RBC Capital Markets Conference
Will Software Take Over the World???
Mike Marcellin, SVP, Strategy & Marketing, Platform Systems, Juniper Networks
December 10, 2013

Operator:
Thank you for standing by. Welcome to the RBC Technology Speaker Series. I would now like to turn the conference over to Mark Sue, Managing Director of RBC. Please go ahead sir.

Mark:
Thank you Melody. Good morning and thank you for joining us. Welcome to another exciting RBC Conference Speaker Series. My name is Mark Sue and I'm one of the data and wireless technology analysts here at RBC. I'm joined by my colleague Jon Atkin who covers Datacenters.

Today's call is part 4 of our ongoing speaker series on software defined networks. A dynamic industry which is in rampant change, not only in the potential benefits, but also the definition itself. So far we've heard from several companies with many more companies to come.

Our featured highlighted company today is Juniper Networks and our featured speaker is Mike Marcellin, Senior Vice President, Platform Systems Division, Juniper Networks. We're also joined by Lisa Hartman, Director of Investor Relations.

There is a slide presentation which will accompany this call and you should have already received a copy, but if you haven't, please send an email to mark.sue@RBCcm.com.

So we've seen the complete resources in the IT industry become virtualized. Now we're seeing that as well in storage. And looking into the traditional networking infrastructure for data centers and for enterprise customers, some of the problems of underutilized assets and high op-ex and inflexibility can potentially be dealt with new technology and new solutions.

And, unlike the past where applications had to get the permission to run through the network, numerous companies are providing flexible solutions to quickly onboard applications as we move to a flexible cloud overlay.

Juniper will discuss today, their approach to software defined networks, the evolving overlay network. Their internal resources to date and also the added benefits of the acquisition of Contrail as we move to a more on demand networking cloud.

After our intro discussion, we will take questions from the audience. So with that I would like to pass the call to Lisa Hartman who will then pass the call over to Mike. Lisa, please go ahead.

Lisa Hartman:
Thanks Mark and good morning. Before we begin I want to refer everyone to our safe harbor also included in the slide deck that Mark referred to. Information, statements and projections

included in this conference call and accompanying presentation slides regarding Juniper Networks' business outlook, economic and market outlook, future financial and operating guidance, and overall future prospects are forward looking statements that involve a number of uncertainties and risks. Actual results could differ materially from those anticipated results in those forward-looking statements. For more information and factors on such risks, please refer to our most recent 10-Q filed with the SEC. All information, statements and projections contained in today's call and accompanying slides speak only as of today. Juniper Networks undertakes no obligation to update the information contained in this call or slides in the event facts or circumstances subsequently change.

With that, I'll turn it over to Mike.

Mike Marcellin:
Thank you Lisa, thanks Mark and good morning everybody. Before I jump into the topic at hand, I thought it might be useful, at lease for some on the call that maybe aren't as familiar with Juniper to spend just a second on slide number 3 which really outlines the businesses that Juniper focuses on.

We really approach two primary customer sectors. Of course Juniper has, for a long time, been leader of the service provider sector, helping to enable the largest service providers around the world – that could include wireline, mobile, cable, even web 2.0 providers. We also focus heavily on the enterprise sector and that can be your global 1000 all the way down to mid-size enterprises.

The businesses that we're in that support both of those sectors are really three – the routing, switching and security. And as we look at how we approach the market, we really want to align with how customers think about their networks. And so they really think about them in various networking domains. You can see the list if you're on the slide. The service provider oriented domains are really Access and Aggregation, EDGE, CORE and Data Center. And then for the enterprises… we also, of course, have data centers that have wide area networks, Campuses, Branches and then the Consumer and Business Devices are really the end points for many of our networking solutions.

So this gives you a flavor of the businesses that we're in. And because we're talking about SDN, we're going to spend a good bit of time today talking about the data center. I do think that SDN is a really compelling way to make the network more agile. And so over time we will see it addressing some of these other networking domains.

But really the first one that it's really started to impact is the data center, so that's where we're going to spend the bulk of our time today.

