07 – MPLS

Hello, and welcome to this “Cisco on Cisco” Seminar which is on Cisco IT's MPLS VPN network in Europe. I'm Rich Gore, Manager of the Cisco at Work team, which is a part of Cisco IT's Cisco on Cisco initiative. Now the theme of our show today is going to be MPLS VPN networking at Cisco. We'll be showing you an overview of Cisco IT's early use of a Service Provider's MPLS network to support our WAN in Europe. You'll get a chance to hear about how and why Cisco IT chose to deploy an MPLS network in Europe and several of the problems that were overcome along the way. Now it's my pleasure to introduce the star of today's show, Craig Huegen, who is Cisco IT's Chief Network Architect and a member of technical staff in the Intelligent 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's networks during that time. Craig, thank you very much for coming today. Well, thank you, Rich. It's my pleasure to be here and to be able to present the best practices and lessons that Cisco IT has experienced with MPLS VPN over the last few years. So sit back and enjoy and we're going to explore Cisco IT's use of MPLS VPN to support Cisco's WAN.

OUTLINE

First, we're going to talk about the Cisco IT IP VPN vision, which is a little bit of background but set up our strategy in 2001 and as we have implemented several networks based upon that. We're going to talk about some of the early hurdles to that adoption that a lot of Enterprises have found and Cisco itself has found. We're going to follow it up with the lessons learned in our Europe, Middle East and Africa implementation. And then we'll go with a few questions I know that we've collected from the past.

THE CISCO IT IP/VPN VISION

So first, let's talk about the Cisco IT IP VPN vision and why it's very important to us.

CISCO’S INTELLIGENT NETWORK INFRASTRUCTURE DELIVERING VALUE THROUGH THE NETWORK

As we take a look at what the underpinnings of Cisco's network mean, we can show the building blocks of the Cisco global converged IP network. Now this is the basic fundamental architecture that defines Cisco's network as a whole. It's not specific to IP VPN, but it's very important to understand in the context of how Cisco IT does its work. The first is optimum latency and this is really important in the world of supporting data, voice and video. Because interactive applications require that we have the lowest amount of latency between locations, from a network component, as well as at the application layer, to provide the best level of productivity to the user. Flexible access is another underpinning. So as we talk about the benefits that we get from all of the methods that we use to attach to Cisco's network, this is an important characteristic. It absolutely has to be highly available, obviously, to run Cisco's business. It's one of the most critical elements of the operations of our network, as well as the architecture and design. It's IP based. And this really talks about the convergence element. Instead of operating many different separate networks over the years Cisco has had a very strong focus on converging to an IP network, both for voice. Then for video and now we're getting into some of the storage and other applications. We talk about advanced services on top of this. Now, these advanced services are intelligence within the network. It's really intended to make sure that we have optimum distribution for our applications. These are technologies like Quality of Service, to support these enhanced applications. Multicast, to optimize the delivery of content and of software throughout the infrastructure. Wide area application services, all of those elements that contribute to the intelligence of the network, and ultimately providing a better foundation for our applications to operate in the environment. And how does MPLS VPN support that intelligence? Is that a well supported set of services on top of an MPLS network? Well, we'll talk about it. It has to be integrated into the network and I'll specifically talk about the importance of that in the coming slides here. Great. Secure access, security is absolutely critical nowadays. I don't know that we need to say anything more about how important it

is. But we need to ensure that we feel a comfortable level of security as we're going forward with these types of technologies. And finally, converge, once again, onto that IP foundation. We're looking to drive as much complexity out of the environment and take advantage of the economies of scale that an IP network offers. Now with that foundation set, that really forms a layer, if you will. That network forms a foundation for supporting data, voice and video and the rest of the application categories that you see there in the red blocks. All of the support for corporate administration, employee collaboration, etcetera, are really built upon that foundation. Ultimately the net result from a business perspective is the productivity, the shareholder value to Cisco and also a showcase. We have to really point out the fact that the Cisco on Cisco element is very important to us.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE DIRECTION: LEVERAGE ENTERPRISE CLASS IP/VPN

