Unlocking the Potential of Virtualization

THE REAL IMPEDIMENTS TO VIRTUALIZATION

ROBB BOYD: Explosive growth in under-utilized, costly servers and storage is stretching the limits of cooling, power and floor space while the cost of capital needed to support this growth is increasing. It's true, we need to do more with less and we have answers for you on how to do this from a data center perspective. Welcome, my name is Robb Boyd and this is Jimmy Ray Purser. Today's topic, unlocking the potential of virtualization. So, if we're going to unlock that potential I think we've got to understand some of these problems, right?

JIMMY RAY PURSER: Yes sir, we do.

ROBB BOYD: So, as it takes off virtualization is big, everybody is doing it. At the same time we're starting to feel like, you know what, we're uncovering new problems that we have to deal with.

JIMMY RAY PURSER: Yes.

ROBB BOYD: You open one door, but yet another one closes on you.

JIMMY RAY PURSER: That paradigm shift, right? Any time you add a new technology kind of the ripple effect of that rolls somewhere else on the network.

ROBB BOYD: Here's one of my concerns, you always think here we are, Cisco, we're a networking company and we're claiming to have some answers to unlock virtualization, which is a server technology because we're talking about server virtualization. Is that a little bit cocky also to say we're going to just start virtualizing our things? Does that sound kind of bandwagon approach-ish?

JIMMY RAY PURSER: Well, you know what, here's the thing though, it depends how you look at that. A good network engineer, the answer is almost always going to be it depends, right?

ROBB BOYD: It depends, there...

JIMMY RAY PURSER: It's like you shook the little magic eight ball and you get the it depends and you're like err, and you throw it against the wall.

ROBB BOYD: There's only one answer in that ball, yes.

JIMMY RAY PURSER: That's the only thing that is there. So, you know, we've been doing virtualization for a long time. VLANs are virtual. You know we've got virtual context for firewalls. I mean there's always some kind of virtualization out there. Applying this to the server is just kind of the next step up the OSI layer. We all kind of memorized that for out various certification exams, but it's also the template we use to design and grow and develop our network. So, if you look at how networks are built out, how we did our technologies, we really followed the OSI model pretty strictly. And so, yes, of course we're going to look at this, because any time you add virtualization to a network it is going to break other things. It's going to break my access control list. It's going to break some of the policies that I have to my network. And it's going to really limit how redundant my networks are on the network because I've got everything virtualized in this one space. So there's got to be a better way of doing things. And to solve the networking problem in either virtualized or not, physical or virtual, to solve that networking problem you've got to go with a networking company.

ROBB BOYD: Yes, okay, fair enough so you talked me into it. But we do have a data center expert with us. I'd like to go ahead and welcome Steve Phillips. Steve, actually, you're going to be our regular data center expert on the show.

JIMMY RAY PURSER: Excellent.

ROBB BOYD: Because we are diving further and further into data center than we ever have before, but 15 years of experience with Cisco. Am I getting that right?

STEVE PHILLIPS: Thirteen years.

ROBB BOYD: Thirteen years, I added two years on you already there.

STEVE PHILLIPS: I don't need that.

ROBB BOYD: CCIE since the 1990s, but you fully focus on data center. This is what you do day in and day out. Let me ask you what do you see as the problems of server virtualization and its growth? And you were showing me some numbers just before that talk about how large and how fast this is growing and organization. Because really, what today's focus is about is virtualization is really about consolidating things so that we're not wasting money on inefficient use of resources basically, right? What types of problems are you seeing from your perspective? Does virtualization work in all instances?

STEVE PHILLIPS: So, virtualization is probably the biggest disrupting technology in the last decade.

ROBB BOYD: Right.

STEVE PHILLIPS: It's breaking everything. It's breaking down organizational boundaries. The server guys now get to manage network elements within components of their domain. The network guys can sometimes see traffic and other times they can't see traffic depending upon where two virtual machines live in one particular instance.

ROBB BOYD: Right.

STEVE PHILLIPS: So there's a hell of a lot of things that are becoming broken now just by virtue of the fact that we've gone to this virtualized server model. But to be very, very clear, the business benefit of virtualization is so great that people will live with those disruptions at this present time.

ROBB BOYD: Isn't that funny?

STEVE PHILLIPS: Yes.

JIMMY RAY PURSER: Yes.

ROBB BOYD: So here's the problem as I've seen it, and I'm curious from a networking perspective and a server perspective because we've got two different audiences that we kind of have to speak to here. But as kind of the clean lay as I was trying to work it out is before, you know, the server guys were asked to deploy applications. And correct me if I get this wrong, because you've got more experience in this. But the idea is we have a new application that supports our business, so we dedicate a server to that task. And so for every application we had a server, because what did every manufacturer or every software developer do is go, well, we want dedicated resources. Well, virtualization does actually a very good job of saying you know what, instead of now three servers to cover three applications, you can have three applications running on the one server. But then what naturally happened is VMware came in and said, well, wait a minute we've got to solve this problem of the fact how do these machines communicate with each other? Well, that's always been a traditional switching function and the switch has always been located as that official demarcation point. So we kind of took something that belonged to the network side and moved it inside the machine now and lost visibility to it.

