

Data Center Technology and LAN Design at Cisco

Hello, and welcome to this “Cisco on Cisco” seminar on Cisco IT Datacenter LAN Design. I'm Rich Gore, Manager of the Cisco at Work team, which is a part of Cisco IT Cisco on Cisco initiative. The theme of our show today is datacenter LAN design at Cisco. It'll be an overview of Cisco IT's planned migration to a whole new datacenter LAN design. You'll get a chance to hear about how Cisco is integrating visionary new networking services into its current server and storage architectures. And also what implication this has for the overall LAN design. It is my real pleasure to introduce the star of today's show, Craig Heegan, who is Cisco IT's Chief Network Architect and a member of technical staff within the information network services technologies team. Craig has been with Cisco IT for ten years now and has been a central figure in the design of many of Cisco IT networks in that time. Craig, thank you very much for coming today. Well, thank you Rich, it's my pleasure. I'm happy to be here to talk about Cisco IT's new datacenter LAN architecture and designs. It's critical for supporting new improvements in storage networking, in datacenter application hosting overall in the utility computing environment. So sit back, enjoy, and we'll get into the material here.

OUTLINE FOR TODAY'S DISCUSSION

The outline for today's discussion is really to give you an idea of the background of Cisco datacenter environment. To talk a little bit about how our network is built today in that datacenter, and then to drive the future evolution of that datacenter network environment. We'll describe that in some detail.

CISCO IT DATA CENTER BACKGROUND

So first let's talk about the datacenter background and why we even have datacenters in the first place.

CISCO PRODUCTION DATA CENTERS

So Cisco's production datacenters throughout the globe are made up of really two classes of datacenters. The first is our corporate production datacenters. And these are the key primary datacenters that we have located in San Jose, in Raleigh, North Carolina, and in Amsterdam, as well as one in Southern California supporting the Linksys acquisition. On top of that we support many development datacenters which number many more around the globe and are intended to cover the engineering requirements here at Cisco for producing our products. Why do we have so many more engineering datacenters than production business datacenters? A lot of those datacenters have to do with reaching more interactive applications close to the users and providing the types of development environments that those users need to be very productive in creating new products and solutions for Cisco. More specialized, more local. Yes.

CISCO DATA CENTER LANDSCAPE

If you look at the datacenter landscape at Cisco we have approximately 7,000 servers in our datacenters around the globe. We have about one system administrator for every eight servers. Both that ratio as well as the number of servers is going to continue to grow over the course of the next few years. We have a heterogeneous environment today that has multiple hardware vendors. So you can see the difference in Solaris servers versus HP servers versus Windows and versus Linux in the chart here. And ultimately that multiple OS environment is something we will have to continue for the future. Now is this the landscape for combined engineering and production datacenters or just for the production datacenters? This is for both.

CISCO DATA CENTER GOALS

So some of the goals that we have within our datacenter are really around three areas. The first is obviously optimizing our cost of ownership. And really, ultimately that means consolidating our datacenters. Making sure that we're operating as few as we need while at the same time supporting the applications and the characteristics of the applications such as latency that the users require or that they desire. Ultimately keep a lifecycle management of our hardware assets as well as our software assets. Make sure that we are properly using the assets that we have and that we are getting the most out of them. And ultimately establishing the standards because standardization drives efficiency, and efficiency is what really affects that TCO equation as we talk. And we're going to get into significantly more detail about that as we talk about the future of the datacenter. The next is business agility. And ultimately in this particular area, this is where a lot of the shifts are taking place. We have a datacenter that has traditionally been a model of one application per server or one server per application. So typically when an application is deployed they call up the infrastructure team asking for a server to deploy it on. Now that's transitioning to a new model where we talk about the on-demand environment and a utility computing environment where we want rapid delivery of services. And we want to enable our business goals. And just the basic idea there is, rather than maintaining one to one relationships between applications and servers, we will migrate to a pool of compute resources. The same that we do for network, the same that we're doing for storage. It's establishing a horizontal layer that can meet the business objectives without doing one to one relationships. And finally we want to support business continuance. This is absolutely critical to making sure that in the case of a disaster we're able to continue in our environment. Now it's made up of a few different areas. First security, right? Everyone knows that security is absolutely critical to the environment. It's gotten more focus lately, and we want to ensure that both from a physical perspective as well as the network as well as the applications, we continue to drive that security mentality throughout the infrastructure. The next one is an active/active architecture. And what we mean by this is a change from many of the traditional models that you've used in the past which were active/passive. That change really talks about having an active server, which is running applications with a standby somewhere in another environment that you can roll to. What we want to do is maximize our investment in there. And typically what ended up happening was the backup server ended up sitting idle most of the time. And it aged without providing any support to additional load. In periods of peak access many times the backup server was sitting unable to help, unable to participate. And so we're evolving that architecture to support that active/active. But doesn't that run the risk if they're both active at the same time, that at peak times if one of them fails the other one remaining won't be able to support the full requirement? That is a danger if you're not properly handling capacity management. I see. So that comes as another part to play. You also have that when you're operating a primary and a backup mechanism where the backup has less capacity than the primary as well. Makes sense. Finally, ultimately, what we want to do is virtualize the operating system and application layers from the servers. Because today as we talk about the servers that go into the datacenter and the applications that ride on top, typically we are taking the operating system and installing it as a custom piece for every application we deploy. So how can we separate that operating system environment? Make it virtual. The application environment, make it virtual so that we can ultimately deploy the application independent and support a patching strategy for the operating systems in the datacenter for security purposes. Independent of any particular server that either the OS or the application is riding on. Correct. Now, in every case that's probably not going to be possible. But in a lot of cases you have very similar hardware, very similar OSs. And you want to make sure that when a security vulnerability is disclosed for an operating system that we're able to get the patches on all infrastructure that's running that particular version. Even if it's embedded and hidden from the user in some specialized applications. All right we'll move on then and talk about the foundation for integrated service management.

