DEFINING THE DATA CENTER SWITCH

PRINT DOCUMENT

ANNOUNCER: Welcome to TechWiseTV on the Cisco Interaction Network. Please remember to submit your technical questions in the field provided at the bottom of the viewing console. You can also email questions, comments and suggestions to techwisetv@cisco.com. Thank you for joining us, and enjoy the show.

JONAS TICHENOR: The data center is the heartbeat of the 21st century business. Today we focus on mission-critical IT computing as we look at Defining the Data Center Class Switch. Our own Robb Boyd sits down with Doug Gourlay, Senior Director, Data Center Solutions, to explore the challenges and demands that are putting data centers under the gun to deliver.

DOUG GOURLAY: When your cell phone has 20 times the storage capacity of your Microsoft Exchange mailbox, you often have to wonder what's going on here.

JONAS TICHENOR: You'll get a first look at the incredible Nexus 7000 and hear the top five requirements for a data center class switch.

JIMMY RAY PURSER: This is good stuff, man, this is really built hard. As a matter of fact, this was designed to support 384 Cat 6a cables, that big, thick, heavy stuff.

JONAS TICHENOR: We'll show you advances in network virtualization and how it's changing today's data center.

ROBB BOYD: How can we deploy capacity in applications quicker right where we need them, when we need them?

JONAS TICHENOR: And we'll reveal the secrets of the new operating system that puts the Nexus 7000 in a class of its own.

JEFF RAYMOND: The NX-OS operating system actually can do self-healing.

JONAS TICHENOR: Well, you've impressed Jimmy Ray, I don't know if you meant to do that.

JEFF RAYMOND: Hey, my work here is done.

JONAS TICHENOR: Exactly, exactly.

JONAS TICHENOR: Restock your pocket protector because you're going to want to take some notes. This is TechWiseTV, technology you can use from geeks you can trust, only on the Cisco Interaction Network. Welcome to TechWiseTV, I'm Jonas Tichenor along with Cisco solutions experts Robb Boyd and Jimmy Ray Purser. And our focus today is Defining the Data Center Class Switch. Of course, what business is not dependent on their IT resources these days?

ROBB BOYD: We know there is no question that the data center has become the lynchpin for any business that depends on their data network. Data centers are under increasing pressure from regulatory demands, higher expectations of service delivery and a not-so-distant history that relegated the IT department to a simple expense to be managed. Add in power and cooling challenges and you can see that the expectations for data centers to deliver are very high.

JONAS TICHENOR: Yes, so talk to me a little bit about what Cisco's experience in this area is.

ROBB BOYD: Well, Cisco has made advancements on many fronts in the data center. Incredible innovation around application networking, WAN acceleration, storage, security of course, and let's not forget about switching.

JONAS TICHENOR: Okay, talk to me about that switching.

ROBB BOYD: Well, Ethernet in particular; you know, our advances with Ethernet have been well placed, it turns out. And it was not always so obvious that this would be the standard that everyone would rally around, but just look at the evolution. It took us eight years to go from 10Mb to 100Mb, just two more years to get to 1Gb -- that was in '99. 10Gb of course is commonplace now, and 40Gb to 100Gb Ethernet is on the horizon.

JONAS TICHENOR: Yes, and obviously you have seen some trends that would kind of show us why the data center is really important to folks right now.

ROBB BOYD: Well, from an architectural perspective, a dramatic re-centralization has taken place, and for a number of good reasons. Just look at the history -- remember the mainframe days?

JONAS TICHENOR: Oh boy, do I ever, Jimmy Ray?

JIMMY RAY PURSER: The HAL 9000, remember that? HAL 9000, big hair.

JONAS TICHENOR: Big hair, tight pants, I remember.

ROBB BOYD: And I wouldn't call that one a dumb terminal necessarily, but we moved from that to the PC age. Push all the intelligence out to the edge, to the user. Now the issue with that was loss of control, what does IT do for me? Yes. Now we're seeing it come back and we're seeing this re-centralization and the benefits we get are control, consistency, cost reduction, you name it. This now defines what's become a data center renaissance; IT is relevant again. And this is largely under the guise of reduced operating expenses as IT now has a business relevance that nobody can challenge.

JONAS TICHENOR: Yes, I like that when you say renaissance, that's very...

ROBB BOYD: Do you like the renaissance?

JIMMY RAY PURSER: Like the Medici renaissance in Florence.

JONAS TICHENOR: Yes absolutely, so what you're saying essentially is that the network is becoming more relevant to a specific business.

ROBB BOYD: Absolutely, and this is what Cisco refers to as Data Center 3.0. And as I sat down with Doug Gourlay to talk about this, he is Senior Director, Data Center Solutions at Cisco, I asked him what does that mean?

ROBB BOYD: Doug, I appreciate you taking the time to sit down. It's obviously a very busy time, so I appreciate that.

DOUG GOURLAY: You're very welcome.

ROBB BOYD: Let's get right to it -- what is Data Center 3.0?

DOUG GOURLAY: Well, Data Center 3.0 is the third evolution of the application architectures that then define how we have to build and operate our data centers. We went from mainframe worlds to client server to now a world of service-oriented architectures and Web 2.0 technologies which really create the most network-centric data centers we have ever seen before. And to respond to that we had to increase server capacity, increase network capacity and increase storage capacity. Yet first we ran into some physical boundaries, boundaries of space and power, thermal capacity of our own data center facilities. So Data Center 3.0 is about understanding the application architectures that define it but then also working within the physics boundaries we have of our facilities. Maximizing efficiency, productivity and there will be more IT workload output for a given power input. Leveraging a series of network-centric, server-centric and storage-centric technologies working together. One of the key ones is virtualization.

ROBB BOYD: Interesting, tell me a little bit more about virtualization.

DOUG GOURLAY: Sure, virtualization generally is seen, starting from the server world where we start breaking a server up into multiple instances. We can have five operating systems running on one physical server. We often joke that the server becomes software. All right. Well, then there becomes a set of network dependencies. So we have to link those different operating systems together. We have to provide segmentation between them and as you compress them, the bandwidth actually increases, it doesn't go down. So again, it becomes yet another driver for increased network capacity and demand.

ROBB BOYD: Interesting. So with these changes afoot and this type of thing, what would you say is the current state of data center networking technology right now?

DOUG GOURLAY: The majority of the market out there and what our customers are doing is Gigabit Ethernet connections to the server. So 10Gb in the aggregation, multiple Ethernet networks, Fibre Channel networks. If you look at a server you often see -- we often call it the server arachnid because you have the two Ethernet front end, two Ethernet back end, Fibre Channel fabric A, Fibre Channel fabric B, lights-out management. VMotion oftentimes, from VMware, will get its own interface; backup and restore networks and this proliferation of interfaces coming out of the server. So it starts demanding a larger and larger server, not because of hard drives or not because of number of CPUs but because of the amount of memory in the I/O interfaces.