JIMMY RAY PURSER: Well, you know and that's the thing, right? I mean that's one of the things that I think has really worked pretty good in the server side, data server side, because we see some problems that are happening. You've got virtual machine sprawl. I mean you have server creep. After we got off the mainframe everybody was going client/server, then you had server creep everywhere and you had racks and racks and racks and heating problems and everything else. And then you started putting virtual machines on. You had virtual machine sprawl and on the storage side you start using Fiber Channel. You have Fiber Channel IDs exploding everywhere and you can't keep track of them. Well, on the Ethernet side we figured this out. This was VLAN IDs. We figured out how to run multiple VLANs, multiple tags on a single link. So why not take that basic concept and kind of port it over to the data center side?

ROBB BOYD: But if Cisco comes in here and we go, hey, we've now virtualized a switch. I mean if we cut to some of the specifics that we're building into, one of the things I see and you and I were talking about this early on, was this notion of, well, VMware put a switch in there to solve some of these problems why should I put another switch in? Are we just trying to say, well, we do switches so you need to put our switches in, Cisco equals switches. But there are others to it and not to get into the specifics there, but staying kind of on what keeps things broken. You had mentioned that the virtualization has broken all the traditional models and blinded us to a lot of things. What kinds of things has it done from that next level of virtualization is around mobility.

STEVE PHILLIPS: Mobility is absolutely huge. The ability to take a virtual machine, literally create it anywhere within a VMware

cluster where you have resources and then to dynamically adapt or to dynamically adjust the amount of resources that particular VM can get by just moving it around within the infrastructure is huge. It gives you better availability characteristics for example, better performance characteristics. But you do pay a penalty because as you move that virtual thing around how do you maintain policies? How do you maintain counters for troubleshooting? How do you provide visibility into the movement of that particular element? And that's where the challenges lie.

ROBB BOYD: How is that being dealt with today? Because I mean with virtualization growing like it is maybe we haven't got the mobility side. But I mean what's been the reaction? People are still virtualizing, so are they literally just flying blind in most operations these days and kind of getting by or has there been some tension? How would you describe it?

STEVE PHILLIPS: There's a lot of tension going on, but, as I said before, the business benefit is so huge, 42% of customers today have production workloads running on virtual machines. So they are doing something, but I think they're living with some of the problems that they have today and they're looking for solutions. How do I get that consistent visibility, the consistent policy definition, as well as separation of roles that I had before in my physical environment so I can map that to my logical environment?

ROBB BOYD: Right.

JIMMY RAY PURSER: And that's the thing, we're looking at trying to get the best of both worlds. There's always a tradeoff in any type of technology, right? I mean there's always something that has to be given up to get, hopefully, a larger benefit. And virtualization always has been kind of a black hole and every hacker on the planet is trying to figure out a way to hack the hypervisor. It's a gold mine I mean, man, if we can get in there and exploit the switch in there, then man we really can get... So I mean, you know, it's tough enough to do a man-in-the-middle attack on a server. But now I can actually get in the middle of a virtual switch, going in that hypervisor itself and start redirecting that traffic, man, I can steal a bunch of data off the network very transparently. It's going to be hard to detect. It's going to be really hard to monitor because I'm doing it right in that shim level, if you will. And so that's a very dangerous tradeoff that as a networking guy myself, a security guy, I would tell my folks I know it saves a lot of money, but you know what, I just can't sign off on it from a security standpoint, the risk is just too great.

ROBB BOYD: Yes. So Cisco worked with VMware to start solving some of these problems, right?

STEVE PHILLIPS: Yes.

ROBB BOYD: This isn't something that we just started doing on our own. Can you describe kind of the results of that? What I'm leading to is help me understand kind of, there's a framework established that began dissecting some of these components.

STEVE PHILLIPS: Yes, so we began working with VMware around about 2 1/2 years ago taking some of the switching concepts and constructs that we've developed within our physical environment.

ROBB BOYD: Right.

STEVE PHILLIPS: Taking those features and putting them within the hypervisor itself. So we worked very, very heavily with VMware to come up with a framework that allowed us to insert an XO ISA based switch within their hypervisor infrastructure. Now, there's a lot of value in that. We're obviously going to add back a lot of features that were missing previously like access control lists, private VLANs, NetFlow accounting. But the real key was observing the administrative boundaries. So essentially what was happening in the past was the server guy was getting to manage a lot of network type functions. He had to understand about VLANs and NIC teaming and EtherChannel and negotiating with the server guy about what type of VLAN he needed to see. And it all becomes a very, very slow and complex and they talk different languages.

ROBB BOYD: But is it not just a matter of server guys learning how to understand networking and therefore can pick this up? I mean why wouldn't that work?

STEVE PHILLIPS: They can do it, there's no real impediment to that. It's just something else that he's got to do on a day-to-day basis and that can just add to the overall workload that he has to go through.

ROBB BOYD: Yes.

STEVE PHILLIPS: Instead of devolving some of that responsibility away to the network guy and saying okay you handle that for me. But provide me with some clues and hooks that allow me to do my job much more efficiently using information that you

can provide to me.

ROBB BOYD: Right.

JIMMY RAY PURSER: Well, see that's the problem right now, is that we can talk about how the technologies tie in this and that but the real issue with virtualization, it kind of is more of a Layer 8 problem.

ROBB BOYD: Yes.