So if you move to the next slide – slide number 4. I just want to put a frame on what's going on in the market. Mark talked about virtualized compute which has really become a reality, for, but now the majority of workflows in the data centers. So essentially now you don't have to put a single application in on a single server, you can use virtual machines that take better advantage

of your server resources which drives up utilization, improves your cost profile and makes the turning up of applications more agile.

This has led to on-demand expectation that enterprises, and even consumers, have of applications. And the way the market has addressed that is by the building of clouds. As you can see from statistics on this slide, talking about really the growth that we have seen and expect to see, these are Gartner statistics, from the cloud.

Significant revenue expected to be driven over the course of the next four years – applications in the cloud, public cloud services. And interestingly, and I think for this conversation, private cloud. What that means is I'm an enterprise and I like the on-demand nature that a cloud can offer, I may not be willing to put any or all of my applications in a public cloud infrastructure because of security, because of geography or whatever the reason may be, but I want that on-demand capability. I want the agility of resources, so I am investing in a private cloud.

So I'm taking cloud techniques, cloud capabilities and I'm putting them in my own data center. So that's really the growth that we're seeing in private cloud.

So those are really important trends to frame this discussion because server virtualization, even to some extent, storage virtualization, have really become pervasive and have enabled the investment in cloud that we've seen thus far.

But there's been one impediment. And this is an impediment that Juniper's actually been talking about for a couple years now, which is that historical data center network has not had the level of agility that is required to keep up with what's happened on the server and storage side.

So if you move to slide number 5, we really tried to capture it in really two different dimensions. What are the challenges with the data center network up to this point? And this is really what the promise of SDN holds.

So the first one is that the network impedes time to value. So what does that mean? What that means is if you're able to spin up an application in a matter of minutes, which is literally the case from a server perspective – a virtualized server perspective – how long does it take for your network to be able to support that application?

Well today most of the network changes that occur, occur via a change ticket. Where the application or server guy goes to the network guy and says, hey I need new network resources or I need a change in my networking construct to support this new application. And depending on how agile your IT infrastructure is, that is days in the best case. But usually it's more measured in weeks.

So if I can spin up an application on the server in a matter of minutes, and I have to wait weeks for my network, there's obviously a challenge there.

But the other thing that impeding value over time really is meant to address is increasingly because applications are really the heart of the enterprise, and because most enterprises are

looking at a variety of different solutions, whether it's across routing, switching, security or even other capabilities they require, often the environment is multi-vendor, often there are emerging protocols every day that could be SDN related, they could be other.

So a variety of tools and vendors in the data center suggest to, historically, a lot of those tools and vendors have done some things that are pretty proprietary and so making the network and all of its component elements work together takes some time just to get it stood up.

So while you can choose best of breed, which is attractive to many, getting it all working together has been historically a challenge.

And that's one of the promises SDN hopes to solve too if you move in the direction of open standards and we'll talk a little bit about that today.

The other important dimension that we have to consider is value over time. You know, server refreshes have gone from five, six years down to three years, sometimes even two years. We're in the midst right now of important move from one gig to ten gig and that's driving a lot of server refreshes. And then associated network refreshes come along with that as well.

But what has historically happened, and in this world of SDN emerging too, what enterprises are at risk of having happen is they invest in the network and then because of a change in their server, because a change in the applications they're running, because of new protocols that may come from SDN, they have to then throw out what they've invested in and buy a whole new round of networking equipment. So the value over time, the overall total cost of ownership over time, has not been all that attractive.

So those are the challenges fundamentally that I think SDN ideally addresses and that Juniper is really trying to address.

So let me take you through how we do it. On slide number 6, it's important that when we talk about the data center, to really understand that often, in most cases, it's not just a single location. Not just a single data center. Most enterprises have what we call an on-premises data center and they probably have multiple ones. Certainly for any enterprise that might be national or global in scope, they're going to have multiple on-premises data centers.

But it's also likely that they have… maybe they're using a managed services provider for some services. Or maybe they outsourced some of their infrastructure to a hosting provider which would be their infrastructure sitting in someone else's location.

