Dawn of the New Data Center

YOU SAY YOU WANTED A REVOLUTION?

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

JIMMY RAY PURSER: Hey Valerie, how are you?

VALERIE ST JOHN: Good, our focus today, the dawn of the new data center. Now in the past you guys have maligned the use of the word of unified as overused, but this time's different, this is huge right?

ROBB BOYD: It really is different. In fact one of the things that I think is going to become very important to this show is that we understand that there's a true paradigm shift, I don't think that's a term that's been abused in a long time. But there's very real technology that backs up why unified, if you were to pick any word to try to describe the level of innovation and such that we're going to be talking about here, that is one that describes it very well. And so I think it's important and there's real technology to back it up.

VALERIE ST JOHN: Just to make sure that I understand this a little bit. Let's take it back, let's talk about the history here. What's been the evolution in the data center, what's changed?

ROBB BOYD: Well a lot has changed in the sense that if you go back far enough, you look at we had mainframe days where everything was centrally contained. Then as intelligence got put into the endpoints, we put those in the hands of users, we lost a lot of control but we gained a lot of things and we took that same PC model that developed out of that and we've attempted to bring it back into the data center. But there's been problems from that because we've got a PC model that's really not built to handle the stresses that we're putting data centers under now. And so you know a lot of innovation has been attempted to bring that under control but fundamentally it's been the same thing attempted over and over again, when in reality perhaps we can look at it a different way and that's just something we'll talk about in this moment, but could we do it differently?

VALERIE ST JOHN: Right, and which leads us to this question of business agility, the whole idea that must respond very quickly to the business challenges right?

ROBB BOYD: And I'm glad you brought that up because I think business agility is exactly what we're talking about here when you talk about the desired end result. So Jimmy Ray and I, we may get wound up about talking about technology and we probably could get carried away with talking about technology for technology's sake, but at the end of the day the reason the IT department exists is to execute on something that the business has said hey we need to do something to take care of the changing market condition. The speed in which you can do that, the speed in which the business uses their intelligence, using this loosely, and says we need to deploy these applications to do this, how fast we can actually make that happen is what we might call business agility and becomes very crucial to being competitive to be able to differentiate. And in fact I'm reminded of a story, customer had made a reference to the Golden Gate Bridge and describing how they rolled out applications. They had, I think in this particular situation it was around 400 locations which is not an atypical situation really. And the problem they've got in rolling out technology is they said it's like painting the Golden Gate Bridge. We start at one end, start with one location, and then we go through each location rolling out the new application. We're doing this to respond to a market condition but it takes 6 to 12 months. You know as soon as you get that one done, what do you do, you start over and start painting again right, because at that point it's time for the next thing. Well how agile is that if it takes that long to roll out an application that the business already decided, in fact...

VALERIE ST JOHN: Not very.

ROBB BOYD: What is the reality behind the scenes? Because you've dealt with this before, what happens? I mean how different people are really involved when it comes to deploying something new?

JIMMY RAY PURSER: Well you know it's a real pain in the tail. I mean the further up the ladder that go, the bigger your business gets the more hands you kind of have involved in here. Smaller business rollout applications. That may not seem like a big deal, I go out and buy the apps, I buy the license, I install on the server, everybody uses it and have a nice day. The larger, the more data center base-type customers and stuff, typically if I'm going to even put a new application on I got to go through some type of change management protocol. I got to actually have backup strategy, I got to have plans, I got to have maintenance windows to actually set all this stuff up, I have to have people standing by. You know it is a humongous task to rollout. I mean a patch for the application.

ROBB BOYD: But aren't we talking about different people handling different levels of expertise?

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: I mean you've got server people doing one thing but then they rely on getting data from a networking person who's got to give them certain pieces of information, then they've got storage. And you have these different elements and it becomes kind of this complex orchestra that's got to happen every single time an application comes out right?

JIMMY RAY PURSER: Absolutely, you got to have NOC people that are watching the stuff, seeing any ripples in the network happen.

Because if I have a small -- let's say I've got an international company and I have a small blip in one data center site, the hysteresis effect of that, that rippling effect going all over the world is huge. And once it starts, it is nearly impossible to stop.

VALERIE ST JOHN: So what you're describing here is not just a major logistical problem, this is a major fiscal problem. What's been the solution up to now?

ROBB BOYD: Well here's the thing. And this is where I say it requires a paradigm shift because we describe this, obviously we're trying to do -- I think, let's be straight up about it -- is a typical presentation fashion, we need to understand the pain that we're here to solve. So that's why we're bringing it up, we just want to make sure everybody understands the area that we're trying to address because things we're talking about are normal, these are status quo for any operational environment and it isn't often questioned about how do you actually change things. Because the stuff -- back to your question what has been done to change it, well if you -- I mean I oversimplify it to a certain extent, but really all we've done in the past as an industry is what we've either sped it up or maybe we've overlaid some management on it to try and solve some of these issues. But what happens when you do a lot of that is you simply either speed up or introduce more complexity to the situation. Now virtualization is probably the most fundamental game changer, if you will, in terms of what it brings to the industry, but virtualization for the benefits it brings to optimization and better utilization of resources also introduced a lot of additional complexity or problems that now had to be answered because we didn't even know that the question needed to be asked until we actually did it. And so addressing how virtualization plays into that is still important, we just need to understand how to deal with it.

VALERIE ST JOHN: Is Cisco doing anything to address the challenges here? Yes we're a company that makes acquisitions and things like this. But if peel back that onion a little bit further and you look at the history and where some of these things have come about, you're doing a disservice to us and yourself if you don't realize it. In this particular situation we acquired a company called Nuova in last -- I think this is the last few years, I'm not good at quoting the dates. But essentially Nuova was a company created by some very smart people at a high level within Cisco that said we're going to go out and solve some problems that haven't been solved before. And they went out -- but they didn't do this alone, they looked at the problems that they wanted to solve in the data center and they brought in experts from the virtualization side, founder from VMware, experts from Sun, you know server experts, virtualization experts, and they got their heads together and gave themselves permission to start over, permission to say hey let's ignore legacy for just a moment, let's look at the problems that are really driving us nuts and say if we could ignore how all this stuff got to where it is and just start over, what could we do differently? Give themselves permission like that and innovation comes out of that kind of thing. And so I know I'm leaving a lot of the details out but I just want the groundwork laid for what we're going to talk about here for this notion of Unified Computing.

VALERIE ST JOHN: Right, which allows us to expand on that a little bit. Why don't you start by just defining that for us in the scope of this conversation? What are we really talking about here?

ROBB BOYD: So as I mentioned before, I think unified is the best single word to say it. Then the next question becomes what are you unifying? And we still laugh at the term, I know.

VALERIE ST JOHN: Still using that unified, can't get beyond it.

ROBB BOYD: But it makes sense. Because what we're doing is, from the definition, back to your question is we're unifying the compute area that's normally handled on servers, and still is really to a certain extent, it's just how we do it, storage, network, virtualization, these different someone siloed areas now being unified under a certain umbrella. I mean does that definition work for you or how would you go about describing it?

JIMMY RAY PURSER: You're half close. I would say if you add stateless to that you're on target.

ROBB BOYD: Okay so you consider stateless a silo that gets introduced to the equation or just becomes part of the solution?

JIMMY RAY PURSER: Stateless is a term that we've used a lot of times in firewalls and security and things of that nature because we need that type of healing aspect when it comes to any application. But stateless is not a term we use too loosely in data center type of products here. And to really get our arms around what this is, this is more than just a unified -- unified computing architecture is good but it doesn't really kind of put your arms around what the product really is. I mean there's a lot of ways I can tie -- say this is unified and that because I'm plugging it into this or I'm using Fibre Channel Ethernet, I'm unifying it that way, this is stateless unified computing. Now we're really changing the game.

ROBB BOYD: Okay well let's go ahead and address the elephant in the room because there's been no shortage of press guessing at what Cisco's coming out with. Now as we talk about it today the product has not been launched, it's no secret this show is taped in advanced and we're going to get some advanced look at some certain things. But by the time people are watching this show the product will have launched, some of these questions will have been answered. But the biggest thing that everyone seems to love to lead with is Cisco's getting in the server business right?

JIMMY RAY PURSER: Yes, right.

ROBB BOYD: Cisco's coming out with a blade server, is Cisco in the blade server business?

JIMMY RAY PURSER: No, and you know what, that was something that I thought too when I very first heard of this. I'm like what

are we doing making a blade server, I mean there's no shortage of those out there, you can get those from all kinds of vendors and stuff. And then I found out who was doing it, the Data Center Group, I'm like I know it's not a blade server if those guys are doing it because I mean you're talking about the Nexus team who put together some of these products. I'm like but what are they doing? You know because knowing them it's not a blade server. What really is this product?

ROBB BOYD: It's obviously a missing in that backyard right?

JIMMY RAY PURSER: Yes, so there's something here that's a little different. And once you start understanding the story it's -- a blade server doesn't even come close to describing what it is, that's not even in the same category.