So let's move now into the reasons that we want to support Enterprise-class IP VPN. And why it's important to us both today as well as what the trends we see in the future. Excuse me, when you say Enterprise-class IP VPN, what other classes are there? Well, the thoughts are that you can find more carrier class elements, carrier supporting carrier type services. There's the wholesale element of it when you have a smaller Service Provider operating on a larger backbone. This is really targeted at what an Enterprise network and how they would use that service. Makes sense, thanks. So the first element is the transport benefits. These were the first direction or the first items that we looked at in terms of building our direction, because they are the most tangible from the Service Provider in our initial deployment. The first thing that it does is really it out tasks the core of the network. If you take a look at how our network has been built over time, to get that optimal latency, to get that flexible axis. We have had to go through a procedure which took our Service Providers, asked them for their fiber backbones, and in order to make sure that you have the optimal latency, we mapped onto those backbones. For instance, if a Service Provider didn't have a hub located in a particular city, say maybe Dallas, Texas or the Bay area, and you established your own hub there, you could end up backhauling a lot of the traffic. And so what we needed to do was study the fiber maps, build the infrastructure in those cities and then interconnect them with circuits. Well, if you look at what the Service Providers do in that particular context, they're doing exactly the same thing. After all, they have to do the planning for their underlying fiber, their intermediate systems as well. So the idea of MPLS VPN as a technology here really leverages a single backbone that the Service Providers build and that the Service Providers operate. It creates a virtual any-to-any on top of that that Enterprises can take advantage of. And so for us, it makes it very easy to add sites. All we do is worry about connecting a site into the cloud. Not having to go through which hub does our site need to connect into for optimal latency and what is the impact on the backbone links. Because the assumption there is that the Service Provider is maintaining that as part of the service. So that adds a lot of flexibility to the service. Something that made me frightened just a little bit was when you said out tasking the core or the heart of the network. How comfortable are we giving the heart of our network to a Service Provider? Well, when I say the core of the network, immediately that draws up the picture in your mind of the center of the world, if you will. Thank you, yes, that was -- And it's really that picture that I think scares a lot of individuals. If you take a look at what Enterprise network architects and engineers are doing today, when they have to analyze the circuits, they have to look after the fiber paths. They have to watch out for redundancy. And in some cases they are actually having to pull design reports for their circuits every thirty days to make sure that they are absolutely redundant. And that's something the Service Providers already do, so why not take advantage of that? That is ultimately the reasoning for driving it. Now there are also a couple of other operational elements to this. The capacity planning now gets reduced from a matter of instead of having to watch capacity for each individual link of the puzzle, it becomes about capacity from one site to the rest of the network. So instead of having to worry about bandwidth between A, B, C, D, E and F on a path, we only have to worry about A to the cloud and F to the cloud, very valuable. The other element that we wanted to really drive IP VPNs with is a foundation for in cloud services. And really the concept behind in cloud services is the use of telephony, for example. In every one of our branch offices with a traditional network connection, we have not only the data connection, but also voice connections to the local Service Provider for voice. We have a connection perhaps to our long distance voice Service Provider using traditional ISDN primary rate interfaces. We want to take that and leverage IP VPN as a foundation to access services like IP public switched telephony, really to connect to our Service Providers, and send our inbound and outbound voice to and from customers over that network instead of establishing gateways at each of these offices. And IP VPN provides a foundation to more easily do that. It's not to say it can't be done in a traditional network. It's a foundation to

make it easier to support and easier to use. Now, that said, there are key service requirements. Because, as you mentioned, we are taking some of, what was in our control previously, and we're really expecting the Service Provider to step up. And there is such an integration here, there are key requirements. The number one is that we have a transparent service. The service needs to be a drop in replacement for an existing WAN. And the reason for this is you don't want to have to rearchitect your network. You don't want to have to reengineer the network. And so in the areas of routing protocol integration, for example, you don't want to have to reengineer your network to use another routing protocol. You'd like to be able to just drop that in and replace what is an existing collection of links with this service. You want it to support Multicast so that you don't have to build really workarounds of tunnels and establish a more complicated scheme for handling that traffic. And finally the QoS needs to support. If you already have QoS within your Enterprise the QoS needs to be supported throughout the cloud as well. And offer the transparency such that your traffic classes can be preserved inside that cloud. In addition, you also need sufficient footprint from our Service Providers worldwide. And the reason for that is as you're talking about the nodes and each of the potential nodes for traffic in the network and where traffic can be redirected; you need to have enough coverage that traffic can for optimal latency purposes go the correct direction. If you have only two, if your Service Provider has only two provider edge routers in its network, for example, one in Los Angeles and one in New York. And an office, let's say Denver, tries to talk to another sales office maybe in Indianapolis. The traffic would likely have to go from Denver all the way to the west coast. Go across the country to that other provider edge router and then come all the way back west to the Indianapolis office. Instead, what we need is the coverage of those provider edge routers in order to address optimum latency. So that Denver, hopefully, connects to a hub close by or a provider edge router close by. And Indianapolis, hopefully, connects to that connection close by and allows for the optimum latency directly between the cities. So the Service Provider basically has to have sites in any of the cities that Cisco has a real interest in. I would say that the target is very close to that and we'd like to have that coverage as close as possible to that. Okay.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE DIRECTION: CONVERGENCE OF PIPES

