

NETEVENTS

APAC PRESS & SERVICE PROVIDER VIP SUMMIT

Draft

Data Centre Interconnect - the perfect solution?

Guest Speaker: Gint Atkinson

**Vice President, Network Strategy & Architecture
KVH Colt**

So I'm going to talk a little bit divergent and off of the slide deck and just use the slide deck really as a use case of what we're doing and our offer and how we've gone from very plain vanilla DCI to, I would like to compare it to Tata and KVH Colt are doing similar to their offer.

We're looking at DCI in a very progressive way as a cloud connectivity service that integrates DC, multiple datacentres, even partner datacentres and puts ideally an entire ecosystem on to that service. So cloud providers on one side, cloud consumers on the other and tons of datacentres all over the world inside that cloud connectivity service.

So the first point, from our experience a couple of years ago, we saw that public cloud is not everything. And in the process of setting strategy and direction, it gets really difficult for a lot of companies to ignore the successes of the public cloud, to ignore the growth rate of Google and Amazon and these other massive players and get back and figure out where do we fit in, what do we offer.

In KVH's experience, we were so focused on niche markets offering high performance, white-gloved service for very demanding industries like high frequency trading that required ultra low latency networking, co-location services at the stock markets themselves. So serving these types of vertically focused industries with very specific and demanding requirements, you end up with very niche services.

This is the complete opposite of public cloud which is meant to serve everybody with a very general requirement that applies to the majority of the market, something that a

small company like KVH can't do and something that I think for a much larger company like Colt is quite challenging to address.

So the point is that public cloud is not everything. It's something and it's even part of the mix for these specialised niche customers that we're serving. They have a need for the public cloud too. So the question is do we have a role in taking them to the public cloud and that's where I say yes.

So we saw that going beyond serving our enterprise customers, who had N number of datacentres -- we've got nine datacentres in APAC, including Singapore, Hong Kong and Japan. And we saw that a few years ago, our customers kept on going to more and more datacentres. Why? Because they had trading partner in those datacentres. So to get closer to their trading partners and to do a cross-connect, they had to take down rack space and put their servers there or get hybrid private cloud environments from us at that datacentre.

We then realised that in the long run, this is not efficient for the customers. It's not in their best interests to go all over the place. We also saw them taking down a lot of public cloud resources. And the way that we saw this was customers were ordering connections from the facilities in our datacentres to get to cloud providers like Amazon, Microsoft, IBM SoftLayer.

And again we saw that this would not scale for the customers. And what the customers needed was a direct connect. They didn't need to be in seven datacentres, they needed to be in two or three.

So we've all heard there's issues with getting at your cloud services via the Internet, that providing private direct connect model connectivity options was the way to go. But still the customers had the problem, they're in too many datacentres, they have too many direct connects.

So how did we solve the problem? What we saw was that instead of offering datacentre connectivity which many hardware vendors and many enterprises felt I've seven datacentres, I need fat pipes between all of them, I need dedicated fat pipes, well, this is still not solving the scalability problem. If they're in one datacentre and they have five trading partners in different datacentres, how do we solve that problem? They should be able to get some kind of a virtual private network service that connects them to all their trading partners.

So what we did was we created DCNet. And with DCNet we're in Asia's top 100 datacentres. We get the cloud providers to come in and if you're in one datacentre that's in DCNet you can get to all of your trading partners, all of the cloud providers. To you as an enterprise it just looks a cross-connect away.

The flipside of that is that for these cloud providers, they could have virtualised pops on DCNet. So they may have one location in Japan or in Asia and if they want to create presence on DCNet, they can create a virtual pop simply by putting their service on to a network instance and putting ports at different DCs on there and then customers would see that for example as an Amazon or a Microsoft or a Google cloud service port.

