NETEVENTS

EMEA PRESS SPOTLIGHT ON 'THE CLOUD'

First Draft

Debate 6:
SDN – Now Let's Talk Benefits

Ian Keene

Vice President, Gartner

Panellists:	
Marc Latouche	Cisco's Consulting Services, Cisco Systems
Derek Granath	Extreme Networks
Mike Banic	VP Global Marketing, HP
Brian Levy	ONF

Good morning. So - Software Defined Networking, what are the benefits? Let's talk about the benefits.

Well, I think we've been talking about the benefits for three years now and think it's three years ago that Software Defined Networking was the promise of really low-cost networking; white boxes and people were predicting the best of Cisco and Juniper, and the whole networking thing's going to change, but we've moved on from that.

Two years ago every vendor, equipment vendor or networking vendor had to have a set of PowerPoint slides on Software Defined Networking. Now, one year ago that actually had to have a migration strategy for their customers on how to move to Software Defined Network architecture and now I see that really things are starting to get serious. People are testing it out, service providers are testing it out. My main interest area is this carriers' of service provider and public networks and there's a huge amount of real interest generated now. In fact I'd go beyond real interest. Basically people are betting their financial targets on improving the operational efficiency to implementing some sort of Software Defined Networking.

But I think as we move on, the great ideal of Software Defined Networking as it [catches] the very peak of its bated expectations in the Gartner hype-cycle has slipped a bit and we're getting to a more reality now in what it actually can deliver.

I'm a last minute replacement for this talk and it's entitled Software Defined Networking, Let's Talk about the Benefits. I suppose we had better talk about the benefits briefly and that's going to be my second question. But really I would like to make this into a short analysis of really where we are right now and what we can expect in reality can happen.

So to help me with this, I've got Marc Latouche again from Cisco from the previous panel, Derek [Granath] from Extreme Networks, Mike Banic from Hewlett Packard and Brian Levy, well, Open Network Forum, is that the hat you've got on today?

Brian Levy - ONF
Yes and also I'll talk with Juniper too.

Ian Keene
I'm going to start off and before I ask you to very, very briefly tell me about the wonderful benefits of Software Defined Networking, where do you see the biggest opportunity for this in networking? Is it in the data centre; is it in enterprise network, enterprise local area networks? Or do you see it in the carried infrastructure? Where's it going to make the biggest impact? Brian, would you like to start for us?

Brian Levy - ONF
Yes, thank you. Before I start let me define a couple of terms, it's important we understand things first. What we mean by SDN it's like cloud, what does it mean, what actually is it? Can I define another term called NFV, Network Function Virtualisation. I actually think the second thing, it's a horrible acronym, the second thing is actually going to be much more important for the industry than the first one. So there's a radical thing to say.

In Juniper and in the ONF we've defined what we mean by Software Defined Networking in more detail. What we mean is about the separation of the planes of the network, these are the different ways that the network communicates. The forwarding plane that carries the actual data, the control plane that tells the data where to go, the management plane that knows what's going on, and we talk about the forwarding plane still being distributed. The control plane that used to tell the data where to go being partly distributed and partly centralised and then the management plane being centralised. The benefit of SDN is having that global intelligence and global vision over the whole of the infrastructure to make really intelligent decisions.

You've got this map, you have one box, you have the view of the whole infrastructure, you can do some really neat and clever things with that. Network Function Virtualisation is about taking the things that we've learned in the IT industry, this idea of virtualisation which came in many years ago to virtualise computing. Where we had machines that were not carrying very efficient workloads, they weren't fully loaded and we wanted to share all of those resources across the infrastructure. Now we're taking that idea and we're applying that idea to network functions, to routers, to EMEA Press Spotlight on 'The Cloud' firewalls, also to other components which used to be boxes and dedicated appliances in the network.

There's a fundamental change that's occurring as we do that. We're moving from the world of vertical scale where you need the new box to do anything, you need the new box to scale, new box to do new features. To a world where we can create thousands of virtual machines and scale, so if you want to add more capacity we add another virtual machine in parallel within another compute environment. This is a complete change of structure; it's tremendously radical to networks. In a few years networks will be nothing like they are today.