And then, of course, increasingly as we saw in the Gartner statistics, more and more applications themselves are moving to the cloud. And so they're putting applications out in someone else's environment.

So when we talk about the data center and solving challenges of the data center, it's really solving the challenges of how do I connect all of my sites and they could be my own sites or places where I put my applications, how do I connect them, how do I secure them, how do I

decide what goes in virtual resources or physical resources. And that really is the entirety of the problem set.

And so back in late October Juniper introduced an architecture that we call metafabric. And that's on slide number 7. And simply put on this single slide, the goal of metafabric is to drive those connections across data centers, optimize within data centers, secure that entire profile and all done with some important architectural principals around being simple. And simplicity drives operational savings, openness, which drives investment protection, and drives inoperability, and being smart.

And we'll talk about why that's so important because as we start to introduce SDN into the network, we need to make sure we still understand the applications that are running on the network and what the physical infrastructure is doing so that we can optimize that network.

So if you go to the next slide and this kind of eludes back to something that Mark mentioned at the kick off of the call. This is not a major right turn, I would say, for Juniper. If you think about our journey in the data center, summer of 2008 when we shipped our first switch, and of course at that time, we already had routing and security capabilities.

So starting in 2008 we really had a full networking portfolio in the data center. Since that time, we have invested in a common management platform to drive automation in the data center to drive improved operational simplicity. We also have introduced switching fabrics in the data center, first with our EX portfolio and then with Qfabric back in 2011. Over that time we've also been building a set of eco system partners and, as Mark eluded to, and I'll talk further about this as well, we acquired a company called Contrail late last year and began shipping our SDN controller called Contrail, just a few months ago.

And so over that time, over this journey, we've built up, I would say, a pretty important business for us – switching, routing and security. A broad base of customers and broadening. And really, as we looked across the industries that we served, it really runs the gamut.

And I'll talk a little bit as we go through here what industries, what types of customers are looking at different capabilities around SDN and who are maybe some of the early adopters that we're seeing.

So if we move so slide number 9 and I'm moving quickly, but I want to make sure we have a lot of time for your questions. I mentioned three architectural principals and I'll read these as we go through the solution in how we address them.

But fundamentally they are simplicity, openness and being smarter or providing analytics and visibility into what's going on to the network. And if we can do these things, it's going to lead to a more agile infrastructure for our customers, faster time to value, a varied network resources, and Juniper's really been known for and will continue to be known for investment protection with our solutions.

So I'll take you through real quickly the components of the solution and you can understand how it all fits together.

The first one I want to talk about on slide number 10 is the switching architecture. That's really kind of the life blood of the data center network. It connects servers and storage to each other. It connects them to typically an edge router on the data center that then will take that network out to the wide area network or to the end user.

And we have made great strides in simplifying and driving high performance in the switching network in the data center.

So just to orient you to this slide. We've done something that I think is far and away the most compelling modular and best investment protection story in the industry around the switching architecture.

You can see in the middle of the slide the QFX5100, that's the new topography switch that we've developed that can serve in all of these different architectures. Many of you may be familiar with QFabric, this is really Junipers lead foray and we really led the industry in the development of fabrics in the data center. Fabrics give you any to any connectivity, give you extremely low latency across the data center and allow you to manage what could be a complex underlying switch architecture as a single logical switch.

So the QFX5100 can be an end point on QFabric. But it also can serve in much smaller configurations which is really important new news for Juniper. QFabric solves a very complex problem for very large data centers. And so we've seen it be successful in Greenfield situations for some of the largest data centers around the world.

But what we're now able to do is to take some of that learning and some of that technology and bring it down to two switches, four switches, 10, 20 switches and you can build your way up to a QFabric, or not. And the QFX5100 is that common building block that can work in all those configurations. Again, extremely compelling investment protection story.

The QFX5100 also has the rich features head of Junos, layer 2 and layer 3. It's got the leading resiliency story with in-service software upgrades that are hitless, which again leads the industry.

So we're really excited about what this can do for our switching business. It's shipping this quarter and we'll bring out additional models as we get into 2014.