ROBB BOYD: So how would you describe it architecturally?

JIMMY RAY PURSER: Well what we're really looking at is truthfully an architecture that is built for optimization -- not built for optimization, of course it is.

ROBB BOYD: For virtualization?

JIMMY RAY PURSER: But for virtualization. I get my -ations mixed up you know. But what we're looking at is how do we optimize, everybody's going to virtualize in data center, how do we build an architecture that supports and takes it to the next level with a virtualized-type environment? And this computing platform does that. How do we make it scalable so that I can actually scale this stuff? And most importantly how do I manage all this stuff? You know blade servers, data center stuff, I got all this crap that I've got to manage individually with all these different little element mangers everywhere. And to change the game I can't be running from manager to manager, I can't buy a whole box of blade servers and have to buy management servers to go in those blade servers to manage those servers. It's like a snake eating its tail, it's like --.

ROBB BOYD: But that's what we're trying right, historically?

JIMMY RAY PURSER: Well historically that's what it's been.

ROBB BOYD: Okay so hold on, I've got a list of things, we're calling this kind of a setup, a tease because the idea is we need to find the real technology that backs this stuff up. But I was keeping a list of the stuff that has been jumping out at you that you wanted to make sure we covered. And so I would say this to our audience to say look to how we prove...

JIMMY RAY PURSER: Keep your checklist.

ROBB BOYD: ...these things exist, keep your checklist. First thing on the checklist, single management platform.

JIMMY RAY PURSER: Yes absolutely, I mean like I said, that is really critical when it comes to managing a data center today. The downfall of data centers is how do I manage those, how do I extrapolate the information that's in those and actually have some usability in there. How do I go beyond just standard MIB pulling and monitoring and just some basic stuff that's built in with a lot of the proprietary management systems. None of the management systems talk to each other, they're all very closed.

ROBB BOYD: Well here's the thing, I'm anxious to get to the meat of this stuff so I'm just going to list it out real quick. We've got the VN-Link technology that we've talked about. Very innovative stuff, we've done a show on that before. We've got the virtual machine policy, the application look out there in terms of how that's done. Comparing kind of traditional architecture with really what is a new architecture here in this unified computing, centralized and uniform switching model, that changes the IO bay reduction. And as the software and where do we locate the management modules now and how does that change? All stuff that I think we should be looking at here.

VALERIE ST JOHN: Alright, lots of promises made here, I will have to keep my eyes on you and make sure you live up to them coming up next.

FOUR INNOVATIONS THAT FOREVER CHANGE THE GAME

VALERIE ST JOHN: Big promises for hope and change are only meaningful when backed up by real technology. Robb and Jimmy Ray have obviously been given some privileged VIP-like access to the development team. So gentlemen, your task now is to begin backing these claims. Go.

ROBB BOYD: You know, and that's the truth, right? We need to back them up. And as we look at, how do we begin this paradigm shift, how do we truly define by what we mean by unified computing. Because I think everyone expects us to be able to tell a networking story, and that's going to come out actually on its own any time we're talking about Cisco. But we looked at different guests because there's quite a few innovators that have been a part of bringing this solution to the point where it is today. And there's one guest that we talked to that we said, we've got to have this guy on the show. So I do want to welcome David Lawler to the show. David, you spent time -- was it... As I was going over your bio and we have talked, you had a lot of development expertise at Sun, you've done a lot of very interesting things on the server side throughout the 90s. You were there for quite some time. You obviously left and went to Nuova at some point and helped these guys. And you're a part of that side that has brought

the market to where it is now, and now you've got the -- I was going to say unenviable task, but actually it's kind of quite fun, as we've been talking about. Saying, hey, there's a new and there's a different way to do things. Right?

DAVID LAWLER: Right.

ROBB BOYD: So welcome.

DAVID LAWLER: Well, thank you.

VALERIE ST JOHN: Robb and Jimmy Ray are obviously extremely excited about what you and your team have brought to Cisco. What's brought you to this point?

DAVID LAWLER: You know, that's a good question. I mean, basically, when I first heard about what we were doing and what we potentially could do at Nuova, right, there was a Cisco angle to it. Looking at a fabric, and frankly, I'm not a networking person, right? I couldn't set up a network if my life depended on it. But what they're actually recognizing and having people from VMware there, Ed Bugnion who was Founder and CTO, looking at and actually recognizing where the data center was going where the server is not going to be something you install in the data center and four years later you go and uninstall it. The server is going to be this thing that actually is fluid and moving around in the data center. So it's really important to customers. It isn't just the server, but it's also the connectivity. Those two things have to be one, right? And that's what we were looking at. Is looking at that and the scale that people were actually getting to in the data center or have to get to, right? It's a new model. Today, to scale, and to scale a data center, what people are doing today is they say, okay, I need more computers. I need bigger data center. That's the answer today.

ROBB BOYD: More and more and more.

DAVID LAWLER: More and more and more, right? That's all, right? You know, and it's more physical and it's also more virtual. A virtual machine still uses resources. It uses storage, it uses network, still have to license the software for it. So that's still there and somebody has to manage all that and it's all moving around now. It's not like it was just fixed in the data center, it's all starting to move around, and there's more of it. And people are going, you know, okay, we'll get some more, right? Wrong answer.

ROBB BOYD: Wait, from server sprawl to virtualization sprawl.

DAVID LAWLER: That's exactly right.

ROBB BOYD: Well, so, you've got your list of, I think it's four game changes. There might be a few more that come out by the time we get done. But I want you to keep us honest here. So Jimmy Ray, where do you start? What are the things you say, hey, to make sure you notice this?

JIMMY RAY PURSER: Well, you know, I think that one of the things that we need to look at is the fabric itself. Now we're talking about data center designs and stuff. We're kind of far beyond the loop data centers designs. I mean, everything now is a fabric, and I think in a fabric type of design we need to -- our entry point needs to be 10G. And especially some of the work going on in the IEEE with 40 gig and 100 gig, we need some scalability to kind of build up to that. But the problem is speed in the enemy in the data center a lot of times, because a lot of times we're not prepared for all that extra speed. It's like having a great car and then driving on a dirt road. Like, man, this thing drives like crap, and then driving on a good road it drives great. Well, you know, you've got to consider the road that you're on and how we're transferring all this data back and forth. And look at how data centers are designed today -- kind of by accident or kind of by happenstance, if you will. The higher the speed the more problems that can cause, and stuff. I'm kind of wondering, how are some of these problems solves in this current architecture? What are some of the speed issues, some of the sprawling issues. You know, because we've really got to look at designing this from scratch. And this is what killed me in understanding this product here, is that we did everything different. I mean, we designed it from -- you designed it, not me. I didn't go out with a soldering iron or anything. But how did you build this up, I guess, starting from scratch?

DAVID LAWLER: You know, if you're starting from scratch, right, and you covered just the multi-faceted team basically, there are guys coming in and looking at it and asking questions. Like you look at the other guys and go, why? You're serious, right, in looking at it, and let's talk about the 10 gig. You know, I'll tell you a secret. You talk to the server guy and they go, I don't need 10 gig.

JIMMY RAY PURSER: Well, yes, of course.

DAVID LAWLER: I don't need 10 gig. I need Fibre Channel, I need Ethernet, but I don't need 10 gig Ethernet. It's just not there, right? And the other thing about it is that 10 gig, you talk to most people and they go, it's 10 times 1 gig, right? That's a lot faster, right? That's not what we did. What we've done basically, and working with Cisco you look at data center Ethernet and evolving Ethernet. Ethernet is a healthy living protocol, it's evolving over time basically. And we've involved it into a lossless, low latency fabric, right? So you can actually get a higher quality fabric out of it. And that was one of the things, was actually taking it and saying, okay, 10 gig has to evolve into being actually a fabric that can cover multiple things. And so it's, you have the fabric that can do a lot of stuff. We build with that in the standards bodies and basically say, okay, we want one fabric to actually address the issues we've got here. You look at the challenge in the data center and people have servers with multiple 1 gig connections coming out of it, and they'll have some Fibre Channel connections coming out of it.

JIMMY RAY PURSER: (Inaudible).

DAVID LAWLER: Exactly, right? You know, and it's not that people say, I just need -- they don't need 10 gig but they need more than 1 gig. Okay, so there's a challenge. So looking at it what you want to do is actually simplify that environment, right? And say, okay, it's not about some Ethernet connectivity, some Fibre Channel, it's like, you need bandwidth out of that box. And that was a simplification on it. It's like, okay, yes, you talk to a customer and say, yes, I need like 8 gig total bandwidth with Fibre Channel and Ethernet and all that stuff together. But if you just talk to them about this, okay, it's just one 10 gig connection, it's like, I'm not there, I don't see that. So it's evolving that into a unified fabric, right? And in doing that basically, and this is the acceleration problem you talk about is, you talk to customers who said -- and everybody's been trying to go for this stateless computing. Great, nice vision.