ROBB BOYD: Well, with all these increased interfaces what are the pain points that come from that, is that one of them?

DOUG GOURLAY: Well, one of the pain points actually, I mean the sheer amount of network technologies and number of networks you have to manage. But in the end it starts creating some boundaries and barriers to the adoption of further virtualization technologies. Like how can we move a virtual machine from one rack to another, one row to another or one data center to another faster than the client can detect that it failed? You start running into addressing problems, storage problems, segmentation problems, all of which networks are poised to solve.

ROBB BOYD: The idea being there that if something fails in the data center but nobody notices it, did it really fail.

DOUG GOURLAY: Exactly.

ROBB BOYD: Something like that?

DOUG GOURLAY: Exactly, because in the end it's about supporting our end users. And we're finding nowadays we have this empowered end user, the type of recent college graduate coming into the workforce who has a higher expectation of IT service delivery than many IT departments are prepared to fulfill. When your cell phone has 20 times the storage capacity of your Microsoft

Exchange mailbox, you often have to wonder what's going on here, why can't IT get it? So we need to make IT empowered. We need to give IT the opportunity to be very responsive to their end users and have an infrastructure that has that perception, and from the network perspective always available. But also being to move dynamically instances of applications from one location to another, balancing thermals, maximizing application responsiveness and shifting workload to where it can best be processed in the most efficient way possible.

ROBB BOYD: Yes, shifting expectations, I think that's a powerful point as well.

DOUG GOURLAY: That's a good way of thinking about it, yes.

ROBB BOYD: Thank you, Doug.

DOUG GOURLAY: You're very welcome.

JONAS TICHENOR: So Robb, we are going to hear more from you and Doug a little bit later in the program but it's clear that a lot of energy is being put into the modernization of data centers.

ROBB BOYD: Well, ultimately it's about solving business problems and how can we make the data center infrastructure last longer. How can we increase stability and reliability to reduce service costs? How can we deploy capacity in applications quicker, right where we need them when we need them?

JONAS TICHENOR: Absolutely, so hang on to your propellers, folks. In segment two, chief geek Jimmy Ray Purser takes us into the world of data center switching and gives us a first look at the new Nexus 7000.

JONAS TICHENOR: Today business needs are driving advances, and we certainly need technology that can keep up.

ROBB BOYD: Increasing dependence on IT systems means that there is zero tolerance for downtime, and the window of opportunity for maintenance is shrinking. Data centers are evolving rapidly just to keep pace.

JONAS TICHENOR: So exactly what makes a data center class switch?

ROBB BOYD: Well, Jimmy Ray Purser put together a list for us, top five requirements for a data center class switch.

JIMMY RAY PURSER: You know it, absolutely.

ROBB BOYD: So I've got this here, let's go through this.

JIMMY RAY PURSER: All right.

ROBB BOYD: The first one on your list -- upgrades with zero service disruption.

JIMMY RAY PURSER: Well you know, that's always been the dream, to be able to upgrade firmware, software in the switch without taking the switch down. And you have to consider how that's actually going to be done. I've got a running switch that is actually sitting there doing processing and stuff, and then I want to take that image that's running, put another image here and then switch them out real quick. And I've got a lot of other processes and a lot of states attached to this image and so what I need to do is kind of pseudo-attach and gracefully put these over. And it's just been a dream for a very long time to do that, and in data center type of environment that's absolutely critical. Downtime in these environments is measured in seconds lost so every second is worth so many thousands of dollars. So yes, in a data class switch I absolutely have to have a zero service disruption. It's when you mentioned maintenance windows decreasing and things of that nature, we're talking about 24 by 7.

ROBB BOYD: That's part of it, right?

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: Okay, well second on your list is lossless fabric, what is that?

JIMMY RAY PURSER: Well, lossless fabric is again, it's kind of one of these dreams, if you will. Having a backplane, a forwarding plane, a control plane that is going to check and make sure it can send data before it sends data. A lot of people look at buffers and they look at speeds. You can see all these performance tasks and so we're going to do some of that stuff too. But a lossless fabric is a fabric that's going to check and make sure that the fabric is available, that if I've got to send 10Gb worth of data, number one, do I have 10Gb available to send on that backplane, and then transfer that over and then make sure it went. A lot of devices today, it's just like I've got 10Gb there, I'm going to send it and if it makes it, well it makes it, and if it doesn't, well it doesn't. Well data, I mean we've talked tons of times before on this show how much data is worth. If I start losing data in a data center switch that's running 10Gb and all that stuff, your losses add up pretty darn quick.

ROBB BOYD: That sounds critical.

JIMMY RAY PURSER: So lossless fabric is absolutely critical.

ROBB BOYD: Moving to the next one, this has got to be huge -- physical and environmental design. There's got to be a number of

things in that category.

JIMMY RAY PURSER: Well, you know, I think every IT person out there today is asked to do more with the time and with the skills that they have. A lot of us in any IT job are hired to do one thing but it always sprouts into multiple things and we feel like we're doing the work of three different people. Ask any networking engineer who's now been assigned to understand how phones work on it and now understand how power works on it, now understand how code works for security. I mean, we all get stretched. This type of design means that I can, as I'm walking by and scanning, I can look at the lights and I can tell what the health of my switch is. I can look and I can see my cable manager, I can look at it while unbundling all these cables, I can look and trace that stuff out visually. That makes sense. And fast, I can troubleshoot right on the spot.

JONAS TICHENOR: That's fast.

ROBB BOYD: So LEDs, intelligent power management, you mentioned cable management, that type of thing.

JIMMY RAY PURSER: Absolutely, yes, those physical attributes are key. I cannot tell you how many times that I've spent chasing gremlins on the network and it's a physical problem or somebody's hit it with a broom and knocked a connector out.

ROBB BOYD: They say always check that first, right?

JIMMY RAY PURSER: Absolutely, absolutely.

ROBB BOYD: Yes, number four on your list is lights-out management. Okay, when I say boom boom, that's it.

JIMMY RAY PURSER: Out go the lights.

ROBB BOYD: That's perfect, well Pat Travers there.

JIMMY RAY PURSER: Well, Pat Travers' action there, yes.

ROBB BOYD: But seriously though, you do manage -- thank you -- lights-out management, what do you mean by that?

JIMMY RAY PURSER: Well, you know, it's kind of funny because lights-out management means a lot of things to different people. A lot of people think that okay, I've got lights-out management means I have a remote console into my data center and then I just turn the lights out. Well, that's the spirit of it. But realistically what that means is that truthfully you don't have to go in this data center to do anything. It's designed and laid out so that even if I have a failure, even if I have the box that dies on me, I can still get access to it and that's a big thing. As long as everything's working, everybody has got a lights-out data center. You can manage it with SNMP, you can actually interrogate it with --

ROBB BOYD: As long as everything's working, that's the key point.

JIMMY RAY PURSER: As long as everything's working and when things fail -- that's what I always tell folks. That's how you know how good of a business and how good of a network you have when things are falling apart. And then you know, lights-out management, when things fall apart, I can still access that gear without having to go in there and actually touch it.

