

Technology Tutorials

LAN DESIGN IN THE DATA CENTER TRANSCRIPT



Program and Presenter Opening

Sidney Morgan:

Hello and welcome to this “Cisco on Cisco” Tutorial on Cisco’s data centers and network environment. I’m Sidney Morgan, and I’m an IT Manager in the IT Cisco on Cisco team.

The theme of our show today is an overview of Cisco’s data center environment and its evolution toward a virtualized, services-oriented model.

It’s my pleasure to introduce the guest today for the show: Jon Woolwine, an IT Architect in the NDCS – Architecture team. Thanks Jon for being here.

Jon Woolwine:

Thank you, Sidney. It’s a real pleasure to be here and to be able to present our data center strategy.

Data Center Networking: Cisco IT Data Center Background

Jon Woolwine:

Okay, let’s get started with the data center networking discussion. Just a quick outline for today’s discussion. We’ll begin to talk about a little bit of background, take a very high level look at the data center environment. We’ll follow that up with more of an in detailed view of the network architecture and design, and we’ll end the discussion of where we see the data center network going and how it will evolve into the future.

Cisco Data Centers

So first to give a little bit of background. We're gonna talk a little bit about the geographic distribution of data centers. We're gonna get a sense for the size of the data centers from a server perspective and we're also gonna talk about some of the fundamental goals that are driving our thinking around data center network design and architecture.

So a very high level view of what our data center environment looks like geographically. In this diagram here you'll see several data centers. We have about 50 total globally and of those 50, we have different types of data centers. We have what we call production data centers and we have what we call development data centers. And the differences between the two are that the production data centers are primarily for hosting business applications and workforce productivity services. These kinds of applications demand high availability and we also have fairly rigorous governance and policy around changes in those data center environments.

Now the development data centers are a little bit different. They support the product development, the product engineering side of Cisco. And here in these data centers, we're looking for a little bit more flexibility without such rigid governance so we physically separate those data centers and host those engineering services out of those data centers.

From a network perspective, traditionally we've designed these networks differently. However, more recently we've begun to converge on a single design for both production data centers as well as development data centers. Altogether, all this data center space adds up to about 180,000 square feet, which is a pretty significant amount of space that Cisco IT manages.

One other key point here I think is worth calling out is that we're gonna have a significant amount of new data center square footage opening up here in the next six to eight months. We have a facility opening up in Richardson, Texas. You may have heard of the PDC program. This is a new production data center that will be build out there, a Tier III facility, very excellent facility for us.

We also have additional capacity coming in in the form of new data center space in WebEx's Mountain View facility, and that will also add to the space coming in in the near future. So together, we see about an additional 50,000 square feet of data center space coming into the picture for us in the near future.

Server Landscape

And now to get a little bit different perspective on the scale of that data center environment. We have about 12,000 physical servers spread across all these data centers today. Now of these physical servers, about 20% of them are virtualized, meaning that we take a physical server and we virtualize many servers within that physical server. And what this is allowing us to do is really optimize resources within the data center. In this case, we're optimizing the server resources. We're also saving quite a bit of money in terms of power savings, as well.

So running on these servers, we have almost 13,000 applications as well as 325 production databases. Operating systems, on the right there in the pie chart you'll see a distribution of operating systems. Primarily we're Linux and Windows based. We do have a fairly large installed base of Unix with Solaris making up the majority of it and HP-UX having a small slice of that pie.

Cisco Data Center Goals

So now to take a look at some of the data center goals. The real focus here for Cisco IT is around optimizing the total cost of ownership, maximizing business agility, and focusing on business continuity. Now the optimization of total cost of ownership we've been achieving through several means. One is through the consolidation of existing data centers, so taking those many data centers that you saw before in the previous diagram and consolidating them into a fewer number of data centers. In addition to consolidating those data centers, we're also undergoing consolidation within the data centers, so taking existing equipment and consolidating it to end up with a relatively smaller amount of equipment to manage.

Now lifecycle management is another area where we've been able to optimize total cost, and what I mean by lifecycle management is hardware lifecycle management or, as some people call it, fleet management. It's a refresh process that allows us to get in and continually refresh equipment within the data center. This includes networking equipment, where we're on what's typically about a four-year lifecycle for routers and switches within the data center.

We're continually going around data center to data center over these four years and refreshing. And what that tends to do is it tends to flatten out that capital expenditure over time and it also tends to introduce new capabilities slowly rather than having a sudden need to introduce new capabilities into the data center and having a big expenditure come along with that.

Standardization is another way where we've been able to optimize total cost of ownership, and with standardization, what we've really been able to do there is narrow down the physical infrastructure in terms of the types of physical infrastructure we have and the standards behind the physical infrastructure. We've been able to define standards around hardware platforms, software platforms, as well as configurations. And what that does is it allows us to minimize the complexity we have in our data center environment. With that minimized complexity, we're seeing some significant gains in terms of operating expenditures in the data center.