ROBB BOYD: It sounds good.

DAVID LAWLER: It sounds good, you know, diskless computers and you boot off the net and there's SAN from wherever you're coming from and it's a great idea.

JIMMY RAY PURSER: Back to VMS 1982.

DAVID LAWLER: Oh, exactly. Nothing's new here, right? And what they're actually seeing is that people focus on making the computer stateless and then this computer's got three Ethernet and two Fibre Channel and this one's got five Ethernet and one Fibre Channel. It's like, okay guys, you left state here.

JIMMY RAY PURSER: Yes, what's that?

DAVID LAWLER: Exactly.

JIMMY RAY PURSER: We're more stateless than that.

DAVID LAWLER: That's exactly right. So that's what we wanted to do. And we looked at the technologies on the market and some of the technologies on the market were saying, okay great. Some of them said, okay, you've got Ethernet, you've got Fibre Channel. What we're going to do is we're going to give you InfiniBand and we're going to roll everything into that. So you have two solutions to make it three, and that wasn't what we wanted to do. We wanted to simplify this environment.

ROBB BOYD: Going the wrong direction.

DAVID LAWLER: Exactly.

JIMMY RAY PURSER: Did that. Been there, done that.

DAVID LAWLER: So we said, okay. What if we can actually, in evolving Ethernet and making it a better thing, making it be able to handle Fibre Channel. So you get the unified fabric, you get a true wire once infrastructure. Take the state out of the wiring in the data center, reduce the cost in the networking, reduce the cost of the adapters, switches -- all those components. That was the first epiphany, sort of like, okay, that's state, we can simplify that, and there's pure cost savings right there and customers can understand that so long as they trust it. And trust is the key.

JIMMY RAY PURSER: There you go. See, that's it right here, because we're talking to the data center crowd here who have been promised a lot of things before in the past and it hasn't been delivered on. And we can go through technology after technology. So how do we build that trust up here, because, I mean, that's something I didn't trust. When we were first talking about this solution I'm like, huh, yes.

DAVID LAWLER: Yes, the way you earn trust basically is, first of all, through your partners, through certifications, right? What's different about the networking world versus the server world is the networking world actually does, because it's defined by definition, interoperability, there are very solid and clear standards. And the server world is like, well, it's either an Intel or an AMD processor, it should run the operating systems, and there you go, that's it. So to get that trust in the fabric level basically is, okay, we can do this, we can drive it to the standards body. But people also say, well, who else? And we came to the market with the QLogic, Emulex, we're doing all that and we've got the EMC -- all these companies...

JIMMY RAY PURSER: Heavy hitters.

DAVID LAWLER: Exactly, the heavy hitters going, yes, this is a fabric.

ROBB BOYD: Well, speaking of that, though, because you bring up the adapters and stuff in there, you've got that on your list as well as some of the innovation that's occurred.

JIMMY RAY PURSER: Oh yes, absolutely. That was something that impressed me, if we just came up with the adapter only in this solution. The adapter was incredibly impressive, how it was designed and how everything was brought into the single Cisco ASIC to bring all this connectivity in. Can you give us a little more background on just the adapters. I think we could do a whole show on

the adapter. That amazed me...

DAVID LAWLER: That one alone, I mean, it's a phenomenal engineering team that did that one. They focused on it and they said, okay, I can cut down on wires, but, you know what, a lot of operating systems they don't do VLAN tagging, for example, in the OS or optimized virtualization. So what we did is we developed a virtual life adapter. And this adapter basically is following the standards, PCI Sig, SRIOV which is called Single Root I/O Virtualization. What that means is a single PCI Express device, physical device actually, to the operating system, to the PCI Express bus can be up to 128 separate devices. All right, so what that means basically...

JIMMY RAY PURSER: Modem adapter.

DAVID LAWLER: To the operating system, one adapter, yes, but the adapter looks like up to 128, either Ethernet or Fibre Channel, right?

JIMMY RAY PURSER: How did you do that, though?

DAVID LAWLER: Good engineering. I mean, it's looking at the problem and trying to optimize -- and trying to look at the problem overall. What we look at, and this is where we start and I'll start to use the term -- I use the term unified computing. We've been talking about unified fabrics, unified this, and I agree, it gets overused. But what we saw is that, you know, configuring the network for all this compute stuff, if I want to optimize, actually the network's got to go all the way into...

ROBB BOYD: There's too much of a hard line between the (inaudible).

DAVID LAWLER: That's exactly right, exactly. So where does it need to stop, and one of the things we want to do basically is, let's optimize. Optimize the performance of this platform, right? So if you're looking at optimizing virtualization one of the biggest challenges you actually have and the road blocks is I/O, right?

JIMMY RAY PURSER: Oh, absolutely, positively.

DAVID LAWLER: Exactly. So we want to say, you know what? The way it works today and everything, all the I/O has to go through the hypervisor and the hypervisor goes, okay, I'll deal with all that, right? What we did basically with this adapter is said, okay, what we can do basically is have the virtual machines communicating directly to the PCI Express bus. Now, you have to do that, you have to work with partners. It's about the trust basically, so we did this with VMware. So this is something that we're coming out with. We've announced it, by the time this has been aired, but it's something that's going to be shipping in a couple of months. And the technology here basically is that, with the single adapter what I can do basically is coordinating within -- and it takes management coordination, and we'll get to that later. I can actually coordinate virtual machines to communicate directly to the PCI Express bus. I create a PCI Express device and then the hypervisor creates the virtual machine, connects it directly up to the PCI Express Bus. And then what that means basically is that as I have I/O, I use less CPU cycles, increase the actual bandwidth, and a lot lower latency from I/O.

ROBB BOYD: Whoa, see, you couldn't have done this kind of innovation within any single company, right? A server company's not going to do that, the networking company is (inaudible).

JIMMY RAY PURSER: Can I build that from the parts bin?

DAVID LAWLER: No, this whole solution is not something that comes out of the parts bin. There's people that's going to say, well, you know, if you do this, this, this and put a wrapper around it and then do that, I can do it.

JIMMY RAY PURSER: So just create a logo and...

DAVID LAWLER: Well, let's look at just the adapter. If I take this adapter up to 128 devices, single machine, but I've got this and in the overall unified computing system I can have up to 320 compute nodes, right? So multiply it out and let's just get -- let's say it's thousands, right? Coordinating virtual interfaces with PCI Express devices with network policies all the way up to the network out. That has to be one unified definitional phase. If you have to coordinate all that, and I talked earlier about moving everything around when there's thousands of objects here. But now, if you want to do it with performance, you want to do it with scale and then simplify it, that's what we looked at and identifying and being able to actually map these components out. That's what we can do with the adapter, but the adapter fits into the overall model. It's part of the solution.

JIMMY RAY PURSER: I guess the thing that kind of gets me too, which is another thing on my list is, is the problem of virtualization is manageable all that mess. And so now you're talking about putting 128 of these on an adapter, and I'm thinking, okay, that's cool for the data sheet. But as the SAN guy, how do I manage this stuff? I mean, are we still talking about calling MIBs and, you know, running a proprietary type management solution. Like all the blade systems out there today, they're closed. I mean, the management is very closed and it's hard to integrate applications into them and stuff.

DAVID LAWLER: Yes, absolutely. I mean, if you look at blades today, basically, blades are a physical optimization at best. They basically took the same parts and said, you know what, to make these smaller we can sort of share cooling and fans, and that's it.

JIMMY RAY PURSER: Yes, it's (inaudible).

DAVID LAWLER: And then because they did that they actually threw more management at the problem, right? And as we know in computer science, every problem's solved by another layer of indirection. That's what they did.

JIMMY RAY PURSER: That is absolutely right.

DAVID LAWLER: So it doesn't make it more reliable or better, it just solves a tactical problem. We looked at this and we looked at the platform as a device -- a large scalable modular device with all of the components being known and defined. And so as you add in componentry, whether it's the adapter, blades, chassis, fabric extenders, everything, they plug into the system and they become resources for the device. What we have is a cohesive management model in the system that's built on a model of understanding it and it actually builds up and creates a finite state model of everything that's in the system. So you're actually manipulating a single device, but that's scalable.

JIMMY RAY PURSER: Yes.

DAVID LAWLER: Up to 320 compute nodes, multiple racks, multiple chassis. A chassis, you know, the overhead of an additional chassis for us is zero, right? It's the same management overhead to manage a single chassis of eight blades as it is 320 blades. It doesn't matter, that's just capacity to us. You know, the computer industry over the past 20 years -- you talked about the evolution, but it's always been this optimization for the current cost master, whatever that is -- the biggest part of the cost. In the mainframe it was the computer. Let's cut that down, right? Now everything has changed. It's actually gone from where it was mostly capital costs with the computer down to operations. What we're doing basically is looking at it and saying, let's get the operations' costs back under control, allow people to actually manipulate it, increase the productivity to the people we have today. At the same time, and this is critical, it's not just about, hey, we'll make them do much more faster, that's great, because you also have to plan on failure. And what I mean by that is, hardware will fail, software will fail, and usually you don't know what failed, to be honest with you, right? And so the industry term we use is mean time to repair. It's a great term, right? Let me ask Jimmy Ray, once you've found the problem, how fast does it take to fix a problem?

