Revolution on the Network Edge

THE QUANTUMFLOW PROCESSOR

VALERIE ST JOHN: From the beautiful industry briefing center here at Cisco, welcome to TechWise TV. I'm Valerie St John along with Cisco Solutions experts Robb Boyd and Jimmy Ray Purser. Hey guys. Our focus today, the ASR 1000, revolution on the network edge. And by the year 2011, global IP traffic is expected to reach 29 Exabytes per month. That's the equivalent, get this, of 144 times the entire world's printed matter, six times every word ever spoken, or 7.25 billion DVDs streaming online. It equates to more than 1,100 times the amount of information supported by the US internet backbone in 2000, that's less than a decade ago, guys, a lot of IP traffic. Here's a softball question for you, Robb. We've been getting along just fine, why do we need the ASR 1000?

ROBB BOYD: Well, I think you mentioned a lot of those same points in your open. The idea here is, obviously, things continue to change. I think the number one thing driving usage on the internet right now, quite frankly, is video and we see that proven over and over and over again. So you need something that's not just about bandwidth though, it's about how do we serve these things up in a really reliable consistent manner because so many companies are building their entire presence connected to the communications and stuff that they can't afford for these things to not be available or not operate at the speeds and with the reliability that they have come to expect. What's really interesting about the box, the ASR 1000 and the ingenuity that went into it and that we're going to dive into, I was just blown away by this. You mentioned some wonderful stats, I've got one for you as well.

VALERIE ST JOHN: Okay.

ROBB BOYD: The 7200 router, which is installed, $10 billion in sales Cisco has made with those boxes with customers. It's the router right now for edge connectivity. This box is the equivalent to 160 of these routers.

VALERIE ST JOHN: Whoa, not shabby.

ROBB BOYD: This is a huge change, but it's not just about speed. There's more, but we'll continue to get into that.

VALERIE ST JOHN: Okay, so it's got greater capacity, obviously. It's greener, uses less power, saves you money, right?

ROBB BOYD: Absolutely, I'm glad you mentioned green as well because the idea's how can we continue to not only do things more but we can't afford for energy use to scale up exponentially. This box actually has won awards. In fact, just recently at Interop it won award from Miercom, being tested as being much more efficient than even many of its predecessors as well. So we're going to continue to see those kind of trends and should expect to see those kind of trends out of all of our network gear.

VALERIE ST JOHN: Excellent, then who better to bring into this discussion than the father of the ASR 1000, Michael Beesley.

MICHAEL BEESLEY: Hi Valerie, thank you for having me.

VALERIE ST JOHN: Good to see you.

ROBB BOYD: Absolutely.

VALERIE ST JOHN: So I feel kind of like I'm in the presence of royalty because you are one of the first engineers to join Cisco Systems. You left in '97 to help found Juniper Networks. But that's okay. You got brought back into the fold by an acquisition back in 2004. We're very happy to have you back.

MICHAEL BEESLEY: That's right, thank you.

VALERIE ST JOHN: So you saw a need and you worked to fulfill it, but tell us what was wrong with the old routers. Why did routers need to change?

MICHAEL BEESLEY: I think really it's a combination of a changing set of requirements as the networks have become more mission critical, they've become faster, the customers required more services and features to be applied to the traffic. That combined with the fact that the previous generation of mid-range and edge routers had been built quite a while ago. Their architectures were a little old, their technologies were a little old. So certainly using more modern software techniques, software architectures and more modern silicon processes, you could revamp the mid-range router, the edge router to satisfy the ongoing needs of mission critical network in terms of speed, features and functionality as well as reliability.

ROBB BOYD: Let me ask you a question on this because you make some points about how those things change. But you were also quoted as saying that the ASR 1000 is a radical -- I don't know if that was actually your word or not because who knows who may have attributed that one to you. But nonetheless, the idea was that it was a radical departure from the way routing has been done in the past. So as we start getting into what this box can do, radical, that's kind of a big term to use and throw around, how would you back that up? What do you mean by radical?

MICHAEL BEESLEY: I don't know whether I'd go quite as far as radical, but certainly some very, very significant changes in the approach to the design of these types of systems are embodied in the ASR 1000. I think the two main aspects would be first the data plane. We worked for many years to produce the Quantum Flow Processor, which is really a next generation network processor that allows the data plane of the box

to be very, very fast, very, very flexible. You can add services and features by just doing software downloads and instant on for many different services and features. So that's the first thing that's significantly different in the ASR 1000.

ROBB BOYD: So the Quantum Flow Processor.

MICHAEL BEESLEY: The Quantum Flow Processor. I think the second aspect is actually the software architecture. In my mind, the ASR 1000 is the first mid-range edge router to embody a full carrier class, modular, redundant software architecture where we continue to leverage and use IOS and its functionality but IOS is an application running on a carrier class operating system with redundancy and scale and performance that really hasn't been found in this class of product previously.

ROBB BOYD: And so you had to reinvent a lot of these things. And that's kind of one of the key points. This is a ground up reinvention of how things could be done. You helped us invent a lot of innovations in routers back in the early days. You then took some of that knowledge to another company, we won't mention their name again because why the unnecessary press. But you broke out again because you've got this bug to continue creation. The first thing you mentioned, however, was the Quantum Flow Processor and that's what I wanted to focus on. Jimmy Ray's lighting up like a Christmas Tree over here. Walk us through what that means because, first of all, it blew my mind and Jimmy Ray, you've actually been talking about this for a while.

JIMMY RAY PURSER: Yes, quite a while.

ROBB BOYD: Because this router's been in customer networks. But you had mentioned this early on, Jimmy Ray, as being quite amazing because it was just different to think of a company like ours using something other than merchant silicon. The fact that we spun our own and I'm like, do you build these things in your garage, what did you do, how does that kind of thing come about and how do you begin explaining that. Do you mind -- I'll let you two walk us through it, but I want to hear more details.

MICHAEL BEESLEY: Okay, well certainly, I think it comes from the realization that modern routers and modern networks need very fast, high bandwidth, high packet per second, high services in their data plane. And certainly it's very expensive to develop this kind of technology.

JIMMY RAY PURSER: Well, let's get down to brass tacks here though. We're talking about a PROC, let's get to the real technical stuff.