JIMMY RAY PURSER: I mean we're talking about how these server folks are looking at this, how the networking folks look at it and it's kind of your job/my job. And there's some fear in there that the further we go down this path the less likely I'll be employed here because I will be technology-ing myself right out of a job. Both sides are feeling this and so there is an ingrained interest to kind of fight this technology, if you will, for job security.

ROBB BOYD: Yes.

JIMMY RAY PURSER: Now, I know it's silly because one thing we know about technology is it's always going to move forward whether you fight it or not. You know I don't care how much you liked token ring, it's just not around. And that's one of the things we're seeing here. But I think some of the concepts we have seen have really kind of split into some good middleware areas that kind of still give both people the same responsibilities but they kind of modify them a little bit I think. And that's the biggest issue with virtualization.

ROBB BOYD: Now, well, these are good points. Anyway, well, so Cisco worked with VMware to help create the framework where Cisco can work natively directly within the hypervisor. This is where we're going next, the promise of reestablishing a clean dividing line between server and network responsibilities. It is starting to take shape. The concepts do make sense, but how does it really work?

UNDERSTANDING VN-LINK AND THE NEXUS 1000V

ROBB BOYD: Hello, this is TechWise TV. I'm Robb Boyd along with Jimmy Ray Purser and our resident data center expert Steve Phillips. Server virtualization has created the ability to run multiple instances of a server on the same physical box and that is a good thing. But the communication between these instances, well, it's a networking function that server folks have to contend with. I mean it's switching, but it's on the servers. So to address this Cisco has been working directly with VMware to create a framework called VN-Link to help us restore these operational models that we need because we know they work. We need to understand how all of this works in terms of our daily job, so Jimmy Ray why don't you get us started man with VN-Link?

JIMMY RAY PURSER: Well, then we need to go to the good old whiteboard here.

ROBB BOYD: Bring out the magic tablets.

JIMMY RAY PURSER: Let's take a look here at the magic tablet. Now, when we're talking about the Nexus 1000V, let me grab the marker here, we're really talking kind of a virtual Ethernet module and that's how we're looking at this. You can see that this is actually embedded inside the VMware ESX server itself. We've got this nice little software switch that's tied here and all of it is managed by the VSM. And this actually scales out to -- when you actually design this out, this actually scales out to a 64-slot switch.

AUDIENCE COMMENT: 64 nodes yes. And that limitation actually is because that's what VMware does, right? So this is actually how this all ties together. You can see how my virtual machines are plugging into here. It really looks no different than kind of almost like a micro network is what it looks like, inside of this whole architecture, if you will.

STEVE PHILLIPS: And the significance of this little switch here is the fact that the VSM can run in two modes. You can run it as a virtual machine in its own right in one of the clusters, in one of the cluster nodes. Or you can run it as a separate appliance depending on the number of nodes and the type of performance characteristics that you've got.

JIMMY RAY PURSER: I think that one of the things, and I want to make sure that we get into this kind of quick, because this is kind of cool, we look at the software features of this and functionality but what really makes this powerful is something we addressed earlier and that's the concept of a port profile. And one of the demos that we'll do is actually show how one of these sets up, how we can move this around and the real power of actually having this VM to manage all these instances here. So let's

take a real quick look.

ROBB BOYD: And so the idea behind the port profiles that we're getting into here is this is the way we actually interface with these boxes, right?

JIMMY RAY PURSER: Yes, that's correct.

ROBB BOYD: So now that that...

JIMMY RAY PURSER: Well, it's how we define a policy to that virtual machine itself so that it can travel with that machine. So anytime we move that virtual machine anywhere else in that compute architecture the policy follows it. And the real important thing is that the policy not only follows it based upon any QoS, any ACL stuff I have set up, but what also follows it also are my stats and counters. When I move that stuff from one to the next I don't reset all those. All that accounting inventory or that accounting information actually comes across with it as well. It's all preserved.

ROBB BOYD: Because the piece that's been missing, and this is what we're going to address with the port profile, it sounds like it is before we would establish that information based on physical connection so when the server guys said, hey networking guys, I need a connection because we've got this new application that needs to go out and they would assign the profile of information to that physical port. Now we have to re-create that in this virtual world so that it can stay with it and also adopt those characteristics as it moves around.

JIMMY RAY PURSER: Yes and you know not being a snoot here, but this is how this kind of works. This is a networking guy and a server guy kind of working in an environment, is how these port profiles work. So what does a server guy like? Well, a server guy likes an interface like this, right? This is via vSwitch to your client. This VSM is operating and so this is actually showing all my virtual interfaces or my port profiles that I've got set up here and you can look. It's color-coded really nice. My virtual ones are blue, my physical ones are green, and I would kind of tell what's going on and who is assigned where. I mean it really makes it pretty easy to sign this up. Now I'm a networking guy I see this stuff and I'm like, man, no, here's what I like right here. I go to the command line and, oh, that's what I'm talking about and so now what I'm doing is, as the networking guy, then we're working together. We've got a server coming in. We've got a virtual interface coming in. We've got something that comes through change control and change control will say we need the following things to happen.

ROBB BOYD: It's got to be a process.