JIMMY RAY PURSER: It's relatively quick. I mean, you can bracket it out, you can jump right to it and fix it and solve it and go.

DAVID LAWLER: Especially, in our case, everything's on a plug so I can pop it out and pop a new one in. It's pretty simple to solve, right? Mean time to diagnose is a different world, right? Finding what the problem is.

JIMMY RAY PURSER: Yes, that's a big problem.

DAVID LAWLER: That's right, but in the industry, MTTR, mean time to repair, includes diagnostics.

ROBB BOYD: I want to hit one more point.

JIMMY RAY PURSER: Wait, wait, stop, stop for a second, because you're saying that everything's hot plugged and I can't let you get away with that. Where's my application layer? I can't pull the server out and put this application on another server. I mean, how am I...

DAVID LAWLER: If you pull out a running blade it's going to die. There's no doubt about that.

JIMMY RAY PURSER: No, what about the application that lives on it? If I've got to replace that, how can I pull that application and assign that to another resource?

DAVID LAWLER: And you're absolutely right. What we do basically is, opposed to actually having to manage down to the physical node, what we do is we take the model that people have today. The model that people have today is you define a certain amount of compute resources, you define the network resources...

JIMMY RAY PURSER: That's a silo.

DAVID LAWLER: Exactly, and each silo defines their component. We take that and make that a logical model. If the company wants to for good organizational reasons or whatever, keep having the networking guys do that and the server guys, they can do that illogically at this point, right? And then once it's defined we call that a service profile. Once that profile is defined you use that and you apply that to the hardware. Now to your point, once I've got -- you know, I want to pull the blade out and there's something running on it, I simply say, I'll take that service profile, apply it to another blade. From every way you look at that blade, from MAC address, worldwide name, Bios...

JIMMY RAY PURSER: Bios.

DAVID LAWLER: Yes, part of the service profile, right. This is a complete system. It's not just an afterthought. It's, as I said, not the parts bin. You're like, oh, what we've done is we've consolidated multiple independent management planes into a single system so that you're looking at it and you're managing a single system. You know, competing systems out there, they're management brands, and I'm not going to pick on any one. I can pick on all of them, frankly, because they're all the same thing, right? I mean,

they're brands, but if you have it you'll have a single stream, and that single stream says, oh, I've got firmware over here and I've got identity over here and I've got this, they're all maybe on the same stream. But for identity to work you've got to make sure that firmware on all the blades is correct, so you'd better check that first or I'm just not going to work, right?

JIMMY RAY PURSER: That's right.

DAVID LAWLER: And then access to the chassis managers, right? You know, the state of the art for chassis today -- and not for ours. Ours, we've changed that model. The state of the art for chassis today is not you have 16 blades with a clustered management model inside of it that does three things.

JIMMY RAY PURSER: Which is pretty silly.

DAVID LAWLER: Right, it does power management, cooling management and access control. And by the way, you have to update the firmware on that separately. And maybe you can buy another piece of software that actually does that and coordinates it across it all.

ROBB BOYD: I'd like to say, once people get it working too they're like, don't touch it.

DAVID LAWLER: Exactly.

JIMMY RAY PURSER: Yes, that's exactly what (inaudible).

ROBB BOYD: We've got everything just right so we'd better not mess with it, and thus your agility goes right out the door.

DAVID LAWLER: Exactly, that's exactly right. I mean, the management models today and the way I like to look at it is almost, it's almost the McDonald's model. They'll sell you a server and go, oh, would you like fries with that? They'll sell you a server and go, would you like management with that? For us it's built into the system, right? A resource you can't manage is a useless resource.

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

DAVID LAWLER: And so server vendors, basically, if you look at a server, a blade from any of the vendors out there, I'd say 80 percent of the value, processors, memory, disk, chipset on it...

JIMMY RAY PURSER: Off the shelf.

DAVID LAWLER: Off the shelf.

JIMMY RAY PURSER: Yes, nothing new there.

DAVID LAWLER: That last 20 percent, well, that's a little bit management glue and then hand-holding. But then they say, oh, would you like management with that, would you like other stuff? What we're doing is saying, you know what, that's not hard. The hard part is actually coordinating it at scale in the network and so everything can move around.

ROBB BOYD: So, before you close us out Valerie, I just make a point that we only hit three of your game changers. We hit the fabric, we hit the adapter and we talked about management. So in the next segment I want to make sure we bring up -- and actually we're going to have you back in segment four. I want to make sure that we come back to some of the innovation you've done around memory because that's a story I don't want to miss.

JIMMY RAY PURSER: Oh yes, that's heavy duty. Oh yes.

ROBB BOYD: That was a great story. That's right, there's a lot of good stuff to talk about.

VALERIE ST JOHN: Thanks guys, so we're not done with David yet. It's time to open that shiny new toolbox and put it to work. We're taking apart the hardware next.

EXPLORING THE NEW HARDWARE

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 looking at the dawn of the new data center. Robb and Jimmy Ray twisted some arms, called in a few favors, maybe made a few threats and as a result we have a first look. And now that we've gotten it all together here, we welcome Brian Schwarz, an engineer with the Unified Computing Center, to help us take it apart.

BRIAN SCHWARZ: Thanks.

VALERIE ST JOHN: Hello Brian.

Dawn of the New Data Center

BRIAN SCHWARZ: We've got some serious toys on our table here. Why don't we just walk you through them? Let's start with the blade design here. In some respects it looks to be a pretty standard x86 blade. It's got two sockets, four cores in each sockets, bunch of DIMM slots. Two very distinct differences about the blade, one is it's stateless blade design which you mentioned a little bit in earlier segments, it's really an anonymous piece of computing capacity that isn't really enabled until you get a service profile associated with it.

ROBB BOYD: It's an awesome piece of computing policy. (Inaudible) an anonymous.

BRIAN SCHWARZ: Anonymous, so it's kind of a scratch piece of...

ROBB BOYD: It is anonymous.

BRIAN SCHWARZ: Yes, it's anonymous because all of the state has been sucked out of it essentially into the network. So it's just kind of bare capacity waiting for you to associate a service profile and turn it into a real server. That's what gives the dynamicism of being able to provision servers a lot faster. So that's one big difference and we'll talk a little bit about that more when we talk about the Unified Computing Manager. The second piece that's very different about it is this adapter here on the back. This is a converged network adapter, so it has that capability, what it shows to the operating system is a 2-port, 4-gig HBA. This one is actually based on an Emulex chip. We also have another one based on QLogic and it also has a 2-port, 10-gig NIC. Both of those are basically muxed onto two unified fabric ports on the back, so what the operating system sees is pretty standard. A lot of good compatibility in the system but it allows us to implement the unified fabric on the back. So you don't have to have multiple mezzanine cards on actual blades themselves.

JIMMY RAY PURSER: Is your port -- so then this is the muxer right here then?

BRIAN SCHWARZ: Yes, that's the special Menlo ASIC right there.

JIMMY RAY PURSER: Whose product is this?

BRIAN SCHWARZ: These are based on the Intel and the Halem processors.

JIMMY RAY PURSER: Oh really?

BRIAN SCHWARZ: Yes, so they're each brand new, hot off the press x86 server designs with four cores in them, so it's a 8-core blade and this blade design has 12 DIMM slots in it. I think you'll talk to David a little bit more in the next segment about some special memory expansion technology we have with (inaudible) DIMM slots.

ROBB BOYD: I want to get your reaction of this too.

JIMMY RAY PURSER: So what's the...

ROBB BOYD: (Inaudible).

JIMMY RAY PURSER: This is the half slot one, right?

BRIAN SCHWARZ: Yes, yes.

JIMMY RAY PURSER: So you've got half and a full.

BRIAN SCHWARZ: Yes.

JIMMY RAY PURSER: This is pretty darn interesting. I mean I like the fact that this is an anonymous piece of hardware. That truly makes it stateless until we actually initiate it. That's pretty sweet but we've got a couple hard drive bays here in front too, right?

BRIAN SCHWARZ: Yes, yes, so you get in two SAS drives there if you want to store the operating system there or use them for space.

JIMMY RAY PURSER: What's the council port for, old school management?

BRIAN SCHWARZ: Yes, it's for debug. You can certainly go, for folks who are actually walking around the data center floor who might need to plug into one of the blades...

JIMMY RAY PURSER: Oh sure.

BRIAN SCHWARZ: You get KVM access, virtual media, so you can mount to a CD ROM or something to it if you need to.

JIMMY RAY PURSER: Very cool, very cool.

BRIAN SCHWARZ: Yes.