ROBB BOYD: Hold on, you know what, I just remembered that we were supposed to have a PROC On set.

JIMMY RAY PURSER: I thought you had it with you. I don't have it, I thought you had it.

ROBB BOYD: No, this is the kind of stuff we need to plan ahead. I think it's coming in now.

JIMMY RAY PURSER: Did it just get found. Thanks, man, appreciate it.

ROBB BOYD: Wow, he's taking that swine flu thing seriously.

JIMMY RAY PURSER: It was still warm, too.

ROBB BOYD: Yes, straight out of the fat (inaudible). All right, all seriousness though, what makes this thing special? How do you begin explaining something like that because it fascinates me.

JIMMY RAY PURSER: Right, because this is the thing that when this first kind of came out there was press about this PROC before it even was announced anywhere else. In Network World it was called the Mystery Processor. Then when I heard it was a 90nm PROC, I was like, I'm like well that's not ours, that's probably Sons of Niagara, it was the very first thing I thought of. But it's not. This is a multi-core processor. You've got what, 40 cores?

MICHAEL BEESLEY: Yes, there's 40 cores inside of the QFP. They're soft, they run regular software that's written in a high level language. Each of those processors is (inaudible) threaded.

JIMMY RAY PURSER: Wait a minute, no, because that's not right because you can't -- any time we write proxy we had to write them in assembly. And you're saying it's wrote in something other than assembly?

MICHAEL BEESLEY: In prior generations to get the speed a lot of the features would be implemented in assembly. The downside of that is the software becomes very complex, it's very hard to add features.

JIMMY RAY PURSER: That's right.

MICHAEL BEESLEY: So on the QFP as a goal we wanted to be able to write the software in a high level language so that it's easier to debug, it's quicker to write, quicker to test. To do that you need a more powerful engine. So from the grounds up we designed into the QFP the silicon power and the processing power to take into account the fact that we were going to use a high level language, we're going to use a C Compiler and a Debugger.

JIMMY RAY PURSER: And CC, huh?

MICHAEL BEESLEY: And CC.

JIMMY RAY PURSER: Very cool.

MICHAEL BEESLEY: So that allows us to implement features quicker, to be able to test them and basically produce a more reliable product and software release to the customer base.

JIMMY RAY PURSER: Now see, here's what you've got to help me with because any time that -- I've been involved with a little bit of ASIC rolling and one of the things that we always had a problem with with multi-processors was cash coherency.

MICHAEL BEESLEY: Yes.

JIMMY RAY PURSER: And what cash coherency means is that basically if I've got PROC 1 and it's writing to the cache in memory location X and PROC 2 and it's also writing to the cache in memory location X but it's down as X2, I don't want this one to read that's data and this one to read another data so we use another protocol called, well we always called it Illinois because it was developed in Urbana. What's MISA, right, MISA is actually the acronym for what it is. How did you guys solve cache coherency with 40 PROCS, I mean 40 cores?

MICHAEL BEESLEY: Certainly there are two aspects to that. The first is that just in the hardware the caches are coherent so the software doesn't have to explicitly worry about accessing stale data or data that's being written by a separate processor or a separate thread actually because each of the 40 processors has 4 concurrent threads of execution inside of it. So that's where the equivalent of 160 7200 routers comes from. Inside the QFP it's actually operating on up to 160 packets at the same time. So the first aspect of that is simply the coherent caches. The second aspect is the flow management logic that is inside the device. Certainly, when you have that many processors working in parallel one of the things you have to worry about is actually getting packets out of order. They have to leave the router in the same order that they came. The QFP has some very sophisticated flow management hardware inside of it where the packets are tagged and stamped so that the device knows which was actually packet one and which was packet two and then on the way out those packets are reordered to make sure they came out in the same order.

JIMMY RAY PURSER: Yes, but as you're passing these packets through, maybe we're going to talk about this on the software side of it, my understanding is you're passing whole packets not just headers.

MICHAEL BEESLEY: That's also true, yes.

JIMMY RAY PURSER: Now I understand this in the CRS because of the way it does tree lookups. How are we doing that in this? This just sounds so --.

MICHAEL BEESLEY: Certainly in previous generations to minimize memory bandwidth and to increase the performance, the devices would only operate on the headers, they'd only have access to the headers, which is a nice thing to make something go fast. The downside is if you've got features that require the whole body of the packet, anything from GET VPN or IPSec, content-based routing, there's a multitude of features that require access to more of the packets at the header, the body, the tail and whatever and as a matter of policy we wanted the QFP and the ASR 1000 to be able to do those services and those features in the data plane without adding service cards, without adding any extra hardware and what have you. So the QFP was designed from the ground up to have enough memory bandwidth and to have enough cache space and whatever such that each of those processors could access all of the packets regardless, so they could do any feature that we wanted and the hardware was designed to provide the performance and the memory bandwidth to support that.

JIMMY RAY PURSER: You know, because here's the thing, you're describing this stuff and I'm like, so which came first, this is like the chicken or the egg. Did you write the software and say to run this software I need to have this type of PROC or did you write the processor pair and you say okay, I wrote this to make this optimized, I need to write this type of software because we're really talking about something completely different here.

MICHAEL BEESLEY: Yes, I would say that we took a system and software feature perspective first. We said, well, we want to be able to build these kinds of systems with these kind of scale points, price performance points and over time we want to be able to run all of that feature set in the data plane, everything from straight forwarding to firewall to NBAR to, obviously, ACLs and all of the security features. And from that starting point we started working on the micro architecture and the design of the QFP and the processors. We decided what kind of memories we need, what kind of processors we need and then what kind of hard-coded silicon accelerators that we need to do some very, very well-defined jobs that we know we have to do like an address look-up for IPv4 or IPv6 or whatever. It's a very invariant thing. We know the processor has to be able to do that. We know it's not going to change. So we were able to put that actually into silicon, hard code that into silicon to get a performance boost, but still leave everything else being executed in software on the 40 processors inside the device so that we can over time as requirements change, customers needs change we can, with just software, adapt the feature set of the system.

JIMMY RAY PURSER: What is it about nine layers, you guys, nine layers of metal to get this to work?

