

VALERIE ST JOHN: Welcome to TechWiseTV, I'm Valerie St John, along with Cisco solution experts Robb Boyd and Jimmy Ray Purser, hey guys.

JIMMY RAY PURSER: Hey, how are you?

ROBB BOYD: Hey.

VALERIE ST JOHN: Our focus today, breaking the application bottleneck. We demand more and more out of our networks today, you know that. The need to regain control over a multitude of business critical applications is leading a wave of recentralization. As with everything in life, the purpose of course is to either avoid pain, or seek benefits. Robb, what's at stake here?

ROBB BOYD: Or as I say, avoid pain, or gain pleasure, or could we put an and in there? Can I gain pleasure while at the same time avoiding pain? I think that would be an ideal situation. I think the issue here becomes we've got users already complaining about application response times and things like this, and at the same time, we're telling a lot of people that there's benefits to centralizing applications, moving them potentially further away from the users. Is that simply going to exasperate the problem? Is that going to create more of an application bottleneck that we need to be aware of?

VALERIE ST JOHN: So if recentralization then is the right way to regain control of our applications, does this maybe challenge the user experience? How do we get ahead of this problem?

ROBB BOYD: Well I think there's a couple things to keep in mind, and I believe this is my whiteboard debut here. So let's take a look. First of all, where a lot of people focus is this notion of optimization. In fact, we focus a lot on optimization, because there's a lot of neat stuff to talk about here, and we'll talk about that a lot today as well. The other thing, though, to consider as an equally important leg, is this notion of availability. We can't forget about that, and then security, of course, is one that we'll often talk about. Forgive the handwriting, I think actually you have to have bad handwriting on these things. The idea is that you've got to balance between these things, so don't overly focus on optimization at the expense of these other things, keep a balance between them. And the best way to go about doing that is think about this as a stool, and what you've got is your guiding principle to keep things balanced is policy. In other words, how do we focus on the fact that business changes? There's dynamics, boy, it's never been more relevant and timely to talk about business change than now. And so if you're overly focused on one area of this at the expense of the others and don't have the balance, you don't have that policy to point to, you may be setting yourselves up for creating more complexity than you actually attempted to solve in the first place, right?

VALERIE ST JOHN: There's a long way we could go here, though, speaking of balance. What's the right level of investment, would you say?

ROBB BOYD: Oh, you know, actually, that actually might even be a fourth leg on the stool here.

VALERIE ST JOHN: Hey, there we go, a stronger stool.

ROBB BOYD: I think that's a good question to ask. And the idea is really understand, focus on what it is you're trying to solve. If you've got a point problem, then perhaps you can solve that with a point product, okay? But if you do that, and you lose sight of your policy and the dynamics of the business, you do run the risk of saying you

know what, I'm spending a lot of money in this one direction, but at the same time it's not flexible, it doesn't support all those legs of the stool. And I'm again creating more complexity. The business needs are going to change, make sure the policy maps to it, don't overly focus on those.

VALERIE ST JOHN: So it seems to be all the more important, then, to start with a judicious plan. Now as we're going to talk about application delivery networking, let's bring in our senior network applications deliveryman, Jimmy Ray Purser, welcome.

JIMMY RAY PURSER: Thank you.

VALERIE ST JOHN: So you probably have a lot of customers listening to this saying okay, this all sounds really good, but how do I know if this is for me, how do you begin the process of guiding customers through thinking about this?

JIMMY RAY PURSER: Well, you know, a lot of times it starts with defining really what it is you want to do. A lot of folks look at their applications, and they're kind of thinking that they're probably working pretty good on the network as they are, or they really don't need much optimization. But as their network starts to grow, the answer has always been inevitably throw more bandwidth at it, right? And there's always going to be bandwidth available. We've got 10 gig right now, we've got 10 gig cards right now that are coming out. I even see 10 gig is available on a lot of motherboards for servers nowadays, so speed is never the problem. But that's always a great answer to throw at it, hey, we're both a field salesman, we love selling higher speed switches and stuff, right, that we're talking about more quota. But realistically, the best thing to do is take a look at what it is you need to do to optimize and make best use of your bandwidth out there, not with QoS and stuff, but how to streamline and compress some of these streams. You seem like you're in mid-question.

ROBB BOYD: Well, this may not be the right time to bring it up, but I want to make sure it gets addressed throughout the show. But if I'm a customer looking at what we're talking about today, I think we can all say yes, I've got slow response times, I've always had complaining users, who doesn't, right? But how do I know that my particular situation warrants these particular solutions? And so you don't necessarily have to answer that now, because we're just getting to the beginning of the show here. But I want to keep that in mind because I think that's going to become critical, as we look at these, we go these are great solutions, how do I know they're right for me? Let's make sure we're kind of readdressing that thread as we go through things.

JIMMY RAY PURSER: Well you know, that is a good question, and it's something that's important to ask not only on today's show, today's segment, but also all throughout your network, is is this important and stuff? A lot of times I see a lot of projects fail, a lot of users complain, because their high value transactions, like some of these higher end databases and stuff, they're not optimizing that traffic. Those databases are designed to optimize internally, not externally. You know, they're not designed to actually handle the traffic on the wire. And so being able to accelerate all these processes and our high value transactions, especially if it's coming in from the outside, like over SSL and stuff, and we need something a little bit better than just SSL optimization. There's a lot of boxes I see today that talk about all this load balancing, all they really are is SSL optimizers. I'd buy a couple RIMBO cards, put in a regular server, and do that myself, that's not a big deal.

ROBB BOYD: All right, well let's get into that. I want to get a feel for exactly where we're going to take the show today, because today's show is about clearing up any confusion some may have with how Cisco is solving this application bottleneck. There's innovation here that may have been missed. Cisco groups a couple different

technologies under this umbrella term ANS, or application networking services. And there's two primary product areas under this umbrella. First of all there's ACE, the application control engine, we'll talk about that one, and then there's WAAS, wide area application services, they're not the same thing. We've talked about WAAS quite a bit, I think there's more that we need to spend time on today, right?

JIMMY RAY PURSER: Oh, yes, absolutely.

VALERIE ST JOHN: But before you explain this further, Rob, Jimmy Ray, you actually did an informal survey of some of the customers, and there's a lot of confusion about this, isn't there?

JIMMY RAY PURSER: Oh, mercy, let me tell you. I'll meet with some folks in the field, talking to them about how they're deploying ACE, or looking to deploy it, or really any type of content switching type of solution, if you will. And I asked a lot of folks, I'm like hey, what is the difference between ACE and WAAS? Where do I need one or the other?

ROBB BOYD: And what did they say?