JIMMY RAY PURSER: All right, let's check out the next piece.

BRIAN SCHWARZ: So that thing will slide right into the front of the...

ROBB BOYD: Am I allowed to touch this?

JIMMY RAY PURSER: Yes.

BRIAN SCHWARZ: Of the chassis here. We've actually got one of the blades already slid in here. So the design is eight of the half-width blades, you can take the dividers out to get to the full-width blade.

JIMMY RAY PURSER: Oh that's pretty cool.

BRIAN SCHWARZ: Yes, yes, yes.

JIMMY RAY PURSER: So if I want to do that, oh that's cool.

BRIAN SCHWARZ: Yes.

JIMMY RAY PURSER: So we just pop the tabs and then I can put the full -- that's pretty cool.

BRIAN SCHWARZ: Yes, you've got the power supplies on the bottom here and in some respects it looks pretty standard. The big differences turn out when you look at the back of the chassis. This is where a lot of the optimizations around the unified fabric come from. A traditional blade chassis, what you would see in the back is really small fans and a lot of space for networking elements, essentially I/O modules. The chassis was designed straight from the unified fabric, so what it has is it has two fabric extenders plugged into this one here at the left and in the center. And these implement the unified fabric, so each of them has four ports of 10-gig unified fabric.

JIMMY RAY PURSER: These are 10-gig apiece?

BRIAN SCHWARZ: They're 10-gig apiece, you can just start with one each for redundancy. So you have 20-gig of capacity to the eight blades and you could scale up to use all four ports on each fabric extender to get you up to 80-gig.

JIMMY RAY PURSER: Hey man, you know what? To be honest with you, if you brought this into my data center and showed me, I would say, this looks like a oven.

BRIAN SCHWARZ: It's certainly got some...

JIMMY RAY PURSER: How do I cool this thing off?

BRIAN SCHWARZ: It's got some -- these big things in the back here, they've got eight fans, so...

JIMMY RAY PURSER: Just pull these out?

BRIAN SCHWARZ: Yes, the size of the fans...

JIMMY RAY PURSER: Goodness sakes man.

BRIAN SCHWARZ: Is optimized to cool the chassis itself.

JIMMY RAY PURSER: Oh that's fans inside of fans.

BRIAN SCHWARZ: Yes and...

JIMMY RAY PURSER: It's a jet engine.

BRIAN SCHWARZ: Yes, it's two fans. It actually gives you better performance. You have two fans front to back like that.

JIMMY RAY PURSER: Very cool.

BRIAN SCHWARZ: Yes, less power consumption and also the backplane which you can't see which is in the middle of the chassis here, because there's only two elements...

JIMMY RAY PURSER: Yes because that doesn't go back that deep.

BRIAN SCHWARZ: Yes, because there's only two networking elements, most of the backplane is actually open, so the air can flow through, easy to cool the system. The fans shouldn't have to work that hard to keep the eight blades cool.

JIMMY RAY PURSER: Is that right?

BRIAN SCHWARZ: Yes.

ROBB BOYD: You actually did not have to follow a model. I mean that's one of the fundamental differences here. You did not have to follow a model that said, how much stuff can we cramp into a small place creating such an incredible oven that that actually stresses the data center even further, right? So there's benefits that continue to accrue from that?

BRIAN SCHWARZ: Yes, yes, unified fabric is a great technology because older blade chassis designs had two physical networks essentially. So they had to have four or eight blade switches in the back of them to go out to two physically distinct networks.

JIMMY RAY PURSER: That's right.

BRIAN SCHWARZ: There's a lot of optimizations that come from just using a unified fabric design with the fabric extenders. The other nice thing about the fabric extender is they're essentially unmanaged. You're essentially going to manage all of it in the switch. So there's no individual management the way you would in a traditional blade chassis where you're going to go into the device manager of the individual blade switch and configure the ports that connect to the servers and the ports that connect to the network.

JIMMY RAY PURSER: Is that right? You don't groom the ports out actually on that?

BRIAN SCHWARZ: No, no, it's all managed in the UCS Manager, in the network essentially, all of the management of the system, the chassis itself. You also notice there's no management modules in the chassis, so there's no separate...

JIMMY RAY PURSER: Yes, no, I don't see any of it.

BRIAN SCHWARZ: Separate extra management servers that are needed. It's all done in the network, essentially in the UCS fabric interconnect here.

JIMMY RAY PURSER: Hey, you know what? Let me -- so I'm trying to look. I don't know if they can see that on camera or not but the depth of this, and then we swing this around to the front and we're looking at the depth of the blades here and stuff. How much airflow free space do you have in this box? Because in a data center, heat is the enemy, right?

BRIAN SCHWARZ: Yes, oh yes, it's hard to cool...

JIMMY RAY PURSER: And so I mean...

BRIAN SCHWARZ: Hard to cool these systems.

JIMMY RAY PURSER: The smaller we get, that's great. But the hotter stuff gets, usually we start (inaudible).

BRIAN SCHWARZ: Yes, yes, traditional blade chassis, the backplane is kind of a little bit under 30% open. Ours is above 60%, so it's dramatically...

JIMMY RAY PURSER: Seriously?

BRIAN SCHWARZ: Huge holes in it that allow the air to flow freely.

JIMMY RAY PURSER: Is that right?

BRIAN SCHWARZ: Yes, yes.

JIMMY RAY PURSER: Holy smoke man.

BRIAN SCHWARZ: It's just if you design a chassis from the ground up based on unified fabric, there's a bunch of physical optimizations you can do around it.

JIMMY RAY PURSER: Uh-huh, okay, so that's cool.

ROBB BOYD: So this is part of that paradigm shift we're talking about here as well. Now I noticed we haven't talked about (inaudible).

BRIAN SCHWARZ: So the fabric extenders will connect up and really the brain as we've said before lives here in the fabric interconnect.

JIMMY RAY PURSER: Now so the fabric extenders are -- for somebody trying to get their hands around the traditional data center, the fabric extender is equal to what?

BRIAN SCHWARZ: The best way to think about it is kind of a remote line card. If you're familiar with like a Catalyst 6500 or a Nexus 7000, they have line cards you plug into it.

JIMMY RAY PURSER: Right.

BRIAN SCHWARZ: They're all managed as one large element. The same kind of concept applies here except they're not physically in one chassis, one physical chassis. Think about taking one of those line cards out and physically putting it into the blade chassis itself. So you can have up to 40 of these blade chassis connected up to the fabric interconnect. That all acts as one big switch essentially. It's all managed from one spot here in the UCS Manager.

JIMMY RAY PURSER: Really?

BRIAN SCHWARZ: And you're not going to configure and manage these as individual switches the way you would've in the past.

JIMMY RAY PURSER: You know, then that really kind of gives us some true N+1 redundancy here that we can do. I mean we talk about that but realistically it's kind of really hard to get to that level without spending a ton of cash.

BRIAN SCHWARZ: Yes, so certainly there's a bunch of optimizations in the design that were designed for HA N+1. You know the two fabric extenders, the reason you need two of them, each one will connect up to a fabric interconnect. So the loss of any one of the components, the servers themselves will still have full network connectivity to both the LAN and the SAN.

JIMMY RAY PURSER: Wow, so any gig capacity is what we're looking at here on the (inaudible).

BRIAN SCHWARZ: Yes, yes and there's also some future proofing in the design to go up to higher levels of capacity, more I/O modules to be slid in the back here, more ports, faster ports. You shouldn't have to replace the chassis or do any of that.

ROBB BOYD: Let me ask you something because something that jumps out at me is we don't have the benefit of seeing like the cable reduction here because we're not actually having this cabled. Obviously there's not too much from a cable perspective here to deal with. I just want to grasp this -- because I don't have the engineering background that you do -- but kind of the before things were versus what we're looking at here. Help me understand. I mean can you draw that comparison, that contrast picture for me?

BRIAN SCHWARZ: The easiest thing would be to imagine instead of two fabric interconnects here in the back or fabric extenders, imagine six or eight blade switches in the back, each of them with six or eight cables coming out of the back of it because it's often 1-gig and things like that. It's really just monstrous amounts of cables coming out of the back and it makes the fans work a lot harder.

JIMMY RAY PURSER: Yes it does.

BRIAN SCHWARZ: In addition to just the challenges of laying up cable...

JIMMY RAY PURSER: It's a lot bigger too at that point too.

BRIAN SCHWARZ: The cost and you know (inaudible).

ROBB BOYD: What kind of a difference does this physically make to the person's life who's managing this? I mean one of the things I know that we're still going to be talking about more when we get deep into the management is obviously we're blending some worlds here. And I want to make sure that we also understand that we've spoken to the fact that I think that can cause some fear because of the uncertainty that we introduce to the situation because we've changed a few things with this. It's disruptive to a certain extent. We've got server people looking at this going, how does this change my job and is my job in danger, which is not a popular topic right now, and the networking guys looking at this and going, does that mean the server guy's about to take my job? You know everybody's looking at it this way. I know you guys have thought through this and...