ROBB BOYD: All right, this last one -- of course not least -- forward investment protection.

JIMMY RAY PURSER: Well, you know, that's a huge one because a lot of times the design of a product looks really good the day that it's purchased, and I've been caught in this. So I buy a product and it looks great and I'm like, holy smoke, look at all this capacity. We've all heard the argument, who needs more than 64k of RAM, how can you ever use that? So it's one of these things where if I buy this switch, I'm not in this endless hardware upgrade cycle all the time. I put it, it stays, it lasts for a long time.

ROBB BOYD: That's a good list.

JONAS TICHENOR: Great list, guys, does a great job defining the challenges that lie ahead for the next data centers. Now Jimmy Ray, despite an increase in security, you managed to find your way into Building 7.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: Building 7, that's the data center networking architecture lab, I believe.

JIMMY RAY PURSER: That sure is, I took Doug Gourlay's ID when he was here.

JONAS TICHENOR: Let's take a look.

JIMMY RAY PURSER: Man, I love this sound. From this I'm obviously in the data center and you can hear all these awesome fans burning. And I'm here to look at one of the most awesome data center products that's ever been designed. This right here is the Nexus 7000 series data center switch. This switch was designed by data center people for data center people. It is revolutionary; it's designed to be serviced easy, it's designed to be maintenanced easy. You can just look at it and see what's going on -- take a look at this. I've got all my status lights, my main status lights that tell me what's going on with the switch right here on the very top, very bright green LEDs. So as I'm walking through the data center and just trying to get that visual queuing of what's going on, I can

actually look and I can see the components and what their status is doing on the back of this switch. I can tell if they're doing bad just by taking a quick look and not ever look at the management console. I walk in here and take a peek around. One quick look tells me I'm all green and everything's ready to go. That's just a little piece of it; let's take a look at some really cool stuff in here, how about this? Got a nice door cover to this, a really cool lock. So it keeps people from having all those physical errors that happen in data centers where people are cleaning up and take a broom and snag a cable or just anything it can cause in here. I can't get in here to the switch unless I open it up, and that's the best part. Let's open this thing up here and let's take a look. These doors come shipped with the product but you know what? We can take these off if we don't like them. In fact, I love them but this is what I'm talking about. You know, in any data center one of the biggest problems is cable management. Look at the cable management built in this product. This is not the cheap stuff. This is good stuff, man, this is really built hard. As a matter of fact, this was designed to support 384 Cat 6a cables, that big, thick, heavy stuff, hard to bend, 384 can fit in here, absolutely no problem and it supports that good, solid stuff. We're using more tie-bands in here so I don't have to replace these by the thousands. I can reuse these, replace them and look, they're strong, man. I can jump on these things and just keep and reuse them -- great stuff. Well let's go down a little bit more and look at this product. Really nice, common fabric or switch layout here. I've got all my ports on the front obviously so I plug everything in. I've got my supervisor cards right in the middle so I can actually see them. And actually even on my supervisor cards, my flash cards that fit right in the center of this, I can slide right in here and if there was somebody popping them out or causing more errors, I could put a cover over them and lock them down. And here they've got these really cool logos, we'll throw those in for free. I've got my 10Gb cards that are sitting right here, we have that dedicated port slots that tell these are my high-performance slots. So if I want to run at line rate, I just plug in my SFPs in these dedicated ports and these are guaranteed line rate ports. And you know what? This is what really amazed me. These are 10Gb SFP+s -- I've never seen a 10Gb SFP, this is awesome. This can address the ability to add tons and tons of 10Gb ports in this switch so I get that higher line rate of data transfer that's imperative in any type of data center. Let's go down a little bit more and look at some of this cooler stuff. How about some of these line cards, no one's seen these before. There's nothing new on this but this is what makes it really cool. These are actually locked so I can't pop them out. Once again nothing new but check this out. If I've got to take a card out and service it, I don't have to go back and do a shutdown command on these ports or turn the line card off. I can actually pop these buttons and the second I pop these, I've got a limiter in here that actually cuts the circuit off and I can pop this board out. It powers it down just by mechanically pressing these two buttons, how cool is that? And then of course, one of the more important features about this product, I've got front-to-back air flow on this product. This is absolutely critical in any data center design, so I've got a cold aisle and a hot aisle. My cold aisle, and man, you can really feel the air coming through this product, it works great. My cold aisle is pulling all my air through the front and it's spitting it all out through the back, and there's a lot more cooler components in the back. Why don't we go round and check this out, let me button this thing back up. Robb, tell me how am I going to get this thing on the plane with me. It certainly doesn't fit in the carry-on that well. Let's go take a look at the back, there's some more cool stuff back there. Y'all, this is where the rubber hits the road, and a lot of switches out there today, when you go to the backup switch, there's nothing there but a blank panel. The Nexus 7000 series, fully functional. Every piece of this product actually has been engineered and designed to be functional, perform and to work perfect in the data center. Man, you can really feel the air pushing out of this product because I actually do have the front-to-back air flow. This would be the hot aisle side of this product. Now let's take a look at one of my favorite features of this product -- there are so many. This is one of my favorite features as far as replacement goes, this is the ID light. The ID light actually allows me to light up a component, and this works all over the front cards, the back. If I discover a problem, anything that needs to be services, I can actually light up a light and dispatch somebody to service that product and not have to worry about did I say the right slot or terminology correct or whatever. I can just say replace the card that's lit up with the ID light and that makes it excellent because I don't have any human errors that could be involved in here. And if I slide on down to this product, this is where I have my fabric in here. These are five fabric cards in here that provide 230Gb per slot for the front slots. That is amazing, 230Gb from day one. That means I can actually upgrade and do things as they roll out, like 40Gb or 100Gb. This box is designed to scale, and check this out. One of my favorite features when it comes to taking switches apart is actually looking at how the cards are designed. Now what we did was we took a good thing and added on to here. We've still got these same type of release levers on here that allow me that same functionality I had in the front, and here I'm going to pop out my fabric card, check this out. On a half board I've got all that intelligence on a single ASIC right underneath on this heatsink. And this is great news because the less components I have on the board, the less problems I have actually to fail, the less components that I have to die on me. And I have (inaudible) scores that are astronomical, and this is excellent for a data center switch. Put this thing back in, there we go. Now let's take a look at the power supplies. These are 6kW power supplies that are actually going to power this. I have got three of them in here and this is great because I can actually design these so I can have redundancy in my data center. Or more importantly, I can actually have a grid A and grid B type of redundancy solution. I can have it so that I can actually lose the entire power grid, and grid B on this example, but actually keep my switch up and powered and delivering data in any type of catastrophic emergency. This is excellent plus it is designed to be serviced. Look how big these handles are, these are not small, I can put my whole fist on these thing and actually pull them out. And what I really like about this product, all my serviceable components are in the back away from my cabling plant. So I don't have to worry about unplugging cables and a failure that I can have replaced in minutes, I don't have to worry about taking cables or taking hours actually to do. I can light up my ID light, send somebody out, pull a fan tray out, put another one in, boom, we're done. I love this switch. I'll tell you what, I'm going out to look for a forklift to get this thing back to my lab in Wisconsin so I can actually start packing some serious data.