JIMMY RAY PURSER: Absolutely, man, you've got to document this stuff out. So change control gets in, networking guy says, okay, well, I need to set up a port profile for another server. It's like, okay, well, let me set this up because this is what the networking guys would do. If we didn't have this then a network guy would go on to a switch port and do some of these same commands. So first he would go in and start setting up his port profile. Well, this would be different, he wouldn't set this up on a switch, but this is how we actually identify a common name that our server administrator is going to see out there, right? So the port profile is actually any common name we want to set. So, on this one we'll call it TechWise TV. That's pretty creative, huh?

ROBB BOYD: Going forward.

JIMMY RAY PURSER: TechWise TV is the port profile that everything you can see on change context, right here, TechWise TV in config mode and then I drop down into this context, no big surprises here for networking people. This is just like configured interface, config, interface config, I drop down, I change my context. Then I go in here and I actually start adding some of the stuff that we know, if I can spell it right, switch port, mode, access, I set that up. Now I need to tell it what VLAN it goes into. So we'll up arrow action here, switch port, access, VLAN, looks like a VLAN 12. I sign that in and now I've now told it what mode it is, what VLAN it is, and, again, it's my network guy doing this stuff. And my next thing I want to do it is the most important command, which is something that I messed up on a lot when I first started doing this demo, is I'd forget to type in state enable. State enable is actually what sets up the state to turn this on and now it looks like I really haven't done anything. But here's the trick, my next command is my VMware port hyphen group command. Once I assign that command you'll look right down in here and you'll see it actually propagate out and this is what gives us the power. Now, this is the software switch. These are all IOS commands. I'm all setting this up. This is just like setting up a Nexus. So I'm configuring all my information into here and the second I put this down I'm going to communicate back through this device and you'll see it push down so watch this. So VMware I'll do a port hyphen group, I'll do tab completion, I'll enter that. You can see that add distributed port group, if I push it down, there's my name, TechWise TV. Boom, as a networking guy that's it.

ROBB BOYD: Well, and you said it, as a networking guy I got really excited. What does this mean to a server person, right?

STEVE PHILLIPS: So all of those commands that Jimmy Ray has just entered in there, the server guy doesn't really need to understand that stuff.

ROBB BOYD: Yes.

STEVE PHILLIPS: All he needs to do is attach his particular virtual servers that belong to TechWise TV.

ROBB BOYD: Right.

STEVE PHILLIPS: Drop them into that port profile. He doesn't need to understand about the VLAN dependencies, access control lists, all of those things have now been taking care of for him. So he can now simply create the virtual machine, pick the right port profile, and add it to the virtual machine. Now, when we create the virtual machine within the ESX hypervisor cluster, we detect that virtual machine has come in. We apply the TechWise TV policy and it's dropped into the right VLAN, apply the access control lists, turn on NetFlow accounting if we need to, turn on EL Span and we're up and running.

JIMMY RAY PURSER: Well, watch this, let me show you how it works. So here, if I click into host and clusters then what I'm doing is I've got my list of machines here. These are all my virtual machines. This is my physical, these are my virtuals here, and I'm going to assign this to the XP client. So we'll say, okay, edit virtual machine settings. Now, this is what my server guy would do. He'd come in and say, networking guys, they probably didn't set up my profile there's supposed to be a TechWise one here. Let me see if it's here. Oh, right there it is look at there. Push it out. So he clicks on that.

ROBB BOYD: So that's server view.

JIMMY RAY PURSER: Yes, server dude is doing this and the networking guy, man, he's off playing cards with the guys chilling. So the server guy assigns this out. He clicks on this. You can see, reconfigure virtual machine. I went ahead and set that up and now this policy is on here. Now, here I am, I'm a networking guy, I may not trust that server guy, the server guy didn't even do that. So I can do a show IMP virtual and see if my XP stuff is assigned. Yes, look at there, it's assigned a port, a virtual interface Ethernet 2, which is right here. The XP client is signed up and it's assigned to this physical server. Now I can say, well, let's just see if everything took. I can do a show port command is show virtual Ethernet 2.

ROBB BOYD: Well, what was happening Steve when it came to on the server side of things though, could this not be done within VMware's vSwitch? I mean where was all this taking place in the vSwitch environment? Was it just the server guy had to learn how to do it or was asking for help from the networking side?

STEVE PHILLIPS: So you have to look at one thing, he either learned how to do it or he simplified his life and just dropped everything into a single VLAN. Obviously, that's going to break quite a lot of policies and some of the security rules. And we have heard of some customers that have run afoul of compliance rules and security audits.

ROBB BOYD: Oh so it's not just a preference that the company would have to worry about, it's a fact that you could actually be in breach of some government issues.

STEVE PHILLIPS: There are some legal ramifications here in certain instances. So this is why people are very, very careful and they want the tools now where they can actually see what's happening within the virtualized environment and why Nexus 1000V is one of the critical components of that.