And the last are where we've been able to make a difference on total cost of ownership is around virtualization. And what that really is to us is about optimizing or fully utilizing the resources we have within the data center and we've found that virtualization has given us a big payback in terms of maximizing that utilization.

Now the next goal I'd like to cover is business agility, and really what we're talking about there is the rapid delivery of services into the data center. So for example, how long does it take for an infrastructure team to provision a new server for an application to run on? A key component of rapid service delivery here is automation and the ability to take what used to be maybe a multi-week process and bring that down to a couple of days or perhaps even a few hours. And ultimately this is all about being an enabler for business schools, for growth, and being able to take advantage of new opportunities very quickly.

Cisco Data Center Goals (Contd...)

And the third big goal for us within the data centers is around business continuance. Security is one way we've been focusing on getting business continuance and from a security perspective, the real focus has been on looking at security across all layers of the data center environment, not just at the data center perimeter and at the application layer, but looking at a more holistic view of security and how we can implement it throughout the various – the layers of the data center.

Availability and recoverability are also big key areas for us.

Recoverability meaning if there's a disaster in a certain location and we have a need to bring up critical services in another location. Virtual servers, virtual operating systems, virtualizing applications, that's bringing us the ability to really make it a lot easier to provide recoverability. So these virtual environments allow us to very easily move services and applications from one data center to another, which can be really, really beneficial in a situation where you have to recover from a disaster.

Data Center Networking: Current Data Center Network Architecture

Okay, now we're gonna make a little transition and move away from that high level background around the data center and focus more on a detailed view of data center network architecture, as well as design. And to begin that discussion, what I thought I'd do is just bring up a few key points around the data center network and give you an idea at a rough high level how we build things and how we design things, and then we'll start really looking into more detail how these individual elements play out.

Today's Data Center Network in Cisco IT

So to begin with, the data center network is what we call a standard Layer 2 based ten Gigabit Ethernet core distribution access architecture, and it's all based on a Cisco Catalyst 6500 platform. Some of you may be familiar with the Cisco Validated Design program. It's an internal program that Cisco has established to provide customers with data center designs that Cisco has actually validated, and tested, and approved, and put them out there for our customers to leverage. Cisco IT's data center environment very closely aligns to the Cisco validated designs in that we do follow that Layer 2 based approach in many of the elements that you'll see in upcoming slides very closely align with where the Validated Design program is going.

So as I said before, the 6500 platform is really a core platform for us within the data center. In addition to the 6500, more recently we've started looking at the Catalyst 4948 ten-gig e-access switches, and we'll talk about some of the drivers for those. These are switches we're beginning to use in the access layer of the network and we have some specific reasons that drove us to look for the adoption of this particular product.

One other point here is that most servers within the data center are connected at gigabit Ethernet speeds and they connect via copper gigabit Ethernet ports. We standardized on server platforms that leverage copper. We do have a handful of servers in other storage platforms out there that use fiber-based interfaces, but the vast majority are copper today.

Now within the last two or three years we've put a big focus on introducing what we call IP network services into our data center environment. And what we're doing there is we're embedding capabilities like local server load balancing, SSL offload, firewall and other security capabilities into the network so that these capabilities can be leveraged by the applications in the data center.

Global load balancing is one of these services that we've introduced, well actually, quite a while back but we've grown more and more reliant on global load balancing over time as we stressed availability and recoverability in our architecture. Traditionally we've used Cisco's Distributed Director product for that. More recently, we're making a transition to Cisco's Global Site Selector product.

And the last point here is around server management and network management. We build what we call an out-of-band network to allow us

to manage servers, to manage network devices. And this out-of-band network leverages the Cisco 2800 and 3800 series routers for async connectivity, that's async serial connectivity, as well as the Catalyst 3750 and 6500 switches for Ethernet management, and this is typically through an integrated Lights Out Management interface on a server.

Cisco IT Data Center Network Layout

Okay, here's an interesting perspective that I thought would be good to look at before we start taking too detailed of a view into the IP network. What you have here in this diagram in the green boxes on the left are servers. What you have on the right side of the diagram are the three different network fabrics that we have within a data center. Now on the right-hand side in the red oval you'll see a management network. I just spoke to that on a previous slide. That's the management network where we provide out-of-band access to servers, as well as network devices for management purposes.

Now in addition to that, we have a storage area network within the data center. Servers can establish connectivity to that storage area network through fiber channel switches. And in addition to the storage area network, we have the traditional data or IP network based on Ethernet within the data center. Servers establish redundant connectivity to Ethernet switches. As I said before, that's 6500s, the 4948s in the data center to get access to that IP network.

Now in addition, we also have what we call heartbeat switches in the data center and these are really purposed around providing what we call cluster private connectivity. The application clustering environments within data center require a private network connection to allow cluster communications to happen across and we provision dedicated switches to provide these cluster heartbeat connectivity types of functions out there in the data center.

Sidney Morgan:

So Jon, are they dedicated switches or are they dedicated VLANs?