JIMMY RAY PURSER: Well the 10 people I asked, 8 of them said that ACE actually replaced WAAS.

ROBB BOYD: Oh, really, did it now? Someone needs to tell our business unit.

JIMMY RAY PURSER: Yes, that WAAS is the older product and ACE is the newer product. Well, you know, I would hate him to call Norm Skyren, or Sirkin Skyren, I'd hate calling him that too. Hate to call the key WAAS, he would have to abdicate his throne. But yes, there is some confusion out there, a lot of folks really don't see where these are deployed at. And it's tough, man, I've said this before. Anytime we start climbing that OSI model, the higher up the stack we get, the more confusing, the more complex the picture gets. And if we're talking about balancing based upon applications, we're really talking about some very unfamiliar territory, and some developer type language to be able to craft that traffic in multiple ways.

ROBB BOYD: Well I want to take my new whiteboard here, and we're going to attempt to do this again, I want to introduce something new.

JIMMY RAY PURSER: I'm going to take that away from you...

ROBB BOYD: Let's see if we can do some of this detail correctly, and we'll pray that the pen works here correctly. But first of all, let's introduce something new, to map this stuff out, make sure there is no confusion, I'm going to introduce the TechWiseTV Cisco unified acronym demystification guide, the TWTVCUADG, just to keep things...

JIMMY RAY PURSER: You know what, give me that pen.

ROBB BOYD: No, no, hold on, I'm not done yet, this is going to make a lot more sense, because once you've accomplished that acronym, all the rest get easier, right? So we've talked about ANS being kind of the umbrella, ANS is not a service, it's application networking services, right? So within that, and the two areas we're going to focus on a lot today, are first of all, we talk about what is WAAS? Man, we've talked a lot about WAAS on the show, you're talking about our friend Noam, we've got new friends coming in today to take it to a different level. But this notion of wide area application services, just to define it, it's really about improving on the inefficiencies that are

present in the way protocols operate across a long distance, right?

JIMMY RAY PURSER: Yes, right, that's right.

ROBB BOYD: And it's amazing to me how big of a difference you can get. And you say, well I can get you 400% improvement without expanding any of your bandwidth, and people's jaws drop, or they just write you off, and think that's not possible. But in truth we see that easily, and much more, and we've got data to prove that. So we'll talk about that one. But think about all the different places that WAAS -- WAAS is about you've got a device that's going to be out at the office, you've got a device out at the remote office, and you've got a device at the home site, and it's about improving that transmission in between. Which has allowed us, with a couple of new innovations that we've talked about, the collaboration we've got with VMware, so we're doing some of the desktop virtualization, so you can minimize the number of desktops that have to be physically deployed as remote office. As well as this Windows on WAAS, or WoW, is the popular acronym that has emerged from that one, but minimizing the number of servers that have to be deployed at a branch office, all good stuff. We'll talk a little bit about more of that in the next segment, but then we also talked about the other key product area in there being ACE, or the application control engine. Completely different, right? These are very complementary, ACE is about taking the complex transactions that are now becoming part and parcel of today's server application experience in the data center, and saying you know what? You're not good at doing all that complex SSL calculations, right, or even more it's these XML translations that have to happen now that are so complex, servers are buckling under on that load. And we're just saying hey, we've got a dedicated appliance, or a module here, that's going to be able to offload that processing so that the server can focus on doing what it needs to do and what it's good at. At the same time, let's make sure that these other transactions are done more quickly, so that the end result is a better experience for whoever, whomever is actually accessing those applications. Would you say that kind of maps it?

JIMMY RAY PURSER: It's true, it's like a preconditioner, because if we're just looking at standard data transfer like we've been doing for 20 years, these type of products don't really play in any network. But if we're talking about the processes that we have today, an XML process like you mentioned, it's four times the processing equivalency of a standard data packet. So you're really talking a higher end workload here, and we're still talking about the same type of servers. So instead of keep beefing up these servers, buying multi-core processors, and more memory, and faster ports and stuff, precondition that traffic.

ROBB BOYD: Well because you're approaching the problem from the wrong direction when you do that, right?

JIMMY RAY PURSER: Yes.

ROBB BOYD: Because it's the fundamental understanding of what's happening here that's actually causing the bottleneck?

JIMMY RAY PURSER: And XML is a base technology, there's a lot of other protocols that stem off this base technology, so it's just going to get a lot bigger, the problem's going to be a lot worse as time goes.

ROBB BOYD: Well it's going to get worse, too, because we're seeing a growth in XML, right, we're seeing 15 to 50% now in terms of XML.

JIMMY RAY PURSER: That's low, I bet.

ROBB BOYD: Businesses have a lot of their money tied up in terms of there's transaction processing happening with XML, right, so those are pretty key transactions that not only need to happen quickly, but they better happen reliably.

JIMMY RAY PURSER: They're high value, you start losing this stuff here, you can directly tack dollars to XML transactions. Other stuff, you kind of have to fudge it a little bit, but XML, you can tie a dollar to every transaction that goes across.

ROBB BOYD: Absolutely.

VALERIE ST JOHN: Good job guys, boring through some of the complexities here. Looks like a lot of interesting angles to pursue, so let's do just that. In segment two, we'll dig a little deeper, and show you how WAAS deployment just got easier.

VALERIE ST JOHN: And now for a super-exclusive moment. Robb and Jimmy Ray have some unreleased footage. Well, not played directly for this audience, anyway. You guys wanted to tackle a misperception about WAAS, right?

ROBB BOYD: Absolutely, well the misperception was everybody thinks WAAS is very difficult to install, very difficult to begin getting the benefits out of.

JIMMY RAY PURSER: Yes, I sure did.

ROBB BOYD: You know, and you did think that, right? So we approached it, we said well let's take a look at this, because there has been a lot of changes made. But before we get into that, because that's what we cover in the video, where do you think that perception comes from, is it legitimate?

JIMMY RAY PURSER: Well yes, you know, I mean any time you climb that OSI model up the stack, the higher up the stack you go the more complex your configuration is going to be. Configuring a Layer 2 switch, some complexity, configuring a router, even more. Configuring a firewall, even more, I mean we start getting more. And so we're talking about optimizing these protocols, we're really talking about a lot of complexity here. I mean I've set a bunch of these types of products up, and I was not looking forward to playing with WAAS, because I know what a pain in the butt it is, and how tough it is to get these circuits up, and get them matched, and set all my stickiness qualities and all that stuff. But you know what, I was pretty darn surprised. I hate to sound like a marketing commercial here, but realistically, the setup wizards, and the guides and everything was phenomenal, I could not believe how simple it was to actually get it...