JONAS TICHENOR: Well, well, well, Jimmy Ray Purser caught on tape. I think the cat is out of the bag officially, what do you think?

ROBB BOYD: It's definitely out of the bag, good stuff.

JIMMY RAY PURSER: Thanks, man, that was happening.

JONAS TICHENOR: Well, in segment three, guys, we are going to actually -- and I didn't know we could do this -- widen the geek net and break down this new beast to see exactly what it's made of. Online, on demand and on the go, this is TechWiseTV, technology you can use from the geeks you can trust. And today we are Defining the Data Center Class Switch and Robb, I'm going to underhand one to you right now, so just get ready, just get ready.

ROBB BOYD: Make it a big, slow one, all right.

JONAS TICHENOR: How important is reliability in the data center?

ROBB BOYD: Thank you for the soft ball.

JONAS TICHENOR: You got it.

ROBB BOYD: I think the obvious answer is it's very important. The not so obvious answer, but common in this world, is well, it depends and it's depends on your needs.

JONAS TICHENOR: Absolutely, so the more mission-critical the application, the more redundancy and security is required.

ROBB BOYD: Exactly, and as data centers can be classified by tiers, and with tier one being the most basic and inexpensive and tier four being the most robust and costly. A tier one data center is not required to have redundant power and cooling infrastructures for example. But it needs only a lock for security and can tolerate up to 28 hours of downtime per year.

JONAS TICHENOR: So what's so different about the other extreme?

ROBB BOYD: Well, tier four is at the opposite extreme; there's a lot more requirements with that one. With that one you've got redundant systems for power and cooling, multiple distribution paths that have to be active and fault-tolerant. Access should be controlled with biometric readers and single-person entryways. Of course gaseous fire suppression. Cabling infrastructure should have a redundant backbone, and a facility can permit only about 0.4 hours of downtime per year.

JONAS TICHENOR: So Robb, how do those tiers apply to different businesses across the spectrum?

ROBB BOYD: Well, tier one or tier two is usually sufficient for enterprise data centers where the users are within a corporate environment but these lines could be blurring. Financial data centers are typically tier three or four because they are critical to economic stability and thus need to meet higher standards set by federal regulators. Public data centers that provide disaster recovery and backup services are also built to higher standards.

JONAS TICHENOR: Right, so heavily regulated businesses already know this, so what is the relevance here?

ROBB BOYD: Well, regardless of industry or regulations, customers are constantly looking to strike that balance and make decisions about how much do I invest in the reliability of my infrastructure? It's that classic economic concept of opportunity loss or, as Jimmy Ray would say, how do I get the most bang for my buck.

JONAS TICHENOR: Yes, Jimmy Ray also says let's quite talking and get the meat and taters. That is better. We also want to welcome our guest, Mark Berly, Product Manager at the Data Center Business Unit and of course Jimmy Ray Purser. Jimmy Ray, we're still talking about the big switch here, and what kind of innovation are you seeing in this bad boy?

JIMMY RAY PURSER: Well, things that really impressed me with the switch, just looking at the outside of it, it was really cool. Some of the thought that went in just to the physical design of this product, very impressive. But you know, I've been in many cars that look good but drive horrible and done have any kick to them and stuff.

JONAS TICHENOR: Like a lot of them.

JIMMY RAY PURSER: So what I really want to see is the guts. Let's talk about the operating system, how this switch really works. And so when I've done some reading on it and research, I always cue on marketing terms and start laughing. And so one was I saw it was a multi-threaded operating system, I'm like okay, right. Here we go again, another multi-threaded operating system that's modular. And I'm yet to see anybody demonstrate that, like run a debug all command and actually start doing regular operational switch because we know what that does to a switch/router, it chokes it off.

MARK BERLY: Well, NX-OS is the next step in network operating systems and it's based on the industry-proven SAN-OS. And it's really focused on feature richness, not any feature but features that are relevant to the data center. One big emphasis for us is continuous systems operation and as you were talking about before, data centers need to be up 24 by 7. This is byproduct of globalization, markets move around the world and they access the same data centers. It's also a byproduct of the proliferation of server virtualization. We have servers from different business units running services from different business units that have different service level agreements. So you can generally never get an agreement on when you can have a downtime. And so basically the rule of thumb now is no downtime, so you have to have an operating system that can support it. So we looked at our operating system portfolio and we looked at MDS and the SAN environment basically has that. We can't drop packets, we can't have downtime.

JIMMY RAY PURSER: Yes, SAN-OS is good, that's good actually, that's good going.

MARK BERLY: So we looked at that, we took that and we added to it. And we added Ethernet extensions and enhancements, and that's what brought us to NX-OS.

JIMMY RAY PURSER: So this is not IOS. This is not just an IOS port and we're kind of adding data center features to it. You're talking about the birth of a new operating system.

MARK BERLY: 100% -- this is the birth of a new operating system. We have new L3 code, new L2 code and like I said, it's tailored, it doesn't have everything. So this isn't the Swiss army knife to everybody, it's very focused on the needs and features of the data center.

ROBB BOYD: When you guys talk about multi-threaded modular, is there something you can show maybe to help explain that? What do you mean by that?

MARK BERLY: I'll tell you what. Instead let's maybe take a little walk through the CLI.

ROBB BOYD: Okay.

MARK BERLY: And we'll get to the multi-threaded portion of it.

ROBB BOYD: Okay, perfect, that's fine, didn't mean to jump the gun.

MARK BERLY: All right, yes, no problem. So it's got an intelligent IOS-like CLI, number one so we're not looking at having to retrain your IT staff.

JIMMY RAY PURSER: That was my concern, I'm like, oh man, this means I've got to learn another interface and stuff.

MARK BERLY: Let's look at a few commands, so we've got a command like showrun, and as you would expect, the showrun is going to go and show us our running configuration, amazingly enough.

JIMMY RAY PURSER: With all those interfaces, can we come back tomorrow and see the rest of this?

MARK BERLY: The interesting thing here is I can actually do Linux-like or Unix-like pipe commands and affect the output so let look at it here. I can do a showrun and I'm telling it to...

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

MARK BERLY: Yes, begin with Ethernet 2.3 and I want you to show me the first 15 lines and I can hit enter and it will do its little computing and whatnot and it will come back. So we've got a bunch of stuff, I'm not going to bore you with thousands. Actually I don't even know how many lines of config but the neat thing with some of these Linux-like commands is I can say hey, what's the word count. Look at the word count, yes. And it will tell me, hey look, showrun and it's going to say others -- 5.397 lines of config here.

JIMMY RAY PURSER: But that's for...