So I also sit on the NFE forum on the [indiscernible] NFE forum and we're looking at this. So, the benefits of SDN are to radically change the economics of networks. But to also radically change the ability of networks to generate revenue, it must be both of those two things and in Juniper and in the ONF we're focused on both of those things.

Ian Keene
So, where do you see the benefit happening? Do you see that therefore in services like the network plainly? Is it a given in data centres?

Brian Levy
SDN began in data centres there is no doubt. Openflow was one of the key protocols that began at the data centre world. We're actually looking at a whole new second generation protocols. We actually see that SDN can be applied in all aspects of the network. I've got applications today that we're putting together with customers actually and they access network for the latest LTA networks. We're looking at applications in the Edge for virtual CPE and we're looking at applications in the data centre for scale. So I think it's gone – my answer to you is it's gone beyond the data centre, it started there but it's radical and it will go right across the whole infrastructure in my view.

Mike Banic - HP
Yes, I agree with you. I think it's going outside the data centre, we at HP agree with you. Those places you describe are in the service provider networks, service provider access, service provider edge, but it's also true in the enterprise network. At each team we're beyond talking about this now, we can actually show you this. So there's applications that we all use everyday, things like Microsoft Lync that the application is actually run on servers in a data centre but we consume that in a campus network, in a branch network or as a mobile user.

We all sort of experience a certain level of what I'd describe as misery with Lync. You're on a call, and at times I'll talk to one of my employees and I'll say, "I can't hear what you're saying, are you on a Lync phone?" It is because the network isn't aligned to the application it's just sort of a road and the road is kind of like my commute to work it's suffering congestion. What I really want is a road that's aligned to the fact that I have to be at work for an 8am meeting. In this case the network's not aligned to the fact that I can't hear Sam when she's speaking to me. But our application that I can show you actually gets information instantaneously from the Lync server when the call's being set up so I can simultaneously programme the enterprise network.

I can use the same application that's running on SDN to do this in a carrier network across those end points and make that scale. So there's things that, to answer your question, Ian,

it's got to be across a whole network because it's got to serve the users consuming the application and it's got to start with the application as being run in a network, whether that's the enterprise network or whether that's the service provider network. Or in some cases an enterprise network that includes service provider connectivity.

Derek Granath – Extreme Network
So I would agree that we'll see it in the data centre, we'll see it in service provider networks and in enterprise networks. In terms of the benefits that it's going to deliver I think it's more in the agility of the network, the speed of provisioning, changing the quality of service that alignment. I'm not sure that I agree with Brian about it decreasing the costs. I can see it on the OpEx side but not necessarily on the CapEx.

As I said in my panel yesterday there were a lot of misconceptions I think six to 12 months ago that SDN was going to commoditise the data forwarding plane switches and I think that's simply not true. To have that user experience we need switches that have buffers, that have big cables that have a robust operating system that's actually running the box itself. So, the centralisation of control you've got to have a bigger controller than what we put in switches.

Brian Levy
Can I answer that? I actually agree that but let me explain what I mean. We're seeing a [dedonic] increase in the volume of bits going through networks. Networks are rising exponentially. That's an equation that just can't be balanced. So what we need to do is to bring down the cost per bits that travel through the network. SDN will allow us to do that in a very, very efficient way.

So, it doesn't mean that the top line, maybe cost of the network is going to go down. But if you look at what's got to happen, if you look at the scale of traffic in the internet, the exponential scale, that just economic value is not there to really carry that scale, so we must make networks much, much more efficient. What SDN will do is allow us to increase network efficiency both at the transmission level and also at the service level, and that is the key point.

Marc Latouche - Cisco
To continue on that maybe I could show some examples. With the customer we've been working on this notion of centralising some of the functions of the super (inaudible) in the control plane. What we have been able to work with that customer is that with the same infrastructure it was capable to support 30% more traffic. So by being able to centralise some functions, some control planes, we are actually able on the same (inaudible) to support more traffic, which was a great case.