ROBB BOYD: Well they've taken that perception head-on, and we cover that in the video. But before we look at that, I want to talk a little bit about the form factors, because at a high level, you're talking about two physical ways in which this thing can be deployed. You've got the appliance which makes a lot of sense, I think people grasp appliance. But surprisingly -- well not surprisingly, I guess, you think about where does a lot of this stuff need to take place? Well, it makes sense for a lot of that to be done in a router, customers who grasp onto that, and as I was talking to the product manager, about where have they seen success with the ISR, he said the number one selling service module for the WAAS is in the ISR, which makes sense.

JIMMY RAY PURSER: Are you serious?

ROBB BOYD: Absolutely.

JIMMY RAY PURSER: It's outselling all the other modules?

ROBB BOYD: No, it's outselling not only other modules, but it's also the fact that it's outselling the notion of doing it as an appliance. I don't want to keep increasing my devices, I want to build it into the router, that makes sense, because it's a router that's going to be putting it across there, let's just build these...

JIMMY RAY PURSER: Like at the branch and stuff, right?

ROBB BOYD: And it makes a lot of sense when you think about it too, because you think about maintaining security from IOS. And you think about when you start deploying external appliances, or you're trying to optimize video and you're doing different things like this, you've got relationships that can be broken, or have to be really tweaked.

JIMMY RAY PURSER: Oh, you're right, because you're not linking in with protocol, you're actually linking it with -- you're tying your security configuration you've got in the ISR, and it's going to inherit those security properties inside that module.

ROBB BOYD: Exactly.

JIMMY RAY PURSER: Good call.

ROBB BOYD: So I think it's a huge point in terms of why it's been working very well there. And then the other notion was a lot of customers were saying you know what, we are latching onto video, that's a huge growing area, but there is a lot of complexity, when they're adding yet another box at a branch office to optimize through video proxies and saying how can we improve the performance of video out of the branch office? They're saying you know what, with WAAS, it's all part and parcel of that package, within your router now, when you deploy that model, that makes a lot of sense as well. And then of course central management makes a lot of sense with that too, because you're not having to go license things separately, configure things separately, you've got them part and parcel of what you're doing with the router.

VALERIE ST JOHN: Well enough talk, you two, let's take a look at what you actually did. Elizabeth McKoin is a consulting systems engineer from Cisco's Channel Organization. She walked you through the install from start to finish, right?

ROBB BOYD: Absolutely.

VALERIE ST JOHN: Let's take a look.

ELIZABETH MCKOIN: What we do with the new WAAS optimization code is really accelerate those uses of the applications themselves across the wide area network, making it feel LAN-like. So the end-user, the customer themselves, they feel like they're actually working on an application that resides local to them.

JIMMY RAY PURSER: Well yes, I mean we've always known those were the benefits of WAAS, but it costs an awful lot of time to get that stuff to work right.

ELIZABETH MCKOIN: Really, actually it doesn't, and I'm going to show you today just how easy it is to deploy a WAAS environment with a new WAAS code.

ROBB BOYD: I like it, so we've got new code, we've got new stuff to look at, that sounds like a recipe for something we enjoy here.

JIMMY RAY PURSER: Yes, absolutely.

ROBB BOYD: Can we get to it?

JIMMY RAY PURSER: Yes, let's put the rubber to the road, let's see what you've got.

ELIZABETH MCKOIN: So you know, I'm going to start off with the central manager, and we have three devices we're going to configure. And so one of the things that we're really interested here with the new WAAS code is to make sure that its ease of use, ease of deployment, so we've created a lot of cool wizards. And actually we're going to go through the CLI, and we're going to use those wizards. So I'm going to start here by actually configuring the central manager, I'll get into the central manager. So everything has actually been set to factory defaults, so there's no configurations, you notice no host name.

JIMMY RAY PURSER: No host name, yes.

ELIZABETH MCKOIN: So we're starting off, and we're going to go ahead and input the wizard here. And the wizard is going to allow us to go ahead and walk through various steps, calling out the steps as we need to actually walk through them. Note too that it makes a recommendation under default configuration, so it suggests hey, do you want to be an application accelerator? Would you like to enable WCCP? So we've thought through all of these processes in terms of what does a typical configuration look like, and it makes a recommendation. Today I'm not going to actually go ahead and apply the default, I'm going to go and walk through the steps so that you can see just how easy it really is.

JIMMY RAY PURSER: Excellent.

ELIZABETH MCKOIN: This is going to be a central manager, so as I configure, it calls out that there is an interface up and running, so I'm going to use that interface. I'm also going to configure in this case I'd like to set it to full 100, just a creature out of habit actually.

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

ELIZABETH MCKOIN: I can't help myself.

JIMMY RAY PURSER: I'm the same way.

ELIZABETH MCKOIN: Right, so no DHCP here, let's go ahead and configure the IP address, looking good, you've got to have your mask, so we've got our mask here, our default gateway, can't go anywhere without it. All righty, and

we also need to configure our domain name, or IP address, cisco.com as well, and we'll name it central manager, kind of intuitive there.

JIMMY RAY PURSER: Get that NTP happening there, it stops all redundancy, right?

ELIZABETH MCKOIN: Yes, we've got to keep the databases synchronized, remember, this is a database. And we're going to go ahead and apply a timestamp so that we can actually synchronize the databases. Now here we've configured an entire WAAS central manager in a matter of probably less than a minute and a half, so let's go ahead and submit that. And this is the only time that we're actually going to reload a device, because we actually changed the identity. We took it from an application accelerator to a central manager.

JIMMY RAY PURSER: Sure, so it's got to boot up with its new identity and its new role, that makes sense.

ELIZABETH MCKOIN: That's right, so let's go ahead and commit it to the config there, reload it.

JIMMY RAY PURSER: Right, man, old school, I like that, that's good stuff.

ELIZABETH MCKOIN: So we'll confirm it, and we'll allow that to go ahead and start the process of rebooting.

JIMMY RAY PURSER: Okay, cool. While you're doing that, why don't we take a look over here, because we see you've got a rack of equipment here, what are we actually looking at, what are we setting up here?

ELIZABETH MCKOIN: So it's a full environment actually, a fully self-contained environment, so we have the data center.

JIMMY RAY PURSER: Use that one.

ELIZABETH MCKOIN: Oh, thank you. So we've got our data center right there, and here's the central manager, the central manager is basically the brains of the operation, however this environment can work without the central manager if for some reason it goes down. You do want to have it there because it is important in that it allows us to monitor the WAAS environment holistically. You could have 2 to 2,500 wide area engines in a single WAAS environment, and that central manager is actually going to manage that entire network.

JIMMY RAY PURSER: Smoking.