MARK BERLY: That's just the non-default.

JIMMY RAY PURSER: For non-default, okay right.

MARK BERLY: So let's look at our showrun all, which is basically all the configuration options on a particular interface or on a particular piece of the CLI. And over 40,000 and so now you wonder why we don't show all of them. This is the reason because yes, I don't think you have enough paper. And being a green-aware company, we want to make sure we don't waste paper.

JIMMY RAY PURSER: Oh, get that plug in there, man. Yes, there you go.

ROBB BOYD: What else do we need to see here?

MARK BERLY: Sure, well we've got virtual devices so we had talked about that. So we have these virtual devices. Consider them sort of a network-based hypervisor so we can take a single physical device, break it up into a number of logical devices very similar to what you're seeing with the virtual machines. One this box I've created three of them so I'm just going to look here. Show virtual device, I've got three of them, one for database and one for web. You can assign ports to these and see the membership. So if you look down, actually Jimmy Ray, if you look down at the bottom part of the screen you've got VDC db and VDC name web. So those are just two different virtual devices and they both have interfaces.

JONAS TICHENOR: So VDC is virtual device context?

MARK BERLY: Context, correct.

JIMMY RAY PURSER: I mean, this is the next step, right? We've talked about virtualizing servers and virtualizing devices, now we're actually virtualizing networking equipment.

MARK BERLY: Correct, and we can switch to them so let's just say switch to virtual device database, so we've got a bunch of database servers here.

JIMMY RAY PURSER: That is awesome.

MARK BERLY: Now these are independent and isolated, they have all their own processes. So if you run spanning tree in one virtual

device, it's not the same spanning tree process that's running in the other virtual device. So if a problem affects one, it doesn't impact the other. So now you can have different service level arrangements or agreements on the same switch. You can truly virtualize your environment. And as we're developing this, we want to make this reach all the way into the server. So we can develop a single, physical topology, arbitrary number of logical topologies built on top of it.

JONAS TICHENOR: Jimmy Ray, how do you think this fits into your concerns about marketing jargon and this sort of thing, just looking past them?

JIMMY RAY PURSER: Yes, you can probably negate that.

ROBB BOYD: What about the inevitable code updates. These things, you're going to have to get them updated. You have to go through a reboot and then once the system is down, or I assume you've addressed that.

MARK BERLY: We have and as I was saying earlier, it's based on SAN-OS and in SAN-OS we have for years provided the ability to do in-service software upgrades and this isn't just us talking. This is something that has been proven in the field and many of the folks out here watching this have done successful ISSU upgrades on the MDS. We took that and we enhanced it. Oh wow. We've looked at the problems we've had and made sure we iron those things out, added some enhancements to it. And what this really allows you to do is perform any type of maintenance. We covered it before. Folks who have dealt with the data center know that our data centers tend to become locked in time. We turn them on and once they go live, once we get services running, you don't change them, they just...

JIMMY RAY PURSER: They freeze.

MARK BERLY: Yes, and so if you want to even just fix a bug, you want to introduce some new cost-reduce hardware or some really cool feature, you can't do it. Either you have to turn on a new piece of the data center or you have to wait for a new data center to come out. So these datacenters will -- 4, 5, 6 years and the will just sit there and then a new data center will open up and that's when you get to add all this new innovation. So we're able to do that with this anytime maintenance concept and the heart of that is ISSU. The really neat thing with the ISSU and the entire system is we have this concept of stateful process restart. So our processes can crash and restart in milliseconds, not seconds and not minutes.

ROBB BOYD: Really?

MARK BERLY: There's no re-convergence, no control plane sending and receiving messages and all this. The adjacent boxes and the services running over that network have no idea that something happened to a process on this box.

JIMMY RAY PURSER: What do you mean, processes? Define that a little bit because I think of processes...

MARK BERLY: Spanning tree, and actually we're going to have a demo a little bit later. Spanning tree could crash and it could restart before the adjacent boxes knew about it. There would be no re-convergence whatsoever, nothing at all.

ROBB BOYD: As far as they're concerned, nothing has changed.

JIMMY RAY PURSER: That is unbelievable.

ROBB BOYD: I think that is something we're going to talk about that, all right.

JIMMY RAY PURSER: I can't even... Hey, you know what though, man, let's be serious. I don't mean to be a wank but you're kind of being an Ethernet attorney here.

MARK BERLY: I am?

JIMMY RAY PURSER: You still haven't answered my debug question so can we really do debug, also do all this then --

ROBB BOYD: Hey, why are you avoiding the question, Mark?

JIMMY RAY PURSER: Well actually I'm not, and it's something I want to point out and I want to talk a little bit, run through some of the CLI. As you can see just looking at the screen, we've been running debug all the entire time we've been talking. You're kidding. I've been issuing the commands, we've been doing stuff, we've been having a conversation. We've got term-on on you haven't seen one message, the prompt's been responsive. And these aren't just showrun commands we're doing; I'm actually doing grep filters on them. I'm filtering for certain commands, looking for certain output and it's still running debug all (inaudible) multi-threading.

JIMMY RAY PURSER: So it's a truly...

ROBB BOYD: Do you recommend we leave debug all on all the time now, is that what you're saying?

MARK BERLY: Generally no, it's got a lot of output. Okay, okay. I just wanted to, simplest...

JIMMY RAY PURSER: That's pretty cool.

ROBB BOYD: That's pretty clear.

JIMMY RAY PURSER: You know what, that's actually pretty cool because that does take obviously a lot of resources. All of us have had routers or switches run away from us; that's pretty clever, that's (inaudible).

JONAS TICHENOR: So Mark, one of the things we hear a lot about is data center operations. What components of the Nexus 7000 simplify what we call remote operations?

MARK BERLY: Well, one of the real keys or one of the things that causes the most downtime in data centers or really in any network are user errors. And so the less time you can have people physically in the data center, physically touching things, is a good thing. We also see data centers being built in more and more remote areas so you may have less operational staff to go to those. I never thought about that, yes. Yes, that's very, very common, especially with our large service providers.

JIMMY RAY PURSER: Colos and stuff like that, absolutely. Oh yes.

MARK BERLY: So you want to be able to touch these devices without physically having to be present. And we have, it's basically a lights-out management processor we call the CMP, connectivity management processor. It's its own CPU with its own memory running its own image. It connects to a separate power bus so it always have power as long as the system has power. You can reset the supervisor, you can reload it, you can do whatever. You still have connectivity of the CMP and the CMP has complete awareness of everything that is going on in the supervisor as well as the line cards. You can use that for troubleshooting, for image recovery. There is a whole multitude of things you can use it for. It really can take the place of the old terminal server networks that we have seen people deploy in conjunction with their data centers.

JIMMY RAY PURSER: So is that in-band or out-of-band?

MARK BERLY: This is completely out-of-band.

JIMMY RAY PURSER: That is sweet stuff there.

MARK BERLY: Yes.