BRIAN SCHWARZ: Yes.

ROBB BOYD: You've got the role-based management because you're dealing with organizational structure, right, when you talk about those things. How do we operate our business? Who does what, allowing a business to make those choices and those changes over time? Are there some operational stuff we need to understand while we're looking at the physical hardware that change?

BRIAN SCHWARZ: Yes, a lot of it comes into the UCS Manager of the software itself in terms of how you manage it. So we still preserve kind of the role specialization. You still need best practices for a network person to define network policies. The server person (inaudible).

ROBB BOYD: Their knowledge isn't going to waste here, right?

BRIAN SCHWARZ: No, no, absolutely.

ROBB BOYD: I mean it's not going away.

BRIAN SCHWARZ: What's going to change though is that they're all going to be operating on a single management system as opposed to using a ticketing system or email or Excel or something. To communicate amongst different teams, they'll actually be looking at the same GUI. They'll have different parts of the GUI which will be kind of their home where they'll do their tasks. It allows them visibility to see what the other groups are doing but kind of preserving the role specialization but allowing the management to be greatly simplified. Just less points of management in an overall system.

JIMMY RAY PURSER: Less stuff to patch, less vulnerabilities out there. I mean you and I are both security backgrounds and one of the things that I noticed in one of the latest security studies is that if you take ten companies that were all broken into, eight of those companies were broken into due to a vulnerability in the software that wasn't patched for over six months. And the more management consoles we have in particular a network manager itself or element managers, they tend to be very vulnerable because they have so much deep access and higher privilege access of the gear. So even limiting that really increases your security posture as well.

BRIAN SCHWARZ: Yes and the policy-based management should really help a lot to get consistent server deployments with the same set of network policies applied for all the servers (inaudible).

JIMMY RAY PURSER: Oh that's a excellent call.

BRIAN SCHWARZ: And it should really have better consistency which not only for security but is good for high availability. You know, there's...

JIMMY RAY PURSER: Sure.

BRIAN SCHWARZ: The new one that's rolled out will look the same as the other ones such that if you need to do a failover like an application-level clustering failover, the new server's going to look the same. It's going to have good connectivity, connect to things, (inaudible).

ROBB BOYD: This really addresses a lot of the issues that virtualization had kind of brought out for us here which was we went from physical server sprawl to reining that in for higher optimization and getting this virtual server sprawl. But you've now given the ability to truly optimize the assets but not lose control of them which has security implications whether you want to call it that or not just because you know where things are and you know that the policies are applied consistently. And you've removed because of these service policies, we've begun seeing this with VN-Link and such as that technology comes out because we've now locked stuff down. We've said, you know what, let's take some of your famous fat finger errors out of that equation and make sure that things are applied the way the networking person identified says, for this set of things, it's all wrapped up. I've made it possible for you. Let's not screw it up now. I like that.

BRIAN SCHWARZ: Yes, yes.

JIMMY RAY PURSER: Okay, we need to find out where the brains are in this stuff before we run out of time because you're talking about no management here and this and that. Where are the brains?

BRIAN SCHWARZ: The brain is actually inside the fabric interconnect. So not only does it perform the traditional data plane, it's got a host of -- this one has 20 fixed ports, all 10-gig, data center Ethernet on it. You've got an expansion module which gives you either more Ethernet ports or native Fibre Channel ports that would connect up to your SAN.

JIMMY RAY PURSER: Oh really, okay.

BRIAN SCHWARZ: So in addition to being the normal switch, this is where the UCS Manager lives, so it's very different than a traditional switch in that it not only is used to configure the ports here, it manages the chassis, the fabric extender, the blades themselves, the adapters. It's all managed from a special piece of software that's written in and resides in the fabric interconnect.

JIMMY RAY PURSER: Really.

BRIAN SCHWARZ: Yes.

JIMMY RAY PURSER: Okay, we're talking about obviously a very important piece of hardware here. How are we keeping this thing alive? What makes it redundant, how are we building some durability into this product so instead of actually -- because a fear of mine be is, okay, all right, I can put all my stuff here, all my brains here, but how am I...

BRIAN SCHWARZ: It better stay up.

JIMMY RAY PURSER: You ain't kidding buddy.

BRIAN SCHWARZ: Yes, yes, so certainly the design was built kind of the same as you saw in the chassis with two fabric extenders. Each fabric extender would go up to essentially one of these switches and the two switches essentially form you can kind of think about as an application cluster. The management is actually running on both of the switches at the same time and there's replication between the two switches. The failure of any one of these switches, yes, the application would essentially failover to the other -- the UCS Manager would failover to the other fabric interconnect.

JIMMY RAY PURSE: Oh, very cool, very cool, very cool. And we've got what looks like redundant fans here. Let me flip this around so we can see this in the camera.

ROBB BOYD: This makes me wonder have we ever had a data center on the desk before?

JIMMY RAY PURSER: Oh, that is an excellent call.

ROBB BOYD: It just occurred to me that I don't think we've ever had an opportunity to put a data center on the lab table. What's -- because something's got to be missing here.

JIMMY RAY PURSER: Excellent call.

BRIAN SCHWARZ: Yes, you know, it's amazing. You could have a huge number of virtual machines in these chassis that the amount of networking capacity, we have a 2U form factor of this fabric interconnect that has a terabit crossbar built into it, so it's just a throughput, just incredible. The one thing that you haven't seen here is the external storage.

JIMMY RAY PURSER: Well, yes, I mean that's...

ROBB BOYD: That makes sense.

BRIAN SCHWARZ: You connect this up to your traditional SAN storage because we have native Fibre Channel ports in it. Or obviously it has screaming fast Ethernet connected up to NAS storage and use that as well.

ROBB BOYD: This just messes with my head about how much stuff we've actually got here and kind of recognizing, don't belittle the fact that -- I don't mean that word but you know there's a lot here in a small form factor.

JIMMY RAY PURSER: Well you're right and that's a lot of computing power to look at this because I think typically if I was watching this show and I'm a data center person, I would look at it and say, huh, is that part of the small business communications solution there?

BRIAN SCHWARZ: Yes, well imagine a large Unified computing system which might have 10 or 15 of these chassis all being managed from a single point essentially, multiple racks.

JIMMY RAY PURSER: I can't. That's hard to believe that we can take multiple of these chassis and I can unify that management into a couple of these, really one of these in essence because that's what we're managing (inaudible).

BRIAN SCHWARZ: Server device management, adapter device management, Fibre Channel connectivity device management, LAN connectivity device management all done in a single point of management.

ROBB BOYD: Were you going to point something out over here?

JIMMY RAY PURSER: Well I was actually just talking about the -- hey, you know what? Hang on a second.

ROBB BOYD: You got something better?

JIMMY RAY PURSER: No, because I'm trying to get my head wrapped around this management. How are you doing that? Because we're not doing this with MIBs, right? I mean how...

BRIAN SCHWARZ: No, there's special firmware that runs in all of the components so that run inside the fabric extender that actually run on the blade itself. We also do some unique things that allow us to boot our own operating system on the blade to do some blade configuration on the blade.

JIMMY RAY PURSER: Seriously?

BRIAN SCHWARZ: So the whole system is designed to work together.

JIMMY RAY PURSER: That is heavy duty but let me just real quick point out a couple things here because it looks like we got -- obviously it's just a data-center-class product. We've got redundant fans, redundant power supplies, all the power and the cables come out through the front just like a good...

BRIAN SCHWARZ: Font-to-back cooling.

JIMMY RAY PURSER: That's a 1U, very sweet man.

VALERIE ST JOHN: Better performance, smaller device.

BRIAN SCHWARZ: Absolutely.

VALERIE ST JOHN: Thanks very much. As with any new technology especially one as disruptive as this is, how it works within your system is absolutely crucial. What do you need? What might you already have? That's what we're going to show you in Segment 4 coming up.

DESIGNING YOUR NEW ARCHITECTURE

VALERIE ST JOHN: Today we're looking at truly groundbreaking innovation for making the data center more agile and more responsive to today's business. Brian Schwarz and David Lawler are back with us in the same studio at the same time and you guys are here to explain how well this technology plays with others, right?

ROBB BOYD: Well yes they are. The whole thing about let's get physical, we did that in the last segment, I want to see how we actually interoperate with this stuff, now understanding of course that we're looking at some of this stuff in advance. So some of the things we see here from a skin or a GUI perspective, probably subject to change, I think that might happen.

BRIAN SCHWARZ: Yes, it's an early beta build of the software but the functionality's pretty much the same as what the released product will have.

ROBB BOYD: Well let's fire it up.

JIMMY RAY PURSER: Let's see what you've got.

BRIAN SCHWARZ: So the first thing you'll notice is that it's a classic embedded device manager, so it comes free with the system. And first it does discovery of all the components, so you can see here that I have a pretty simple system with just a simple chassis, one chassis, one fabric extender, shows you the server itself and the interface card. It also...