ELIZABETH MCKOIN: Yes. And the other thing that it does is it allows us to see alerts. So we can see if there's a CPU utilization issue, we can see disk issues. We can also see whether maybe an application didn't turn up properly, so all of that is being monitored by the central manager across the entire WAAS environment. And I think what's even more important is that we've added application monitoring, we had it before, but we've actually enhanced it here, so we have canned reporting, reporting we can actually schedule and email out to individuals who need to know how their applications are performing.

ROBB BOYD: Right out of the box, you're saying?

ELIZABETH MCKOIN: Right out of the box, exactly. We also have the wide area engine that is setup as a device in the data center. And we are using the inline module here, very quick and easy, basically it forwards packets,

optimizes, so there's no WCCP configuration in this environment.

JIMMY RAY PURSER: Although you could, we're just not using it in this one.

ELIZABETH MCKOIN: Right, exactly.

JIMMY RAY PURSER: And then over here?

ELIZABETH MCKOIN: This is the branch, right? And basically the same configuration, a wide area engine supporting the branch, inline, using the inline module, so again very quick and easy to get up and running. And that's really the most important point of this, ease of deployment.

JIMMY RAY PURSER: Well let me ask you this, so in the newer code, is that all we really did here, we just added ease of deployment, and fixed a few bugs, and there you go?

ELIZABETH MCKOIN: Oh, no, no, it's a rich environment now. The new WAAS code includes the application specific optimization. We've actually gone to vendors and validated the application's performance and optimization, so we've worked very, very closely with them. There's no reverse engineering.

JIMMY RAY PURSER: You went to vendors and worked with them on their own applications?

ELIZABETH MCKOIN: Yes, with Microsoft, we licensed the code from Microsoft. And we did this for Windows Media Services, and for MAPI, and so what we're doing is we're trying to be lockstep with the vendors, no reverse engineering, and hence the customers don't absorb any kind of risk.

JIMMY RAY PURSER: That's very cool, what else have we added to the code to make it pretty smooth? I know we've got to kind of wrap this up pretty quick, but I just...

ELIZABETH MCKOIN: In terms of application optimizers, we've included SSL, HTTP, NFS version 3, we've got Windows Media Services and the virtual blade. And the virtual blade is very cool, because we can literally have Windows 2008 running on a WAAS device in a virtual instance, and we can support up to three virtual instances. So what we're trying to do again is work with those vendors, look to see how we can extend the reach, and allow them the flexibility into the branch with their applications, through the virtual blade.

JIMMY RAY PURSER: No kidding.

ELIZABETH MCKOIN: Yes, it's wonderful.

JIMMY RAY PURSER: Can you support Active Directory on that?

ELIZABETH MCKOIN: Absolutely. And read-only mode, security is a big deal, we have DNS and DHCP servers.

ROBB BOYD: So we're collapsing services at the branch. Those necessary things sometimes do benefit by being a little bit closer to the user, so it's all about increasing that user experience, right?

ELIZABETH MCKOIN: Exactly, and providing the optimization that they want.

JIMMY RAY PURSER: Hey, you know, I'm a little selfish here, because I've setup WAAS before, and I really want to get into setting this back up, but I've just got one last question to ask you here on some of the features. I also heard that you guys made some changes in CIFS and print.

ELIZABETH MCKOIN: That's right. So CIFS and print services, but CIFS is now transparent. No more core/edge, and so no more connectivity directives. Literally you just enable it, and it's up and running, and you're going to actually see where that happens, it's done automatically now.

JIMMY RAY PURSER: That is so awesome.

ELIZABETH MCKOIN: It is.

JIMMY RAY PURSER: Now what do we need to configure here on the CM, is there anything else we need to configure here, or are we ready to jump straight to the...

ROBB BOYD: Yes, you guys better boogie.

ELIZABETH MCKOIN: Actually I'm going to go ahead and enable the CMS, so I'm going to load up the database now. And then we're going to move to the data center wide area engine. So let's do a CMS enable. And while that's actually going on its merry way, why don't we go ahead and start the configuration?

JIMMY RAY PURSER: You got it, let me change this here real quick.

ELIZABETH MCKOIN: On the data center, wide area engine.

JIMMY RAY PURSER: All right, you're in.

ELIZABETH MCKOIN: Excellent, thank you, beautiful, makes my life so much easier.

ROBB BOYD: Nice to have a handy assistant like Jimmy Ray.

ELIZABETH MCKOIN: It's nice. Okay, so we're now in the wide area engine, and this is the data center side, and we typically are going to have these in a cluster, for resiliency, right?

JIMMY RAY PURSER: Of course, right.

ELIZABETH MCKOIN: So again, it finds the inline module, it knows that it's available, it's asking do I want to apply this basic configuration with Autosense enabled, but because my configuration is slightly different, I've got this nice little kit that I carry around with me, I'm going to say no to it. And I'm going to configure it as an application accelerator. I'm going to specify the central manager's IP address, look, it detected inline, I'm going to accept that. It's also giving me the interfaces that are up and running, and so what I want to do is I want to configure my management interface, on option 3. I'm going to go ahead and configure the duplex, because I'm a creature of habit.

JIMMY RAY PURSER: That's just how it is, yes, exactly.

ELIZABETH MCKOIN: All righty. And so let's get that straight. All right.

JIMMY RAY PURSER: That can cause some problems, huh?

ELIZABETH MCKOIN: Yes, we don't want to give it the wrong IP address here. So let's go ahead and configure the default gateway. Oops, there we go again. And specify the name server, really easy stuff, cisco.com, all right, and simple name. And that NTP server, we've got to synchronize the clocks, and set our time zone, and let's see what the offset is, 06. And I'm going to do everything in line, so I'm going to choose my inline module. And again, creature of habit, I like to set it. And now, if you notice we've got all these licenses?

JIMMY RAY PURSER: Yes, what's that?

ELIZABETH MCKOIN: So these are the new licenses.

JIMMY RAY PURSER: That's new.

ELIZABETH MCKOIN: Yes, well, there was transport and enterprise, but now you notice that we have video and the virtual blade. And you do have to purchase licenses to support the Windows Media services, Microsoft likes to get paid.

JIMMY RAY PURSER: Well, yes, because you actually had to engineer with them, right, of course.

ELIZABETH MCKOIN: Yes, right, and so we pay for the privilege. And the virtual blade as well, so there's a fee for the virtual blade. And you can always upgrade to the virtual blade and add the video services later, so that's very...

JIMMY RAY PURSER: Do I have to come back and walk through this whole config again to do that?

ELIZABETH MCKOIN: No, very easy installation through the device manager.

JIMMY RAY PURSER: Excellent.