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

MARK BERLY: So it's really designed to let you manage your boxes remotely. You don't have to touch them, you don't have to physically get into the data center once you get it set up and running.

JONAS TICHENOR: Good stuff, thank you very much, Mark, for that fascinating stuff. Jimmy Ray, thank you as well. Robb, let me go back to you here. What kind of stands out in your mind?

ROBB BOYD: Well, a couple of points about what we just heard. First of all this idea of a modern, multi-threaded operating system with a zero service disruption architecture. There is no single point of failure during system upgrades, I mean that's huge. This light-out processor management, I've never seen this in networking gear before, I think that's fascinating. And the idea of virtualization, providing fault containment as well as an allocation of resources and stuff like this, there's just a lot of innovation in this box.

JONAS TICHENOR: Absolutely, and in segment four we're actually going to do something a little bit different and that is called tweaking the geek. We're going to tweak the geek.

ROBB BOYD: That's going to leave a mark.

JONAS TICHENOR: We're going to dive in even further, even deeper into the Nexus family. There is nothing more exciting than being witness to innovation that literally redefines the game and sets new standards, and that is what we have here with the Nexus 7000 series. We are joined by Jeff Raymond, Technical Marketing Engineer with the Data Center Business Unit. Robb, Jimmy Ray and Jeff, it's time to consider more of the Nexus 7000 design elements here.

ROBB BOYD: Jeff, what can you tell us about the overall fabric architecture?

JEFF RAYMOND: Absolutely, the fabric architecture is kind of like the plumbing in your house. You've got to get it right the first time because you don't want to have to change it. So we put a lot of work and investment and intelligence into our fabric architecture. Instead of trying to explain it, let me go to the whiteboard because a picture is worth a thousand words and this is a complicated picture.

ROBB BOYD: All right, let's do it.

JEFF RAYMOND: So when we talk about the fabric we have the main purpose is to connect the input module to the output module, so let's first draw those two modules. We'll call those I/O modules.

JIMMY RAY PURSER: So this is ingress and egress?

JEFF RAYMOND: Yes, so we'll even put that up there, we'll say in and out.

ROBB BOYD: Good call.

JEFF RAYMOND: And now as we saw in the earlier part of the show, we have five fabric modules on the system. And so if we draw each of these fabric modules as entities in between here, these five fabric modules actually work in conjunction to provide aggregate throughput to the system. And the connection from a single I/O module, from a single slot, actually provides connection to all five fabric modules. Each of those connections is 46Gb/s. (Inaudible). So if you have all five fabric modules in the system, you work this together and you have your 230Gb/s of bandwidth that that fabric provides day one. So if we finish off this diagram right here and talk about it, there's a load balancing algorithm. We do a per frame load balancing across all available fabric modules so that you can utilize all five if they're there.

JIMMY RAY PURSER: Per frame?

JEFF RAYMOND: Per frame so it's very granular and you get a very good distribution of traffic across those fabric modules.

JIMMY RAY PURSER: Jiminy Christmas, I mean when you consider these are 10Gb ports, per frame load balancing, that's unreal.

JEFF RAYMOND: Yes, and so the other thing is, if a particular fabric module needed to be serviced or removed from the system, we have a very graceful operation for reducing the actual overall capacity. But simply what we do here is, as you see, we just load balance across the rest of the available modules in the system, whether it's four or three, two, down to even one.

ROBB BOYD: So one can go away, no big deal?

JEFF RAYMOND: No big deal. There is a bandwidth hit but you can just replace that and there is no traffic hit.

JIMMY RAY PURSER: What's the time differential between the time I lose one and between this (inaudible).

JEFF RAYMOND: We're talking milliseconds here.

ROBB BOYD: Wow, well that's what I'm wondering is you start going through all those. How much time does something like this take?

JEFF RAYMOND: Yes, that's important in the data center, right? The amount of downtime is measured in those milliseconds these days.

JIMMY RAY PURSER: Yes.

JEFF RAYMOND: The other key point of the fabric that we put a lot of investment in is the intelligence of how we handle multiple traffic types. We talked about in the data center we have Ethernet fabrics, we have storage fabrics and both of them have different requirements. Specifically, storage actually requires zero packet drops, it can't handle packet drops, whereas Ethernet is a little bit more flexible. So we built this fabric to be able to handle both of those types, what we call the unified fabric. And the mechanisms we have at work here are a couple of different things. The first thing is we have on the supervisor something we call the arbiter. This is a mechanism that actually communicates with all the modules in the system so that if a particular ingress module has a frame that it needs to send, it sends a request to the arbiter. The arbiter in the meantime has received information from all the output modules of the available buffer, the available bandwidth. So it actually is a traffic cop that knows what's going on in the fabric at any one point in time. When the arbiter gets this request, he returns the grant so now the I/O module can send it through the fabric uncongested.

ROBB BOYD: So tell me...

JIMMY RAY PURSER: Guys, wait a minute -- I didn't mean to cut you off, Robb, I'm sorry -- but that's why it's a lossless fabric then.

JEFF RAYMOND: That's effectively why it's a lossless fabric and guess what? We're just taking a page of what we already do on the SAN-OS or the MDS product line for carrying Fibre Channel. Now we're bringing it to the Ethernet world and providing higher quality of service and more intelligent congestion management.

JIMMY RAY PURSER: That's pretty cool.

ROBB BOYD: What's the total amount of bandwidth we're getting out of this thing?

JEFF RAYMOND: Well, the total amount of bandwidth per slot is 230Gb per slot. But the overall series, the Nexus series that we've launched, has the future capabilities of going to over 15Tb.

JIMMY RAY PURSER: Jiminy Christmas, man.

ROBB BOYD: Well, I guess.

JEFF RAYMOND: So actually there's one other component I want to talk about here, something we call VOQ or virtual output queuing. This works with the arbiter to actually prioritize traffic so you have your best-effort class, your cattle class, all these different classes of traffic going through the switch. You want to make sure your high-priority traffic gets through first. And so in addition to that arbitration scheme, we use the VOQ to provide prioritization and to prevent any blocking within that switch fabric as well.

ROBB BOYD: So we're running out of time, is there something we can see in action here?

JEFF RAYMOND: Absolutely, it's always going to be better to see in action so why don't we come over to the demonstration here? And what I have is a traffic generator generating an effective 80Gb/s of traffic.

JIMMY RAY PURSER: Can you guys see this? Check this out.

JEFF RAYMOND: An effective. So basically there's eight 10Gb ports all running at line rate and that's what these transmit rate counters are showing right here.

ROBB BOYD: Oh wow.

JIMMY RAY PURSER: Yes, that's pretty cool.

JEFF RAYMOND: Okay, so this is the traffic generator.

JIMMY RAY PURSER: You can see them transmit.

JEFF RAYMOND: Yes, this is the traffic generator; let's jump over to the switch itself. And in this case, looking at the switch, what we're going to see here is that we're generating 80Gb/S. How fast did I tell you the switch fabric was per slot?