Jon Woolwine:

So in this case, the heartbeat switches are actually dedicated switches. And we've taken that approach historically just to maximize stability and to not create any kind of dependencies on our existing Ethernet switches that typically provide data connectivity. The thinking there is if we have some kind of a soft failure on the Ethernet side, we can at least maintain that heartbeat network and avoid some potential problems around data corruption should they happen.

Going forward we are looking at opportunities to possibly consolidate that heartbeat connections onto the Ethernet switches. We'll just have to carefully evaluate the pros, and cons, and risks of doing so.

Sidney Morgan:

Thanks.

Cisco IT Data Center Network Layout

Jon Woolwine:

Okay, so here's that detailed look at the IP network and our network architecture and design. And as I mentioned before, we take a layered approach in our architecture. The various layers are core, which you see at the top, distribution or what some folks call aggregation, there in the middle. We have a services layer which is hanging out on each side of the distribution layer and at the bottom we have what we call an access layer. So some of the key differences between each layer, I'll start at the top.

The core layer, the primary function for the core, is all about high speed switching of packets. We're not doing a lot of fancy policy kinds of things in the data center core. It's basically packets in, packets out. It's 100% routed.

Now you go down to the distribution layer and its primary purpose is to provide high-density ten-gigabit aggregation for the access switches. What's important for that distribution layer is that it have a high density of ten-Gigabit Ethernet ports. Obviously there's gonna be a lot of access switches that get aggregated on those distribution switches.

Now the service switches which you see hanging off the side there, their primary functionality is to host those IP network services that I spoke about earlier. So they're hosting the SSL offload capabilities, content switching, server load balancing, and security services, as well.

Sidney Morgan:

Jon, what are the advantages of building a separate Layer 2 network services switch?

Jon Woolwine:

That's a great question and I actually am gonna address that here in a couple of slides. I can point out the pros and cons there. Don't let me forget to do that.

So going down to the access layer, this is obviously where server connectivity is established. At the bottom you'll see individual servers. Most of the servers in our environment establish redundant connectivity at the access layer, meaning they have Ethernet connections that go to two different switches. Those two different connections are redundant through NIC teaming so on the servers the network interfaces are teamed where one interface is active, another is standby. Should the active interface go down, a standby interface can come up? That allows us to withstand failures of access switches, as well as network interfaces within servers.

Which brings up another point. We have redundancy from the top core layer all the way down to the access layer in this network and that's a key, key characteristic here and a key requirement for us on delivering the availability that the business expects from our data center environment.

Now as I said before, and I'll say it again, the primary product here that's playing in this space is the Catalyst 6500 and we have seen recently an introduction of the Catalyst 4948 switches at the access layer here. Today we're primarily still Catalyst 6500 but we're beginning to introduce those Catalyst 4948 switches and I'll actually tell you what the situations are that kind of drove us to that kind of thinking and have led us to introduce the Catalyst 4948 switch. If we could just hold that for a couple of slides I'll hit on it eventually.

I'm gonna back up and take us back to the distribution layer because I think that's an important place to really take a view on the data center. The distribution layer serves as a logical boundary for us. The distribution layer is where we draw the line between the Layer 2 server access network and the Layer 3 high speed switching network which goes up to the core. So connectivity about the distribution layer to the core is 100% routed. It's all Layer 3. We use equal cost multipath routing to provide multiple paths from each of the distribution switches up into the core.

Now from the distribution switch down through the access switches and to the service switches, it's a complete Layer 2 environment. And what that means there is that the distribution switches are the default gateways for the servers. They are the first Layer 3 hop for the servers.

Okay, another important point here that I don't want to miss is the fact that the network all the way from the access layer up through the core is based on ten Gigabit Ethernet. So as server connectivity has moved from 10 to 100 to 1,000 megabits per second, we've been upgrading this data center network at the various aggregation points to be able to

support those higher speeds that servers are demanding.

And one other point here that I'd like to make is around the Layer 2 nature of this network. Before I stressed that this is a Layer 2 network from the distribution switch down and being a Layer 2 network, it gives us a lot of flexibility around where we can physically locate servers. Because we can put a server in any rack connected to any switch and extend a VLAN or a subnet to that server, as opposed to being forced to stick a server in a single rack associated with a single switch to make sure it gets in a specific VLAN.

So having that Layer 2 design gives us a lot of flexibility around where we can physically locate servers and that leads to one other point here, too. The ability to actually extend any VLAN and its associated subnet to any access switch means that we are trunking all VLANs from the distribution layer down to the access switches.

The last point I'm gonna make before we leave this slide is around server virtualization. So as we said before, Cisco IT has been moving rapidly in the direction of server virtualization and one of the techniques we've used to support server virtualization, which is not shown here but which I'd like to call out, is by extending that trunking using 802.1Q from the access switch actually down to the server. And what that lets us do is it allows us to provision different virtual machines, different virtual servers, within the same physical server but have each virtual server be on a different VLAN or subnet. So good flexibility there.