So what can we do with this? And I covered this a little bit earlier, really converging to common facilities for the services that we have in the environment. You know, as we take a look at long distance PSTN connectivity and the local PSTN connectivity, along with IP video conferencing capabilities, and content distribution. All of these types of services that we have seen delivered via different connection types in the past, we really want to take those and aggregate those. Use IP VPN as that foundation. Leverage the core services and the economies of scale that are built by that, in order to provide that on an operational infrastructure that is less expensive to maintain -- a single pipe taking advantage of a pool of bandwidth, as opposed to individual services with individual bandwidths that we have to manage. So it simplifies management a good deal. One worry I would have about converging everything onto one pipe and I'm sure you didn't mean that exactly is what happens if that one pipe goes down. So there needs to be better backup. When I say one pipe, I actually don't mean one physical pipe, but I mean one logical pipe that reaches that particular office. Physically, we look at actually deploying two different circuits that provide that service. But it is IP VPN services, as a whole, provided redundantly, as opposed to or with resiliency, I should say, as opposed to having many independent connections that each need their own redundancy. Makes sense. Also taking advantage of the resiliency characteristics of aggregation. Now the challenges that we face there are around economies of scale for those integrated services. For example, most of the Service Providers today are just now beginning to develop their IP PSTN service offerings to customers, so that you can connect with the public switched telephony network through their services on top of their IP VPNs. The other element that we observed very early on and it continues to some extent today is the idea of bandwidth mentality. Many of the Service Providers have really focused on what has been their core business. They've offered bandwidth. And so the IP VPN offerings are very much a basic service that they offer. What they haven't done very effectively to this point and are really making strides toward now are taking the other services that they offer, the voice services, some of the content distribution, etcetera and really marrying that with the IP VPN vision to offer that to Enterprise customers. And so much of our focus in the earlier days of implementation was around just the transport aspects. But the services and the services environment is something that we definitely want to continue pursuing, as we go forward with those offerings from Service Providers.

HURDLES TO IP/VPN ADOPTION

So now let's talk a little bit about the hurdles that we have with IP VPN adoption. And there are a handful of them that have really made it difficult for us to go very fast with this particular approach.

IP/VPN ROADBLOCK 1: PRICING

The first one is road block number one was pricing. Now, some of these that I talk about today have largely been mitigated and we'll talk about how they've been mitigated as we go through them. The pricing for the early IP VPN services were fairly high. And the reasoning for this is that the Service Providers looked at the MPLS VPN as a full mesh network. And said, comparatively speaking I'm going to charge the Enterprise customer the same price or near the same price that I would charge them, for building a full mesh of ATM or Frame Relay PVCs, the virtual circuits, throughout their network. Now, as we approached it, as we talked with the Service Providers we said, well, Cisco IT and almost every other Enterprise, when you really look at it, they have some any to any needs for optimal latency. But the vast majority of our applications still come back to Data Centers. We still have the main traffic grooves, if you will, are based from the facilities where we have data and that we support them. And what we're really trying to enable is more collaboration, is more telephony. So when you look at the pricing, we told them it's not about a full mesh network. It's actually about providing enhanced or more enhanced capabilities in that particular fabric to enable optimum latency. And so that really is changing. Most Service Providers are restructuring their pricing for the purpose of adoption. And, in fact, the recent bids that we've seen are demonstrating significant savings compared with building your own traditional network. And that I contribute to a direct push in leveraging these types of services in the intelligent network in the Enterprise. And a realization that it isn't about an any to any network, but about providing optimal latency for those applications and enhancing the network.