So that's really the quick run through of our approach to simplifying the switching architecture.

But if you go to slide number 11, what's equally exciting for us is we're able to bring our MX3D to full position of prominence in the data center. And in this world of SDN, what is happening as people are starting to adopt SDN is, and it might be working with VMware in a given data center.

And that being for some time and they're going down the path of VMware, SDN solution. They might be doing open stack in another data center. Maybe their e-commerce data center wants to utilize open stack, so they've got that sitting somewhere. And maybe they're using openflow in another data center.

So what's happening is they're trialing SDN, they're going down different approaches for a variety of reasons and islands of SDN are popping up. And the reason why that's happening is because different protocols can't necessarily communicate with each other.

So what the NX can do for us, and this really is founded by Junos and our trio silicon kit set is really connect these islands of SDN. We utilize open state as protocols to do that and they can do all the encapsulation and decapsulation between these areas.

So whether you have some bare metal resources and some virtualized resources, whether they're using NSX, Contrail, Open Stack, whatever it may be the NX can connect that and it can do it in a high performance way.

So now the promise of SDN spanning all those different data centers that I mentioned really becomes real.

If you look at slide number 12 an important part of SDN is really around automation analytics of the infrastructure. That's both the network infrastructure, physical and virtual, as well as having visibility into applications and what's running on the network.

And our approach there is with network director. Network director is actually our solution that spans the entire enterprise. So it's not just for data centers it's also for your campus environment, your wireless LAN environment. But in the context of the data center, it allows you to visualize the network infrastructure giving you really easy view to go in and manage that and for the first time ever, really manage the inside of the fabric itself. So in a Qfabric or in a virtual chassis fabric.

But it also allows you to do some pretty rich analytics to make you and your network much smarter. And so when I talk a little bit about Contrail, I'll talk about this notion of an overlay network and an underlay network.

But one of the risks you have with an overlay solution is you lose visibility into what's going on in the network. So if you need to troubleshoot, if you need to optimize – if you don't have this level of visibility and the analytics to back it up then you could have a challenge. So that's what network director brings to us.

And also, again, allows for automation which really drives op-ex savings for anyone managing the data center.

Just a couple more slides real quick to get to the entirety of the solution. Security is a critical element here and as you saw with all the different resources that are now required, you've got some virtual machines and you need security down into the hypervisor, you've got certainly a

physical network that needs to be secured. You've got a cloud infrastructure where you need to ensure that connection, especially if it's with public cloud resources, need to be secured.

So we take a very holistic approach to securing the data center environment with metafabric. Obviously performance is key there as more and more gets put into the data center, but giving you a simple platform to manage and a simple tool using Juno space to manage your security profile is critical.

And we've got some really unique approaches to managing intruders in the data center that we've said are intrusion deception which really puts a different spin on managing it. We don't have time to go into it today, but we certainly are taking a different approach to this to really both fend off attackers in a different way and then also provide visibility to those attackers via fingerprinting them and exposing that to our broad customer base.

And then the final part of the solution is around Contrail. And let me break down our approach to this in two ways. As we think about the entire data center solution, more and more of their customers especially in the service provider, cloud provider space, as well as very large enterprises who are building private clouds that are really looking at open stack as their orchestration answer. That's because it is open and because it's agile, it's been built for the cloud.

So Contrail is an ideal SDN controller in an open stack environment. And I think that's the sweet spot of this solution. And it can provide true network virtualization so your network now acts as a common pool or resources just as your service and storage do. We do have broad hypervisor support across KVM, Xen and ESX coming in the first half. And ESX is the VMware hypervisor.

And, as you may have seen, when we made Contrail generally available, we also announced OpenContrail which is the open source version of Contrail allowing the developer community to take advantage of it, allowing our customers to bring it into their labs, to test it out, to write add on code to it. And so Contrail really becomes the way that the network becomes as agile as the cloud has been.

Now the other part of our strategy is that there are a number of enterprises that are VMware shops, VMware users. They have known VMware for many years. VMware has supported them and has grown with them. And VMware now, with their ESX solution, has a compelling SDN network virtualization solution.