UTILITY DATA CENTER: FOUNDATION FOR ISM

Cisco's view of services provided to the internal users within Cisco is around services. It's not about an app. It's not about just a server or just a network element. Because services are really what our consumers use. And so as part of the foundation that provides the network servers and the applications, and the applications themselves, management is absolutely critical for this. So on the slide here we talk briefly about client services on the left, about calling cards, about pagers, about VPN access and the types of services that we offer. Now all of these are made up of a lot of individual components. The servers that provide the networking infrastructure, the identity management - all

of those capabilities, we have to wrap up into a client service. We build an infrastructure to support those client services. That infrastructure also has to turn around and provide management services for the infrastructure itself, right? The network management infrastructure and systems monitoring, trending, capacity management, all those elements. And then in addition to that we have the basic infrastructure services. Now this goes to focus on the applications that we know and use everyday which are critical to infrastructure. E-mail is a good example of that. Enterprise paging this is the identity services that help validate your password every time you log in. All of these services are native to the infrastructure themselves, and we have to operate those in that framework as well.

CISCO IT DATA CENTER NETWORK OVERVIEW

All right, now we'll talk about the Cisco IT datacenter network architecture. And this is the specific components that make up the network infrastructure at the first layers, the bottom layers of networking, as well as some of the services that we offer at Layer 4 through 7 which includes content networking and such. And briefly cover the existing network environment before we talk about the future.

TODAY'S DATA CENTER NETWORK IN CISCO IT

Today's production datacenter network is fairly standard network architecture that you'll find in most datacenters. It's built as a Layer 2 switching domain inside the datacenter. That is provided by Catalyst 6500 switches as part of the core distribution and access hierarchy. Not entirely Layer 2. I mean, you couldn't have the whole datacenter floor be on one single Layer 2 LAN, could you? Depending upon the size of the datacenter we provide a Layer 2 domain. In some cases our very large datacenters are segmented into multiple chunks based upon where the servers sit on the floor. So we only segment by VLAN within the datacenter floor, for the most part? We will use VLANs to provide the services required in one of those sections that I mentioned inside the datacenter. The IP network services that we provide inside that architecture are provided using what we call the server switch model. And basically what this server switch is it is a pair of Catalyst 6500s that provide the Layer 4 to Layer 7 services such as firewalling, content networking, etcetera, on a separate set of chassis. Now in the hierarchy of core distribution and access, these really look like access switches. They sit off to the side like many other switches. But ultimately it's the centralized area where we bring some of the special infrastructure services such as firewalling and content switching. And these are all done with blades on the 6500 chassis? Correct these are all done using blades in the Catalyst chassis. The out of band access for this network, because we need to be able to provide access in case there is a problem with the infrastructure, is provided by serial consoles terminated on Cisco routers. The 2800 and the 3800 series, as well as Ethernet lights-out management. The access from the servers. The specialized cards that allow you to boot and manage a server remotely. Typically on Catalyst 3750, the rack switches. And this infrastructure is kept independent so that we have access to that management plane should we have a failure of the production environment. Most of our servers today are connected at 100 megabits. The primary reason for this is that many of our servers do not reach that level of bandwidth where they need between 100-Meg and 1-Gigabit. Now this is rapidly changing. The availability of Gigabit ports on our servers, almost every server we buy today has a Gigabit NIC copper put onto the server's board. And so we are rapidly migrating toward Gigabit connections to our servers in order to take advantage of that. Now the question is do you immediately need that? That's a question a lot of people have asked. And in our particular case what we're talking about for the future, virtualization, really driving our utilization up, and standardization driving efficiency. Many of these add up together to form a higher bandwidth need as well as a higher CPU need and ultimately a higher resource need on these devices. Because it is no longer a one to one, now it's really an on-demand infrastructure where many applications can run on a single platform. And drive utilization up? And drive utilization up. So that's why we are focused on Gigabit connections to the servers, and even 10-Gigabit for the future when we start talking about blade centers and things like that. Finally we have global load balancing provided by distributed directors as well as an investigation into the global site selector as a successor to the distributed director going forward. And these are used for our key services such as active directory lookups, as our directory, our internal phone directory. All of these services use that capability to find the local instance of our infrastructure when you perform a lookup.

CISCO IT DATA CENTER NETWORK LAYOUT

Roughly, what you see here on the slide is our architecture for supporting the datacenter in terms of networking. As I mentioned it is a Layer 2 domain. Our service switches enables that Layer 4 to Layer 7 services. And you can see those two switches sitting off to the side providing those types of services. And we'll get into a little more detail later about how logically this looks from a network perspective. This is really a physical depiction. Servers can be either singly connected. In the case of the development datacenter this is fairly common where many servers can perform the same tasks. But many of them are only connected using a single link and the failovers will take place across multiple hosts as opposed to a single highly available host, which is the most common we find on the production or corporate side. So we might have, for some hosts, dual linkage from a single host to two switches for instance? To two separate switches, that's correct.