JIMMY RAY PURSER: You know what though? Well wait a minute though, let me kind of interrupt you here because you're showing me this stuff here. And I'm thinking, if I'm watching this show and let's say that I'm a SAN guy and I've got a certain level of insight into my network, this is not really that impressive.

BRIAN SCHWARZ: You know this is a lot different than traditional device managers because what you had in the past was you would have different device managers for each of these components and each administrator would login and manage a different component. So there's never one view where could see the access layer switches for both the LAN ports here and the SAN ports and all the things about the servers and you know I/O elements in the chassis.

ROBB BOYD: So you're saying in the normal world that we're accustomed to and the shift we're going through on today's show...

BRIAN SCHWARZ: It would different theories.

ROBB BOYD: You couldn't have seen all these things, it would've been individual device managers, so don't overlook the fact that we're looking at it all together here now for the first time ever.

BRIAN SCHWARZ: Yes, yes, and imagine the system at full scale, 40 chassis, 80 fabric extenders, 320 blades and adapters all in a single pane of glass.

ROBB BOYD: Does that answer your questions though?

JIMMY RAY PURSER: Yes, it really does because that is really -- the difference here is that by integrating the network -- I won't to say network layer because that obviously doesn't map too well -- but it's the network portion of this that actually makes the difference.

BRIAN SCHWARZ: Yes, the fact that the access layer of networking and the servers themselves have essentially been fused together from a management perspective into a single system. It's all one device essentially, so you could see in all these things it's kind of still the -- to the question you asked before Robb about, how does this play with existing specialties? There's kind of -- envision all the --

ROBB BOYD: (Inaudible) management stuff?

BRIAN SCHWARZ: Yes, there's kind of tabs here that you'll see. There's a place for a server admin to do most of their work, a LAN admin to do most of their work, and storage guy to go in and it can configure...

JIMMY RAY PURSER: Hey that's sweet, that's cool.

BRIAN SCHWARZ: Yes, so there's -- the logical administration...

ROBB BOYD: (Inaudible) because I don't know if everybody caught that, so what are you talking about?

BRIAN SCHWARZ: So you can see these three big tabs across the top. So you'll still have role specialization. There's a lot of knowledge that gets built up in individual server teams in the LAN networking team and SAN team to have a place for them to do their traditional jobs, to go in and do the same types of tasks they've done before. But now they're putting it into a single system as opposed to a bunch of discreet device managers.

JIMMY RAY PURSER: That's very cool, that is very cool.

BRIAN SCHWARZ: So let me just quickly show you, if we wanted to create a service profile.

JIMMY RAY PURSER: All right, first of all, now what is a service profile?

BRIAN SCHWARZ: A service profile is the way that we essentially implement the stateless computing concept, so you create essentially what you want to be, the end state of the server and the access layer connectivity in the actual system. You create it in software, it's essentially stored in software and the three specialties kind of put in what the configuration should be. And then you could choose to deploy it out to an individual server in the large kind of compute that's connected to the fabric interconnects. And when you deploy it, it configures not only the server but the adapter, the fabric extender and both the server facing port and uplink port on the fabric interconnect all in one shot.

JIMMY RAY PURSER: Let's see it man, let's go.

BRIAN SCHWARZ: So let's go in here and do a quick TechWiseTV service profile here. And so one of the things you'd first do is you would create essentially what you want your NIC to look like. So you're going to go and create a NIC. And not only are you going to say, I want a NIC, I want it to be on a certain VLAN that's already been defined by say one of the network architects who's gone in to find all the VLANs. I'm going to choose this one here, TechWise VLAN.

ROBB BOYD: Oh, thank you.

JIMMY RAY PURSER: Love that.

BRIAN SCHWARZ: Yes, go in and --

ROBB BOYD: That's the party VLAN.

BRIAN SCHWARZ: Go ahead and connect it and here's some of the -- you can quickly see some of the policy-based nature of the systems. So we're pretty big on policies because the data center's becoming a lot more dynamic and you want to have some of the experts define a set of policies and then some other admins, maybe some of the more junior admins --

JIMMY RAY PURSER: What's a pin group?

BRIAN SCHWARZ: Just go in and select essentially a bunch of drop-down boxes here that a pre-defined policy. And of course from the expert in the back, I've already created a certain pin group which is, which uplink port do you want to be configured to for this service profile? Just go ahead and select that and again when you deploy this out, it doesn't matter which chassis, which blade you define. This service profile, the LAN traffic for that is always going to come out of that uplink port.

JIMMY RAY PURSER: Groovy man.

BRIAN SCHWARZ: Yes and you basically go through and do the same thing for -- no, let's not do that, I want to cancel it and just finished. You basically go through and do the same thing for the HBA, determine which VSAN you want to be on. You can select which server you want, pick an individual server or select on a pool that's based on certain hardware characteristic. Give me one with the memory expansion, give me one with the fancy virtual interface controller in it. Same thing, set up all your boot order.

JIMMY RAY PURSER: Set the boot order, how cool is that?

BRIAN SCHWARZ: Yes, yes, so you can go in and say, do I want to do local boot, do I want to do SAN boot, do I want to do PICC boot?

JIMMY RAY PURSER: Awesome.

BRIAN SCHWARZ: And again when you deploy it, it goes in and programs the BIOS directly.

JIMMY RAY PURSER: Unreal.

ROBB BOYD: I'm kind of curious because with your background from a server perspective and stuff, is it even possible -- because I know you've been involved in this deeply in the development David, -- but as you pull back and you think in the old world about the way you used to think and the development you were doing at Sun and stuff, what's this kind of stuff mean to you as you look at this? And you're seeing kind of a baby having the birth here if you will. Here's a child that's come out. Any thoughts on what this means.

DAVID LAWLER: It's solves the adolescent problems in the child we have basically. You know what? You've got all these different systems and, as Brian's saying, different tools for managing them, different tools for taking care of firmware, different tool for taking care of the adapters, different tool for taking care of identity, all that. And people have to coordinate that. What this does is it actually makes it single, right? It makes it actually working with the infrastructure. It's a addressing a problem that we had that frankly the only way that people ever should be able to look at it because of the way they've been doing business basically is additional tools, more tool. I'll give you another tool for that, wrap it in another tool and --

JIMMY RAY PURSER: Yes, that's right.

DAVID LAWLER: And that's what the answer is. So this basically says, okay, you know what? Let's make the device itself as a whole smart and then give you an interface into it.

JIMMY RAY PURSER: You know there's so much here obviously. That's incredible and I think that as I was learning more about this, the thing that kind of just absolutely blew me away was some of the problems you guys solved with some of the memory mapping. And that's such a big topic because when it comes to adding memory in machines, a lot of server folks out there today, they're buying four-way systems, only putting a couple of products in just so they can get the extra memory boards.

DAVID LAWLER: Right.

JIMMY RAY PURSER: And so I mean we're really talking about a new way here to really cut back on costs just because of hardware and where we're deploying this stuff. But how did you guys solve some of these problems with this mass amount of memory here? I mean if I saw that on a data sheet, I'd be like, huh, there's Marketing at work right there.

DAVID LAWLER: Yes.

JIMMY RAY PURSER: There's no way they can do this.

DAVID LAWLER: Yes, so what we looked at, I mean we looked at the end of the blade basically and we looked at the capabilities of what we could do with Nehalem and we talked to customers. And first of all it was identifying the problems, right?

JIMMY RAY PURSER: Right.

DAVID LAWLER: You know, is memory the problem or what is the problem? It was looking at virtualization. People were running out of compute cycles before they were running out of memory -- vice versa, excuse me. They were running out of memory before they ran out of compute cycles, sorry about that. And so how do we solve that? And you're right, people were looking at four sockets, so it was extremely insightful. And we work with an absolute gifted engineering team who looked at the problem and saw a simple answer, rock stars. I mean a simple answer that basically was in hindsight looks, well that's simple. But like anything, you look at the system. And this is a Nehalem system, it's got 12 DIMMs, right? And so each one of these DIMMs -- and I have no idea what the size of this -- let's say these are 8-gig DIMMs. Each DIMM is a physical, single physical DIMM. That's what the memory controller's looking for. Now in Nehalem the memory controller's actually on the processor, so the memory controller...

JIMMY RAY PURSER: That's different.

DAVID LAWLER: Yes, that's different, that's new to this new architecture. They just announced it and basically what that does basically is it actually gives them much higher performance, more bandwidth, a whole bunch of pluses to it, so it puts them back in the game on that standpoint. But you look at the memory, and I've got on this KISS system 12 DIMMS and that's at full performance and so at maximum memory bandwidth. Well we looked at it and as I said just absolute ingenious insight into this basically was, well if I've got one DIMM that's 8-gig, can I represent that as four DIMMs at 20-gig? Sounds simple and it's a simple -- in networking it's like I'm addressing a 8-gig DIMM. I've got the addressing mechanisms, let's see if I can do it. And you know what? You can't. So what we're doing basically with that system is you'll take effectively the exact same blade as this from a compute standpoint, same processor, same capability in a disk subsystem, same service processor. And then because we want to have more DIMMs, effectively it's double (inaudible), so we have up to 48 DIMMs on that system basically and you can use less-expensive DIMMs. So instead of a bunch of 8-gig DIMMs, I can use a bunch of 2-gig DIMMs.