IP/VPN ROADBLOCK 2: TRANSPARENCY

The second roadblock that we ran into was transparency. Some of these have changed and some of these still continue to be a concern for us. The first one from a QoS perspective that we hit very early on, was that Cisco had a seven, at the time a five class of service model, but now a seven class of service model, that we've establish end-to-end for our network. And so with these levels, these classes of service for voice, for video, for interactive applications. What we found was that the Service Providers had many times only three or four classes of service. And even within those classes of service we could not recover our own markings at the opposite end. So if we had to compress five classes of service into three, at the other end what could happen is video and voice could look like the same traffic. And we really needed the ability to separate it back out into its components. We worked with Cisco's development teams as well as the Service Providers to establish a method to do that. And almost every Service Provider today can offer QoS transparency to the Enterprise. The second one is around Enterprise IP routing integration. And what this really speaks to is that the service needs to integrate with the network. In our particular case, EIGRP is the routing protocol that we use between all of our sites. And many of the Service Providers only offer the option for static routing or to use BGP, an external routing protocol. Now it's not so much the need to learn about BGP, as we have BGP in some portions of our network. However, what it is about is the need to architect the network to support it, to do the redesign. And the complexity of having multiple protocols in the environment, because BGP still needs an underlying routing protocol to operate. And so for our network, we have determined that we really wanted to interface with the Service Provider network. And that other Enterprises would want to do the same with traditional Enterprise routing protocols, such as EIGRP, OSPF and RIP, Routing Information Protocol. Finally, we wanted to address the implications of IP Multicasting. Cisco uses Multicast to transmit audio and video, using a single stream worldwide for key events such as Cisco's company meetings. Cisco also uses Multicast for content distribution and rolling out software to our content networking out in the field. As we took a look at Multicast, we really needed to ensure that that capability still existed in the MPLS VPN network. Now a lot of Service Providers have seen that this is very important and it is also, even more important than to Cisco. It's important in the financial vertical as well. Most of the financial Enterprises use Multicast to distribute market information in real time. And so Multicast is a critical element. Most Service Providers have seen this and it's a matter of doing the testing and supporting Multicast within their network, and determining very early on what the model for Multicast as a value added service was. Some Service Providers

wanted to charge for it. Others included it as a base part of the package. And our particular position was Multicast actually saves some bandwidth. It actually saves bandwidth in the core and for the most part it should be an integral part of the network both, from an Enterprise perspective as well as a Service Provider perspective. Makes sense. So of these three different difficulties in transparency, which of them are still current today? Well, Multicast to some extent, but I would say that that is moving away. Most Service Providers have said we're going to support Multicast as part of our base offering. The routing protocol is the one that really is an outstanding. And we're finding that a number of Service Providers, a good number of them, are supporting and they understand the concepts behind needing to have that level of integration. And others have not really progressed to that point. They're looking to establish a basic level of standards that they can operate their network with. And then customers should interface with their network using those standards. And so a lot of it is just a matter of where the Service Provider is looking to offer that level of integration to the Enterprise customer. And we think it should be very heavily on the integration side. So ultimately, the partnerships that we have with Cisco's Research & Development Departments, as well as the Service Providers themselves, are really driving this through new feature development, as well as implementation. I'm sure a lot of customers are happy that Cisco IT is going first. I think in many cases that it is very true, because we've heard many other Enterprise networks give us the very same feedback, both in customer councils with Service Providers as well as in our own Enterprise advisory boards.

IP/VPN ROADBLOCK 3: PE COVERAGE

The third roadblock that I talked about was provider edge router coverage. And I mentioned this a little bit earlier. Specifically in the early services that we looked at, they relied upon a small number of routers for a trial perspective. And the number of those and the locations of those really didn't lend themselves to our network very well. So in many of our trials what we found was that there was only five or there were only five PE nodes n the United States, where Cisco has nine hubs located throughout the US. Actually moving to that service would have increased our latency, on average, by about 10% to 20% across the network. Ouch. Now, the services had developed today to the point where coverage is almost a non issue except in emerging portions of the world. So in the United States throughout Europe, western Europe, as well as the core of Asia Pacific, we are seeing that the coverage is very good, and is not too much of an issue anymore when compared with what it used to be. And so the service evolution there to the IP services network, really a lot of the transitions that Service Providers are making towards MPLS as a core of their network for carrying all applications, not just VPNs for the Enterprise, that transition is enabling them to build their footprint and to establish that as a key portion of the network. So this is really not a giant concern for us anymore. We still ask it as part of the procurement process, but largely we expect that most Service Providers that come to us will have that capability.