ELIZABETH MCKOIN: So nice GUI, configuration. So we're going to choose enterprise, I don't have the Windows Media service running, and we're going to accept the configuration, we're going to apply the configuration. And so while that's happening, why don't we go ahead and move over to the branch?

JIMMY RAY PURSER: You don't have to register this with the CM?

ELIZABETH MCKOIN: No, actually it does it automatically. I don't have to reload this.

JIMMY RAY PURSER: Wow, that is really cool.

ELIZABETH MCKOIN: So sweet and easy.

JIMMY RAY PURSER: Well let me move you down to this other one here.

ELIZABETH MCKOIN: All righty, so we're back to the default configuration here, and I'm going to go ahead and do setup again. Now again, if a customer had a configuration in which any one of these applied, all he'd have to do is answer yes. But because I have a unique configuration in my little box here, I'm going to say no. All right, application acceleration, tell it where to go to get the central manager. It detected inline, and look, it knows that all my interfaces are up and running. Pretty much the same configuration, I like consistency, creature of habit here. All righty, no DHCP, let's configure an IP address, and our mask, there we go, our default gateway, pretty easy stuff. And we've got to know our -- there we go, cisco.com. And let's give it a real unique name like branch.

JIMMY RAY PURSER: There you go, ain't nothing like uniqueness.

ROBB BOYD: Won't confuse anything there.

ELIZABETH MCKOIN: Okay, and we'll set the clock or the time zone, just making sure our databases are synchronizes, and the clocks are set equal, enable that inline module. I'm going to say no here, all righty. Again, enterprise license is what I'm selecting, and there's my configuration. Done, enter, enter, I am ready. And so now it's going to register with the CMS, and so as that actually communicates back to the GUI, we can go in and we can start managing the entire device, this is the only time we're ever going to touch the CLI.

JIMMY RAY PURSER: I cannot believe you setup...

ROBB BOYD: And these are pretty much standard configuration steps too, I mean those are things any networking person would probably be doing on any configuration, because you can't get past some...

JIMMY RAY PURSER: I can't believe you setup three of these in like 10 minutes.

ROBB BOYD: So this is interesting, but how do we know that it took?

ELIZABETH MCKOIN: Well let's go ahead and go through the central manager. Now you can see that all of the devices are online, they join, status is green, we're looking really good here, right? So let's go ahead, and what I'd like to do is show you our validation page in terms of the device groups. So we're going into the device groups, I'd like to show you where we actually have those application optimizers and accelerators already enabled by default. So literally out of this configuration that we spent only minutes on, we're able to actually bring up a system, have it fully configured and enabled, again, CIFS configuration is easy as hey, up and running...

ROBB BOYD: Yes, well literally I think this was under 10 minutes. I mean we talked a lot through this, kind of understand some of the background, and some of Jimmy Ray's uh-huh moments, which are always enjoyable for everybody.

JIMMY RAY PURSER: Yes, but your comment though on how do we know this works also reminds me, you guys also added a lot of new reporting to this thing too, right?

ELIZABETH MCKOIN: Absolutely.

JIMMY RAY PURSER: We don't have time to cover it all, but you've really added a whole bunch of other stuff.

ELIZABETH MCKOIN: Just a tremendous amount of reporting, about 45 reports, and probably still growing, you can

customize your reports. And then you can schedule them and email them out to others, so that's what makes it so rich.

ROBB BOYD: Well I think what I like here too is if there was any smaller businesses that maybe didn't have the depth of IT talent that was familiar with this kind of thing I think we've removed a lot of the fear, because the steps that can be automated are automated, a lot of things are self-discovering.

ELIZABETH MCKOIN: Really, you don't have to have trained individuals installing at the CLI because they literally can walk through that CLI.

VALERIE ST JOHN: Great job, guys, and thanks to Elizabeth McKoin, Consulting Systems Engineer for Cisco's Channel Organization, very nicely done. In segment three, we're dealing all aces with James Schultz, our ace on ACE, coming up.

VALERIE ST JOHN: Online, on demand, and on the go, this is TechWiseTV, technology you can use from geeks you can trust. Today we're focusing on application delivery with WAAS and ACE. We've seen how cool WAAS is, now let's turn to ACE. Robb, you mentioned earlier that ACE is really good at providing assistance at the host location, right?

ROBB BOYD: Right, there in the data center, the idea is to have something in front of your key applications that are offloading and optimizing the processes that have to occur but maybe aren't best done at a server, can be better done on an appliance. Which is really not a new concept, this notion of an application front end, we used that stuff back in the 60s and 70s, with these AFEs with IBM mainframes, right? So we've optimized these processes; that is really all ACE is doing these days, which is saying hey, this stuff is better done over here, man, leave the focus where the focus needs to be, and improve the performance.

VALERIE ST JOHN: So ACE is really good at relieving the server of these important tasks that the server's not really good at accomplishing anyway.

ROBB BOYD: Absolutely.

VALERIE ST JOHN: And the result?

ROBB BOYD: The result is better performance, wouldn't you say, Jimmy Ray?

JIMMY RAY PURSER: Faster than an in-laws phone call after winning the lottery.

ROBB BOYD: Oh wow, that's a good one.

JIMMY RAY PURSER: They cook, boy.

VALERIE ST JOHN: So let's welcome our guest James Schultz, engineer with the Application Delivery business unit. How would businesses start using something like ACE, James?

JAMES SCHULTZ: Well to start off with, the industry recognizes ACE and the product group that it fits into as application delivery controllers. And the way I like to describe that or explain it is it's kind of the next generation of server load balancer. So going back to the original question, the first role that businesses would typically use an ACE for would be as a typical server load balancer. So smaller businesses might need additional servers to add redundancy for their website, larger businesses just need to be able to scale, so they need to add more servers. And the ACE as a traditional server load balancer gives them that functionality. The second main role that they can deploy ACE for is application acceleration. And what we've done here is we've taken the patents and everything we acquired from the FineGround acquisition, which ultimately became the AVS3110 and 3120 products, and we've ported those features into ACE. So you can now take the ACE, put it in front of your applications, and do some interesting things with the web pages as they're being sent back to the clients.

JIMMY RAY PURSER: Like what?

JAMES SCHULTZ: So basically we can do things like look at the HTML as it's being sent back, we can alter the HTML to make the client more aggressively cache objects that are in the web pages, so they don't come in and request them. So instead of doing things like compressing the data which we also do, we can actually make it so the client doesn't request the data at all. We don't even have to worry about sending 304s back, or any sort of that type of thing. It saves businesses a lot of bandwidth.