OPERATING SYSTEM ARCHITECTURE EXISTING SERVER ENVIRONMENT

As we talk about content switching in a clustered environment and distribution of applications across multiple hosts in the datacenter, we've made a progression from the traditional, the first implementations of content switching that we used. The good old local directors, all the way through the CSS or the content services switch, and now into integrated content switching modules that provide the primary load balancing service that we use today amongst our servers. So you can see that during our peak we had a fairly significant number of local directors present. Fifty-five pairs, actually, of these devices. And this is because the local directors operated in that same one to one fashion, that one to one ratio. That each application really needed its own local director to do distribution. As we moved into the content switching area with the CSSs and later the CSMs, we were able to do some consolidation. As the power of those devices became much, much higher we were able to move multiple applications onto single instances. And so at this particular point most of our applications are moving onto common pairs of CSMs as a horizontally offered service rather than a per application service in the datacenter. And this is common for every type of application in our environment J2EE servers, our Oracle 11i ERP environment, Exchange servers, Lotus Sametime. All of these -- and just to name a few -- all of these are powered by that content networking infrastructure supplying the connections to the servers. So a couple of questions. It looks like the number of actual things being managed is decreasing as you move towards the CSM. The CSM offers more features than just simple load balancing between boxes. Are we in the process of using those higher level features or planning to move in that direction? Absolutely, and that's a great point. We are, while the local director did traditional Layer 4 load balancing, just simply splitting connections among the servers. We are now using URL rewrite capabilities, we are using the other services that the CSM provides to inspect content, to rewrite some content, do redirection. Ultimately providing to the users what looks like a common name space for web applications, but at the backend splitting it up and providing the services across multiple infrastructure elements. That's a very good point.

IT CONTENT SWITCHING SERVICES

As we talk about content switching in a clustered environment and distribution of applications across multiple hosts in the datacenter, we've made a progression from the traditional, the first implementations of content switching that we used. The good old local directors, all the way through the CSS or the content services switch. And now into integrated content switching modules that provide the primary load balancing service that we use today amongst our servers. So you can see that during our peak we had a fairly significant number of local directors present. Fifty-five pairs actually of these devices. And this is because the local directors operated in that same one to one fashion, that one to one ratio. That each application really needed its own local director to do distribution. As we moved into the content switching area with the CSSs and later the CSMs, we were able to do some consolidation. As the power of those devices became much, much higher we were able to move multiple applications onto single instances. And so at this particular point most of our applications are moving onto common pairs of CSMs as a horizontally offered service rather than a per application service in the datacenter. And this is common for every type of application in our environment J2EE servers, our Oracle 11i ERP environment, Exchange servers, Lotus Sametime. All of these -- and just to name a few -- all of these are powered by that content networking infrastructure supplying the connections to the servers. So a couple of questions. It looks like the number of actual things being managed is decreasing as you move towards the CSM. The CSM offers more features than just simple load balancing between boxes. Are we in the process of using those higher level features or



planning to move in that direction? Absolutely, and that's a great point. We are, while the local director did traditional Layer 4 load balancing, just simply splitting connections among the servers. We are now using URL rewrite capabilities, we are using the other services that the CSM provides to inspect content, to rewrite some content, do redirection. Ultimately providing to the users what looks like a common name space for web applications, but at the backend splitting it up and providing the services across multiple infrastructure elements. That's a very good point.

CONTENT NETWORKING PRODUCT EVOLUTION

As I mentioned, the product evolution of the content networking platforms for Cisco have moved forward. And from the local director to the content services switch to the content switching module we have really made a lot of progress. The original local director was a separate standalone device. It migrated to the content services switch which was still a standalone device but could perform higher layer operations, as I mentioned a little earlier the URL rewrite capabilities. Then that technology was taken and integrated into the switching platform. So that with integrated within Cisco IOS we now have the capabilities in our server switches to offer these services independent of what used to be separate implementations.

CONTENT SWITCHING: BRIDGED MODE

Now the way that we deploy this, and this gets back to the conversation we had earlier about the L2 domain. We have deployed the content switches in bridged mode. And the primary reason for this is from a logical perspective, there are a number of positives that we attribute to that. The first is that we don't have to make any major changes with hosts in order to add this type of functionality. We simply logically move the servers providing load balancing, or needing load balance services, behind the CSM module. We simply configure it up on a separate VLAN. And logically it looks like a bridged segment through the CSM and that the servers are on the same subnet. I'm confused now which one of those three pictures is the bridged mode that we're talking about? The bridged mode is actually all three of those. Okay, there's my confusion. So from the Layer 1 view this is how they're physically attached on the right hand side. I see okay, thank you. The Layer 2 view shows from a subnetting perspective, from a bridge perspective how the devices look to that layer. And then Layer 3 to the network, it actually looks as if the hosts are on the same segment as the load balancer, as well as any other resources available there. So it's really invisible. Now another positive that's really key to mention here is that the IP addresses that are used by clients to access these services and these applications, are preserved. Whereas in a Layer 3 model where we redirected them, where we directed them through a proxy, they would normally have to change. And so in essence we have addressed the problem of web logs on the web servers with them. Now the downside to this, and I think you might have been leading this way with the discussion of Layer 2, is that there's the potential, if misconfigured, to create a spanning tree loop within the environment. And Layer 2 networks can be susceptible to this. And the key to preventing that is to ensure that the network is properly configured. That there are no errant links that are created between the two creating that loop, and ultimately watching the configurations to make sure that they are properly configured. So Cisco IT has some method of making sure that, for instance, if I go in and reconfigure one of the switches and I do it badly, there's some way of catching that using a standard config monitor? Rich, we wouldn't let you touch our switches. I approve of that policy, yes. But in general we have change management put into place and specific controls on an agreement on what content teams are permitted to change such as URL rewrite rules. And what the networking teams are allowed to change, at the lower level. And some capabilities in Cisco IOS can help us enforce that.

CISCO IT FUTURE DATA CENTER NETWORK EVOLUTION

So now we'll talk about the datacenter network evolution and where we're taking the current datacenter environment. Where we're going to transform it into that on-demand utility computing environment.