So the way that we tackled the problem, you know not everyone can put everything into a datacentre, even one datacentre. You still have branch offices. You still have trading partners that aren't in your datacentre. So how do you take care of that? Really easy. You do a service mash-up. You take your DCNet and then you simply hang a tail using our typical Carrier Ethernet service and you can attach your branch office. Or if you want to, you could go straight to Amazon Direct Connect or to SoftLayer.

So you can either come across this cloud connectivity network or you can still get your dedicated connection, different price points of course. Or you can mix these two together. You can blend these two connectivity services and this is where orchestration comes together.

Who's going to do the stitching? That's where lifecycle service orchestration comes in. And that's why I say SDN and NFV are past the leading edge, we're at the leading edge point.

Orchestration is technically -- I don't want to say it's a lot easier, but we just don't have as many complexities to deal with for example that SDN does, legacy network interfaces, all of this other stuff. Orchestration is going to be able to progressively rely on the SDN controllers, on the NFV managers, on the existing element network management systems. So there's a layer of abstraction that's going to allow orchestration to get down to business a lot of faster and with less difficulty than maybe we're seeing with SDN and NFV.

So ideally, this is how it should look. It should be an ecosystem and our enterprise customers come in and they can take down space and power in the datacentre and put their servers there and cross connect into DCNet. Or they can get hybrid private cloud services from us in one of our datacentres, get a cross connect there on DCNet. Once they're on DCNet, they can reach everybody that's on DCNet. Today Amazon, Azure and SoftLayer.

We're also going to be bringing in other cloud providers. I'm not saying these are cloud providers that we're currently partnering with, but good example, the big names, salesforce.com and all of these other applications putting them on to DCNet is going to allow the enterprise customers to jump in and connect to their cloud services.

Now one of the things that we can do beyond just offering this connectivity and where orchestration comes in, DCNet should not really just stop at our physical port that then attaches to Amazon. The customer is getting a VPC and they have a virtual switch inside there and a virtual router. With orchestration, today we call the Amazon APIs to provision a direct connect service and hook it up to the customer's VPC.

But we can do more. We could configure the network environment inside the customer's instance. This would guarantee that the experience is appropriate, that it matches the SLAs of their DCNet service, match all the way up to their virtual router inside Amazon Web Services. And that's just a few API calls away.

You can then take it even further. What if we inventory the services and the price catalogue from Amazon, SoftLayer and everyone else? Then we can really start

doing service brokering and we can look at their network consumption, look at their compute, look at their application utilization and say you know, SoftLayer is having a promotion; you could move X amount of compute over here and you could save on your connectivity and you could save on something else.

So as we put orchestration into the mix with the right kind of cloud connectivity service between DCs and DC players, for the service provider it enables becoming a very new kind of systems integrator when you have orchestration tying everything together.

So just quickly today the way that it looks, very heavily deployed in Tokyo, Osaka, Hong Kong and Singapore. Specifically in Hong Kong and Singapore, we're also building our own metro. So DCNet is going to be running on top of our own managed fibre [plan].

So for the customer experience, the benefit is that with DCNet they don't need to get all these different point-to-point connections from different connectivity providers to get to their cloud service provider. One experience, they take down one port on DCNet and then there's simple logical cross-connect away from their content provider or their cloud service provider. They can literally spin it up in minutes. That's why the orchestration is there.

If you place an order, if you don't have Amazon Direct Connect right now, but you're on DCNet, you go to the portal, you spin it up and minutes later it's done. How is it done? Because the orchestration engine comes in, calls the Amazon APIs, orders it, pays for it and hooks everything up. The hooking up that happens is mapping VLAN tags. Amazon has a VLAN tag for that customer, we have a VLAN tag for the customer's DCNet instance. That's the hook up because Amazon is already on DCNet.

So ready access, stable performance, dramatically reduced cost, usage based charging, increased security. So the usage based charging is interesting. Both etherXEN and DCNet offer a mixture of CIR, committed information rate and excess information rate. So they subscribe to CIR, it might be 10 meg or 100 meg and then they can set their level of EIR that they're willing to burst to. So this is kind of the elastic part. So they can limit their burst and they pay for EIR on a usage basis.