ROBB BOYD: I want to talk about how that kind of thing is deployed, because I think when you start talking about delivering that kind of performance improvement, there's two things you brought up there that I'm curious about. Maybe I've been overly hyper-focused, if you will, on this notion of improving the application performance, but you mentioned load balancing as well. Where does load balancing fit in that equation, because that obviously is still an important part of understanding the value provided, right?

JAMES SCHULTZ: Well load balancing really is the fundamental portion of ACE, so almost every customer that uses ACE is going to be using the load balancing functionality. And layered on top of that is where they add on application acceleration or security, which is the third role I was going to get to.

ROBB BOYD: Oh, the whole security side of things, okay.

JAMES SCHULTZ: Yes, so the third part, security, basically we've added in a stateful firewall, so we call it the last line of defense. Typically we still recommend you have a full-blown firewall, an ASA, the gateway to your data center. And we have the ACE in place on the backend in front of your servers. So if an attacker does happen to break into your data center infrastructure, you still have a last line of defense between any of the systems they might have compromised and your server sitting on the backend.

JIMMY RAY PURSER: You don't want to accelerate them, do you?

JAMES SCHULTZ: No, you definitely don't want to accelerate attacks.

ROBB BOYD: Well what are we looking at from an installation perspective, when you talk about how hard is that kind of thing to get started? Because one of the things that I was understanding is the whole notion of transparency in terms of being able to do these things -- because the first thing you always think about is we can do this, but oh, by the way, you're going to have to deploy all this software on your end users' computers to get these kind of benefits. But that's not the case here, right?

JIMMY RAY PURSER: Yes, exactly, I mean a lot of folks don't really look at setting stuff like this up, load balancer stuff that I've setup, I mean they typically do require what you're talking about, at least software on the servers and stuff, and they're harder than a brick bath actually to configure and get going and stuff. So what is the deal?

JAMES SCHULTZ: So there's a couple things that can help here, and it's actually one of the key points that we can mention on how ACE differs from WAAS, is that ACE doesn't rely on any sort of software on the client side, ACE does everything within the HTTP protocol in terms of optimization. So the acceleration that happens is done through standards based protocols, and the client web browser knows how to interpret this data. Whereas with WAAS you need to have a WAAS device on both the server side and the client side.

JIMMY RAY PURSER: Oh, right, of course, you have to have terminating points.

JAMES SCHULTZ: Right, so it's got to decrypt the traffic, and unencapsulate it, and all that stuff. Whereas with ACE, we just send it directly to the client. And in terms of ease of use, we do everything through, we have a product called the ANM, and there's also an embedded version called the device manager, which is built into our appliance. And what those do is provide a full web GUI for configuring the load balancing, security, everything in one pane of glass.

JIMMY RAY PURSER: Well you mentioned appliance, what other form factors do we have on this product as well? We've got appliances, we've got...

JAMES SCHULTZ: So currently we have two form factors, we have the ACE module which fits into the Catalyst 6000 chassis, it's great for enterprises that already have Cat 6000s deployed, because they're all over the place. And that thing scales from 4Gb per second up to 16Gb per second. And the good thing is you can scale it using a license. So we can sell you a 4Gb per second module upfront...

JIMMY RAY PURSER: Oh, that's right, it's not a hardware upgrade, oh, that's great.

JAMES SCHULTZ: Right, basically we call it pay as you grow. So as your needs increase, as your business grows, you just buy a new license, and we give you more throughput.

JIMMY RAY PURSER: Man, I like that idea, because once you set these things up, if you've got to actually replace these, and redo the hardware, it doesn't make sense to even buy them, because the configuration time and stuff. So set it up once...

ROBB BOYD: It takes a forklift upgrade.

JIMMY RAY PURSER: Yeah, that's really cool, I'm glad we do that actually.

JAMES SCHULTZ: And the other good thing is we have the other form factor which is the 1U ACE 4710 appliance, and it scales from 500Mb up to 4Gb per second in the same manner, you just use a license. But the configurations are the same, so you can just move them directly over from your appliance to a module if you do move.

ROBB BOYD: That's relatively new, right? I mean by coming out with that, we've opened up those capabilities to be experienced in what Cisco might call a commercial marketplace, or kind of the smaller customers who are always

like well, I love the benefits, but those are big company solutions, I'm not going to necessarily, if I don't already have a 6500 chassis, that may not make sense for me. But now we've got the 4710, allowing people to reap some of those benefits. And this may be an overly obvious question, I always see your whiteboard sitting here, and I'm thinking you've got to draw something, does it make sense to kind of make sure we understand where this sits in the data center? Or is it just as simple as it sits in front of the applications, Robb, come on?

JIMMY RAY PURSER: Yes, and before we do that too, as we're doing that, the discussion here, because one of the things that I thought was the coolest about this, because in the world of server load balancers and stuff, there's a lot of datasheet points, that ours does this, you know what I mean, we get into so much minutiae, the higher level we go, the more crap we can pull out of the air and say well we can do this and they can't, and a lot of people say well I'm not a developer, I don't even know, and do I even care? One of the things that we do that I think is very important, that is really critical, really sets this product apart, is the virtualization aspect. Is that we talk about form factors and stuff, and that's a very tiny picture up here, the mere fact that we can start virtualizing this, and start growing this device, so I don't have to put a whole bunch of them in there, means a bunch of things, right? I've got great security processes, I've got great granularity, I also save a ton of money in heating and cooling because I can virtualize this stuff. Let's get over here and draw some stuff out, show us actually how you would put this in your network, and anything else cool you want to show us.

JAMES SCHULTZ: Okay, so start out with your basic internet cloud here.

JIMMY RAY PURSER: Which is required by law.

JAMES SCHULTZ: And sitting above it you have your data center. So within your data center you're going to have your gateway router, and beyond that we're going to stick an ACE. So all your traffic is going to go through the ACE here, and above that you might have your Layer 2 switch.

JIMMY RAY PURSER: Your LAN switch.

JAMES SCHULTZ: That's a switch, I can't draw diagrams very well. And behind that you've got your servers. So basically you're using the ACE here to load balance the different servers. Across the backend we can take them in and out of service using health probes, if a server fails or you need to do maintenance.

JIMMY RAY PURSER: So the ACE is sitting in line?

JAMES SCHULTZ: Yes, the ACE is sitting inline.

JIMMY RAY PURSER: How is it determining this traffic? How is it inspecting to know that this is traffic it needs to optimize?

JAMES SCHULTZ: So there's actually a virtual address that you configure on the ACE, which is what the clients are going to, so when they type in www.cisco.com, it resolves to a virtual address that lives on the ACE. And the ACE then accepts that traffic, changes the destination to one of the server's addresses, and passes it back on the wire.