And so for those customers that are VMware users, we want to make sure that we have the best networking solution to work with VMware. So we've invested pretty heavily in that partnership with both our switching and routing and security portfolio to make sure that we have the right connectors and APIs in the VMware, that we have the right analytics that can support a VMware environment.

And one of the things that we can help VMware with is really bridging between the virtual resources where they're still heavily embedded and the physical resources. And we can do that with both our switches and our routers.

So we've got really two primary architectural approaches depending on the orchestration and hypervisor approach that the customer is using.

So slide number 15 really ties it all together. It's a little better if we could do the bills. But just the net of it is, as applications get launched, you don't want to be constrained by the network. So with this solution, upon the launch of an application, you an automatically determine the right policies for that application, it's security policies, it's how the network needs to respond, it's what QOS you need across the network.

And then you automate the provisioning of those policies, you can provide analytics all the way down to the physical network along with the virtual network, back into network director. And then if applications need to move to a different virtual location or even a physical location, you can do that in an automated way.

So that's really the promise of having the application be the centerpiece of what's going on in the data center.

So with that, the last slide is just if you want to learn more… when we announced Metafabric back on October 29, we walked through some of this information in more detail, but we also had a customer panel where customers like Northwestern Mutual, Tesla, IBM and Bloomberg talked about what they're looking for in the data center, talked about how a solution like Metafabric really delivers what they're looking for.

So if you want to hear from some customers and really get a sense of the use case it would be worthwhile, 30 minutes or so, watching that webcast at this place.

So with that, I think Mark I'll turn it back over to you. I'll take a breadth and then we can open it up.

Mark:
Okay great. Thank you Mike. So if we look at the evolving world of SDNs we're seeing many different approaches – overlays, underlays, hybrids, management orchestration – each one has their pros and cons. Overlays for example you can lose some visibility. And at the other end, you want to allow for increased automation.

You have a lot of assets, so the question becomes which way is the world moving as we build SDNs? Or rather, is it really about many islands of SDNs and trying to piece it all together. What are your customers saying in terms of their approach?

Mike:
That's really the operative question. I think in the grand scheme of things, we're pretty early in this journey. I mean we've been talking about SDN and Juniper was actually one of the original

contributors to the openflow spec very early on because we realized that the network needed to be more agile. It needed to be more open.

So I think this is something that the industry has been talking about for some time, but from the standpoint of adoption, we're pretty early on.

And part of the reason is what you just said which is, there are a variety of different approaches. So before taking the plunge most of our customers are trying to just get information. Virtually all of our customer conversations nowadays, SDN is a topic. And some of them that are further along are doing trials or even doing some implementations.

So I think the answer to your question really, at this point, varies depending on the type of customer. You certainly have a pretty small handful of very large companies for whom the data center is their business. So maybe we'll call them the 2.0 properties who have big staffs of software developers, big staffs of people running their data centers, and they staff it up with application developers and they can control a lot of resources at this.

So I call them the master chefs. You give them a few sharp knives and you give them a light recipe book and they can create some things with that.

But at this point I think most of the enterprises that we talked to are probably looking for a little bit more of a prescriptive approach. And so really the operative question to ask them is what are the biggest pain points that you have? Is it the deployment of new applications? Is it scaling up and scaling down at a moments notice? For some businesses that is the reality and for some it's actually not.

Is it you're trying to better manage your op-ex because you want to put resources in another part of your business to address your core competencies or your core business. So really honing in on that will help us understand the best approach.

And, I think the other important question is, what have you invested in thus far? Maybe it's got nothing to do with SDN, but unless you're building a Greenfield data center you have a network, you have probably virtualized at least some of your server infrastructure, you have some tools to manage orchestration, maybe you're using some automation, maybe you're not. So you've made some investments and so it's about how you get from here to there.