Sidney Morgan: Any helpful hints, since you're trunking down to the server, about spanning tree?

Jon Woolwine: So what we found at this point is in testing we haven't seen any problems with spanning tree. VMWare, which we use, does not internally bridge traffic so we haven't seen spanning tree concerns there. The primary concern or the real focus of effort in introducing server or virtualization and extending those VLANs down to the access layer is coordination amongst teams.

So we've made a decision to have a separate team responsible for managing virtual servers take care of the virtual servers and provide all configuration for the virtual servers separate from the network operations team that does the configuration of the switches. What that means is the folks managing the access switches in the network need to be communicating with the operations team that manages the virtual machines to say, "Hey, this server needs to be on this VLAN and be associated with this subnet."

So organizationally, some little hops to get over but nothing too significant that's gotten in our way. Just a little bit of coordination is required there.

Sidney Morgan: Okay, thanks.

Cisco IT Data Center Network Layout: Scaling the Data Center

Jon Woolwine: So this next slide gives you an idea about how we scale this data center network and what you see here is the single pair of core switches that we saw before in the previous diagram, but underneath this pair of core switches we have two different sets of distribution, service, and access layer devices. So what we've done is we've essentially created a modular approach from the distribution layer down. We replicate that distribution layer and its associated services layer and access layer multiple times throughout a data center to give us sort of a modular architecture that can scale with the capacity of the data center. And oftentimes we refer to these modules as pods within Cisco IT. So if you

ever hear that term pod this is what people are referring to.

Cisco IT Data Center Network Layout: Service Switches

Okay. Our next focus is on the services switches and a key point around these that I'd like to reinforce is Cisco IT made a conscious decision to separate out services in their own dedicated switches apart from the distribution switches. Many folks think these services logically fit within the distribution layer and while that is a good potential place to put these services, you also have to keep in mind that the distribution layer had requirements around ten-gig density. We have a lot of uplinks that come from the access layers into these distribution switches.

Cisco IT made the decision to maximize the number of ten Gigabit Ethernet ports available for this uplinking, and by doing that what we decided was that we would take these services and essentially locate them separate from the distribution devices in their own dedicated chassis which we call services switches. So what's running in these chassis? Well, its things like content switching, server load balancing, SSL offload, as well as security services. These are all implemented through modules so we're talking here about the CSM module, the SSLM module, as well as the ACE module. And I'll talk a little bit more about the evolution of those modules in upcoming slides.

You can tell that with all these services there's a significant amount of complexity that we're introducing here into the network and to help manage that complexity we've introduced Cisco's Network Analysis Module into these service switches. And what the Network Analysis Module does for us is it gives us visibility into what's happening on a packet-by-packet basis within that environment to allow easy troubleshooting and the ability to work through problems fairly quickly.

Sidney Morgan:

Does this architecture improve the resiliency of the overall data center architecture?

Jon Woolwine:

Yes. Another key point here is these services that are provided here on the service switches, these are redundant services. We have a pair of service switches. You'll see one on the left, one on the right hanging off there to the side. One's in an active state. The other's in a standby state.

We have state information replicating from the active switch to the standby switch such that should there be a failure of the active switch or the links that connect it to the distribution layer; we can automatically field traffic to that standby switch and make that the active switch.

That's an automated process that happens for us. The convergence within the network after such a failure will happen via spanning tree to reroute traffic over to that other server switch.

Content Switching: Bridged Mode

Okay, let's move on. So how do these services get integrated into the network? Well, there are three common ways this gets done. One of the ways is by integrating these services at Layer 2, which is commonly called bridge mode or sometimes transparent mode. The other options are Layer 3 integration where we basically make these modules part of the Layer 3 infrastructure and have them terminate a Layer 3 domain, or we do what we call a one arm mode implementation where we use network address translation and make modifications to Layer 3 headers and packets to make things work.

What Cisco IT has done is made a decision to go with the bridge mode and, Sidney, I think you asked before a question about the pros and cons of this integration and really what the bridge mode gives us, or

what I should say what the bridge mode gives application teams, is a lot of transparency. There are no Layer 3 modifications to the data center network at all. There are no Layer 3 modifications on the server side. There's no Layer 3 modifications within the network. So from a Layer 3 perspective, our routing domain can remain unchanged and very stable. We like it that way.

The server folks also like it because we maintain the source IP address of incoming packets which they commonly use for various purposes. So from an application perspective, the bridge mode integration is a much more transparent way of integrating these services for them.

Now from the network's perspective it's a little different. There is some complexity and some challenge here and those are shown here on the cons listed. Primarily there's some risk here to bridging traffic through services modules. You have the risk to spanning tree loops being created, accidentally of course. You also have some complex troubleshooting in here that may have to take place. And really what we've had to do is come up to speed operationally to make sure we have the skills necessary and the processes and policies in place to avoid these kinds of problems and to maximize our ability to troubleshoot.

Sidney Morgan: Have we developed any processes to improve the troubleshooting ability?