ROBB BOYD: And that's part of how we achieve that transparency, right, in terms of not having to change anything from the client side? Before we run out of time, I want to make sure we understand, you mentioned three unique attributes. You wanted to make sure we hit fast, green, and secure. Fast, what examples would you cite for that?

JAMES SCHULTZ: So examples, from the beginning, we basically designed ACE to be extremely efficient in handling this really deep inspection of the different protocols. So we built it on a multi-core architecture, so the ACE appliance uses 16 processor cores, and the ACE module uses 32 processor cores.

JIMMY RAY PURSER: Jimminy Christmas, man.

JAMES SCHULTZ: Versus the competition, which uses somewhere around four processor cores on their highest end boxes. We've also bundled in hardware compression on the ACE appliance, and hardware SSL on both form factors, and those are both options on a lot of the competitors' products.

ROBB BOYD: Okay, so when I hear us doing that many more processor cores, it makes me wonder why you bring green into this as your second point, because it sounds like it would use a lot more energy. So someone in marketing has told you that you have to talk about green, so fess up.

JAMES SCHULTZ: Right, and so from the green perspective, there's a couple things we can talk about. Basically we have the device virtualization that was a topic we mentioned a minute ago. And what we can do there is we can carve up a physical ACE into multiple virtual contexts. So on the 4710 appliance, we can have up to 20 virtual contexts.

JIMMY RAY PURSER: Jimminy Christmas, man.

JAMES SCHULTZ: And on the module we can do 250.

ROBB BOYD: Okay, so you take that increased performance, and then by virtue of virtualizing, the idea is that you don't have to buy a new appliance or new module for every application or group of applications you need to optimize. Also opens up the management benefits too.

JAMES SCHULTZ: Right, and also in the service provider market, it works very well for service providers. Instead of having to buy a separate ACE for each customer, they just carve it up into different contexts, and hand them out to the different customers.

JIMMY RAY PURSER: Oh, that makes a lot of sense, because they're doing the hosting and stuff like that. Also works great for the eighth layer of the OSI model too.

ROBB BOYD: Oh, the whole political side of things?

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: You get role based access control and things like that, so people focus on just their virtual instance that they need to manage and apply their policies to, that works pretty well. Security, near and dear to our heart, was your third option as well?

JAMES SCHULTZ: Right, so security, as I mentioned, we have a stateful firewall that's built into the ACE, and since it's virtualizeable, you can actually put that stateful firewall kind of wherever you want it on the network. If you have a multi-tiered application, like you might have a web tier, an application tier, and a database tier, you can

actually put firewalls between each of those tiers if you wanted to.

ROBB BOYD: Wow, that's a lot of granularity.

JIMMY RAY PURSER: Wow.

VALERIE ST JOHN: So it's fast, green, and secure, attributes we can all keep in mind, thank you, gentlemen.

ROBB BOYD: Good summary.

VALERIE ST JOHN: I like to keep it pithy. Up next, ACE meets WAAS. They said the union would never work, we'll show you why it does, coming up.

VALERIE ST JOHN: So we've looked at WAAS and ACE separately, let's look at how they work together. Seems there may be some confusion that they have overlapping functions initially, which is why we invited Gabriel Dixon, from the Enterprise Solutions Engineering team, or ESE. A cool acronym, because you can pronounce it easy, thanks for joining us. So Gabe, tell us what the ESE team does.

GABRIEL DIXON: Sure, thank you, first, for having me here, Valerie. ESE stands for Enterprise Solutions Engineering, and we are a technical group within Cisco's central marketing organization. And what we do, we build complex networks and application solutions for enterprise and midmarket customers. We focus on what we call PINs, or places in the network, those being campus, WAN, branch, and data center. My particular focus area is in data center applications. And the program that we're working on is DCAP, data center assurance program for applications. And what we do is in that program we work very closely with third party application vendors such as Microsoft, IBM, Oracle, you name the vendors, both large and small. And we work closely with them to jointly test, jointly validate, and collaborate at every step of the way, so that our customers could more rapidly deploy and competently deploy these solutions in the network.

ROBB BOYD: And this is why I'm glad you're here, because I think this is a unique point that a lot of customers may not get in terms of an investment that Cisco's made, to be blunt about it, to our customers' benefit. But the idea is that we say our boxes and applications and things that we want you to buy from us will do these things. And what your team does is say okay, you say you make this application work better, well we're going to test it out. We're going to put it in a real world environment, we're going to benchmark these results, we're going to see what stuff comes out of it. And each one of you guys focus on very unique areas to do that. So that's where my first question comes in, we're talking about ACE and WAAS, we're talking about getting rid of this application bottleneck, and increasing performance from the user and that endpoint perspective. What type of unique results are you guys seeing in that area since that's where you focus? Anything you've learned that you're passing onto customers there?

GABRIEL DIXON: Sure, sure, we tested a slew of applications, and all of them yielded very good results. One example that we tested was Microsoft SharePoint, which is a web-based collaboration tool.

JIMMY RAY PURSER: Oh, right, that's big.

GABRIEL DIXON: And with that application, with both WAAS and ACE in the picture, we saw fabulous results, and I've got a few notes on some of the results.

ROBB BOYD: Share with us a few details.

GABRIEL DIXON: And with WAAS we saw a 34 times faster downloads, 50% faster site navigation, and 88% reduction in bandwidth going over the wide area interconnect.

JIMMY RAY PURSER: Yes, but you're doing this in a lab right, I mean couldn't you kind of expect, you're going to get good results? I mean I get good results in my lab too, I'm a little skeptical of your numbers here. I don't mean to be a jerk, but...

GABRIEL DIXON: No, that's a good question, Jimmy Ray.

ROBB BOYD: But that's your role on the show, Jimmy Ray.

GABRIEL DIXON: Yes, and a fair question, we get that all along. But remember, we are working side by side with the application vendors, at every step of the way, and they have to certify this. So what we're doing...

JIMMY RAY PURSER: Wait, your application vendors certify the stuff we do?

GABRIEL DIXON: Yes.

JIMMY RAY PURSER: Oh, I didn't know that.

GABRIEL DIXON: Yes, and it's posted on their websites, on the partner proxy.

JIMMY RAY PURSER: Is that right?

GABRIEL DIXON: Absolutely, absolutely.

ROBB BOYD: Well it makes sense, right, because they go we make them look better. Why wouldn't they want to be involved?

JIMMY RAY PURSER: Well you know that, that's interesting because I was out at an Oracle deployment about a couple months ago, and I was introduced to your config guides by one of the installers out there who was saying that yes, since we're doing this on Cisco, they did the joint testing. And Oracle recommends that you do not set the sticky persistence, but Cisco does because of the way this architecture works. And sure enough, we did on and off, and it made a huge difference. And it's like those are little tricks you learn along the way kind of trial and error, and it's great that you don't have that much trial and error with the stuff you guys make.