CISCO DATA CENTER TRENDS

From a trending perspective, if we look at the environment today we have a very heterogeneous environment, and it's a rigid environment, that one to one application type environment that I talked about earlier. We want to migrate that toward the flexible environment. We want to use a standards-based approach, probably take on commodity hardware, and really drive our applications so that they can be easily deployed across any of those assets. We want to move from the application specific security to a self-defending network which monitors itself, watches out for the threats that are in the datacenter. But also I think one of the keys is that we cannot afford to lose focus on security across the entire stack. So not only do you have to take application specific security and individual login prompts and transition those to more of an identity management type environment. But you still want to have some of the application level security's authorization checks, etcetera, in there. That's absolutely key. From a multi-services network, which is providing voice, data and video, we want to enhance that. We want to continue driving that. Define, really driving intelligent network services. This speaks to more intelligent application networking, application optimization services that the network automatically handles for us. Having the ability to even take the storage and integrate that with the IP network is really an important characteristic, and where we want to define that going forward. We'll take SAN and NAS storage to really create network virtualized storage. So today many of the hosts, we have to make a decision whether they're going to access storage services on the storage area network, or whether we're going to access storage on the IP network through network attached storage. We want to remove that barrier and really make storage as a whole a pool. And have perhaps differentiated access mechanisms but ultimately a single pool of storage that we're maintaining. The primary benefits of that are the efficiency and the economies of scale that we get out of managing a common environment. As opposed to maintaining a bunch of separate filers, a separate SAN, or even islands with multiple SANs that we have to maintain in that environment. So just to probe a little bit on this. We're going to have storage in one datacenter in San Jose, another datacenter in Research Triangle Park or someplace else. And also storage as in network attached storage that could be in hub sites or in branch offices. And we're going to manage them all as if they were all in one virtual pool? Well, I think the exact approach that you use and how you take a look at it will be different. There could certainly be separate pools of storage located in each datacenter. But I think the key is the convergence of the primary technologies here from separate direct attached storage, from separate network detached storage, and from separate SANs. And move those into a common pool of storage, at least within a given datacenter and proceeding then to perhaps a cross datacenter as we look going forward. And it's going to be a progression. We want to take dedicated compute resources -- this is really that one to one map -- this is precisely that point -- and move that to dynamically allocated compute resources. So that we can take applications and get them deployed onto available servers. This also helps us as we upgrade server components, as we continue the concept of fleet management where our hardware is upgraded on a fairly regular basis to keep it fresh, to keep it very efficient. We can move applications around on the servers in order to do that. Whereas in the past you had to prepare the new server, copy the application over, follow it to migration, and then finally turn down the other host to automate this. So could you describe this for me, this new server environment? Just what would it look like if I walked into the datacenter physically what would the boxes look like? Well, they look like many of the servers that you would find today. It's just that many of the transitions will be made from some of the previous hosts that were very large in nature. Probably the key difference that you'll see is a move toward more commodity hardware. We definitely know that that is a trend within Cisco IT and with other enterprises. The move to supporting commodity hardware, the move to supporting Linux in the datacenter, etcetera, is really critical in order to level out the platforms that we have in the environment. If you have many different operating systems that run on custom hardware for those operating systems you'll have a fairly difficult time doing some of that consolidation. And what we're looking to do is to, over time, consolidate the operating systems, consolidate the underlying hardware, and bring that onto the network such that it's that on-demand pool. So then from an application developer's point of view, suppose I've written a wonderful application that requires this particular OS and I say, here... I'll stop you right there. So in general the application developers should not be developing to a particular OS. Instead we establish the standard application development platforms, right? So we standardize on the J2EE platform, for example, for offering most applications. That won't cover 100% of applications. We know that there will be some custom ones that come into place and that we'll have to handle. And we will handle those as necessary. But by and large the majority of the environment will move to a standardized approach so that we can move these applications from host to host. So then I have a non custom application but it's really huge and it requires a lot of storage and a lot of processing power that might spike from time to time

because of user periodicity of use. And I hand it to you, and you're the datacenter person. How do you allocate resources to that? Well, there are a couple of ways to do that. The first is that if we have an infrastructure that can adjust on demand. If it's very, very spiky based on time of day, it's likely that you will have excess capacity allocated to it throughout other hours of the day, unless you're in an emergency situation. If you know that there are periods of months or weeks or even days where you're going to experience that, we can adjust accordingly for it. The basic idea of this is to have an infrastructure that based on some capability that you have to move applications around and what the weight is in terms of the difficulties of moving an application. We can make adjustments accordingly. Obviously if an application is too difficult to move on a very regular basis we won't adjust in terms of minutes. And when you say move, you mean move from one server to another? From one server to another dynamically. So if my application takes three and a half servers to run, you'd just toss it into a pool of server resources and let it find its own server? That's the objective. So there are tools that can help us. And in reality, when the rubber meets the road, there are some limitations that come into play like how fast you can make the change. We're looking to address those as we go forward. Finally, we want to move from an application specific management environment to a policy based management environment. And this really controls -- it really means setting up a standardized environment for expressing policy, where data can flow, which applications can execute those types of things. So that we are not building and rebuilding and rebuilding all of the logic to support the management of these applications. Really driving it as a horizontal because management is a fundamental element of everything that we do. And once again, that is the service oriented environment that we're looking to drive.

DATA CENTER CONSOLIDATION

So one of the areas that we're focused on is datacenter consolidation. And really, we'll talk briefly about the three areas. Datacenter consolidation, the server and application consolidation, what we were just talking about, as well as the storage and storage area network consolidation and what that means for the network. So what we try to do with datacenter consolidation is to leverage a smaller number of datacenters where possible to operate the same number of applications within the confines of the application's performance. So what this really means for us that we'll have a higher number of servers in a given datacenter that we have to account for, a higher density. And ultimately a larger network infrastructure that we'll have to deal with. In a server and application consolidation space we really talked about taking the applications and putting more of them onto a fewer number of servers. And this could be virtualization of just server resources themselves. But it could be actually placing five, ten applications on a single host that we know will not conflict with each other and that would be able to leverage less hardware there overall. You're multiplexing applications? Multiplexing applications. Interesting. And so we want to support fewer application architectures as part of this. And ultimately it means a higher performance server platform that is capable of doing some of this application consolidation. And finally we talk about the storage consolidation, and I already briefly mentioned it. But ultimately it means improving storage allocation. It's really operating it as a pool of resources, much like a pool of bandwidth, how we manage the network, right? We don't do segmentation and individual bandwidth handouts to servers. It's considered a resource that's consumed on demand by many of the servers and many of the infrastructure elements. So ultimately, what does that mean? What does the network need in order to support this? Well, flexibility is obviously key. You have to have deployment of new application environments in this infrastructure while paying attention to standardizing it. That's still key because that's what allows you to reduce your TCL, having that standardization. It makes it very much more efficient. In terms of scalability we absolutely have to have support for a higher density environment. Better aggregate performance and service scaling ability. Because we are looking at bottling more and more and more into a smaller number of components, which is really the value of the trends that we're seeing toward consolidation and virtualization. And ultimately high availability. Now this was always a requirement of the network before. But as we talk about more consolidation of end components they become even more critical. So the focus on more predictability, higher resiliency, and the type of features and functionality that you would expect in order to ensure that the applications got the support that they needed.