Jon Woolwine: Yeah, so one of the processes is going back to the network analysis module, the NAM blade. We've standardized on NAM blades everywhere on all these server switches because we've found that we just rely on them heavily and we've done a good job of bringing the network operations teams up to speed on how to use these network analysis modules to troubleshoot.

Sidney Morgan: Great.

IT Adoption Status-Services

Jon Woolwine: So a little bit more information behind these services. Going back in time, looking several years back, in fact, to the days when we had the Cisco Local Director product providing server load balancing, we had about 100 virtual servers that were distributed across 50 Local Director pairs. Now contrast that today with where we are. We have over 1,500 virtual servers on 25 CSM pairs, so what does that mean to us? That's huge convergence for us. We've gone from managing a relatively large number of devices providing a relatively small amount of services to a small number of devices providing services over a huge domain, and it's that kind of consolidation that we've really found good TCO and benefits from over the long term.

And to give you an idea about some of the applications that leverage these services for various purposes, whether they be content switching, server load balancing, SSL offload, we have Cisco.com. We have same time instant messaging. We have our Oracle 11i applications, Livelink, our J2EE environment. These services serve as a front end for Microsoft Exchange and they also serve as a front end for MeetingPlace. So over time, we've found value across a lot of the different application and services environments out there in our data center.

Sidney Morgan: So we have exciting new network virtualization modules, the application control engine, ACE. Can you elaborate on our ACE roadmap within IT?

IT Adoption Status-Hardware

Jon Woolwine:

I'd be glad to. So before I jump to this next slide, a quick idea on what that roadmap looks like. I said before going back, going way back in time, back to the days of Local Director, we had roughly 50 Local Directors providing service for 100 virtual ITs. Shortly after the introduction of Local Director the content services switch came along. We saw some spotty implementation of that here and there but in general, implementation of both products was sort of on an as needed basis. We'd see an application that had a need for this, we deployed these products.

A couple of years down the road going back two, three years, maybe four years from now, the content switching module came out. That largely has changed the data center environment for us. With the content switching model we changed our mentality from sort of providing these services application by application when needed and providing sort of a box per application to providing services horizontally across the data center using CSM. And that's really where the services switch architecture came into place.

Now the next evolution beyond CSM and beyond the SSL module is the ACE module, the application control engine. And Cisco IT is in the process of validating and certifying the ACE for production in our data center environment. I would anticipate that we're probably looking at implementation in the four to six month time range once that certification process completes. And one of the big benefits for us is additional consolidation.

So we had CSM modules and SSLM modules before. Those services provided by those two separate modules are gonna be converged into ACE modules. So we're gonna go from say I think we're in the neighborhood of 20 SSLMs and 50 CSMs. We're gonna take that down to 30 ACE modules. So another good consolidation opportunity there and another good opportunity for lowering that total cost of ownership.

Sidney Morgan:

Great.

High-Density Attachment to the DC Network

Jon Woolwine:

Okay, the last detailed view I'd like to have of the data center network architecture is around the access layer. Traditionally within the access layer we've had servers in a switch which connect to access switches, traditionally 6500s, at the end of a row. So picture if you will a bunch of racks with servers in them and cabling extending from each rack down to the end of a row where we have redundant Catalyst 6500 switches.

Now over time things have changed in the server world. What we've seen is higher densities of servers, in other words, more servers per rack. At the same time we've seen more network connections per server. That's created some challenges for us. What has it done?

No. 1, it's created some cabling cost implications behind it. More cabling to run from every single rack to the end of the row. One of the problems that comes from that is that cabling all goes underneath the floor tiles in the space where the air flows. So as a result of bringing all that cabling under the floor, we're obstructing air flow and not cooling our high-density data center.

So because of these challenges we've done some thinking around our access switch design and we've begun to introduce what we call a top of rack access switch design. And that top of rack access switch design is shown on the top right-hand side of this diagram. Essentially we have servers within a rack that all connect up to redundant Catalyst 4948 switches within the rack. Now the only connectivity that leaves the rack are the uplinks from the Catalyst 4948s to the distribution gateways. So that's helping us manage some of these physical challenges in the data

center that we're faced with.

Okay, so now we're gonna make a transition and move away from talking about the current state network architecture and design and we're gonna take a more forward looking view on where the data center network is headed and what that evolution looks like. So to start the discussion I'm gonna introduce what we call four of Cisco IT's key intelligent network themes

Cisco IT's Intelligent Information Network Key Cisco IT Intelligent Network Themes

The four themes are No. 1, to simplify, standardize, and converge the production environment. And I've spoken a lot previously about complexity, and standardization, and convergence. Everything I've said before speaks directly to this theme here.

The second theme, provide greater business agility for new technology and applications. We spoke earlier about IT goals around business agility and rapid provisioning, on demand provisioning, to meet those needs around business agility. That's largely what this theme is speaking to here.

The third theme, enhancing IP network security and its integration into a security pervasive infrastructure. What that's all about is going back and really looking at a holistic view of security within the data center. So instead of just looking at security at the network perimeter and at the application, we're really looking at implementing security holistically across all layers of the data center.