JIMMY RAY PURSER: Well now let me show you this, this is what kind of really makes this pretty happening is that you can see my policy is assigned here. I'll look at my show virtual Ethernet. But let's say -- because now you've seen that this is really just a little cool, I mean I assigned that policy, I configured something yes, it's a little cool. But what makes this happening is that this is my physical cluster right here. Right here now I want to drag it down to this server, I want to reassign that, I just grab this, I'm the server guy and I'm like, oh man, I've got some problems here. I want to drag that instance across, VMotion is kicking in which is what you want to happen. I do a next, say yes, that's what I want to do, here's my migrate all my resources. Let's go ahead and it's going to migrate this entire machine including the port profile. Everything is going to be reassigned over here and all of a sudden the stats, all my interfaces, everything that I configured on this machine is now plopped into this server 100% complete. It's over there. Nobody realized this even happened and I can go up and I can do up arrow here, show Ethernet virtual on that same machine look, now see my client has moved. It was one Server 1, now it's on Server 2, shows how many modifications I did and there we go, the port profile has traveled with that device. I don't have to reconfigure it. Once a

network guy does this once, then the server guy can kind of take it from there. The networking guy can go in and mark their stats (Inaudible) application. No matter how many times this networking guy can switch this a thousand times, the stats will be the same. My accounting is still the same, so my networking people still keep everything right in check like they're supposed to. We've got some good harmony here happening, which we never have been able to have before.

ROBB BOYD: You've created a container for them to latch into, so the server person just goes that's all I need. If anything, this actually starts to sound like now virtualization becomes even easier for the server person than it was before. Not just because we got rid of something that was kind of a headache, but I mean now you've even eliminated the phone call in the old days or however that transpired. You know, where I need my physical network link. You're like, hey, you know where your profiles are. I've created the different ones for the types of applications. So unless you've suddenly broken into something completely new then our communication is done, it is nice and clean.

STEVE PHILLIPS: Using this abstraction, what we've actually done is increase the efficiency of not just the technology and the hardware, we've also increased the efficiency of the people and the processes. So now, instead of the boring stuff whereby you're typing things at a command line all the time, just do it once, set it and forget and then keep on leveraging it multiple times. The server guy doesn't need to call everybody. The network guy doesn't have to keep typing things, very much a simplified environment.

ROBB BOYD: I'm glad you brought that up because I didn't think about that, but efficiency obviously being a core theme throughout the show.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: That's very good information. So hold on to those thoughts because we need to pick up with more of that next as we get deeper into the Nexus 1000V and take these concepts even further.

NETWORK INTERFACE VIRTUALIZATION

ROBB BOYD: Welcome to TechWise TV my name is Robb Boyd. This is our chief geek Jimmy Ray Purser and our resident data center expert Steve Phillips. Today we are unlocking the potential of virtualization and I've got a big question on my mind as we take this even further down here. Here's my question Steve, is I don't think everyone is totally on board with running all of that powerful switching, but it's great that we added that stuff in. But is it necessary that we always do all of that in the hypervisor or don't we have some more options potentially to play with?

STEVE PHILLIPS: There are some more options there. One of the key things to remember though is in most modern servers they've got two sockets, four core processes. There's more than enough processing power to run a software-based switches such as the Nexus 1000V.

JIMMY RAY PURSER: Yes, but every time you do that though Steve -- I'm sorry, I didn't mean to cut you off. But every time you do that though you keep adding overhead, right?

STEVE PHILLIPS: Absolutely.

JIMMY RAY PURSER: I mean you're saying there's enough power, but the great thing about virtualization is that I can keep doing a whole bunch of them. So the more I add the more overhead penalty I'm putting on here. You know, I love the Nexus 4000. I'm telling you, man, I was pumped about it, sitting there during the demo, but in the back of my mind I've got to start planning this stuff out in the data center because I've got to factor an overhead budget in here as well. But there's a different way of doing it too.

STEVE PHILLIPS: Yes, which is VN Tag, which we introduced with the Unified computing system.

JIMMY RAY PURSER: Yes.

STEVE PHILLIPS: VN Tag effectively bypasses the hypervisor, bypasses the software switch. So essentially what we do is we attach the virtual interface NIC of a virtual machine to an I/O virtualization capable NIC card and then we push the traffic straight out onto the wire to the downstream switches to handle.

ROBB BOYD: This is where the tagging comes in obviously.

STEVE PHILLIPS: This is where the tagging comes.

ROBB BOYD: Somehow it's got to travel with it.

STEVE PHILLIPS: Absolutely.

JIMMY RAY PURSER: So you have to have a specialized CNA to do this then?

STEVE PHILLIPS: Yes, you do it's called the Palo adapter.

JIMMY RAY PURSER: Oh, right, okay Palo right, of course.

STEVE PHILLIPS: So essentially if we look at the whiteboard here we can see a virtual machine here with this virtual NIC and, as you can see here, there is no network switch.

ROBB BOYD: Right.

STEVE PHILLIPS: It's directly attached to a virtual Ethernet interface on the NIC card here, the Palo adapter. What we do at this point here is we insert our tag, VN tag, that's very similar in concept to what we're doing with port profiles. What it does is it adds a small piece of information to say this is a TechWise TV VMware -- VM rather.

ROBB BOYD: Assuming we'd identified the virtual machine, obviously in that last example it was tagged with TechWise TV being the name of the server.

JIMMY RAY PURSER: Yes, where are you adding that information at?

STEVE PHILLIPS: We're adding that at the NIC card here, at the Palo adapter.

ROBB BOYD: So it is sitting -- but is it in the header of the (Inaudible)?

JIMMY RAY PURSER: Yes, I'm not following you because how does it know? I mean that tag, is it a tag like an 802.1d, a Q tag?