MICHAEL BEESLEY: Yes, it's a big chip. As you said, it's done in 90nm. I think its gate count is well over a billion. I think it's like got eight layers of metal, so it's -- I mean, to put it in perspective, it's a silicon engineering effort that rivals the industry. It rivals what Sun does, it rivals what Intel does.

JIMMY RAY PURSER: Oh, yes, definitely.

MICHAEL BEESLEY: So it's very, very advanced.

ROBB BOYD: Here's what I don't want to miss though and I want to make sure because the producer's going to make me the bad guy in just a moment because we're running out of time to focus on the QFP thing here. But I want to make sure we're not missing, what are the primary points you want to make sure we get across here before we run out of time and lose Michael's access to this topic because I don't want to miss segment two, we're going to get into more of the IOS side of things. What would be the primary things you want to make sure we cover here before?

JIMMY RAY PURSER: A couple of things, I've got this slide up here and stuff of some of the basic logical architecture here. I'm just blown away at what we can do inside of the QFP. What's a typical packet flow through this device look like? We've got a bunch of lines here and stuff, but if somebody's wanting to know how does packet flow run through here, how would you describe that?

MICHAEL BEESLEY: Well, the packets come in and leave the system through the SPAs, the Shared Port Adapters. Those modules are actually common across the ASR 1000 and several other platforms, the Cisco 10000, 7600, CRS. So the customer gets great investment protection, they're modular I/O cards. Packets enter the system through the SPA. They flow through the carrier card and then all packets go up to the embedded services processor for either forwarding or if they're control plane packets, if they're TCP or SNMP or whatever, they will be forwarded up to the route processor for that --.

JIMMY RAY PURSER: To keep state, right?

MICHAEL BEESLEY: Exactly, to keep state and whatever. The nice part about the design is that since all packets will flow through the QFP and through the QoS subsystem, the hardware that we have in the QFP actually serves to protect the control plane as well. So it really helps to minimize the change of a denial of service or whatever if someone's using control plane traffic.

JIMMY RAY PURSER: Oh, yes, definitely, definitely.

MICHAEL BEESLEY: For transit traffic, all of the forwarding and processing is done on the QFP. The ingress interface is chosen and then the packet flows back through the back plane and then back out to the carrier card and then out the egress SPA towards its actual final destination.

JIMMY RAY PURSER: I really do like the idea of having a crypto assist over here as well because that, obviously, gives us, especially in the Service Provider space, where we can actually offer some pretty intense services and a lot of pretty heavy-duty cryptography across those services independently without really dragging the box down, which crypto always has the problem of really slowing everything down.

MICHAEL BEESLEY: Yes, certainly the crypto assist makes the scale of the box very good in terms of key generation and number of GET VPN or IPSec sessions that the system can support and then the actual bulk encryption/decryption is done by a dedicated coprocessor that allows us to go, I think our fastest is up to 8Gb throughput of the box for IPSec traffic with no impact on anything else, no other service, no other features of the box slows down. So that's a very nice extension to these kinds of systems where normally when you turn on an encryption feature the box will slow down.

JIMMY RAY PURSER: Yes, definitely.

MICHAEL BEESLEY: And the ASR 1000, that's not the case.

JIMMY RAY PURSER: Unbelievable, so let me kind of put you on the spot here a second, I know I probably shouldn't do this but if we can't then we'll just do a cut point and then just re-edit it. But in a back alley fist fight who would win, the M120, ASR 1000?

MICHAEL BEESLEY: The ASR 1000, hands down, hands down.

JIMMY RAY PURSER: Why do you say that because man, that's a good box.

MICHAEL BEESLEY: It is a good box, but the architecture and design of that box is as you want to add services and features, which everyone does --.

JIMMY RAY PURSER: That's why you want this kind of box.

MICHAEL BEESLEY: Exactly, there's no vanilla simply transport networks anymore.

JIMMY RAY PURSER: No, no way.

MICHAEL BEESLEY: Everyone is concerned and worried about security, everyone wants their NetFlow statistics. Everyone wants various forms of bonding, whether it be Ethernet or through TDM interfaces or whatever.

JIMMY RAY PURSER: Well billing and accounting, all of that stuff.

MICHAEL BEESLEY: Absolutely.

JIMMY RAY PURSER: If I'm going to charge I need to be able to bill back.

MICHAEL BEESLEY: Absolutely, and as we move forward everyone's going to want things like WebEx acceleration and collaboration features and whatever and with the competitor's box to get all of that you have to add so many dedicated cards that you consume the system, you've no room left in the system for your Ethernet ports.

ROBB BOYD: It defeats the purpose.

MICHAEL BEESLEY: Which defeats your purpose. It also makes the box extremely expensive. The service modules that our competitors offer they come with a price tag, they don't come free.

JIMMY RAY PURSER: No, that's true.

MICHAEL BEESLEY: So it ends up being very, very expensive so for my dollar I'd go with the ASR 1000.

VALERIE ST JOHN: Well clearly, teams of almost freakishly brilliant engineers poured their intellectual resources into this thing. It sounds like one (inaudible) tricked out.

ROBB BOYD: I think that is a compliment.

VALERIE ST JOHN: In my way I have conveyed a compliment.

JIMMY RAY PURSER: That's a Valerie compliment, you'll get used to it.

ROBB BOYD: You do represent the team and that's a good point though, that's a very good point.

VALERIE ST JOHN: The work speaks for itself. Thank you, gentlemen.

ROBB BOYD: Thank you.

VALERIE ST JOHN: All right, for more information check out the show notes or just go straight to our block at techwisetv.com.

INSIDE THE OS: CISCO IOS XE SOFTWARE

VALERIE ST JOHN: And we're talking about the ASR 1000 with Robb, Jimmy Ray and our honored guest, Michael Beesley, who shakes his head every time I call him an honored guest but I just can't help it.

ROBB BOYD: Humble guy.

VALERIE ST JOHN: As much as one operating system for everything sounds like the right thing, the one OS philosophy doesn't address specialized needs, so help me out here, Robb, what do we do?