In the future, we're going to be adding scheduling and calendaring. So this is where the MEF standards come in. The way that the scheduling and the calendaring is going to look is all being driven with the evolution of the MEF standard so customers will be able to schedule more bandwidth. They'll be able to schedule changes in class of services. So maybe at night they want the class of service to change, they want a ton more bandwidth going down the cheapest route with a certain amount of latency, high latency, for back-ups. During the trading day for example, they need very little bandwidth, but it needs to be ultra low latency on the fastest route. So this kind of customer experience we're going to be able to schedule it and allow them to add these elastic characteristics on to this service.

So I touched on most of this, the difference between traditional DC and DCNet or a cloud connectivity service like DCNet. The first thing is that everything is wrapped

up with one service provider, one contract. When you need to go to Amazon in the case of Amazon Direct Connect, we take care of all of that if they want to. They can also come the other way. They can order from Amazon and come through to DCNet if they wish. But the big point is one stop shop.

Very fast. If you're in a datacentre, it's as fast as two days. If you're not on DCNet, you need a cross-connect, two days is the typical cross-connect processing time for our datacentre partners. So if you're in that datacentre, two days, everything else is fast. If you're on net we're talking about minutes or hours for things to happen.

So how does this look a little bit architecturally? So we're really looking at doing cloud federation and lifecycle orchestration. When a customer comes in, they see a portal. And underneath we have a lot of different service experiences they can go after. They can come into the KVH private cloud, the [inaudible] cloud, [B Centre] Amazon Web Services, SoftLayer and then Microsoft Azure which I haven't put on here.

And the big area where we're really putting in a lot of work and it is complicated is WAN SDN which we got from Cyan. So we worked closely for years with Cyan to get their SDN controller to truly turn WAN services into a network as a service.

Then we're currently working with [CRS] to do all of the end-to-end service orchestration from inside the datacentre services and all the way across the LAN services. So the big thing that we really have running here is network as a service is DCNet and etherXEN.

DCNet and etherXEN are all one platform, the modular MSP platform, all built on Cyan and this is one platform across EMEA, APAC and North America. So this has been the case for several years already. Colt and KVH have the same platform with an [inaudible] in between.

Now what we're doing is we're looking at trying to do true lifecycle orchestration end to end and with Cyan and with CRS, DCNet and etherXEN are just simply virtual networks sitting on top of the platform. So DCNet is just another modular MSP service that has -- that's restricted in terms of coverage to datacentres, whereas etherXEN is anywhere and everywhere.

That's our last slide, with more animation just to see a little bit of details, a lot of infrastructure that we're using to glue the stuff together. This is kind of an interesting point. We learned a long time ago when we wanted to do SDN in the datacentre and we talked to our potential provisioning vendor and they told us they needed a couple of years to develop the blade and millions of dollars. Did I say blade? Anyway. And Arista said don't worry about it, no problem, it's really easy. It shouldn't be more than a page or two of code if not less.

So we took our smartest software developer who's in the space and he said yeah, give me a week or two. He took a weekend. Everything we wanted that blade to do was done in about a page's worth of code.

So this with things like YANG and NETCONF and the ability to model network resources and the mapping to network services and then network state and actions

with environments like NETCONF, it really gets a lot easier to do a lot of these things. So a lot of the legacy OSS environment from a price point is just -- and timing development timing -- completely irrelevant.

So we don't need open flow for this. And sometimes the interfaces are proprietary, but these proprietary interfaces can be so easy to use. We can use them now and maybe developing a standard will involve a lot of time. So there's a lot of balance. Some of the vendors are developing tools that are consistent or based on abstract standards, but then they end up being a little bit different. You end up writing scripts for different platforms, but these are ten line scripts not a 100 pages.

So a lot of things are happening in the network environment in terms of tools, development capabilities along with standards that are helping reduce resistance between integrating different platforms and products.

[End]