I could use another case here, which is this hotel they have Wi-Fi. Probably it has only a couple of tourists that use that and then suddenly there is this forum happening, I'm sure you've all tried the Wi-Fi with some success. Wouldn't it be nice if the hotel could just create on the short virtual nets to authorise to have more connectivity? That today is possible but we haven't got to deploy that yet.

So there's a number of functions and clear benefits that will come from functions. I would agree with the comment that that doesn't mean that everything from the control plane becomes centralised, because you still have specific functions that are best placed close to the data planes. So, that's what we call a collaborative control plane, where some functions will be closer to the data plane, because that's where the control plane provides function for that data plane. And you'll have some functions that will be centralised mainly for aggregation for abstraction and to be passed then to orchestration system that will control over all devices.

Ian Keene
Right, so let me get this right, the equipment's going to cost more, you can carry more bits and there's still going to be proprietary hooks on the local level in switching and rounting equipment? Therefore people will still have to go out and they'll have to buy Cisco, they'll have to buy Juniper, they'll have to buy Hewlett Packard, they'll have to buy Extreme they can't go and buy some white box and run it all off a nice big centralised server?

Brian Levy
No, I think there's a mix in that question, there's a balance in that question. There are some things that are really well done by appliances and by ASICS and it will take years before standard computing can do them. And there's some functions that can be really done well in standard computing and we will see a migration of those functions to standard computing environments.

Now, when you look at it from a Juniper perspective you might say, "Well, isn't that a challenge to your business?" I don't think it is because we're embracing this technology and we're using this technology to build efficiency between the equipment today and the vision of tomorrow. It will take many, many years for us to get there and there are many new value chains that are going to be great in this environment.

So I believe that you can devise – that you can use SDN as a way just as we spoke before, of driving up network efficiency, of creating new value chains for our service providers, for enterprise businesses and those sorts of value chains will mean that they're able to pay for new equipment and value. So it's about doing more with the money you've got, not about doing less, it's about doing more and making more money and revenue in the process of doing that.

Marc Latouche
Actually I would agree with what you said. Thanks to SDN actually a customer could buy from Juniper, from Cisco and from HP. While today you might argue that he has to buy a certain vendor if he wants to use ATI because HP has a (inaudible) for that vendor and that's where we want to deploy openflow in all switches to make sure that actually it could be integrated for instance at the vendors. So we have to be careful that there is a switching industry, the difference is in the implementation. Do we think that there's a [magic] control system centralised or there's still some [lights, portion] control that still happen at the data plane and that's [opposition].

Mike Banic
Yes, that's right if we look at SDN enabling the network to better align to an application and it involves the enterprise network as well as one or more service provider networks, then the equipment's coming from multiple vendors. Having this ability for programmability in a uniform way through a standard mechanism like openflow is what enables that. So it's not going to mean vendor lock-in for the customer and that's really critical. It means that the customer has the ability for a heterogeneous environment as well as flexibility of choice.

But the other thing that is really critical is that there's a standardisation around north bound API so that these applications that people build have a level of portability that doesn't become this necessity of a completely vertically integrated environment. Where it becomes a much more open environment where people can start to develop different sets of applications that deliver network services that can be brought together to enrich the user experience and improve how an application runs across a network so that the business runs smoother. So then we're not talking to the British Airways employee who says, "Pardon me for the delay, my computer's acting up." Because it's not the computer, they're using a web browser that's talking across a wide area network to a data centre and it's any number of elements that could be causing that slowdown. The bottom line is you're not getting the answer you need, you're not being serviced and the application is not aligned to the underlying infrastructure.

SDN allows that to happen, that's really what the vision we have at HP is because we're an IT company. We deliver complete infrastructures for people, the compute the storage, the software that brings it all together. The thing that's really holding things up for most people is the network and then we can take that out of the way, we can actually create operational efficiency. That's where the cost saving comes from, it's improving OpEx.