ROBB BOYD: It's funny you bring that up because that's a common question that comes up which is okay, well, Cisco, obviously, prides itself on IOS for quite some time. It's a very venerable operating system, people are very used to working with that, a lot of engineers have invested a lot of their careers into getting very good with these. And then we started talking about data center on this show and what our Data Center 3.0 team has been consistently putting out there, they're always talking about Nexus OS or the NX-OS. And so we've spent a lot of time learning NX-OS and showing the relationships and the places where it makes sense in comparison to what our audience knows from IOS. Well, now we're talking about IOS-XE and so this is a different operating system for the ASR because of the fact that this box has very specialized things that it does better than, quite frankly, any other box, any other technology because customers have special things that they need done at rates and speeds and reliability that they've never had before, both on the Enterprise side as well as on the Service Provider side and so it does take something different. But let me throw this to you guys. You've got a different operating system. This is highlighted. Jimmy Ray, you brought this up as being something as special to you. You brought it up as one of your top three things that made this box radical. Let's cover that one in more detail. Tell me about the operating system. Why change and what makes it so different?

MICHAEL BEESLEY: I think to address that why change first, certainly as networking has advanced I would say the biggest thing that's changed is the mission critical nature of the network, both for Enterprise customers and Service Provider customers. And that really comes back to the reliability and the modularity of your software. So that's certainly one of the big drivers to advance IOS from the architectures that we've used on previous product lines to IOS XE where we've added a significant level of modularity to the system, as well as the ability to have software redundancy in the systems, both in the bigger chassis that have hardware redundancy as well, but also in the smaller chassis where there's no hardware redundancy but we still provide the option for customers to run redundant software. So I would say that's the biggest driver.

JIMMY RAY PURSER: Why not this?

ROBB BOYD: Right off the bat. He's going to improve on it all of a sudden. IOS XR?

JIMMY RAY PURSER: Why just XR?

MICHAEL BEESLEY: Certainly, it's one of the things that we considered when we started the development for the ASR 1000 is whether we should use the XR operating system from the core routers.

JIMMY RAY PURSER: Right, right.

MICHAEL BEESLEY: Certainly there are some benefits to that with regard to a liability and modularity, for sure. The problem comes down to actually a feature set. Certainly the IOS XR was architected and designed from the ground up and has a feature set that really supports a core router. For products like the ASR 1000 that go onto the Service Provider edge, they go into Enterprise networks, the required feature set is so much broader. Literally these products have thousands and thousands of features so we really wanted to start with IOS because it actually has the control plane already implemented and tested and proven in decades of production use that the control plane for all of those features is there. So we really waned to start there but somehow come up with a software architecture that delivers all the benefits of a modular system, such as IOS XR.

ROBB BOYD: Well let's get into how that's broken down. I want you to erase that again. I don't want anybody getting confused --.

JIMMY RAY PURSER: Yes, that's a good idea. So just walking through this real quick, when I first saw this I'm like, hey what's this, kernel, what are you talking about kernel.

MICHAEL BEESLEY: Certainly that was one of the first things that we did was to understand or to appreciate that IOS is a very valuable application. It's a very large code-base with a lot of features in it. But to make it reliable and to be able to allow the system to run two instances of IOS we needed to promote it to be a pure application. So we did, we took a carrier class operating system, a variant of Linux and actually run Linux on each of the CPUs in the box as the operating system and that frees up IOS to be an application. IOS no longer has to worry about managing memory and managing the CPU and to an extent on the ASR 1000 even managing the system because we have external software components that manage the system that handle tricky operations like hot swap and redundancy switchover and whatever. Really leaving IOS to do what it's very good at, which is run control plane, run the CLI, SNMP, routing protocols and what have you. So the first step of that was to actually underpin a carrier class operating system underneath IOS and to elevate it to being a user end application.

JIMMY RAY PURSER: So we're talking really about an RTOS, a real-time operating system that's running right here then basically?

MICHAEL BEESLEY: Well, it's Linux, which is traditionally not an RTOS, but it's got some extensions to it --.

JIMMY RAY PURSER: To make it, yes, right.

MICHAEL BEESLEY: To make it very real time responsive. Although on the ASR 1000 because we have the QFP in the data plane, anything that has to be very, very real time, things like firewall session setup.

JIMMY RAY PURSER: Deep packet inspection.

MICHAEL BEESLEY: Deep packet inspection. We actually do that on the QFP so the general purpose CPUs inside of the box--.

JIMMY RAY PURSER: Oh, I get that.

MICHAEL BEESLEY: -- don't do that work at all.

JIMMY RAY PURSER: Sure, sure.

MICHAEL BEESLEY: Things that have to happen in microseconds or a small number of milliseconds, we actually do in the QFP because we've got an enormous processing power there and the QFP sees the packet first. So that really makes the job of engineering the control plane software that much easier. You don't have to worry about real time in milliseconds where other products and previous products did have to worry about that on a general purpose CPU, which is a very challenging thing to do.

JIMMY RAY PURSER: Well interrupt base processing always slows everything down. So that makes a lot of sense. Let's walk through some of your XE favs, your favorite hits on the XE stuff. What would you like to dig in?

MICHAEL BEESLEY: Well certainly, the first thing to note is the ability to run two instances of IOS, one in active and one in a standby capacity. That really, that gives customers two big benefits. It gives a reliability benefit to the machine. If the active IOS was the fail in a small number of milliseconds the system can switch over to the standby and continue full stateful operation.

JIMMY RAY PURSER: Milliseconds.

MICHAEL BEESLEY: Milliseconds, which is an amazing advance for this kind of product category. I actually think it's the first time in this space. Usually you're up on the big million dollar routers to get that kind of functionality.

JIMMY RAY PURSER: Absolutely.

MICHAEL BEESLEY: And we're able to offer that on the systems that have hardware redundancy as well as the systems like this one that actually don't have hardware redundancy, but you can still protect your software. So that's the first thing of note. The second thing to note is that it is a modular system. On the control plane we have other system processes, such as the Interface Manager, the Forwarding Manager, the Chassis Manager, that actually do work on behalf of IOS, leaving IOS to really do what it's being designed to do, which is run routing protocols in control plane. This gives a level of modularity to the system. You can change each of these components without changing the other component in the system. And as general purpose CPUs move more and more to multi-core, two cores, four cores and whatever, we'll be able to incorporate those types of CPUs into the product and gel that much performance and scale improvements of the system. So they're the first couple of things I'd highlight. The last point I would make is that the software is split into three main buckets that match the hardware layout of the machine. We've got the CPUs on the SPA carrier cards, CPU on the forwarding card, as well as the main CPU on the RP card. So we've relegated a lot of the software to run down on those embedded CPUs. As an example, the SPA drivers that run the I/O hardware are down on the SPA carrier card and its CPU. The software that runs the QFP and manages the state for the data plane runs down on the forwarding card CPU. And that really offloads a lot of the hardware management from the route processor CPU which allow sit to do what it should do, which is process the control plane, run routing protocols and whatever, giving, again, a scale advantage as well as a modularity and reliability advantage to the system.