STEVE PHILLIPS: It's a small piece of information that we add to the tag. We have taken this to the standards bodies with VMware so we are standardizing this as a function within the IEEE standards bodies with VMware. What it enables us to do is insert that little bit of information into the stream such that when it hits the switch we can inspect the tag. So we look at the VN tag value here. See, that's TechWise TV and then we know which VLAN to drop that in, which access controllers to add to that, whether we need NetFlow accounting or not. Some other things that we've had to change just a little bit within IEEE rules of course is the ability to loop packets back out the same port we received them on.

JIMMY RAY PURSER: The same port that received, right, right.

STEVE PHILLIPS: Based upon forwarding rules.

ROBB BOYD: But that's a different port the way you drew it though isn't it?

STEVE PHILLIPS: It's the same physical wire.

ROBB BOYD: Oh, I see what you're saying, okay.

STEVE PHILLIPS: So this is the same physical interface.

JIMMY RAY PURSER: What's the benefit of this then?

STEVE PHILLIPS: Essentially what I can do is bypass the hypervisor entirely.

JIMMY RAY PURSER: Yes, I've got that, okay, I'm cool with that.

STEVE PHILLIPS: So you get around about 10% to 12% CPU efficiency back. Again, we can give that back to transactional processing on the server. So we bypass the hypervisor and leverage a hardware path to provide switching and policy enforcement using exactly the same constructs that we've discussed for the Nexus 1000V.

JIMMY RAY PURSER: Okay.

ROBB BOYD: But does the switch have to be something special on that? I mean does any of our switches work with that because you're talking about a physical external switch at this point, right?

STEVE PHILLIPS: So, the switch is here that supports this concept on the Nexus 5000 series.

ROBB BOYD: Of course, okay.

STEVE PHILLIPS: As well as the UCS 6100 series fabric interconnect. So both of those...

ROBB BOYD: Specifically data center switches that we've recently come out with.

STEVE PHILLIPS: Absolutely.

JIMMY RAY PURSER: Absolutely.

ROBB BOYD: Now, why wouldn't that work for you Jimmy Ray? I mean doesn't break us out of the hypervisor and give you some comfort level there?

JIMMY RAY PURSER: Yes, it does, you know, I'll tell you, you know what, I get really concerned when I start putting tags on my packets and stuff. I'm worried about my consistency checking, I'm working out having a bunch of CRC errors, plus you're having stuff here. I mean any time I start messing with packets themselves I get worried. Okay, I get some back in performance, but if I've got to mark every one of these packets coming out then I'm still marking every one of these packets. And I'm assuming I'm rerunning my share of CBAC on this again so it seems like I'm picking up some overhead somewhere.

STEVE PHILLIPS: You're picking up overhead, yes, but everything is implemented in hardware.

JIMMY RAY PURSER: Yes, okay.

STEVE PHILLIPS: So it's fully transparent to any of the overlying protocols. We just add a small piece of information on that particular link, when it hits the other end, inspect that particular value, apply my policy rules, strip the tag and forward it according to the policy rules that I've defined. A very, very simple model.

JIMMY RAY PURSER: What's the preferred method then? I mean so if I'm looking at this, I'm trying to figure out what my dividing line is here. I've got the Nexus 1000 series, which is simple. It is relatively easy, but I have overhead and I have cost, because I've got to buy the software, I've got to have something to host it on and all that good stuff. I mean you got all that baggage that comes with it, but boy it's nice, the interface looks great and it's pretty slick and cool. This is implemented in the switches, but my concern is that it almost still cuts the server guys out.

STEVE PHILLIPS: It's no real difference to what we just explained with a Nexus 1000V if you take a step back.

JIMMY RAY PURSER: Okay.

STEVE PHILLIPS: With a Nexus 1000V the server guy creates a virtual machine. He applies a profile to it and pushes that virtual machine down onto a server.

JIMMY RAY PURSER: Okay.

STEVE PHILLIPS: He will get exactly the same view on his console using VN Tag. Because when we create the profile with a VN Tag definition we provide the abstracted name, TechWise TV in our example to V Center. So when he does create that he selects the same profiles. There's no change to him. It could be running Nexus 1000V, it can be running VN Tag, it really doesn't matter.

JIMMY RAY PURSER: Where is my -- based on your expertise here, and you're looking at doing one or the other, what's kind of the qualifiers?

STEVE PHILLIPS: So one of the qualifiers is going to be do you have an I/O virtualization capable network interface card? If you haven't got VN Tag capability and I/O virtualization capability Nexus 1000V is the best way to go. That means you can use Nexus 1000V on pretty much any server that's available on the market today.

JIMMY RAY PURSER: I guess that's true, isn't it? I mean it's like a universal adapter, that's true.

STEVE PHILLIPS: If you are going with the latest server models, such as a Unified Compute System that have VN Tag a capable hardware and you do want to leverage hypervisor bypass because you want that transactional performance to be as efficient as possible, VN Tag is probably the right way to go for you in that case.

JIMMY RAY PURSER: Because you'll definitely pick a ton of performance here. I can definitely see, as I'm sitting here trying to work this out in my head I'm like now I see why they use this on Unified Compute, because you really could pick up quite a bit of steam on here. And especially if I'm offloading this on a Palo adapter, then you really could actually have and probably, and I don't know if it's better to say this or not, but I mean hardware is always faster than software that's just life, I mean that's just how it is. But also is there more scalability here too? Because it seems like to me that on the hardware side I'd have a little bit better scalability or is it about equal?