JIMMY RAY PURSER: Two hundred and thirty gig.

JEFF RAYMOND: Two hundred and thirty gig so if I'm generating 80 and we have 230 capacity, what is my utilization going to be?

JIMMY RAY PURSER: Round about 30%.

JEFF RAYMOND: Round about 30% so let's see if that's in fact true.

ROBB BOYD: Quick math.

JEFF RAYMOND: As we look here we're looking at the fabric utilization for module 7, which is the module in question, and for all the fabric channels that we're looking at, ingress/egress, we're hitting 30 right on the nose.

ROBB BOYD: Proof's in the pudding.

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

JEFF RAYMOND: So what this is telling you is that we've got a lot of headroom.

JIMMY RAY PURSER: Jiminy Christmas, man.

JEFF RAYMOND: Even driving an 80Gb/s per slot there's a lot of headroom in this box to grow.

JONAS TICHENOR: Well, everybody is extremely excited, I appreciate this.

JIMMY RAY PURSER: But this is also across -- I'm sorry, I've been cutting you off.

JONAS TICHENOR: Jimmy Ray, you've got to calm down, man.

JIMMY RAY PURSER: But this is also across slots, this is not all the ports on one slot, right? This is slot one, slot two, slot three, and this is where all this.

JONAS TICHENOR: You look like a kid in a candy store, dude.

JIMMY RAY PURSER: Yes, I mean when it transfers data across the switch, a lot of times you run these tests, you'll just do it on a single slot, I'm like ooh wow. But this is actually slot-to-slot traffic transfer.

JONAS TICHENOR: Well you've impressed Jimmy Ray, I don't know if you meant to do that.

JEFF RAYMOND: Hey, my work here is done.

JONAS TICHENOR: Exactly, exactly. Thank you, guys, for that and let's now head to segment five. And we're going to hear more of Robb's interview with Doug Gourlay and dig deeper, even deeper, into the innovation that places the Nexus 7000 a way ahead of the pack. This is TechWiseTV, technology you can use from geeks you can trust. Jeff Raymond is still with us. Sorry to implicate you as a geek, Jeff.

JEFF RAYMOND: No problem. I've been called much worse.

JONAS TICHENOR: It's a badge of honor. We are going to go beyond the nuts and bolts here and takes a look at what puts the Nexus 7000 in a class all of its own. Okay Robb, we only really have time for one more demo here.

ROBB BOYD: All right, let's make a good one, right?

JONAS TICHENOR: Absolutely.

ROBB BOYD: So Mark brought up something that I want to come back and I want to see this in action. The bottom line is there is a lot of innovation within this box, answering a lot of challenges that continue to pervade today's data center. One of the most interesting, I think, comes from the area of reliability and availability. Now for a data center, uptime is something you're always seeking to improve upon. Multiple protocols are in to provide this. Layer 2 you've got spanning tree, layer 3 has got your various routing protocols but the central issue with all of these is the time penalty. What is the effect on convergence, what happens when things change?

JONAS TICHENOR: Absolutely, so Jimmy Ray, you and Jeff I think have a bit of what we like to call innovation to share so share away.

JIMMY RAY PURSER: Absolutely, you know there are so many cool things on this box, asking what's my favorite feature on here is like asking what's my favorite fishing hole. There's so many you know, but dude, you've got to show me stateful process restart, that's happening.

JEFF RAYMOND: All right, let's do it, stateful process restart. What we have here is the NX-OS operating system actually can do self-healing. So we actually keep state within the box of any single process that's running. And so if a process fails or it needs to restart, we can restart from that state without even having any of the protocols realize it. So let's go to the whiteboard and I'll show you our quick demo. If we take a look at a typical network architecture, we many times find ourselves building triangles for redundancy. Oh yes. And so if we look at this as the ag layer of a data center and this is the access layer of a data center. Ag, aggregation? Aggregation layer, yes, and here's our Nexus boxes up here in the ag and we have a 4948, one of the top-of-rack switches that Cisco sells here. Obviously we've built a loop and spanning tree is going to break that loop, so this is what's going to be blocking. What we're going to do is actually initiate a failure on the spanning tree process of this box right here, which happens to be the root bridge, ordinarily a pretty catastrophic failure. Yes. Spanning tree re-converges and you're dealing with some sort of outage.

JIMMY RAY PURSER: Oh yes, re-elect another root and...

ROBB BOYD: All your cascading updates have to occur, right?

JEFF RAYMOND: Right, so what we're going to do here is we're going to initiate that here and we're going to monitor it on the 4948, which ordinarily would have to route around this failure, and we'll actually see how it responds. See if it even recognizes this self-heal.

ROBB BOYD: Let's do it.

JEFF RAYMOND: All right.

JIMMY RAY PURSER: I've got to see this.

JEFF RAYMOND: So let's jump over to the command line here. The first thing we're going to do is we're going to look at the 4948. So if we do a show spanning tree on the 4948 we see everything's running. We're running basically standards-based spanning tree, per VLAN spanning tree, the usual stuff. There's nothing too interesting here. But one of the things we're doing is we're going to turn on debug so if we turn on debug and we start seeing the BPDUs come by. If you look at the time here every two seconds, like a good protocol should, we're sending BPDUs. And so the 4948 is receiving these BPDUs from our Nexus root bridge just like it normally should. So now let's jump over to the Nexus and see how it's doing. Again the show spanning tree says this is the root bridge, as you see right here, all ports are forwarding and so forth. The other thing we're going to show is the actual spanning tree process. So if we look here, we look at all the different processes. We're talking about a modular operating system, many different processes. Process 4652 happens to be the spanning tree process.

JIMMY RAY PURSER: That's not too bad.

JEFF RAYMOND: It's a randomly selected number.

JIMMY RAY PURSER: Yes.

JEFF RAYMOND: And it's only restarted once so now what we're going to do is actually initiate a failure of this process and we're going to see what happens. So right now, oh there it went, we just had a failure of the spanning tree process. So you see service forced exit for the STP process. Now what we're going to do is go look at show spanning tree. The state is still there, the ports are still forwarding. If we go look at the process, hey look, it's not...

JIMMY RAY PURSER: It's spawned a new process.

JEFF RAYMOND: It's spawned a new process so now it's 18196 is the process ID and it's had two restarts. So that's the key, nothing actually happened on the Nexus. Let's go over and look at the 4948. The 4948 is still pumping away; every two seconds we're receiving BPDUs.

ROBB BOYD: You didn't see a thing.

JEFF RAYMOND: No stopping.

JIMMY RAY PURSER: Well, you don't see any, what you would see is what's called a TCN, a topology change notification, clearly the topology changed.

JEFF RAYMOND: So it's effectively healed itself, there was no re-convergence, we restarted the actual process, maintained state. This is going to bring your higher availability during any type of outage in your data center.

ROBB BOYD: I don't know if this digs in a little bit too deep but how do you do something like that? Is that secret sauce stuff we can't get into? That's a good question.