JIMMY RAY PURSER: Very cool. So we're going to pop a board out here real quick.

ROBB BOYD: Waiting for you to open that (inaudible).

JIMMY RAY PURSER: Absolutely, go ahead and pop this thing out. Now one of the things that is pretty amazing on this is really kind of the great board layout design. Right here is the QFPs right here, right?

MICHAEL BEESLEY: That's right, yes.

JIMMY RAY PURSER: So that's what we're talking about.

ROBB BOYD: Underneath there.

MICHAEL BEESLEY: Yes, it's underneath. This is the heat sink. The QFP is actually a two-chip chipset. One chip has the processor cores, the 40 processor cores. And the second chip is a dedicated hardware QoS engine that can run up to 128,000 hardware queues to support very accurate multi-hierarchy egress Quality of Service.

ROBB BOYD: Without even breaking a sweat.

MICHAEL BEESLEY: And that's all done in hardware so the use of QoS features has no performance impact to the system whatsoever. And because the QoS hierarchy is actually drive through a hardware engine, it's got very, very accurate and reliable scheduling, which becomes more and more important.

JIMMY RAY PURSER: That's super important, absolutely.

MICHAEL BEESLEY: These days as networks are running many different applications across congested links, your QoS abilities are becoming more and more important as we go forward.

JIMMY RAY PURSER: Sure.

MICHAEL BEESLEY: The rest of the components on the forwarding card, obviously, the QFP is the main component.

JIMMY RAY PURSER: How many output lines do we have on these QFPs?

MICHAEL BEESLEY: It can support up to 64,000 virtual interfaces across hundreds and hundreds of physical interfaces.

ROBB BOYD: Oodles.

MICHAEL BEESLEY: Oodles is a very technical term to describe that. So along with the QFP on the forwarding card we have the Crypto Accelerator that offloads key processing and bulk encrypt and decrypt for the system. That allows features like IPSec, and GET VPN and what have you to be turned on and, again, no performance impact to the rest of the system. We have a CPU onboard to manage and monitor all of the hardware, which really relieves the route processor in the system from that kind of work, leaving more CPU cycles and more memory for the control plane. And really that's in essence what's on the card. We've got a couple of other chips that just manage the high speed links across the backplane to connect the forwarding card to the other parts of the system.

JIMMY RAY PURSER: Very cool, that was cool.

ROBB BOYD: These, obviously, are the interface cards to the shared port adapters?

MICHAEL BEESLEY: Yes.

ROBB BOYD: Before they go in there.

JIMMY RAY PURSER: This SPA's actually the WebEx SPA.

MICHAEL BEESLEY: Yes.

ROBB BOYD: Oh there we go, don't get ahead of ourselves.

MICHAEL BEESLEY: Yes, that's the shared port adapter. Again, from a form and mechanism fit they are shared across a number of the Cisco product lines which really is a huge investment protection (inaudible).

ROBB BOYD: Yes, absolutely.

JIMMY RAY PURSER: And there's a QFP on here, too then?

MICHAEL BEESLEY: Actually, no, this processor on this SPA for the WebEx is actually the Cavium Octeon processor.

JIMMY RAY PURSER: Oh, really, okay, okay.

MICHAEL BEESLEY: Yes, yes. One of the things that allowed us to basically migrate the WebEx stack and that application from normally it runs on a server in a data center, we're able to migrate it onto a SPA by using the Octeon processor and we run the Linux operating system on that processor and then the WebEx stack and application on top of that. Which really it's a tremendous ability to think that you can shrink what runs onto a server, on a server you shrink it down into a SPA, you're able to put it into your edge router and it gives you a performance advantage as well as, obviously, protecting your very precious WAN bandwidth as you don't have to go back and forth to the data center.

ROBB BOYD: I hate to kind of do this because we have to decide where to end this conversation. I'm blown away, I don't even know how to summarize because the whole idea was I was going to come back to you with a really cool summary of everything that he had just said and it was going to be tight and beautiful and you were all going to be in awe of my oratory skills.

VALERIE ST JOHN: And you're saying we can't have that now.

ROBB BOYD: I don't have it.

VALERIE ST JOHN: No.

ROBB BOYD: Do you have anything?

JIMMY RAY PURSER: Yes, but --.

ROBB BOYD: I'm so happy to have you on the show, by the way.

MICHAEL BEESLEY: Thank you for having me.

JIMMY RAY PURSER: I think this is one of the most advanced products I've ever seen ever.

ROBB BOYD: Until you see the next one.

JIMMY RAY PURSER: Well when I saw the Nexus stuff I was really impressed, right?

ROBB BOYD: Yes.

JIMMY RAY PURSER: I thought that was incredibly advanced. But I'm blown away. I cannot believe what is done in 9nm of silicon here. I'm floored, I don't even know -- this is unreal.

VALERIE ST JOHN: And it's actually pretty on the inside.

ROBB BOYD: See, you've been coming around here too often, mighty nice architecture.

VALERIE ST JOHN: Michael Beesley, thanks so much for joining us. Thanks for giving us a tour of the ASR 1000 and for being a good building representative for Building 20.

MICHAEL BEESLEY: Thank you.

VALERIE ST JOHN: Nice to have somebody from Building 20 over here. Okay, for more information check out the show notes or just go straight to our blog at techwisetv.com.

WEBEX NODE

VALERIE ST JOHN: Online, on-demand and on the go from the industry briefing center on the Cisco campus, this is TechWise TV, technology you can use from geeks you can trust. As travel budgets are cut, naturally online collaboration increases, right. We're all getting better at creatively addressing the workload through collaborative applications. WebEx, Cisco's hosted conferencing service continues to explode as more people are doing audio, video and application sharing inside and outside of the company. So tell me, where does the ASR 1000 fit into all of this?