And the fourth theme here, driving autonomous and utility computing architecture. By autonomous, what we mean is self managing and what we're looking for is a self managing network that can help us deal with some of the complexities that the network has been bringing onto the operation teams recently. And the utility computing architecture is all about taking what we call pools of resources, pools of network resources, computing resources, storage resources and being able to pull from that pool on demand as needed in our data center network environment.

Sidney Morgan: So the utility computing architecture, is that gonna put any extra pressure on our access architecture for the service switches?

Jon Woolwine: With that utility computing architecture, I expect additional uptake of these services within our network. And one of the challenges we will be faced with in the near future, there's no doubt about it, is how to scale that service switch architecture and that's starting to creep into our thinking, especially at the design level. And that's actually a point I was gonna touch upon later in the slide that, yeah, you're definitely right. There's a scalability to these services that we have to be conscious of as we move further and further in this direction of utility computing.

Sidney Morgan: Jon, what are IT network themes?

Jon Woolwine: Yeah, good question, Sidney. I'll try and shed some light on that. So these themes, what they do for us is they serve as architectural principles for us and they guide us in architecture. And these themes typically they don't change much over time. These themes typically stand for several years. What they do is they really provide that guidance for us and direction for us within our data center network.

Sidney Morgan: Thanks.

Data Center Network Architecture Impacts

Jon Woolwine: Okay, moving on from themes, let's take a look at some of the

architecture impacts that we see coming in the near future. The first is around the continuing bandwidth growth with convergence of storage onto IP networks. Here what we're seeing is an eventual uptake of ten gigabit Ethernet at the access layer, meaning servers are going to be accessing the network at ten gig speeds.

Now with this we also are beginning to see the possibilities of converging the storage area network onto the IP network. These are some significant architectural challenges for us but they're also some big opportunities for us, as well. As part of that kind of convergence, one of the things we're gonna have to go back in and take another look at it is our quality of service architecture and the ability of that architecture to support the storage services over the IP network and over the data center Ethernet network. I would anticipate we're probably gonna have to make some changes there.

Scaling the switching architecture to support high density utility computing environment. What that's about is a little bit about what I spoke about before. So we talked about high density computing, more servers in a rack, more network interfaces per server. We see that trend continuing on and we see that density. We see things becoming more and more dense at the server access layer so we'll have to continually revalidate that we have an architecture that's appropriate for accommodating that kind of high-density environment.

There's a required alignment between architecture resiliency characteristics among the network computing resources and applications teams, and that particular impact goes all the way back to one of the Cisco IT goals around business continuity. In order for us to build a highly available and recoverable data center environment, we need to have architecture alignment between the network, the server folks, and the teams running the applications. We can't have application architectures that do not integrate with network architectures and the network providing capabilities that give us high availability but the applications or server is not leveraging those capabilities at the same time. So alignment there is gonna be key for us going forward architecturally.

Enhancing service delivery automation for application deployment or movement among computer resources. So service delivery automation is something we've already spoken to and I'm not gonna dive into that. Really simply, this is about rapid service delivery from an on demand resource pool.

Now the movement of computer resources or the movement among computer resources talks about the ability to support our recoverability objectives and to support business continuance. With virtualization, with automation, we have some opportunities to dynamically take applications and services and the workload associated with them and move them within a data center or move them from one data center to another data center. A lot of architectural implications behind that, and that will take some significant investigation on Cisco IT's part.

And the last architectural impact here is around virtual network segmentation in the data center. The business driver for that has to do with the multiple tenants that we see existing with our data center networks. So today, we don't just have say a Cisco IT data center network. We have a Cisco IT data center network that supports the traditional business applications.

We also have subsidiaries and other organizations within Cisco that leverage data center facilities and our goal here is to virtualize that network so that we don't have to physically create redundant parallel networks for each of those environments. We can virtualize the network and provide data center services for each of these groups and virtually

segment them to meet their business needs.

Data Center Network Design Considerations

Okay, moving down from the architectural impacts to more design level considerations, jumbo frames. Jumbo frames are definitely something we're beginning to look at in support of storage services over IP and over Ethernet. High availability is an area of design focus for us and where that comes into play is in a couple of areas. One is in device redundancy but also around figuring out ways to do operational maintenance on the network and yet at the same time not impact availability of the network. Techniques for doing that are things like in service system upgrades and leveraging the modular IOS capabilities that are available now.

The next design consideration? The number of network devices to manage. As I said before, we are beginning to introduce the Catalyst 4948 switches into our environment. One of the challenges that that's bringing with it is a greatly increased number of access switches to manage. So we went from a handful of Catalyst 6500s or we're going from a handful of Catalyst 6500s to a whole lot of 4948 switches, and managing those switches, we're gonna have to make some changes. We're gonna have to find ways to automate our network management to make that workload a little bit easier.