Ian Keene
That brings me to my next question is in classic networks of today if a router goes down then it's a bit of a problem. Part of the network is okay and okay you might actually have to send an engineer out to go and fix it or replace it, it's a product that can be individually fixed, repaired or replaced. When we have a control plane that's organising the whole network, when that goes down, the whole network is dead.

Everyone shakes their head. Everyone assumes that this technology is going to work fine, it's going to be extremely robust, it's going to offer (inaudible) plus reliability. It really is going to reduce operating expenditure and cost because it's going to be not only highly functional, highly flexible and be able to cope with much higher traffic flows but it's going to be incredibly reliable. But where's the proof of that?

Brian Levy
When we talk about centralisation it doesn't mean having a single point of failure, it doesn't mean having a single thing. So if you look at the SDN architecture, I've looked at others as well, but let me take the Juniper one for example. We use BGP between our controllers, we pass state between the controllers, that's both internal BGP and external BGP. So what that means is we can have within a single autonomous domain within one customer we can have resilience in the controller and we can have two customers communicating together also with resiliency.

We adopted BGP because that is the protocol of the internet, that's how networks of the internet were connected together. So it doesn't mean – don't make the assumption that when we say a centralised controller it's a single controller. And when you talk about resilience, because this is where there's a change in the whole architecture. You used to think about resiliency as box resiliency. You build this box, we'd have another box in parallel, we'd have a very secure architecture between the two boxes, maintain state.

But in the world of SDN and in the world of Network Function Virtualisation we can have thousands of instances of an element that come together to create that element. If one of those instances fail you lose some capacity you don't lose the whole element and you can recreate that instance, so we're in a new world, in a new architecture. That's why I said there's a radical change we're creating things called service chains.

A service chain is a set of virtual machines that are aligned together that create a service. That could be a firewall, that can be an intrusion detection system, that can be a whole lot of network elements, and by the way that can include the compute elements and the application elements of the servers in that service chain. That becomes an overlay to the network and that's controlled by a very resilient controller.

So we are thinking about, as we design this architecture, how we can create the same sort of resilience, even more resiliency than exists in today's networks as we develop this new world of SDN.

Ian Keene
One question I have for all of you is when do you expect we will see a significant tier one service provider running important, not new experimental, but important, common applications over an SDN network?

Mike Banic
Well, if you consider Google a tier one service provider, they're already doing it. They've had an openflow based network for a year or more.

Ian Keene
So you would give that as the best example?

Mike Banic
I think it's the biggest example that exists today. We have some pilots going on in enterprise customers today. They tend to be overlay type applications today because I think people want to dip their toe in the water instead of jumping in with both feet. That's where kind of a hybrid of classic Ethernet and SDN I think will be necessary to accelerate the adoption of it. But those are the types of applications that extreme customers are doing today are overlap applications like network monitoring (inaudible) application.

Derek Granath
We have service providers today that are actually planning to use products that will provide programmability with SDN in just the CPE environment. Where they can automate the provisioning at that part of the network which is the thing they have the most of and if they can get operational efficiency that's where they can actually get the biggest cost reduction, so that's something that we're seeing happening in valuation tests and if there's plan for rollout.

Brian Levy
So for about a year and a half we've had a live customer with many, many enterprises connected to it running virtualised CPE in the UK, I won't say who it is but we've had that going. We're developing new technology that will allow us to change the whole CPE market, that allows us to create – just use very, very simple Ethernet switches as the CPE in enterprises instead of the complex devices that are there today. And to take the intelligence that was in those devices and put them into the actual network. That gives us the ability to deliver all sorts of new services and it creates a new cost model for the service providers, and greater functionality for the enterprise market. That's just one area of SDN that we're looking at, there are many, many others.