ROBB BOYD: Well I'm going to get to that. The idea here is that we're talking about bandwidth consumption in a way that some people, I don't think, have always think of or if they've thought of it they didn't realize that there might be some creative solutions because that's actually what I went through with this one here. Pull up that slide that you've got it here.

JIMMY RAY PURSER: You've got it, man.

ROBB BOYD: Because we've got a demonstration here that I'm going to go through. So the idea, Jimmy Ray, this is what struck me first of all and we've got two drawings here and if we look here we're talking about the WebEx Node, which just came out for the ASR 1000. And the idea here is that if you're a company representative over here you've got multiple connections potentially attending, being driven through meeting services provided from the cloud, the WebEx based kind of a cloud -- that terms gets misused a lot. But the idea is it's delivered as a service. So if you've got, let's just say, 500 people in a given building all attending the same meeting, they each have individual streams as represented here. And so you have this illogical consumption of bandwidth that doesn't make a whole lot of sense that people may not always think of. And so the genius side of this, and I don't even think it's a hard concept to understand, is that you add a bit of hardware, to Valerie's question, in the ASR 1000 in the form of this shared port adapter, SPA.

JIMMY RAY PURSER: The SPA, right, right.

ROBB BOYD: Let me make sure I'm spelling out my acronyms correctly.

JIMMY RAY PURSER: This would be like the multiple sessions that you have like when people attend out WebEx workshop demo that are two weeks following the show as far as shameless plugs go.

ROBB BOYD: It was a shameless plug. Yes, well backup here one second because as you look right here the idea is we now add this shared port adapter within the ASR 1000 and I don't know if this is technically correct, but the way I like to think of, is you essentially almost become a member or a node on the cloud itself.

JIMMY RAY PURSER: That's right. You're talking --.

ROBB BOYD: You okay with that?

JIMMY RAY PURSER: Yes, no, because it's true. Because we're looking at Software as a Service here and you are connecting up and being a member of that MediaTone Network, you're plugging in. And the cool thing about it, it's a simple SPA module so we're really talking about it's very low form factor and for what this does it is so unbelievably easy to configure. There are like four commands to set this thing up.

ROBB BOYD: And I want to say that, so here's the main point that we get across. The idea is multiple streams because of everybody that connects to a meeting and this is multi-media meetings these days, right? I mean, how many of us are turning our cameras on so we can see the other person, make sure that you're paying attention because that's always been good for me to understand on these meetings when you're actually paying attention or not or I can tell you hey they're about to come to you, they're about to ask you a question, Jimmy Ray, come on. Now we get it down to one stream because this connects into the cloud directly with one. It reminds me kind of Unicast, Multicast.

JIMMY RAY PURSER: Yes, that's exactly what it's like, very good, very good.

ROBB BOYD: But it's LAN versus WAN, right, and we pay a lot of money for that WAN bandwidth. And the idea here is let's drive these services, we're all traveling less. This stuff's not going to go down and that traffic's only going to go up, what kind of effect are we seeing on it now in terms of the traffic change. I know you've got that ready to go.

JIMMY RAY PURSER: I do have. Funny you should ask, I have a demo all queued up. If you look at this presentation what we're looking at is we do have a WebEx meeting going on right now and this is our graph of what our traffic is as it's going across the LAN in blue and the WAN in red. And so you can see it just updated again, and it's spiking. So the more people that start sessions, like if it's a company meeting, it's kind of silly to fly everybody out for a company meeting that nobody pays attention to, they could easily not pay attention to it at their house without paying for all the airfare.

ROBB BOYD: Exactly. But the point here though is that LAN and WAN are roughly the same here thought, right?

JIMMY RAY PURSER: (Inaudible).

ROBB BOYD: But we pay a lot of different -- these cost different.

JIMMY RAY PURSER: So much, so, so much more.

ROBB BOYD: Make sure I make my sales points.

JIMMY RAY PURSER: Now let's take a look. Now if we go out here I've got a demo running here and this is the WebEx host with the SPA. I'm going to go ahead and start this off.

ROBB BOYD: Explain what we're seeing here though too because it's confusing. We see the waterfall maybe all of a sudden but that's the traffic you're generating?

JIMMY RAY PURSER: What you're seeing is this is the person -- the SPA host is actually up here. This is where I'm at. This is who's running the meeting. So I'm sharing this full desktop out to these two clients. This is only two people on this network. And so they're all seeing me as I pass this out. And they're seeing different frame rates, passing all that across. And so all these clients are now passing that out. But now they're going through the SPA and look what just happened. So let's take a look at the difference between while this graph is updating. Do you see how we fail? Our traffic really dropped here because in essence what we're doing here is that we're taking all our meetings here and we're consolidating them into the SPA blade, the WebEx blade on the ASR. Now the real nice effect is that instead of having multiple services connect out, everything is being aggregated here. I'm locally serving this up and, truthfully, because of the way this does accounting, I can actually sell those services back to somebody else. So I pay for WebEx services I can actually now bill in locally.

ROBB BOYD: I did not know that.

JIMMY RAY PURSER: Absolutely, because the accounting's so good you could actually resell the stuff back, use it for billing and accounting and actually charge people for how they're actually eating this bandwidth up.

ROBB BOYD: So you just paid for the box.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: You paid for the card anyway. We'll work on the box next.

JIMMY RAY PURSER: Now here's something really cool though, this is transparent. So it doesn't require any user configuration, I don't have to go up to my user site, set this up, anything.

ROBB BOYD: And assuming users were getting good connectivity on a regular basis, they'll continue to get good connectivity. Company saves money but they don't have to do anything different.

JIMMY RAY PURSER: Most definitely, most definitely. Now the sweet thing about this is let's say that it's a TechWise show that people really want to see and you've got everybody (inaudible).

ROBB BOYD: Imagine that.

JIMMY RAY PURSER: Like man, we've got to see it. And they overrun this. I mean, there's a connection limit, right? And so let's say that there's 10,000 connections that hit this booger and they start to overrun this connection. What happens? Well, this is not an inline device that's going to start dropping connections on there. It's transparent so the connections that don't go through here will go directly back out to the cloud.