The network component lifetime versus server component lifetime, the design challenge there is to make sure that from a network design perspective and from a server design perspective we're in sync. Obviously there is tight integration between the server and the network, and these components have different hardware lifecycle periods. In other words, their refresh rate is different so we need to be thinking about what kinds of capabilities, for example, need to be supported in the network to allow that server environment to evolve as needed. And what that typically means is finding a modular and flexible network design.

Moving to the top right, spanning tree scalability. The design considerations around spanning tree scalability are primarily around finding ways to increase the size of that Layer 2 domain that sits within the distribution layer of our data center and yet at the same time maintain a highly available and stable network. So spanning tree scalability is obviously one of those design factors that you have to hit if you're gonna maintain that simple network and allow it to scale at Layer 2.

Data center space, power, and cooling considerations, this falls right in line with the high density computing discussions we had earlier. So continuing to be very conscious of the power constraints and the cooling constraints within the data centers is a top area of design focus for us. And the item just below that, cabling, same goes for that, really being conscious of cabling runs within the data centers in light of the high density computing.

The last item here, a granular incremental deployment. This is largely speaking to a need to come up with a data center design that, again, is modular in nature and that can be incremented sort of piece by piece such that capacity can be built out slowly over time as needed, versus say one big bang network that has to go in place all at the beginning and will sit in a largely underutilized state for an extended period of time.

Sidney Morgan:

For the granular deployment in an architecture picture, do you have set periods of time where you reevaluate the architecture? Because technology will change in a phased type rollout.

Jon Woolwine:

Yeah, actually we're continually reevaluating it, so architecturally we

have a body of architects that bring new capabilities that they like to introduce into the data center. They bring those into a forum where various architects and engineers can take a look at these capabilities and understand how their particular technical domain is impacted. So what we tend to see is one group might bring in a capability, an enabling capability, that forces some architectural adjustment. That doesn't happen on a regular basis. That happens more on an ad hoc basis as new capabilities arise for us.

Sidney Morgan: Thank you.

Data Center Strategies

Jon Woolwine: Now a look at one of the big fundamental data center strategies that we see. We see a natural evolution within Cisco IT's data center and that evolution I may have hinted on before but I'll just put it out there in plain sight here. That evolution starts with consolidation and standardization being a key component of that consolidation. The next building block beyond that consolidation is virtualization, the ability to take a consolidated environment and to build virtual pools of resources that can be provisioned on demand.

The next logical stepping stone beyond virtualization is automation and it's with that automation piece that you get the business agility benefits we talked about before, rapid provisioning. You also with automation take the human element out of it which has some significant benefits in terms of consistently configured environment and making sure things are configured correctly, helping you to maintain availability.

Data Center Evolution

Now this data center evolution is shown a little bit differently here in this slide. This is still speaking to the consolidation virtualization and automation evolution we spoke to in the previous slide but this slide gives you some sense of where we currently sit in that evolution. And today, for the most part, largely we've done a lot of consolidation. That's not to say we still don't have consolidation to do but we've been making some good progress in consolidation and standardization.

Recently, we've been making some – really pushing the virtualization side of this evolution, so we've been virtualizing these network services. We've been virtualizing the server that we talked about earlier, server virtualization. We're talking about virtualizing the network in the near future.

The next stepping stone, automation, we're just starting to tap and we're just getting some early experience in with the Cisco VFrame product and automated provisioning, and the benefits we're seeing there are huge as far as the provisioning lead times. We're taking a provisioning process that traditionally has taken sometimes weeks and bringing that down to days and sometimes even hours for us.

Continuous Availability with Disaster Recovery

All right, this slide's a little bit busy but let me see if I can boil it down to its essential elements. What we're talking about here is availability and recoverability, and we're talking about it in light of the Cisco IT goal of business continuity. Availability is something that's been high on our list lately. As the business has evolved, we've seen requirements and demand for higher and higher availability of applications and services within our data center. So at one point while three nines of availability, 99.9% availability, may have been acceptable, that's no longer the case. We are seeing requirements for four nines and sometimes up to five nines of availability from these data center environments.

And what that's forcing us to do is really rethink architecturally how we architect, design, and build these data centers across every layer, at the physical layer, the network layer, the storage and server layers, as well as the application layers. And right now we've been developing a concept which we call the distributed virtual data center. That distributed virtual data center is embodied here in these two blue rectangles, Data Center I and Data Center II. Now physically these are two separate data centers that would exist within the same city or same metropolitan area.

These two separate data centers would actually be logically operated as one virtual data center and as operating them as one virtual data center, we would be doing things like spanning high availability clusters across these two data centers for our applications and data center environments. We would be synchronously replicating data from one data center facility to another. If we had some kind of an outage event that impacted one data center but not another, we would instantly or nearly instantly be able to bring that other data center, the applications and services within it, online and make that transparent to the users of those applications and services.