JEFF RAYMOND: It is the secret sauce.

ROBB BOYD: Okay.

JEFF RAYMOND: We use something called PSS or persistent storage services to actually maintain that status, kind of like a relational database within the operating system and we just can restart and pull that back up.

ROBB BOYD: You have to do that very quickly, right?

JEFF RAYMOND: Yes.

ROBB BOYD: Because it won't take long to notice it's not changing.

JEFF RAYMOND: It's like in milliseconds, yes.

JIMMY RAY PURSER: Wow, very cool. That is an excellently cool demo. Thank you. That was really something else.

JONAS TICHENOR: Well, thank you guys very much for that. Robb, I think we're going to go back over and take a look at the interview that you had with Doug Gourlay, Senior Director, Data Center Solutions at Cisco. Let's take a look.

ROBB BOYD: Doug, I appreciate you sticking around for a few more questions. I think there's an obvious one though that we need to tackle which is customers, as they look at this innovation and what's being rolled out here, are going to ask about their existing data center investments from Cisco, namely big investments in the Cisco 6500 and the MDS line. What are you telling them?

DOUG GOURLAY: The first thing to understand is how Cisco are organized. When we created the Nexus family of products, what we did was we created new business units leveraging some existing silicon investments. And we reorganized in a way that allowed us to continue our active R&D of the Catalyst 6500, the Cisco MDS and at the same time create the Nexus without making sacrifices in the R&D roadmap of any of those three platforms.

ROBB BOYD: So no changes to the investments of the 6500 or the MDS. This is a complementary addition is what you're saying as well.

DOUG GOURLAY: Absolutely, I'm like passed one through sort of dovetailed one product into another and effected a transition. We don't see this as a transition because the requirements of the data center are very much bifurcated from what you traditionally saw in a campus and LAN branch wide-area network and Metro environment to what we're seeing in the data center. I don't need power over Ethernet. There's a tremendous amount of development in some of these areas that aren't necessary in the data center, yet to bring storage and Ethernet more tightly together there's a tremendous amount of investments I need in the data center that couldn't be leveraged anywhere else.

ROBB BOYD: So it's about filling a hole versus making a replacement.

DOUG GOURLAY: Absolutely.

ROBB BOYD: Okay, that makes sense. Well, to talk a little bit more about the relative positioning here, where does this fit exactly to the other data center products and how do you see that playing out?

DOUG GOURLAY: Certainly the main positioning of the Nexus portfolio is for 10Gb Ethernet-attached servers whereas we'll see the Catalyst lines focusing on the bulk of the market today, which is Gigabit Ethernet-attached servers. Now what's going to drive that 10Gb server? Server virtualization and the unification or consolidation of the I/O. If we take those multiple server interfaces and condense them down to one or two ports coming out of a server, the bandwidth increases. As we bring multiple virtual machines, the bandwidth increases. And we've designed the Nexus family specifically to address the requirements of consolidating the I/O and supporting dense server virtualization environments.

ROBB BOYD: Interesting, so with so much innovation in the Nexus family, what do you think are the main takeaways? What are the key points that you like to leave people with with regards to that?

DOUG GOURLAY: Certainly, when you look at the Nexus family, one thing we tried to do that I think is an absolute first in this industry was understand customer operational best practices. So we look at how do you boot a switch, how do you shut it down and

how do you want to. And what we found is customers have 20-page documents describing how to take a line card out of service. If you think about the continuum of how a product runs, you have the default system behaviors get you here. And then you've got tribal knowledge, best practice, define procedures to fill into where you want it to be. Well, we looked at those best practices and operational procedures and made those the system defaults in as many cases as possible. We added configuration rollback, we added a graceful shutdown procedure, we added a lights-out management port. You can shut down the Nexus and watch it reload and never lose connectivity. A corrupted image you can fix 99.95% faster than you could with an external terminal server so a significant time to restoration improvement. We also looked at the biggest outages we had are human-created because somebody types a command like reload, which has this amazing effect of taking the switch out of service. It went wow, that was painful. We said why don't we make software upgrades hitless, zero service disruption. And design the system architecture from the operating system, the silicon, the software, the hardware, so that everything works seamlessly without packet drops as we do upgrades. And we can manage the system remotely, designing it for the data center environments that our customers are building now so it's purpose-built for it.

ROBB BOYD: It sounds like a good combination of Cisco innovation looking at what customers are saying. These are our biggest pain points, we'd like for you to invent something here to solve these problems. And a little bit of leading the way there as well and saying, you know what? You didn't know you were dealing with this problem as much as you are, but if we solve this would you find it valuable?

DOUG GOURLAY: And to that second point, it's also understanding where the industry and market's going. By leading in the standards bodies and chairing many of these workgroups, we know how 40Gb and 100Gb are getting built. We're on those committees and we're leading them. We understand how Ethernet is evolving to be able to consolidate the I/O. So understanding where the market is going lets us build a platform that not only will last, with an active R&D lifecycle, for a decade and be a viable capital asset for long duration. But it also means that the product is going to be that platform for innovation. So in addition to the usability, it consolidates the I/O, that saves power and improves efficiencies. In addition to that, the platform will be able to have a tremendous amount of flexibility in transport decisions for 10Gb, 40Gb, 100Gb and onward.

ROBB BOYD: So what would you say customers need to be doing today, what's that final point? Say you've heard about this, this sounds interesting, now what do I go do?

DOUG GOURLAY: I'd start testing and qualifying the platform. And I would spend some time on the websites we have and going through some of the videos and some of the tools, understanding it better. And then I would attend a training class. Okay, very good.

ROBB BOYD: Well Doug, I appreciate you taking the time to explain this for us. I congratulate you on what's been a lot of work coming out with this and wish you the best of luck with it.

DOUG GOURLAY: Thank you very much.

ROBB BOYD: Thank you.

JONAS TICHENOR: Great interview by the way, Robb.

ROBB BOYD: Oh thanks.

JONAS TICHENOR: We have covered a tremendous amount of information here today.

ROBB BOYD: We usually don't focus on a single product on TechWiseTV so that's a little bit of a break.

JONAS TICHENOR: Well, this is a big product.

ROBB BOYD: It is a big product, thank you for recognizing that. (Inaudible). It's a big deal and I hope we've provided at least a taste for a number of incredible things that can be accomplished with this data center class switch. I encourage anyone that needs to understand more, certainly talk to your account manager of course but you can also go to cisco.com/go/datacenter. We have these links and more information in the show notes also, you can download this as well.

JONAS TICHENOR: Very good and well defined.

ROBB BOYD: Thank you.

JONAS TICHENOR: Thank you, Robb. Well that is it for this edition. For more information about future episodes as well as a complete archive of our past shows, be sure to visit the Cisco Interactive Network at -- write this down -- cisco.com/go/interact. For Robb Boyd, Jimmy Ray Purser, I'm Jonas Tichenor, thank you for joining us here on TechWiseTV.