ROBB BOYD: Oh, so in other words --.

JIMMY RAY PURSER: If this device fails, that's okay, everybody still connects back out to the cloud. This does not give away --.

ROBB BOYD: No single point of a failure, all of a sudden all the same scalability you had before doesn't go away.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: Oh, okay.

JIMMY RAY PURSER: The only thing it does is it really enhances your connectivity experience and everything that's dropping this is good.

ROBB BOYD: And look now, now the WAN's broken completely away from the LAN.

JIMMY RAY PURSER: Absolutely, so now our WAN usage is lower than our WAN usage.

ROBB BOYD: Which is what we want.

JIMMY RAY PURSER: Which is exactly where we want to be. This is getting up. So this is the really nice thing that makes this pretty efficient. Now let me show you something really cool just real quick. I know we've got to wrap here, but here's the administration panel, pretty darn quick and easy. It shows -- now this is active, I really can't mess with anything here because it actually is connected up. But it'll show all my SPAs connected up, who's using, how they're using, and I can use it for bill back and use it for troubleshooting, piece of cake.

ROBB BOYD: Well see, I like this too because also with my sales background, the idea is that you've got a way to go look and say did we spend our money in a wise fashion, are we actually -- because it's so transparent you're like how do we show somebody that it was good money spent, right?

JIMMY RAY PURSER: It's very cool.

VALERIE ST JOHN: Now more meetings, great. But I am comforted by the fact that we'll be using network resources much more efficiently and that's definitely a good thing.

ROBB BOYD: Now you're thinking.

VALERIE ST JOHN: So for more information check out the show notes or just go straight to our blog at techwisetv.com.

ASR 1000 DESIGN TIPS

VALERIE ST JOHN: For live programs only you may submit your questions to our team of online experts. Simply type your question, click submit and refresh regularly. Okay, let's talk about the design of the ASR 1000. The router's design is a clever combination really of ground up innovation and thousands of hours of R&D. It's built upon years of prior research. It focused on the unique pressure an edge router faces in today's dynamic multi-dimensional networks, right?

ROBB BOYD: It does absolutely and this is probably one of our favorite things to do kind of towards the end of a show like this is really talk about design considerations, which is this notion of, obviously, if something is that fundamentally different, that radical as I kept forcing into Michael's mouth there, we really don't know if he actually said that, but I like the word, the idea here is what do we need to know different from a design perspective. It's not that we have a ton of time here so maybe just focusing on the specifics that make this different. What do we need to be aware of?

JIMMY RAY PURSER: Yes, we'll look at a couple of unique usage cases here as probably the right thing to do. To talk about design with a product like this we're talking about a pretty big product we can really go into deep. So we'll probably do some of that in a WebEx workshop that follows, just kind of dig in there. But let's look at a couple of very specific use cases here. The first one is actually in the Service Provider SP here is a really nice service that Service Providers can provide to their, in this case, like your residential users. So down here I can actually offer a service, a per user firewall that is actually allowing me to charge a service for firewalling my users on the backend here to make sure that they have IPSec or I can actually put a firewall segment for that virtual firewall for every subscriber assigned to (inaudible).

ROBB BOYD: That kind of things been done before. I know it's been offered as a service but --.

JIMMY RAY PURSER: No, not like this because when I wasn't here --.

ROBB BOYD: (Inaudible) a box that time, right?

JIMMY RAY PURSER: Well, here I'm controlling everything at this end, at the head end is where I'm assigning all these firewall services and now since I can actually run this firewall between 10 and 20Gig, I can now run this, an IOS-based firewall with full packet inspection, I'm not talking this little you know take a look at the heard and oh, it's okay. I'm talking about actually digging in, disassembling that packet apart. And I'm offering it at this rate of speed so I can actually have massive amounts of scalability for my customers out there today. Nobody's ever did this before and it really extend what Service Providers --.

ROBB BOYD: -- looking at the full packet and being able to process the things at hardware based --.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: Oh, that's awesome, okay.

JIMMY RAY PURSER: And we talked before about putting duce addresses, ACLS and duce out your ACLs. We can do that here, still build this in to the ASR or if that's part of the firewall service. And since our ACLs run in the TCAMs they're accessed in nanoseconds so you don't even see it. This is a blip on the radar.

ROBB BOYD: Yes.

JIMMY RAY PURSER: So typically you can't do that for that type of folks. Here's another really cool Service Provider type of feature, a Session Border Controller. We do have a CUBE that actually can allow you to give control and signaling services out there for more rich voice services out there today, SIP and things like --.

ROBB BOYD: I was going to say, so you could use it for like SIP Trunking and things of that nature, right?

JIMMY RAY PURSER: Absolutely.

JIMMY RAY PURSER: Absolutely, I can provide that out. I can hold about 32,000 sessions on this and really scale this very long range out of a single box. I'm not talking about a group of boxes, I'm talking about a single box.

ROBB BOYD: And with a Session Border Controller, is that something that's going to be -- is that a SPA card?

JIMMY RAY PURSER: Yes, absolutely.

ROBB BOYD: Okay, so different because you didn't need that for the firewall, right?

JIMY RAY PURSER: No, uh-uh.

ROBB BOYD: So make sure we have that distinction.

JIMMY RAY PURSER: The firewall's built -- oh that's a very good call.

ROBB BOYD: A little bit of additional hardware.

JIMMY RAY PURSER: Very, very good call. And her, of course --.

ROBB BOYD: But that's nice, because we're seeing services go that way, right? SIP's taking off and that's saving companies so much more money to take that level up so they're not waiting money on those extra trunks.

JIMMY RAY PURSER: It's hard to justify money savings on Voice over IP unless I'm doing a lot of collaboration and I start doing SIP trunking. Then I could actually have some money.

ROBB BOYD: I thought you were going to say something negative about voice for a second.

JIMMY RAY PURSER: Oh no, no, no, not after having to deal with Tina. Here's something really cool on the Enterprise space.

ROBB BOYD: Okay, so now we've switched around.

JIMMY RAY PURSER: Now we're going to the enterprise side.