IP/VPN ROADBLOCK 4: UNMANAGED CE SERVICE

The fourth roadblock to this IP VPN deployment that we found was an unmanaged service for the customer edge router. Cisco, as an early adopting Enterprise, typically has needs to run features, and especially with the Cisco on Cisco element, has a need to run features and functionality at the edge of its network. In addition, we're concerned about the loss of visibility into the Wide Area Network. When we have control, or at least access and visibility to the routers on our side of the connection, many times we're able to defeat a -- I shouldn't say defeat, but perhaps warn off a truck roll by a Service Provider by looking at the status of the connection. And in some cases, if the Service Provider can't reach the customer edge router, the first thing they do is they send the truck out to take a look at what may be wrong. By having visibility from our side and an unmanaged CE service, we're able to look into the condition of the Wide Area Network circuit. We're able to provide that information back as we're calling in the trouble ticket. And it can reduce the cycle, the repair time cycle for that particular outage. In addition, that features and functionality that I spoke of a little bit earlier. Those types of features are really applicable, like QoS, for example, right at the WAN edge. Normally we would take a look at perhaps setting the features back one hop within the network in a router that we control. But in many cases the features are most applicable where you have that speed bottleneck, where you go from the LAN to the WAN which is that router right there. And that's changing through customer demands. We're finding most Service Providers are now offering an unmanaged CE offering. Many of the Enterprises have spoken up and said we want control of that edge

router. And we actually use a form of co-management, where we allow Service Providers the ability to see certain parameters on that router. But we hold the control to the configurations and have a process by which we manage that. So they have read access, but ... Generally read access, but very limited read access to those statistics necessary to manage the network from their side, performance, and to identify potential problems. We don't want to take anything away from them. We want to supplement the process. And that's exactly what we've provided as a result of our use of unmanaged CE services. How comfortable are Service Providers doing that these days? These days, it's actually fairly well accepted because of customer demand. In the early days, it was viewed as strictly a managed service opportunity, so that the Service Provider would come in, offer you the router and would hand you an Ethernet cable and say, here's your connection to the network. Throughout the discussions that we've had with Service Providers and other Enterprises, that view has really changed. That said, there is still a good portion of the customer base that still wants to have that type of capability, or I should say that managed service, for the Service Provider to come in, provide the connection and just connect it to my simple switch.

IP/VPN ROADBLOCK 5: INTER-PROVIDER VPN

The fifth roadblock is inter-provider VPN. And this is really the question of resiliency. In prior networks using leased lines, ATM or Frame Relay, it was much easier to purchase a service from one Service Provider, a link between two locations, and from a second Service Provider buy another link. And you could alternate your Service Providers and build a network that, as a composite of all of those elements, was very resilient, and was able to address some of the concerns of failure of a Service Provider. In moving to IP VPN, the very benefit of IP VPN is derived from having that buying an integrated network service, as opposed to individual links in building your own. And so perhaps the discussion of resiliency in that environment had to be readdressed. In our particular case, we said deploying IP VPNs involved taking slightly more of a risk. Enterprises can no longer combine those multiple Service Providers to form that network on a link by link basis. So instead, we looked at using multiple Service Providers who offered MPLS VPN services. This one we can do very easily by buying a primary network from one Service Provider that, say, reaches 100% of our offices. And our second Service Provider, which reaches a small number of offices. By number it's actually about 10% to 15%, but they typically make up 80% of our installed user base throughout those areas. So all the large offices. And then we back that up with traditional internet VPN access, remote access type technologies in case of a failure of an entire Service Provider. In case of that failure, our users can use their home connections, they can use hot spot connections in common places, to get access to the network. For our primary Service Provider we still deploy multiple circuits. We still provide that level of resiliency necessary to protect the network. But largely, this is a matter of protection against an entire Service Provider's infrastructure failing. So I heard four different layers of redundancy there. There was the primary network with, of course, redundant pathways. There was the secondary network that supported the larger sites. There was VPN, internet VPN connections between sites to back that up. And then, of course, there's the resiliency of individual users who can find public internet access at home or in public places. Right, wow, that's a lot. I'd call this careful distinction. Redundancy versus resiliency. Redundancy is provided in the primary Service Provider but we viewed the other options really as resiliency and business continuation practice. And it really goes hand in hand with the mobility aspect of our network, as well. We want the user to be able to access our network from any device anywhere in the world. And so what we're doing is for resiliency purposes we're leveraging those other elements. So we're just not putting all of those elements in place for the sake of backing up that office. They're actually part of a larger strategy to deliver access to our users. Interesting, okay. Now, where we can take that, you'll notice that I said we manage multiple Service Providers. We have them connected to each one of our offices. The difficulty is, of course, is that they are not interconnected inside their cores, the Service Providers. Okay, two separate cores, yes. Two separate cores. And so when you look at two separate cores, and there is a failure of a portion of the network or an individual element like a circuit, how do you provide? The question that the Enterprise has to solve is how do you provide the optimal path through that infrastructure without watching how traffic is bouncing around. So for example, let's say that offices one, two and three are all connected to the primary Service Provider. But only offices one and two are connected to the secondary provider. And the primary link from office two fails, okay? Ah, okay, how does two get to three? How does two get to three, and ultimately bouncing -- you have to see how that -- You have to project how that ends up looking within the network and the possibility of all the different paths that you could do, especially as offices one through three turn into offices one through 100. So one thing that we're really driving with the Service Providers is the concept