Marc Latouche
On the Cisco side we do have clearly today SDN products. The way we approach this is on four elements. The first is in a one platform kit which is a toolkit to help application developers to use programming (inaudible) and APIs to actually interact with Cisco network gears. We have a controller which is supporting openflow that we make available through our open daylight association. Then we have products, the actual switching and a whole lot of functions that could be virtualised, so today customers could run network, virtual routers, virtual firewalls, virtual load balancers which is quite critical when you want to create more end-to-end services in a data centre. To be able from the same operation to create virtual machines attached to a virtual firewall and so forth.

Ian Keene
We're running late but I do need to make some time to see if there are any questions from the floor before you wash up and have your coffee. Do we have any questions, I'm not going to push you?

Questions and Answers

Unidentified Participant
I just understood that there was a big difference between Juniper, HP about BGP. Could you clarify maybe on HP's side, because yesterday I understood that BGP was not any more a big interesting protocol?

Mike Banic
We actually didn't speak about BGP in our session. Actually the things we talked about were mainly focused at enterprise deployments. So I want to make sure that everybody knows that doesn't mean that we at HP think that BGP is no longer relevant. However I think that there is a point of view that's developing in the market about whether openflow – well SDN and openflow as an enabler starts to create opportunities for us to solve problems we've solved with BGP in the past in a different way.

One of the biggest limitations to SDN is our past. What we've done in the past. We've done things in some cases the same way for two or more decades and it's worked, but that doesn't mean it's the way that it always works in the future as it did in the past. Henry Ford, he said, "If I asked people what they needed they would say a faster horse." And he would have never have built the car that the average person could purchase.

Steve Jobs said, "I didn't interview anybody to ask them what they wanted to make the Mac because they didn't know what they needed." And I think we have the same potential limitation with SDN. There's people who will bring creativity and imagination to solve problems in a new way and maybe BGP is right for that. But HP doesn't feel that BGP is so wrong.

Brian Levy
So let me answer this, because there's a bit of a non-(inaudible) there. For example, Juniper fully supports openflow, we were one of the first to support openflow on our switches, openflow 1.0, we're just about, we've got the B2s coming out for 1.3. We were talking about BGP between the controllers as a way of passing state, that's a different thing. That's from controller to controller not between north and south in the network actually controlling the switches.

What we've found is there were different protocols to do different things. We use openflow to actually control flows in the network. But if you want to pass state you want to make sure that one controller is perfectly aligned to another controller. Openflow as a protocol doesn't have the capability to do that today you need something else for the state protocol. And there's a lot of thinking that BGP is a very good way of doing that particular problem, controller to controller, not between the controller and the network which is where openflow goes.

We're actually looking beyond openflow, because we're now looking at some new technologies called proactive network overlays where we think there are other protocols. For example XNPP as another alternative to controlling the actual little virtual routers that go inside the hypervisor where we need the different sort of messaging protocol from the controller to those hypervisors.

If you look at what VM Ware are doing they are also adopting a different protocol for that control between the hypervisor and the controller, because they need a richer protocol than openflow to go and do that. Yet in some cases openflow's absolutely brilliant and gives us the standardisation.

So what we believe is there's going to be new protocols that we're going to need. We talked about the innovation, there's going to be new protocols that we need in the world of SDN and openflow is going to be one of the very important protocols but it's not going to be the only one. If we can find a protocol that exists that will go through the existing networks and give them – I talked about joining existing networks with the new stuff, if we can find ones that will bridge that gap that will join these things together that let's us take a very evolutionary approach, and the best approach to get the best out of both things and that's really important.

So I believe that the use of BGP in the right place in the network is a very good thing to do. I believe there are new protocols like XNPP for example for controlling and passing information to say virtual routers in hypervisors. Net Comp is another very important protocol that we're pushing and other people are pushing for the management protocol because we'll need protocols for actually the traffic to actually create using like VX LAN or MPLS GOE to carry the actual traffic. You need protocols to send the control information that's where openflow sits and where XNPP sits. We need protocols to do the management that's where Net Comp sits and we need protocols to connect the controllers and that's where BGP sits. So we're not talking about the same thing that's important to realise.

Ian Keene
We've run out of time and I'm sure you want coffee. But rest assured SDN is going to save the world.

[End]