ROBB BOYD: Service Providers to Enterprise, okay.

JIMMY RAY PURSER: Absolutely. This is actually the data center interconnect, DCI. And here I'm actually what you're showing is I'm connecting two data centers together in this network at Layer 2. So you can see these are actually my queue links. And what I have here is I have my pseudowire that's connecting these Layer 2 services up. And, again, supporting all of the services and all of the features here at Layer 2, I'm actually able to run stuff like VMotion across this cloud back and forth.

ROBB BOYD: So actually moving, theoretically moving servers or applications can change physical locations?

JIMMY RAY PURSER: Absolutely. Yes, pretty groovy because now I'm actually -- like this can be the great state of Tennessee and this can be, I don't know, Germany and I'm actually transferring --.

ROBB BOYD: Tennessee to Germany.

JIMMY RAY PURSER: -- applications, yes.

ROBB BOYD: Makes complete sense, I know.

JIMMY RAY PURSER: Why not?

ROBB BOYD: I think this happens a lot actually.

JIMMY RAY PURSER: And now I've got these in Layer 2 domains, plus I can encrypt them, I can secure the, I can put all my crypto here and I can actually have a very secure services here that are more secure in my cloud than they are in my local area network. Yes, very groovy.

ROBB BOYD: And that brings up some of the virtualization things that we're going to talk about in our next data center focus show too, right, because mobile -- moving those things around, obviously, implies some policy changes and some abilities, okay. I like it. I like it.

JIMMY RAY PURSER: This is probably my favorite use of the product, believe it or not. I like this one better than any of the other three. And, obviously, I stole that slide because there's no way I could figure out how to make those things do that.

ROBB BOYD: These slides are a lot prettier than yours anyway.

JIMMY RAY PURSER: Yes, I stole them. But this is another enterprise service here and actually what I'm doing here is I'm integrating my threat defense system inside the product because it can do all my denial of services stuff, which is no surprise. Heck, we can do that on a switch.

ROBB BOYD: That's DDoS?

JIMMY RAY PURSER: Yes.

ROBB BOYD: So that sells, okay.

JIMMY RAY PURSER: So what I'm actually looking at doing here is actually because I am running at such a high speed and being able to define all my services and run all my ACLs into a deeper level and do my full packet inspect now I can offer this out at the head end as a service to my branches and really firewall them out accordingly and do deep packet inspection, things of that nature on their end, and really protect my network from that stuff. This is a really cool, cool feature because typically I'm deploying firewalls all over the place, set up all this stuff. I can do this at the head end, there you go.

ROBB BOYD: A lot cleaner.

JIMMY RAY PURSER: Really, really clean.

ROBB BOYD: A lot easier to keep up with, absolutely, okay, very good.

VALERIE ST JOHN: All right, thanks, guys. And for the super serious routing and switching buffs in our audience we have a special guest, Lora O'Haver is with us from Learning at Cisco. Welcome, Lora. LORA O'

HAVER: Hi Valerie.

VALERIE ST JOHN: So you're here to tell us about an updated version of the new CCIE Routing and Switching Certification. Is that right? LORA O'

HAVER: That's right.

VALERIE ST JOHN: What does it cover? LORA O'

HAVER: Well updated version of CCIE Routing and Switching covers all of the important skills that expert level network engineers use on their jobs. This announcement was made on May 5th and it expands coverage of the skills associated with the configuration of security and wireless access. It also expands the coverage of skills associated with the planning and monitoring of the network and also the preparing of the network for the Advanced Services like video and voice. It also adds hands-on testing of troubleshooting skills to the very difficult eight-hour CCIE lab exam. So that will be added into the exam.

VALERIE ST JOHN: Okay, so that's not for me, that's for somebody who's super, super smart. LORA O'

HAVER: Right, the expert level engineers, those who've probably been in the business for at least five years, upwards of ten years, really worked themselves up to the Pinnacle of their career.

VALERIE ST JOHN: And is everything ready now to go or is this something that's happening way down the pike? LORA O'

HAVER: The exams will be released actually on October 18th of this year. The standards and the topics for those exams are already available on the web for candidates to prepare and study because, of course, this is a month's long study process to prepare for this exam.

VALERIE ST JOHN: Excellent, Lora, thanks very much for joining us. LORA O'

HAVER: Okay, great.

VALERIE ST JOHN: All right, Robb, how much did we change the world today? What are the key takeaways on the ASR 1000?

ROBB BOYD: I think the key takeaways is what was really fun about putting this particular show together, and I think all of have enjoyed the fact that we get to meet people like Michael Beesley and he represents, and I thank you so much for making sure that you brought this out consistently because he was hyper-concerned about making sure it was understood that he represents a bunch of very smart people. And we credit him because it's easier to credit one individual and we couldn't fit everybody in here. But the idea is there's smart people that come up with ideas and they execute on these ideas to create innovation that's driving things based on what our customer needs as we take both what our customers are asking for as well as what our customers are building towards. And as every study we look at, it's not hard to imagine these days we're seeing so much video and so much interaction and we're seeing the network take on so much more responsibility for everything

that revolves around it, the ante gets upped, we've got to keep up from a product perspective, we've got to continue doing this kind of thing, we're never done, so we're going to expect more of these type of things. Hopefully, we'll be the ones to cover it here. We're thinking in terms of Exabytes as we look forward, Petabytes, Exabytes, things that don't make any sense to us in today's years. But if you go back simply ten years and say wow, some of that stuff sounded really foreign to me then and now we're already surpassing even our most aggressive predictions. There's nothing on the horizon that says any of that's going to slow down, we had a long way to go. These kind of products are what's getting us there. People like Michael Beesley and who he represents and the many others behind the scenes at Cisco, hopefully, we can keep bringing them. This stuff's exciting to me, I'm so glad we got to talk about it.

VALERIE ST JOHN: Excellent, yes, the upping of the ante is a continual process.

ROBB BOYD: Easy for you to say, absolutely.

VALERIE ST JOHN: Thank you, I appreciate that. And for Robb Boyd, Jimmy Ray Purser, I'm Valerie St John. Thanks for joining us on TechWise TV. To find out about future episodes or check out a complete archive of previous shows, be sure to visit Cisco interaction network at cisco.com/go/interact.