nicolas Fischbach

That's the [hybrid] cloud aspect with the combination of their own datacentres. But I wouldn't – it's also the way they want to have more applications. Their core BSS application is what gives them nightmares. It's really how do they build this new experience using a deluxe model, agile, [updated world] which is completely different from the historical part.

The other thing we see is that there's this trend of – we didn't want to do managed networking for many years. They've brought a lot of stuff back in house. Now they want [out] again, but also in hybrid model, they want to keep managing part of the environment and the rest they want to give it to a sales provider so you see a lot of that. There's many more discussions again about use of dark fibre. As we said, dark fibre's expensive. We don't want to give it away for free. We'd rather sell a merged service. But we see also this push, cloud express want to push it also.

Camille Mendler

On that point, datacentre-to-datacentre interconnect is a function of availability of dark fibre. And if enterprises are looking for this hybrid environment, that tells us implicitly that more dark fibre is available in Europe's cities. Is that the case? Can you characterise which cities are easier ones for that to be a reality or not?

Nicholas Fischbach

I think all the larger ones. If you pick out all the larger ones, just the investment, look what [Zion] did, all the acquisitions [Zion's] done in Europe. And they are trying to change the open market when it comes to dark fibre. And they're also pushing us, to be honest, to rethink our approach to dark fibre. But I think in all the larger cities, if you exclude fibre to the home, fibre to the premises, I'm thinking of Paris, London, Frankfurt, all those bigger cities, there's a lot of fibre assets in the ground.

And the other thing that's changing this pace is the regulations. So a lot of the European countries now, there's fibre ULL. So we have unbundling of DSL, copper cables. Now you have unbundling of fibre. Switzerland is pretty much active there. France, Paris, in Nantes, in Lyon.

Camille Mendler

And just to be clear, unbundling means that—?

Nicolas Fischbach

That the incumbent has to give you access to the assets they have in the ground.

Camille Mendler

And [docks] and/or fibre at a regulated price.

Nicolas Fischbach

Depending on the country. So fibre ULL is not at the stage yet where DSL is. It's still very fragmented. And every country is at a different stage. And what you see is that it's very regional. So in Paris it can only be [inaudible], you know where La Defense is, not the rest of Paris because that's Orange and you have dense areas and so on. But that's what's going to drive—

Camille Mendler

The market conditions for datacentre-to-datacentre interconnect—

Nicolas Fischbach

Is changing.

Camille Mendler

In Europe are improving.

Nicolas Fischbach

They are improving plus changing because you have a lot of new entrants that introduce more complexity, which is going to be very interesting to see. If you as a customer, a large customer start – they want to talk to one service provider. They don't want to talk to ten service providers. So are we going to see the return of this

mobile – virtual operators that address multiple service providers in the back end, provide delivery and assurance, I don't know.

Camille Mendler

Isn't that the function of the Equinix increasingly, to manage that interconnect in their own way?

Nicolas Fischbach

They do. But honestly I think they have pretty good growth in the US, I think. But in Europe, I'm not too sure. We see them more really on the datacentre to cloud, like using AWS and direct connect and things like that.

Camille Mendler

Any other questions from the audience? Bob, another one.

Bob Emmerson

It's not really another one. How does the announcement we saw yesterday about the carrier Ethernet and the third way of networking, how does that play in this setup?

Mike Capuano

Well one of the things that Infinera has is we are compliant today with MEF 2.0 and so if you look at that packet switching model that goes into the transport layer today, you'll have the transport layer and you'll typically have a layer of routers or switches. And the switches will often deliver an Ethernet private LAN service or a virtual Ethernet private LAN service. With this packet capability into the transport network you can fundamentally eliminate that switch and deliver those private LAN services through those Ethernet private LAN services directly from transport. So that's how it impacts us and that becomes more important as we see the rise of 10 gigabit to the enterprise and higher, because once you get to those kind of capacities then delivering it directly from the transport layer makes a lot of sense.

I'll toss it over to you if you want to comment on the announcement.

Nicolas Fischbach

So we've been a longstanding MEF member, longstanding actual active MEF member, so this talks to us a lot. I think for us carrier Ethernet is I think nearly two-thirds of our data revenue. So for us it is very, very important. To look back at the picture, how are datacentres before, how are datacentres now? We used to treat datacentres as customer premises. A lot of [inaudible], no integration. Nowadays, with carrier Ethernet the transport of choice in the metro, Ethernet services, some of them mentioned are E-LAN, E-line, [E-3], EVPL becoming the norm. Customers want to have the same [expense] into the datacentre. So for us the integration between the datacentre internal architecture and the external one is key.

Camille Mendler

Does that mean you're training your people differently?

Nicolas Fischbach

Yes, so let me come to that. So that's very important, but that's what 2.0, [inaudible], what isn't there is very important to us and I don't want to go into the [NFV] part that I mentioned because I think that's them fighting against the ONF, so that's going to be interesting.

I think people is key. I think I mentioned that yesterday. But if you don't upscale your people, change your operating models at the same time, you're going to fail because there's a dramatic change going on there.

Camille Mendler

But then we go back to the context of datacentre interconnect, really to succeed, you need the kit, you need the equipment that can do the job, which is what you're providing, but you also need some knowledge of the internal guts of the datacentre.

Nicolas Fischbach

Yes. Well it's still two different worlds. You cannot tell somebody who's an optical expert that understands fibre characterisation and how you run such an environment, even though the [DNA] platform is pretty good at abstracting, that's the management system, everything happened below. You're going to turn this guy into a VMware expert. You need those guys in between that understand enough of both worlds that they can connect the dots in case of issues, but you're never going to find somebody who can address from layer zero to layer seven, right. If you find one, [how do you find him] because they–

Camille Mendler

With a very nice telephone-number salary.

Nicolas Fischbach

Okay. Because those guys are unicorns, honestly.

Camille Mendler

Okay. Any other questions or comments from the audience? Okay. Well, guys, thank you very much. But essentially it's a growth market. There's going to be more of it and Europe is a good place. And enterprises will be better served. Yes?

Nicolas Fischbach

Yes.

Mike Capuano

And consumers.

Camille Mendler

And consumers. Thank you.

[End]