GABRIEL DIXON: Absolutely, we're putting the customer hat on, and we're vetting this out, putting it through its paces, and we're trying to take the guesswork out for the customers, so that they won't hit some of these gotchas.

ROBB BOYD: Well I think that's what Jimmy Ray just hit on, but what is that deliverable that the customer sees? So a customer watching this goes okay, sounds interesting, what do I do now?

GABRIEL DIXON: Right, right, generally there are two deliverables. One is more technical, and it's prescriptive, if you will, in the form of a design guide. So that goes out to our field and the customers, and it has very step by step guidance on how to configure your network devices, how to configure your optimization devices, your ACE and WAAS, and even application. And that's critical in this test. And the other set of deliverables are more business oriented, and those are aimed at the business decision maker and C-level execs.

JIMMY RAY PURSER: Yes, kind of justifiable clause, and stuff.

ROBB BOYD: Instead of a generic manual of hey, in most situations you would configure these things, it's specific to the applications. You guys are probably focusing on the applications customers use most, you've mentioned Oracle, SharePoint, some others. But what do we see going forward? So we see virtualization, and some other technologies really coming into play, a lot of customers having to work out how do I deal with these new aspects of technology moving forward? I assume you guys are keeping your eye on those trends, right?

GABRIEL DIXON: Oh, absolutely, absolutely. We are currently in the throes of testing Microsoft Exchange 2007 in a completely virtualized environment.

JIMMY RAY PURSER: Are you serious?

GABRIEL DIXON: Absolutely.

JIMMY RAY PURSER: That's a big no-no, I mean you never virtualize Exchange.

GABRIEL DIXON: Well that was their initial position, but with HyperV, when that came out, they sort of softened their position on virtualization.

JIMMY RAY PURSER: Really?

GABRIEL DIXON: Yes.

JIMMY RAY PURSER: Yes, I guess it's their own product and stuff. But how does that work, though? I mean I've tested HyperV too.

GABRIEL DIXON: Right, well we're testing VMware, we're partnering with VMware, they are the market leader in virtualization, and we're partnering also with EMC. So all of the Microsoft Exchange roles are virtualized, and we're bolting on some of our ACE and WAAS products with the latest version of code to make sure that entire environment can be virtualized.

JIMMY RAY PURSER: Well I know it's probably early, but what do you think so far? How's it going?

GABRIEL DIXON: So far the preliminary results are very positive, especially the replication traffic going from one data center to a hot or standby data center. We're able to optimize that replication traffic going over the WAN.

JIMMY RAY PURSER: Seriously?

GABRIEL DIXON: Yes, yes.

JIMMY RAY PURSER: That's huge, I mean that's massive, because a lot of people do co-lo like that, and getting that cold data center back online, that's a pain in the tail.

GABRIEL DIXON: Right, and you're saving money on circuit costs.

JIMMY RAY PURSER: Yes, that's true.

ROBB BOYD: Making room for more applications to work across there. Well this has been very good information, Gabriel Dixon with the Enterprise Solutions Engineering team, thank you so much. Guys, before we go, if you need more information on that, of course please check out cisco.com/go/dcap to get more information on that. And to get the latest on what they've been testing on stuff, we will be keeping the show notes updated as well of course.

VALERIE ST JOHN: Thanks, Guys. Next we're happy to welcome Dave Major from Learning@Cisco, hello, Dave.

DAVE MAJOR: How are you doing, Valerie?

VALERIE ST JOHN: Good, so you've got some big news today?

DAVE MAJOR: We do indeed.

VALERIE ST JOHN: Cisco is introducing wireless certification at the expert level?

DAVE MAJOR: The brand new CCIE wireless, yes, we're really excited about it.

VALERIE ST JOHN: What's the motivation behind it?

DAVE MAJOR: The motivation about it is we see the technologies expand, we see the network, as you could say, converge, we want to have certifications at every level. We introduced earlier this year at the associate level CCNA wireless, CCNA security, and CCNA voice, and so now we're expanding it also at the expert level.

VALERIE ST JOHN: Excellent, when is it expected to rollout?

DAVE MAJOR: So actually the written test will be available on February 17th, 2009, and then the lab exam will be available in March.

VALERIE ST JOHN: Are there any formal prerequisites, or can I do it?

DAVE MAJOR: Well you have to have some kind of a background, absolutely, I mean there is no formal prerequisite for any of the CCIE certifications, but you should know your stuff. I mean you should be at the expert level in terms of networking technologies.

VALERIE ST JOHN: So if I'm interested in finding recommended reading and blueprints, where do I go, how do I find information?

DAVE MAJOR: So that's another great question. We actually have a brand new, we also launched this a little bit earlier this year, which is the Cisco Learning Network, and you can find recommended training, you can find lab blueprints, you can find everything there. The URL is www.cisco.com/go/learnnetspace, and you'll find everything you need right there/

VALERIE ST JOHN: Excellent information, Dave Major with Learning@Cisco, thanks for joining us.

DAVE MAJOR: Thanks so much, Valerie.

VALERIE ST JOHN: And of course Robb, both ACE and WAAS are supported by our friends in Advanced Services, please expound.

ROBB BOYD: You know it's all about implementation. Data centers are complex environments that can often benefit greatly from the expertise offered by Cisco partners, as well as Cisco's Advanced Services group which can serve in a very complementary fashion in that event where you've got a favorite partner that you want to continue working with, but they may not have got all the expertise built out exactly where you want it for these particular applications. The Cisco planning, design, and optimization services for ACE and WAAS technologies can really help accelerate deployment by providing deep expertise and best practices to help you define the IT and the business requirements, as well as implement that customized solution that makes sense for you. These services really round out the parts of the product that can be challenging. So think about implementation, is it new or is it migration? The architectural focus, how do we built out a blueprint? Availability we talked about, load balancing, security, and how that plays into your risk. How do we reduce potential issues both before, during, and after? At the end of the day it's about getting a dynamic, policy-based model, we talked about that, right? You could live with post implementation, deliver continuous load balancing, security updates to your operating environments, incorporate the ongoing changes you need to grow your infrastructure, and respond to the dynamic business requirements. For more information on the Advanced Services stuff, please go check out cisco.com/go/dcservices.

VALERIE ST JOHN: And for Robb Boyd and Jimmy Ray Purser, I'm Valerie St John, thanks for joining us on TechWiseTV. To find out about future episodes, or to check out a complete archive of previous shows, be sure to visit the Cisco Interaction Network at cisco.com/go/interact.