So in this diagram, at the top you see an end user there. If Data Center I were to fail, we would leverage global server load balancing, specifically the GSS product, to redirect application requests from Data Center I to Data Center II. Our applications would fail over from Data Center I to Data Center II, or I should say our active instance of the application would move from Data Center I to Data Center II through clustering technologies and all our data would be synchronously replicated between the two data centers. So it would all be available for that kind of a high availability scenario.

Now in addition to high availability, there is also considerations that need to be made for disaster recovery. These disaster scenarios include situations where both of these data centers may be taken down by a single disaster event, in which case we need to get critical applications and services back up and running and that will be done in a remote data center. The way we do that is by keeping standby systems in this remote data center and by asynchronously replicating the data from these primary data centers to these disaster recovery data centers so that the data and the applications can be made available within a reasonable amount of time and be up and running.

Blade Servers and Integrated Switches

Okay. The final slide here that I'd like to talk about is related to blade switching within our data centers and related to the introduction of blade servers within our data centers. Traditionally Cisco IT has gone with rack mount servers. We've looked at blade servers a couple of times in the past few years and from a value perspective we haven't found that they've added significant value to our data center environment. In a recent investigation we actually did see a strong value case start to emerge for blade servers and we've seen from a power, a cooling, performance, and price point, that blade servers present a pretty good case for us. So we're just starting to look at blade servers within our data centers.

Now one of the considerations here for blade servers is around network connectivity. How do you connect the blade server to a network? I should say we will be evaluating different designs to do this and included in this evaluation are – or I should say will be Cisco's own blade switches and the possibility of using blade switches versus pass through connections to a traditional access switch.

Sidney Morgan:

Jon, can you elaborate on Cisco's blade server roadmap and some of the key decision points that might be made in the future?

- Jon Woolwine:* Yeah, so I think we're at a key decision point right now and that decision point is do we want to standardize on blade servers or do we want to continue going down the road of rack mount servers? We're just in the middle of that decision right now. We don't have any clear direction on whether we'll make that big leap to blade servers but it's definitely under consideration and I would expect to see blade servers become more prevalent within our data centers if not become the dominant form factor in the future.
- So that was all I had for you, Sidney. Any questions?
- Sidney Morgan:* Yes. The first question, what are the top business drivers for automating the data center?
- Jon Woolwine:* Okay. The top business drivers are No. 1, agility, so the time it takes us to provision services in the data center. By automating this provisioning process we can bring that time down from weeks and many days down to a single day or perhaps even a number of hours. So there's a lot to be said there from an agility perspective to be able to provision on demand and quickly like that.
- The second area is around consistency, having a consistent environment. So with automation, we don't have humans, engineers in there creating configurations by hand. We automate that process and with that automation we get better guarantees on a consistent environment and with that consistency we can typically get better availability. So those are what I would highlight as being the two key benefits that we've seen, at least from a networking perspective, around automation.
- Sidney Morgan:* Great, and what challenges have you solved within Cisco IT by deploying a services oriented data center?
- Jon Woolwine:* Yeah, so the challenges that we have been able to overcome with a services oriented data center architecture are primarily going back to standardization, virtualization, and automation. Automation we just hit on but the standardization part's big for us, so not having a complex environment that consists of many different products and that are kind of added ad hoc as needed. We have a very highly standardized network environment and a very consistent environment from one data center to the other. So on that services oriented model and the standardization that it's really pushed has helped us to create stability in the network and it's also helped us to drive down costs in the network through this consistency.
- Some of the other areas in terms of the services oriented data center, the virtualization aspect of it. Obviously the resource utilization part of that's been big for us, so I spoke before about the resource utilization benefits that we've seen through virtualizing services in our network. We talked about some of the resource utilization benefits we got from CSM and we're about to get from ACE and the huge consolidation opportunities there with that consolidation. That's all been enabled by virtualization, which is a key tenet of the services oriented data center. The automation side of it, I'll leave that alone since we just touched on that one.
- Sidney Morgan:* Okay, great. Well, thank you very much and thank you for spending time with the Cisco and Cisco team and sharing your expertise on the data center.
- Jon Woolwine:* I'm happy to do it, Sidney.
- Sidney Morgan:* Thanks.

Outro

Sidney Morgan: Thank you, Jon, I'm afraid that's about all the time we have for questions today.

Sidney Morgan: And for more information about technologies and solutions deployed at Cisco, you can go to the Cisco on Cisco site where you can find Case Studies with information about: what we deploy, what benefits we've gained, what lessons we've learned, and some operational practices and presentations to help you learn more.

Below that, you'll see a toll-free number you can call for more information or to place an order; and you can order Cisco resources on the web from the URL at the bottom of this page.

I'd like to thank those of you watching for spending this time with us, and for being interested in what the Cisco on Cisco Technology Tutorials are all about. We hope that you've enjoyed this seminar and that it has helped answer some of your questions about data centers at Cisco.

And thank you, Jon, for spending this time with us and sharing with us your expertise and your enthusiasm for the Cisco data centers.

Jon Woolwine: It was a pleasure, Sidney. Thanks for joining and see you soon!



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, IQ Expertise, the IQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)