And that's where… we talk about the metafabric architecture, but really it's not a monoless in the sense that we can come in and, you know what, if you're going from 1 gig to 10 gig servers, you're ready for refreshing the servers. You will need a network refresh there so that's a good opportunity to talk about what's the right switching architecture that's going to give you investment protection, that's going to give you the level of agility that you need in an SDN world.

If you are building another data center and maybe you've acquired a company you need to bridge islands of data centers or islands of SDN, we can talk to you about how the MX can help you do that.

So we've taken, I think, a pretty broad approach all aligned to the principals of simplicity and openness, but providing a few different options for our customers. And it's a consultation at this point because we're so early that it's so important.

Mark:
That's helpful. To your last example, Mike, for example, if there's a large data center let's call it Data Center A, using one SDN method, or solution, and another data center – Data Center B – which is using a different solution, what would be the challenge in terms of kind of connecting these islands of SDNs? Or is there a translation there that occurs? And is that really what Contrail will work and openstack open type of environment where the translation does occur so it can be seen as one larger SDN?

Mike:
Good question. Basically networks are built using communications protocols, right? And with the growth of SDN, there've been a variety of different approaches to that. You certainly have things like openflow, you have thinks like VXland, which is what VMware has developed, you have OBSDB which is one of the open communications protocols.

You also have a set of, I would call them, tried and true protocols that have been the foundation of the growth of IP networking over the past 15 years – things like BGP. So you've got a variety of different ways to communicate.

And depending on how those protocols are used and, in some cases, some of the early versions of these protocols have been built with proprietary extensions so they're not interoperable even though at some point over time they might be if they become more standardized across the board.

And so if you use one approach in one data center and another approach in another data center or it can even happen within a data center you can have multiple approaches with different PODS or different areas where you're using different applications, you need to communicate between those two.

So, technically, what's required is you need to de-encapsulate that protocol. You take that protocol and translate it and then re-encapsulate it and send it to its next destination. So there is that translation layer that has to happen.

So that's really where the MX can shine because not only does it have the level of performance that's required for high capacity data centers, but we built the trio chip set a few years ago with this type of need for programmability in mind.

So what that means is unlike most other silican that you would invest in where if you need a new protocol you're probably going to have to get rid of that device and bring in a new one. We can actually program the trio silican using Junos to bring on new protocols.

That's how we've been able to bring openflow under our switches and routers as quickly as we have. That's why we've been able to build the VMware oriented protocols into our switches and routers. And so that's absolutely essential.

Now over time, five years from now, once a given enterprise has maybe said this is my answer for SDN, then maybe those islands start to dissipate. Although we also may find that different approaches are going to be better for different types of workloads and applications.

So it may actually settle out to be two or three approaches in which case you will need that translation.

And so the last piece of the answer to that… so what I talked about is what's happening actually at the network layer and that type of translation and connection has to happen. But then you mentioned Contrail and that's a great point because on top of that how do you control and virtualize this pool of resources in a way that abstracts the complexity from the customer, the end user of the data center?

And that's really where Contrail can come into play. You still need to have the underlying network. It has to be connected, it has to be able to talk to each other, but the way that Contrail then allows you to virtualize that infrastructure, segment your resources against different applications, against different departments, against different workflows and manage that in a way that the end user can use a simple interface and doesn't have to worry about what's going on under the covers, that's an important way to bridge that as well.

Again you still need the network connected and working across these different islands, but Contrail can make it much simpler to manage and orchestrate.

Mark:
That's helpful Mike. And just on connectivity, since there's many ways to do that, for example, proprietary ways to do it and then there are open and new ways to do it. Is there some efficiency and performance that is considered from an openwave versus traditional routing protocol handshakes and those type of things? Or is that coming up on the discussion for customers as well?

Mike:
Well I think proprietary protocols were written to try to address a problem. It tried to address it in a certain way. It tried to address it as efficiently as possible. So an open protocol isn't inherently better from a performance perspective, we believe it's better from a long term evolution perspective because increasingly you've got multiple vendors in the data center environment, you are seeing rapid evolution of SDN and just data center technology in general.

So if you are more open and more interoperable, that can ensure that you can keep up with the changes that are happening.