SERVICE-ORIENTED DATA CENTER MODEL

And ultimately, if you distilled this into a single message for the services oriented datacenter and what IT is doing in that. The primary piece is to create a pool of supply that the application demands from a business goals perspective leverages. So SODC provides that



brokerage, if you will, between the supply of infrastructure and the demand of applications. It's ultimately what we're trying to do in the most efficient way that we can do it. It sounds like a grand vision. It's a good goal. How close are we to doing some of these things? Well, I think some of the pieces are indeed in place to support this. Server virtualization using Cisco's VFrame technologies. Using some external technologies as well are very usable in this environment. We have many networking capabilities to do automation of application deployment. And so much of it is putting this together inside a framework which can support these applications in that environment. Impressive. Now there's a lot of development work to go as well, right? So we paint the vision, we have a lot of capabilities that we've been able to demonstrate. But ultimately as we move forward we will refine more and more of these capabilities.

SERVICE-ORIENTED DATA CENTER ROADMAP

So what does this mean for Cisco in terms of our service oriented datacenter? And you see here a timeline from 2004 to 2008. And we're really making a migration from what is the legacy datacenter that has these individual silos, these one to one applications mappings. All the way through in the steps required to support early consolidation, standardized OS, and eventually getting us to the end goal of infrastructure that's aligned to our services, automated application deployment. So instead of having to singly map an operating -- to install an operating system then take an application and put it on top on the individual server that was assigned. We now just allow the application developers to take their application archive, upload it to the infrastructure, and it's automatically distributed. That's the vision that we're shooting for around 2008, 2009. Ultimately that creates the utility model that we're after.

DATA CENTER NETWORK DRIVERS

So what does this mean for the network? And what are the network drivers that we see from a datacenter perspective? There are hints of all sorts of different areas in here, the first being security. Security's focus has really demanded roles based access to critical areas of the network and all access authenticated. Rules based access? Roles based access. So as we have in the past only provided application access based at the application layer in this area, we see more and more capabilities being pushed down into the network to provide the front gate. So if we set up a business policy that says only our CFO and his staff should have access to financial data we want the network to begin enforcing that as well. Now that has traditionally been a difficult thing to do. Over time we really see that becoming an important shift where the network acts as the front gate and will back that up with application oriented security. Well, right now it's done by passwords on the application itself. So what part of the network is going to support that? Well, in general what will happen is identity at the edge, you will authenticate yourself to the network when you connect. Instead of just plugging in and getting access you'll now be authenticating yourself to the network which will provide an authorization as well. Do you have the access required to reach this portion of the infrastructure? And so we'll begin seeing more and more of that. That also plays into managing the environment as well. So one of the things that has been critical for Cisco IT in leveraging the content switching environment has been the roles based access to manage the environment. Should the network guys be able to touch the content roles? Should the content roles guys be able to touch the network configurations? Those are the types of capabilities that we continue to drive into a collapsed environment. The next one that we're really focused on is integrating the storage transport into the datacenter IP network. And this is really around consolidating infrastructure. So wherever possible we are working toward the ability to leverage IP based storage on a single communications fabric. And over time we really see that being integrated into the IP network for most of the requirements in our environment. And this would be using technologies like Fiber Channel over IP to transport storage connectivity. As well as iSCSI for direct IP communication. We want to automate our service provisioning and delivery. You know the traditional model is the application guy needs a server or multiple servers. Lets say multiple servers. He calls up the server guys and says I need four servers. The server guys call up the network guys and go I need patches for those four servers plus load balancing. So could you guys configure me a load balancer across these four servers? All of that took individual trouble tickets, they took individual interfaces, and took a lot of time. And we want to get to the point where we are able to provide that type of service on just across the entire infrastructure as a service. So this is a utility oriented environment where the datacenter is the utility. And finally, we want to optimize our applications. And we'll do that by using intelligent capabilities in the network. A couple of examples here are the application oriented networking, the AONs capability, which will allow us to take some of the business to

business messaging that we have. This will allow us to take other messaging in our environment and really do some transformation and some intelligent message routing to provide access between disparate environments. And really integrate the environment a little bit more tightly. We'll also use wide area file servers and wide area application services in order to optimize the handling of this type of traffic over the wide area network. So we started at the top of the discussion on why there were so many development datacenters providing that. And a lot of it is because the protocols that are used are not efficient over the wide area network. I think we're all used to application developers who develop an application on the LAN and don't have the resources or don't have the ability to predict what the performance of that application will be over the wide area network. And by using application optimization techniques we'll be able to provide a little bit more insurance. Quite a bit more insurance actually. That when that application is taken to the WAN it will not have adverse performance problems. So could you just give me a quick one-sentence definition of wide area file services, what does that do? And wide area application services. Well, the wide area file services, primarily it is caching and protocol optimization for NFS and for the CIFS or Windows shares, in our environment. And wide area application service sort of extends that paradigm into optimization of all protocols or most of the protocols traveling over the network. Intelligent caching and compression of application traffic. So it speeds up the -- it looks like, to the end user, it looks like it's speeding up the traffic going across the WAN and speeds up my access to files that I may be needing at my PC? Correct. Now we can't beat the speed of light, but certainly it provides that type of acceleration for commonly accessed files and things like that.