JIMMY RAY PURSER: Which is tons a lot cheap in cost.

DAVID LAWLER: Yes, but the other side is that, you know you asked that question earlier about is this something out of the parts

bin? This is ASIC-level, picosecond basis of computing, I mean just basically just picosecond networking effectively on the memory bus, so it's ASIC development from the ground up to be able to do this and that's what we've done. I mean basically doing it, the net basically is that you save just on memory costs alone 50%, right?

JIMMY RAY PURSER: Oh yes.

DAVID LAWLER: Compared to other systems out there. Some people say, well I don't need 384GB, because that's our maximum. On a 2-socket machine we can go up to 384GB of memory.

JIMMY RAY PURSER: That is unbelievable.

DAVID LAWLER: I know, you know, you look back...

JIMMY RAY PURSER: That's, that's...

ROBB BOYD: Well it was previously considered impossible, right?

JIMMY RAY PURSER: Yes.

DAVID LAWLER: Oh yes, exactly. I mean you look back even when 64-bit computings just came out in 2004 timeframe, at that time an x84 box had if it was lucky 4GB.

ROBB BOYD: That was a lot.

DAVID LAWLER: That was a lot, that was a lot.

JIMMY RAY PURSER: That was a really (inaudible).

DAVID LAWLER: That's exactly right. You know, yes, people figured out that they can use the memory now. And virtualization's eating it up, transaction processing is eating it up, search things are eating it up. Well what you can do then is say, okay, I can go to larger memory footprints. And people are doing that and so we've got customers who have said they looked at it and we talked to them about virtualization. They said, man, I can use this over here and I'd like to use this over here. And they're looking at new business opportunities, everybody's been analyzing their data better. Other people are going, you know, normally it would be 96GB but this will save me in that case 60% of the cost.

BRIAN SCHWARZ: You should be able to have higher consolidation, more VMs on a same physical box.

DAVID LAWLWER: That's exactly right.

ROBB BOYD: And this boils down because of the way software is charged for, right?

DAVID LAWLER: Right.

ROBB BOYD: I mean you talk about huge cost savings which contribute very quickly to the ROI when you talk about deploying a system like this in terms of the benefits it can give you. And what I'm worried about is that we've got a lot of revolutionary topics to discuss. One of the things I want to make sure we cover just before we leave, kind of drawing in and talking about the architecture while we've still got both of you here, benefit of 20/20 hindsight now. We've been blown away at the shows that we've done of the various Nexus products but we were looking at them from a networking-only perspective and we were still very impressed, don't get me wrong. But now all of a sudden you go, you guys were holding out on us because now we begin to see, oh, that's why the ports were in the back. That's why the airflow was done the way it was and such because there was a bigger purpose in mind. Can you tell me as briefly as possible because I know that may not be a good thing to ask but tell me how that stuff ties in? How do you normally speak to it so to make sure we don't miss those points?

DAVID LAWLER: When we talk to people about this, there is so much innovation in this. There's too much at one time and you look at unified fabric, you look at vnlink, you look at the fabric extender technology just in the networking side alone. Then you look at the memory, you look at the management, you look at the virtualized adapter. When we talked to people basically is we made sure we're bringing all these other technologies, as many technologies as we can to market separately. So unified fabric's available, fabric extender's available as a product, the vnlink technology is being made available as well in the next generation of VMware.

ROBB BOYD: Right.

DAVID LAWLER: So all those technologies, bring them to market so people can understand them and then it's not all new. But what we do basically in this whole system, you know, it relies on those technologies, people, and then it scales.

BRIAN SCHWARZ: Wraps the management around the whole system.

JIMMY RAY PURSER: I'm amazed at the -- because I'm seeing, I'm thinking like, you know, the management's really the unsung

hero of this story. I'm like they're waiting to see an adapter card, they're waiting to -- it's really the memory when it comes to it. No, wait a minute, it's really how the chassis's built, like all the airflow in there. And I don't even know if you can point to it. How do you even summarize this?

DAVID LAWLER: I wasn't even going to try. I was going to hand it back over to Valerie.

JIMMY RAY PURSER: It's just, it's unbelievable. Yes, take it away Val.

VALERIE ST JOHN: And he was handling ever so nicely.

BRIAN SCHWARZ: Will you summarize it for us Valerie?

VALERIE ST JOHN: Thanks very much guys.

DAVID LAWLER: You're welcome.

VALERIE ST JOHN: David Lawler and Brian Schwarz.

BRIAN SCHWARZ: Thank you.

VALERIE ST JOHN: Glad to have you and next we welcome Dave Major from Learning at Cisco.

DAVE MAJOR: Hi Valerie, how are you?

DAVE MAJOR: It's very good to see you.

VALERIE ST JOHN: In black on black, the very hip, chic San Francisco uniform.

DAVE MAJOR: I had to keep it cool because I'm here on TechWiseTV. I know how cool you guys are.

VALERIE ST JOHN: We are tragically hip on TechWiseTV. What have you got for us today?

DAVE MAJOR: So I have a couple of brand new things. We were talking about data specialization, so we have two brand new specialist certifications that are going to come out.

VALERIE ST JOHN: Huh, do tell.

DAVE MAJOR: They're Data Center Virtual Specialization for Design and for Support. And those are going to be coming out probably a little bit later in the year. We'll have Marcello Hunter back on sometime a little bit later on to give you a brief update on that. And then we also want to tell you about some brand new courses that are going to be coming out on the Cisco Learning Network Store.

VALERIE ST JOHN: Excellent.

DAVE MAJOR: We actually have a brand new store that's going to be coming out right around the time that we're here on TV, here on TechWise, and tell you all about that. We have brand new eLearning products that are going to be coming out. I think we talked about the Mind Share Game on the last TechWiseTV. You can actually purchase that on the CLN Store as well as the brand new CCIE 360 routing and switching product. So those are all available, you can actually go onto the Cisco Learning Network and purchase those.

VALERIE ST JOHN: Oh good.

DAVE MAJOR: So for the first time you can actually go online and purchase Cisco education products from Cisco.

VALERIE ST JOHN: Excellent, Dave Major.

DAVE MAJOR: And I'll give you a quick update just to tag it, www.cisco.com/go/learnnetspace.

VALERIE ST JOHN: Excellent, thanks so much Dave.

DAVE MAJOR: Thank you.

VALERIE ST JOHN: Okay Robb, you did a good job of summing up for us. You kept your promises, you hit the discussion points, sum it up.

ROBB BOYD: That's the big question though. Did we really? You know, because the big fear is that when you talk about all the intelligence that's been put into this and the engineering talent that's led up to this because as we've worked with this team, oh my gosh. Not only are these guys highly intelligent, they're very busy, they've been leading up to these points today, and we're blessed

to get to work with them. At the same time I hope we're doing them justice with regards to how we're covering it on the show. I was trying to take some notes here because some things were hitting me and some points that I think we need to remember when we look at this. First and foremost is that if anybody's watching this show today, we've just been through a full hour of content just about. If anybody thinks that this isn't a fundamental change to the way things were done and they see those points that we've been trying to do a good job of drawing out here, rewind, go back, watch it again because you missed something. Now I apologize if it was us and the way we explained it because that was bad.

VALERIE ST JOHN: Yes, I was just going to say we probably missed something too. There's just so much we could've covered.

ROBB BOYD: But if you think about it, there's this systems orientation, this ability that says, you know what? We're not going to force you to integrate everything on your own anymore. That's what's really jumping out at me. It's saying there's a lot of things that can be done together. This is not about a networking company trying to reach out and act like we know what servers are. We've got server people that have proven their abilities in the market. We've brought them in to help solve what were really looked at as being problems you couldn't solve before. No one even thought to potentially address that. That really jumped out at me but the ability to remove operational complexity, the integration of these things, the memory breakthroughs -- I forgot to ask how many times we probably have on all this stuff -- but it's amazing. It's fun to be a part of it.

VALERIE ST JOHN: Oh, no kidding.

ROBB BOYD: Bottom line and we talked about this at the beginning, it all boils down to this is technology stuff. And when you talk to the business folks and we talk about the reason all of this exists, it's about agility. It's about, how can we improve our ability to respond to the changing marketing conditions which continue to happen much faster? There's many more dynamics we have to deal with. These are the kind of solutions that allow us to focus on the business at hand so we can add value to the things that are going to make a difference. And I sure hope that everybody got that point clear.

VALERIE ST JOHN: Sure hope so, that's the final word. For Robb Boyd, 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.