But I'll give you one example. We were just talking about connecting these islands of SDN. So Juniper's approach there is to use technology called EVPN – Ethernet VPN – it's NPL

encapsulation protocol that is a standard. Now there's another market leader in data center switching uses a couple proprietary protocols to connect data centers that they've created – OTV and Lisp.

In this case, we actually believe that EVPN is a better answer. Forget about it being open for a second, it's just a better technical answer. But it also happens to be open. So what that means is that anyone that sues EVPN or adopts EVPN over time can operate with it and so customers not locked into a single vendor and, in this case, they get a superior technology as well. And that technology can allow you to bridge those islands of SDN, just as one example.

Mark:
That's helpful Mike. If I look at it from the solutions that are being offered now, so I look at the Juniper Qfabric, the metafabric, and Contrail and also the ones that you've done with NPLS encapsulation of the EVPN, and I compare and contrast that to maybe a Nexis 9000 and also Cisco's upcoming controller, how would a customer decide or is it really about leveraging what's installed? What solution will work better in one instance and which solution might work better in another instance? Or is it really head to head? How should investors think about that?

Mike:
That's a great question. The market is moving fast. I walked through the slide that talked about Juniper's march in the data center. So we've built up over time. We've taken market share in switching and routing and security space at the data center. But we're still, certainly from a switching perspective, a relatively distant number 3. And Cisco is clearly the leader there.

And so as we think about… I walked through our architectural principals, I think what a customer is going to have to do is look at the track record of the vendors and how they approach solving problems, what their approach is to open standards.

Frankly, some customers will say, you know what, I don't need it to be open, this vendor or that vendor will take care of me and they'll go with their solution. But I think as we saw some of the announcements that Cisco had made last month, there are definitely some important differences in architectural approach with our solution.

Contrail – our SDN controller – will work with anyone's hardware. That wasn't what I had heard from the announcement that Cisco had made. Our hardware is built to have openflow, to have OBSDV, to have the Xland and to work with any standard space controller. That's not what I heard from the announcement last month.

So I guess the broader market question is, as things are rapidly changing in the data center, will we stick with the status quo or is there enough opportunity and enough benefit to looking at alternative solutions?

We've seen that, I think, somewhat incrementally over time. And we will continue… we can talk all day long about our hardware versus anyone else's hardware and why I think ours is better. So we'll continue with that, but as you can see as we talked about with metafrabric, we've taken a pretty broad look at this problem – not just at a switching level, not just at an FDN

controller level, but also around security, around connecting data centers, around orchestration and management.

So we've got a pretty, I'd say, complete solution and the newest part of it – Contrail – allows us to go in in the least disruptive way possible because it does work with anyone else's hardware as long as they speak BGP, which pretty much any box does nowadays and allows us to come in and help solve some of the immediate pain points of needing more agility and more control over the network – whoever's network it might be.

And then, over time, we think that puts us in an even better position as a hardware refresh might happen or other decisions made in the data center.

So I really feel good about our approach. But look, it's a very competitive market and the best thing that we can do is just talk about how we're approaching these problems. Talk about our guiding principals and then back it up with best products on the market.

Mark:
That's helpful. Mike, I'll ask one more question and then we'll take a question from the audience. So if a customer comes to you and says, these are my pain points, these are the things that I want to do. I want to gets apps on board faster, I want to save money on cap ex, I want to save money on op ex.

In your vision of this open SDN network and the solutions that you provide, what are the range of benefits that I might get if I want to move to a Juniper solution versus an existing solution, for example? What are some of the numbers, a range of numbers, in terms of what I can actually save? Not only in terms of the equipment costs, but the operating costs as well?

Mike:
I think by far the biggest benefits are going to be on the operating side. I mean clearly depending on… we talked about architectural simplicity from a switching perspective and as we look at some of our fabric solutions, you can probably cut the number of devices that you may need in the data center moving from a more complex tiered architecture. You might be able to cut the number of devices 25% or so by moving to a Juniper solution. But the real benefits, the most material benefits, are going to be on the op ex side.