STEVE PHILLIPS: You'd probably get greater flexibility within software, but you're going to suffer a little bit of a performance hit.

ROBB BOYD: There's always going to be something.

STEVE PHILLIPS: With the hardware you're going to get raw performance and even more efficiencies than you can gain today.

JIMMY RAY PURSER: Yes, see, that's what I like.

STEVE PHILLIPS: Again, we're pushing the boundaries of the efficiency to reduce power and cooling costs, to really drive as much expense out of the infrastructure as we can.

JIMMY RAY PURSER: So then for the competitors that are watching this show and they're saying, well, they're just Cisco putting some more gold handcuffs on us with this whole VN Tag thing. This has been submitted as an 802.3 standard along with VMware, right?

STEVE PHILLIPS: Absolutely.

JIMMY RAY PURSER: So this could be used by many vendors.

STEVE PHILLIPS: Yes.

JIMMY RAY PURSER: In the next (Inaudible) time, because we are actually trying to standardize this.

STEVE PHILLIPS: Yes, this is yet another example where Cisco innovate something and then gives that back to the industry. So we're always giving intellectual property back, promote standards, interoperable devices.

ROBB BOYD: Well, we have options for server, network folks for saving money.

JIMMY RAY PURSER: Yes, that's cool, I like that that's a cool technology.

ROBB BOYD: So to make sure we got this straight, following our structure here, we've been talking about the framework around VN Link. We talked to the Nexus 1000V, first native hypervisor, 100% within the hypervisor switch running within the hypervisor. We have the VN Tag to help us kind of break out some options for taking that same port profile and mapping it in and out of the physical infrastructure without losing that information that allows it to become mobile but we need to take it a little bit further next. And next is where we're going to talk about storage. That's right, that's exactly where we're going to go next, is what does virtualization mean for your server environment? We cover it for you in segment four.

VIRTUALIZATION AND YOUR STORAGE ENVIRONMENT

Unlocking the Potential of Virtualization

ROBB BOYD: Cisco has introduced a number of innovations that unlock the real potential of virtualization in the data center. The question now is where can virtualization positively affect another huge area in the data center? Storage. It's all been about LAN so far, so where do things come into play in storage? Did we break just as many things when we started virtualizing on the server side as we did on storage? Because it does feel like the network obviously is becoming a much more critical role when things are virtualized now perhaps than it may have been playing in the old world, so to speak. Can you help us understand what changes and what can we do about it?

STEVE PHILLIPS: So yes, in the SAN there are a lot of changes, as you just said. As we virtualize there's more and more pressure to put the data on networked storage. That's what enables you to move the virtual machines around without having to copy gob-loads of data, for want of a better term.

JIMMY RAY PURSER: Right.

STEVE PHILLIPS: But again, there's challenge is there. What tends to happen is within a single server you'll have a whole group of virtual machines and a host bus adapter connected to my SAN fabric and then that's connected to my storage arrays and life is good.

ROBB BOYD: Looks straightforward.

STEVE PHILLIPS: My problem is, again, I've moved from that physical model to the virtual model.

ROBB BOYD: Right.

STEVE PHILLIPS: Now, what I've done is I've abstracted all of my virtual machines behind one single identity. So in the past I would have an HBA for each physical server here.

ROBB BOYD: Right.

STEVE PHILLIPS: And now I've moved to a model where I've got one HBA that masks everybody. It's abstracted the identities away.

ROBB BOYD: So it's that blindness issue all over again?

STEVE PHILLIPS: So we're blind to whatever's happening from a virtual machine perspective. Now, the good thing is within VMWare's implementation they present a local SCSI disk to each virtual machine. The virtual machine never sees the SAN.

JIMMY RAY PURSER: Right.

STEVE PHILLIPS: So, it is secure. The problem is I can't apply policies on a per VM basis.

ROBB BOYD: Because it's the same across because everything happens to be on that same physical box, okay.

STEVE PHILLIPS: Now, one particular way you can get around this is to use RDM or raw device mappings. What raw device mode does, it basically allows the hypervisor here to present a virtual HBA per VM. And now I can get per virtual machine policies back, so now the SAN administrator can now see every single virtual machine and apply the correct policies so that he can prove that he's compliant with any regulatory compliance rules. Or he can apply the same security rules even as that virtual machine moves around the infrastructure.

JIMMY RAY PURSER: Okay.

ROBB BOYD: Does that work the exact same way that it's been working, the way we have applied tags and stuff before? Are there some other things we need to understand there?

STEVE PHILLIPS: So, within the SAN, they use a completely different technology which is NPIV within the SAN itself. When a host attaches to the SAN it performs a login, so it actually announces the fact that it needs to attach to the SAN. The SAN then gives it an address back from the fabric and then that's given to a name server and then we can apply the particular policies against that login. So a fundamentally different model to what the LAN is, whereby you simply attach unannounced and start

communicating. Within Fiber Channel it's slightly different. You have to login, announce yourself, and then policies are applied. What we use with raw device mode here is the concept of NPIV or network port virtualization, interface virtualization.

ROBB BOYD: Interface virtualization okay.