NETWORK SEGMENTATION

Network segmentation is another area in the datacenter. And our focus has really been around three areas of segmentation. The first is separation of access, and this is where Cisco is adapting to a change in its acquisition model. Cisco has historically acquired companies and done a full integration of them at one time. And where we are migrating toward is an environment where we have longer integration times or we have what are called persistent divisions. If you look at some of the acquisitions that we've made like NetSolve, for example. These specialty networks interface with the customer network. Or in the case of Linksys there is a need to keep their networking separate from Cisco but still use the same common infrastructure for the efficiency that that generates in the economies to scale. So ultimately we want to provide a separation of access on top of that infrastructure that ultimately provides that separation but uses the same infrastructure, provides the security that's necessary. So we're consolidating things in the datacenter, the servers and the storage? We're consolidating datacenters, but at the same time we're allowing for a separation of separate networks? Right, for security reasons. The second area is once you've established that separation of access you still have some elements that you need to protect. So we're looking at ultimately protection of critical datacenter assets. Protecting our intellectual property, source code, financial data, and human resources data. Those types of elements and those types of servers you only want accessed by a specific group of people. And in all of those cases we want to limit the access that others may have toward that in order to guarantee the security, or at least increase the security of that environment. And finally something, with the self-defending network, is the quarantine capabilities that we have for unhealthy or untrusted systems. This is really assessing the posture of our systems and providing a quarantine capability that if your host does not conform to our standards we will certainly push it off to the side where it can be remediated. Where it has no chance of affecting the production infrastructure. And you said it's the host I was thinking of MAC as being something used for laptops and PCs, but it goes beyond that? Well, when I say host I speak in general terms of laptops, of mobile computing devices. Of really any device that's accessing the network. So it could even be, in some cases, building environmental controllers. So if you have a PC that's running an embedded operating system, monitoring the temperature sensors, you may want to validate to make sure that the underlying operating system of that is not compromiseable. Or at least is not subject to common vulnerabilities that are out there that would permit a hacker to come in to leverage that infrastructure, and to attack other infrastructure using it. Again, protecting the integrity of the system.

DATA CENTER NETWORK ARCHITECTURE IMPACTS

So what does that mean ultimately in terms of impact to network architecture? And there are six areas that we've identified that are critical elements in our datacenter architecture. Now we see continuing bandwidth growth with the convergence of storage onto the IP network.

This is in addition to the existing application growth that we have. Cisco is very much a champion of deploying IT capabilities to increase our productivity. In addition to that, as we take and consolidate some more of these infrastructures we have additional bandwidth growth. So we have to continue to think about how that's going to continue to scale going forward, both in the local area network, as well as providing in the wide area network the services between datacenters. Ultimately we want to scale our switching architecture to support higher density utility computing. So if you think about the traditional methods that are used to host applications, many times these servers were fairly large. Maybe six of them fit in a rack, or less. And they didn't have that great of a port density. And a few of our compute farm architectures are now beginning to use one-rack unit servers, appliances, which can be stacked up to 30 to 35 in a rack. And so having a changing switching architecture that can adapt to that, that can handle the spanning tree concerns. That can provide the bandwidth necessary to serve those, I think is critical and we're working on that and thinking about that. I'll talk about some of the architecture we're looking at using later. We want to enhance QoS to support storage traffic on the IP network. Storage has some of the greatest QoS requirements, even beyond the voice and the video which we experienced a couple of years ago. To provide the necessary application requirements that storage has, and that is very low latency. Performance is certainly a key there. But also the reliability of the infrastructure. And so supporting that is one of the critical elements that we're considering as we migrate that storage onto the IP network. The architecture resiliency, and the slide talks it's fairly complex -- it says required alignment of architecture resiliency characteristics. But really what this talks about is, as you begin to consider the reliability of the system, of the service as a whole, in the past we've seen our own organization and many other customers struggle with where do you place the redundancy? Do you make it at a system level? Do you make it at an application level? And as we talk about an on-demand environment where we can take some hardware resources like servers add or delete them at will, we certainly have to consider the resiliency of it. And so we certainly need to consider whether you're going to place very highly available servers in a very highly available network operating that. Or do you place some of that in the applications? And in many cases by having the applications talk to each other and having them across multiple different hosts, sort of in the RAID paradigm where you have many low availability type infrastructures. And you just do more of them with brute force. It turns out it's probably easier to maintain a consolidated number of components. But this is something that we're continuing to explore. And right now we're focused on providing high availability infrastructure, high availability compute nodes to support those applications. Because we're relatively sure that that consolidation and that type of availability is required there. We want to enhance service delivery automation. This is what I talked about a little bit earlier, in automating the complete picture and not just being one org sending a trouble ticket to another org to go do this. This is really around taking the services and providing them in such a way that if I want to add another server to the mix I don't have to generate trouble tickets. It automatically configures the network, adds it to the load balancer for that fifth server. Boots the server off of the SAN, takes the operating system image and puts the application infrastructure on top of it. It certainly is possible today. That's amazing. I mean, just the concept of racking a new server and having the network automatically pick it up and configure the network itself to respond to that. That's absolutely what we're doing in the environment. And finally we want to continue our focus on security. The point I can't stress enough is that as we talk about more consolidated environments, as we talk a more utility oriented environment, and as we talk about protecting critical intellectual property. We absolutely have to make sure that we have an increased security focus.