And we've done a lot of things… the low hanging fruit is often automation. I mean automation… SDN can help with automation, but automation is not a subset of SDN. I mean automation in it's own right can yield significant benefits.

So the first place that we often start is we'll talk about the future vision of your data center, but what can you automate right now? Because it's going to make things move more quickly, it's going to decrease errors, and we've done some nice things integrating with some of the server and application automation tools or having plug ins for things like Puppet and Chef which are automation tools that assist admins use.

Having those plus-ins on our hardware so they can orchestrate and automate provisioning of the network equipment as well. So there is low hanging fruit like that. When you start thinking about what a Contrail might do for you, it's pretty significant.

I mean we talked about the fact that you can spit out applications in minutes on a server. If Contrail was implemented, you can provision the network resources in those same timescales.

So my example of… I think openstack, the stack that they quote is 45 days to provision the network for applications. Now I think some enterprises have done better than that. Made themselves more efficient, or if it's a small change, it can be more quickly.

But let's just say a week. Let's say it's one week rather than 45 days which is what openstack says. So it's 40 hours, let's say, and you move that down to an hour, that's less 15 minutes, but let's just say an hour, then you suddenly have 97.5% improvement. So significant. And so you can be much more agile.

And frankly some enterprises, that's nice, but they don't need to be that much more agile. They much more rather cut down on user errors and so let's go with automation. But I'll tell you, for some enterprises who are doing business in web time who have a lot of their revenue associated with web applications and being able to turn those up very quickly or make in-service changes to bring the latest capabilities onto their web properties, that's extremely attractive to have a network be as agile as the application.

Mark:
That's very helpful. Melody, do we have any questions in the audience.

Melody:
Ladies and gentlemen, if you would like to register a question, please press the 1 followed by the 4 on your telephone. Mark, there are no questions at this time on the phones.

Mark:
Okay, thank you Melody. Mike, as we wrap up, maybe one quick one on scale. What are some of the scale requirements that your customers are talking about? And how should we think about scale in terms of number of new technology for your matrix can actually control? Is it right now in the hundreds of or thousands or potentially higher?

Mike:
That's important and a probably a few dimensions of scale. There's certainly scale in number of network nodes, number of server racks or whatever it is that you're connecting. So we really led the industry as we announced Qfabric a couple years ago and we continue to maintain high levels of scale.

In fact, what we're seeing from some of the largest customers that are looking for significant scale even beyond that are looking at actually layer 3 routed data centers. So their levels of scale are taking them to a point where they're building their data centers like they're building their

wide-area networks or even like the internet is being built. So Juniper knows scale and I feel good about our ability there.

But the other way to think about it is more scaling out and virtualization of network resources and even of the functions that a network performs are virtualizing those, you can scale out across many different user domains, many application domains, many different physical sites as well. We talked about connecting data centers to each others, connecting data centers to clouds.

And that's fundamentally what Metafiber is about. I don't want to leave the impression that it's only about the most massive data centers and the companies with 100 data centers around the world. Because we really think that if this is done… if Metafiber has done anything for us, it's actually supported that level of scale, but it's also allowed us to broaden our market, down market, as well where Qfabric solved a somewhat high end problem, now it's brought our fabric technologies down to the individual switch so that literally two switches connected together can take advantage of them. We've actually addressed a pretty broad array of the market.

Those smaller guys maybe aren't doing much with SDN today because the benefits don't accrue to them as well, but they're still concerned about their applications performing at the level they need. They're concerned about security. And they're concerned about maybe how do their owned assets work with their assets in the cloud.

So still some challenges for the smaller guys to solve. And we'll always be there for the big guys to support their eye-popping scale numbers and making sure that environment can be controlled and be as agile as they need it to be.

Mark:
That's helpful. This is a very exciting time and it's a fascinating discussion. Thank you Mike and thank you Lisa and thanks to our audience for joining us. And please join us next time for our ongoing discussions on the fast evolving world of SDNs. Have a great day everyone.

Mike:
Thank you Mark.

Melody:
Ladies and gentlemen that does conclude the conference call for today. We thank you for your participation.