STEVE PHILLIPS: What then happens here is as the virtual machine is created we simply login using the virtual IP, the virtual HBA address here. We log into the fabric, that responds to say, yes, we accept you onto the fabric. Here's the policy rules that we need to adhere to and then communication to its particular home storage array is allowed. In this way we can segment traffic according to policies that you had in the past in the physical world.

ROBB BOYD: Does that work from a mobility perspective as well?

STEVE PHILLIPS: Yes.

ROBB BOYD: Because the idea there is not have you freed it up to move, except, well --. So I think if it moves around then you've got the issue of consistency among machines having that their physical connectivity is consistent across the board. STEVE PHILLIPS Yes, so I've hooked up my switch here and I move this virtual machine, the virtual machine will move as well as its virtual HBA address.

JIMMY RAY PURSER: Right.

STEVE PHILLIPS: Its identity moves with it. As it moves, it performs another login and we simply respond. All of the policies are enforced within...

JIMMY RAY PURSER: So that flogi is the triggering event then, is that the deal?

STEVE PHILLIPS: That's the triggering event.

JIMMY RAY PURSER: Okay, oh, well, that's pretty cool actually.

ROBB BOYD: You called it flogging?

JIMMY RAY PURSER: Flogi, f-l-o-g-i.

ROBB BOYD: Oh, flogi.

JIMMY RAY PURSER: A fiber login.

ROBB BOYD: Okay, I thought somebody had been flogged.

JIMMY RAY PURSER: After the show.

ROBB BOYD: Got it.

STEVE PHILLIPS: Even more acronyms you need to learn.

ROBB BOYD: Yes, well who would we be without our acronyms?

JIMMY RAY PURSER: Absolutely, absolutely.

ROBB BOYD: So, is that all there is to a Jimmy, I mean do you trust that, does that make sense?

JIMMY RAY PURSER: Well, here's the thing because in a typical network today we're talking about you having all these Fiber Channel IDs that are actually sprawling all over the place. And there's only so much of a limit we can have on these and so NPIV actually kind of helps gather these together, right?

STEVE PHILLIPS: Absolutely.

JIMMY RAY PURSER: And then on top of that in kind of a hierarchical format, NPV is what manages the NPIV, is that correct or no?

STEVE PHILLIPS: It's a related topic if you like. What we see is as we scale SAN fabrics out, if I've got my two central directors, MDS 9500 class machines, as I add edge switches, especially using top-of-rack designs or blade designs, I've got a limited number of domain IDs.

JIMMY RAY PURSER: Right.

STEVE PHILLIPS: Each edge switch uses a domain ID, one per switch. These are a finite set of resources.

JIMMY RAY PURSER: Right.

STEVE PHILLIPS: Although there's 239 within the standard, most vendors only implement a very small subset, so most of our competitors are around about 40 to 50. We support something like 80.

JIMMY RAY PURSER: Okay.

STEVE PHILLIPS: To help better scale these types of environments, NPV enables me to essentially turn these edge switches here into HBA aggregators if you like.

JIMMY RAY PURSER: Sure enough, okay.

STEVE PHILLIPS: We turn off most of the switching functionality. HBAs don't need to talk to HBAs in most instances. They're server ports, so they just need to talk to their storage arrays. So we simply pass through or we actually arbitrate between the login here, but we don't need domain IDs at the edge, which means I can scale this out far farther. It means I can get more servers attached to my Fiber Channel domain and I get better efficiencies out of the entire fabric yet again.

JIMMY RAY PURSER: That's a great way to look at it. These really are just HB aggregators at the edge, aren't they?

STEVE PHILLIPS: Yes.

JIMMY RAY PURSER: Son of a gun okay, alright, that makes sense.

STEVE PHILLIPS: A very simple concept, but a very powerful one when it comes to scaling the environment to meet the demand for more network storage.

JIMMY RAY PURSER: Oh, absolutely or even if I was going to try and start doing some DCI, I think it would also play in there too, right? For my data center interconnect, if I'm starting to sort of bond all these things together, being able to manage this at this level is pretty darn important.

STEVE PHILLIPS: Yes, for DCI we'd probably look at some other techniques like Inter-VSAN routing as well to provide good separation of fault domains between the two data centers.

ROBB BOYD: Well, good guys I think we covered storage, covered a bunch of different ways to extend virtualization into it. I think it's funny because I really didn't think storage was going to fit into that concept, but it's a lot of the same concepts. But we can't ignore the fact that once you understand those concepts, take them out to the storage side because that becomes you've now broken the code in that key that unlocks the blindness that we're experiencing through virtualization so that not only can it be used to restore back the sanity that we've been looking for in the past, some of the stuff that we're used to but also kind of helping us achieve some of the stuff that we were always promised in virtualization around mobility and how to do that still in a secure manner and things like that. Does that sound...?

STEVE PHILLIPS: Yes, it all comes back to that, efficiency, responsiveness, as well as resilience.

ROBB BOYD: Yes, for Jimmy Ray Purser, Steve Phillips, and the rest of our crew here, my name is Robb Boyd. Thank you so much for joining us. For more details and links from today show or to find out about future episodes and check out a complete archive of previous shows you can find it all on our blog at TechWise TV.com.

Unlocking the Potential of Virtualization