DATA CENTER NETWORK DESIGN CONSIDERATIONS

So ultimately the question becomes what are the design considerations that you have to have here. Ultimately we see jumbo frames becoming more and more of a need throughout the entire environment. Before you might have enabled it in a couple of custom areas. But now having that access across the environment is key for enhanced data transfer as well as storage. Storage on IP is likely to require large frames as well. Higher availability, as we talked a little bit about that already. The number of network devices to manage. As we talk about a higher density of server devices, that typically means a higher number of network devices to manage. And so this is something that has to be considered in the context of your designs and your operational capabilities. Also the next topic really talks more about management of the fleet. And as we talk about what the server component lifetime is and the network component lifetime. And as you look at evolving over time to higher performance platforms, to more functional platforms, really aligning, because you're looking at a service now. You're not looking at the network versus the systems versus the load balancers. You want to look at aligning the times. And making sure that you

have a good understanding of what a section of your datacenter is going to look like as a whole as opposed to individual components. We talked a little bit already about the scalability of spanning tree as things get more dense. This is likely to require a couple more changes in the way that we segment portions. That large L2 domain that we talked about earlier may not be able to be as large in terms of physical floor space covered as we go forward, just because of the density. We have to consider that. Space considerations inside the datacenter are certainly another one that you have to consider. And as you talk about more and more consolidation you have to look at the space considerations. And this also ties back to power and heat requirements. When you take a jelly and peanut butter sandwich and you squish it, stuff tends to come out the sides, right? Nice image. And so as you take a look at compressing more and more of the datacenter environment and doing that consolidation. You can see that you'll have to tend to a lot of the sides of the sandwich in terms of the power requirements, the space considerations, the cooling considerations. And those are elements that we have to consider as part of our designs going forward. Now you mentioned datacenter space. I was getting this image of smaller and smaller servers being more and more utilized. So it looked to me like the datacenter was actually reducing its space requirements for this new architecture. Well, I think if you looked at a static application picture you'd find that. But as we continue to grow applications, as demand continues to increase and we continue to build more business logic and analysis, you're going to find that that continued growth in our applications is going to use up some of that space. So I think, net, you're going to find some space consolidation just due to the efficiencies. But you can't forget the growth that continues to happen. Cabling is an issue along with that cooling and power element. How you do rack to rack cabling, and from your distribution frame to the rack. This is critical in that high density environment. Traditionally in the old days we used to do home cabling, home runs to a central server distribution frame. It's not possible anymore. So as we take a look at more dense environments we add more layers of consolidation or aggregation. And we'll continue to drive that in our architecture as well. So we're thinking of a switch or a rack of switches per rack? It's typically moving the access switches closer and closer. Because of the higher density, moving them in the racks themselves is actually serving a very small number of racks with an aggregation switch going forward. And finally granular incremental deployment. And this really speaks to the traditional method for a lot of datacenters was to do a massive build out of a corner or a section of the datacenter. And what we want to migrate to is the ability to roll in almost a rack of servers at a time, or a unit of three or four. And really this is around addressing the enhanced and consolidated utility. So as your application load starts to vary you can roll in a new rack of servers. Before you had to spec what server what going to be applied to an application and you found a way to rack that server. No more. Now we just add to the pool on the backend and continue to use it. See, my mind was trying to grasp adding one server and having the network respond. And now you're talking about bringing in a whole rack of servers. It's actually a rack at a time, that's correct. Amazing.

PHYSICAL SODC ARCHITECTURE

As we talk about what the service oriented datacenter environment looks like, the picture that you saw a little bit earlier starts to morph into something a little bit more in terms of what we're offering. Server fabric switches, like the recent acquisition of Topspin from Cisco. You can see addressed in providing some server interconnect. You see continued deployment of storage area networks. And ultimately we have an infrastructure here which is still made up of core distribution and access. But you also look at the fabrics that are used for the high performance clustering environments, like the server fabric switches, as well as the storage area networks. And the different access methods, different compute type nodes that we use. Traditionally you looked at a server as a server as opposed to a high density one-rack unit type server versus a traditional BigIron versus a blade center that's coming. And we'll talk in the next three slides about that.

CISCO SFS FOR I/O CONSOLIDATION AND CLUSTERING

So the first area that is really on the near term horizon for us is really using the Cisco server fabric switches as a method for consolidating I/O for high performance clustering applications. This will address a lot of the database or high compute needs. And really it is consolidating high bandwidth, Gigabit and 10-Gigabit Ethernet, as well as the storage connections to a server. And providing server to server communications using that fabric, using the InfiniBand protocol. This is going to be used in a fairly specific application context for those high performance clustering apps. But it is something we'll find where there are significant needs for a lot of I/O, and for server to server clustering type apps. So it's performing the same function, this SFS is performing the same function for servers as the MDS is

performing for storage racks or storage frames. Well, actually beyond that. So what the SFS does is it combines the storage and the IP network and a server to server connection for RDMA-type communications in order to provide that low latency, high performance communications fabric. And the applications that we're likely to use in this area are where we have very high I/O needs. I mean, we're talking a lot of SAN connections and a lot of IP network connections. And over time this type of capability, we expect, will be integrated into the infrastructure so that we'll have a common infrastructure for supporting those types of clustering needs.

HIGH-DENSITY ATTACHMENT TO THE DC NETWORK

The next area that we're looking at is the high density attachment to the datacenter network. And this is really around driving more of a server rack switching type model for these server farms where we have 25 to 30 one-rack unit servers. At the top of the rack we'll bring some of our switches in like the Catalyst 4948 10-GigE models. And typically what we'll do with these is aggregate all of the servers from a single rack into those switches and then bring those back to the datacenter aggregation layer. This is really to deal with the density of these servers to address some of the cabling concerns that you have with doing 25 to 35 cable runs times two for redundant NIC interfaces back to aggregation switches. And just using a point of aggregation for that infrastructure. So this is a really silly question, but in the previous slide we had all the servers kind of home-running right back to the SFS server or switch thing, switch if you will. And now we have all these rack servers going into the what is it, 4548? Remember that the 4948. But remember that this is an environment for very high density computing nodes, one-rack unit servers. Where the high I/O requirements are typically used on a different compute platform. These are the multi-rack unit servers that have traditionally had a number of storage area network connections. They have 8 processors, 16 processors, etcetera. It's the SMP type environment. Okay, so it's that slide or this slide and this isn't a combination? Right. Now, a third is when we start looking at blade servers to do these aggregations.

BLADE SERVER HIGH AVAILABILITY USING INTEGRATED SWITCHES

Because blade servers try to provide an aggregation of power, of network connectivity, many times they integrate another switch into the blade center itself. Integrating that is something that the network architect has to understand in both the architecture and the design. Spanning tree scalability for one instance. And in many cases we're going to look at a blade center as a device. The switch within it looks just like another access switch to us. So in general we're going to take those blade centers and connect those into our aggregation layer. We're going to use the features and the functionality in Cisco's blade center modules like the Gigabit Ethernet service module for IBM and HP blade centers. And connect those into our infrastructure like we would a Catalyst 4948 or a 6500. And take advantage of the features and functionalities that are on them. Make it an integrated portion of our network.

SLIDE 28

And that's the overview. That's impressive, thank you. Yes, well as a matter of fact we had collected some questions from previous showings of a similar presentation. And so we wanted to ask you some questions now, if that would be okay. That's absolutely fine. So let's start with the first question here. And let's go to that question. Great, okay, thank you. You mentioned WAFS and WAAS and AON, and those were new concepts for me that I'm still working on. How do these technologies, WAFS and WAAS and AON impact the design of the datacenter LAN? In general I think the technologies that we look at in this area are not going to impact the specific LAN design as much as one might think. In some cases it's likely that they would be provided as a transparent capability inline with the network or integrated with the design itself. But in a lot of other cases the technology which enables this may not even be in the datacenter. It's likely to be more in the wide area network portion. Where in the branch routers or sitting very close to the branch routers we have acceleration capabilities. And that's where we enable that as opposed to in the datacenter specifically. So as we look at these types of technologies that's where we would expect that these types of devices would affect our designs. That said, even within the context of those other places in the network it's very likely that they would be offered as transparent services connected in inspecting traffic very intelligently. So I'm not sure that there are many design changes that need to be made. Nice, okay. Another question was it looks like there are different types of connections for storage switches and for server switches. Are they all physically running over a copper or a fiber Gigabit or a 10-Gigabit

connections or are there special connections for servers and for storage that are required in this architecture? Well, in general in our environment there are two physical layer connections that you can have. There's fiber optics with a varied number of connectors. Although you can always translate very easily just by using a separate patch cord. Or you have twisted pair copper, traditional unshielded twisted pair. The use of SFS brings in a third type of connectivity right now which is an InfiniBand cable, which is a special copper interface that's used. Now in the standards bodies there are developments that are currently underway to standardize the use of InfiniBand over copper cabling as well as fiber cabling. And there have been discussions for quite some time. So over time we'll begin to consolidate. Most of the data center managers are focused on UTP copper and fiber inside the datacenter. Copper being used for the shorter runs and the fiber being used for the longer runs and the aggregation type runs. So what we expect is that regardless of the technology going forward -- actually, let me step back. In the storage area fiber is used almost exclusively. But as we look going forward, as we look at a common communications fabric. To the extent that we can continue to use copper at the speeds that are required, we'll use that for the short runs. To the extent that, mainly because the server manufacturers have copper ports built onto their motherboards. A very inexpensive technology to build onto the motherboards of our servers. We'll take that and use that, provide that aggregation and then use fiber to provide the backend communications where required. So ultimately, regardless of whether it's storage, regardless of whether it's server interconnect, Ethernet, etcetera, we're going to continue to use those and probably that very form and fashion as we look at it going forward. And one last question was we were talking about adding different types of new gear. The SFS switches, MDS switches and some blades within the network. And doesn't adding more types of equipment in the datacenter make it harder to manage and maintain? Certainly you have to consider the number of different types of devices you have in the infrastructure just as well as you have to consider the number of devices that you manage in the datacenter. Good point. At the same time I think where we're looking to make a significant change is, remember, this service environment where we have really focused on the provisioning and the management systems of the entire service. As opposed to managing individual devices. So our attention is really focused on automated provisioning and the scalability associated with making sure we can manage the number of devices. But largely we're trying to get away from a paradigm that has manual configuration of individual devices. Because that's where a lot of the human error exists, right? By building it into the system, by using management access for those devices in the case of troubleshooting as opposed to the individual trouble ticket method. That's where we want to take that. Okay, thank you. Well, that's about all the time we have and that's about all the questions we have.

MORE DATA CENTER RESOURCES

So we're now looking at another slide that should give you information about Cisco IT deployments. And you can go to the Cisco IT at Work website to find case studies about what technologies we've deployed, what benefits we've gained from them, what lessons we've learned. And also there are some operational practices and other presentations to help you learn more. Now you can also find out more information on some of our other IP networks, design guides, operational practices, and several other documents, white papers and presentations on cisco.com. Below that you'll see a toll-free number that you can call for more information or to place an order. And in addition you can also order Cisco resources on the web from the URL at the bottom of the page.

SLIDE 30

So I'd like to thank all of you for watching and spending this time with us. And for being interested in the Global Technology Seminar series. We hope that you've enjoyed this show and that it's helped you to answer some of your own questions about datacenter technology and about LAN design in the datacenter. And most of all I wanted to thank Craig for spending his time with us and sharing with us your expertise and your enthusiasm for datacenter technologies and datacenter LAN design. It was certainly a pleasure, Rich. And we hope you've enjoyed this show, and we'll see you soon.



Corporate 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: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
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

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, 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, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, ScriptShare, SlideCast, SMARTnet, 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. (0601R)