of having interconnects between the Service Providers themselves, to offer this type of resiliency. There are technical methods to do this, but largely the challenges are business challenges. The Service Providers are saying, we have the ability and we have the processes to offer you the network on a global basis. We have the ability to be your sole provider and you don't need to worry about other connections. And, of course, things do happen and we need to protect our business. So convincing them of the need to provide that type of interconnection and that type of capability has been very hard to do. Although we've seen recent movements in setting some of that up through alliances, through partnerships, and really leveraging some of the relationships that they already have. Pretty good.

CISCO IT EMEA IP/VPN IMPLEMENTATION

So now let's talk specifically about the implementation in EMEA.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE EUROPE WAN FY03/04 MPLS VPN

This took place in the old days, in the early days of MPLS VPN. So FY, fiscal year 2003. Correct, this happened in our fiscal year 2003, 2004, which was the end of calendar year 2002 through the end of calendar years 2003. Our implementation, as I mentioned a little earlier, was based on a primary IP VPN Service Provider to connect most Europe, Middle East and Africa locations together. There were some exceptions there. And those exceptions, those hard to reach areas, because there wasn't sufficient coverage to areas like South Africa, and parts of the Middle East, we had to use traditional methods to connect them in. But that is changing, of course, as most of the Service Providers are able to offer that. So we are slowly rolling those sites onto our IP VPN backbone, as well. Very good. We established a secondary MPLS VPN Service Provider to reach ten of our offices, which represent about 85% of the user base as a secondary Service Provider, as I mentioned earlier. Some of the key features of this were the notion of the unmanaged service, as well as the transparent service. So for unmanaged, we reached a co-management agreement with our Service Provider, that we would control the configurations, as I mentioned earlier. With that element of Service Provider being able to reach our device and ask it questions and receive information about the status. From a transparency perspective, all three of those elements I talked about earlier, the routing protocol, the Multicast and the QoS were addressed, such that we connect to the network using the same routing protocol we used in ours, EIGRP. We used native Multicast to connect to them. In fact, we were the first Enterprise large scale network to use Multicast VPN service in a production capacity. As well as the QoS, we were able to preserve the transparency of not only our five class model, but the seven class model that followed onto it. And so the individual offerings from the Service Providers did not have to be customized. Our network does not have to be customized to those Service Provider's offerings.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE EUROPE WAN FY03/04 MPLS VPN

This is a map that you see in the slide here that talks about how our network is connected in Europe. And what you'll see around the outside of this cloud are our individual offices that are connected using the Cisco ISRs, the 3700 series routers and 3800 series routers, to our primary Service Provider, Service Provider 1. And some of our offices as you can see connect to Service Provider 2 in the center of that network. So each router on the edge is Cisco owned Cisco managed 3700s or 3800s? And the routers in the ring touching the edge of the outer cloud are Service Provider? Correct, those are the Service Provider edge routers as well as the inner ring of them is our second Service Providers. I see, thank you.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE MPLS IP/VPN BENEFITS REALIZED – EMEA

So let's look at the benefits that we realized from this. First, ultimately it came down to TCO reduction. We had originally five dedicated WAN staff, the Wide Area Network operation staff, dedicated to watching after our network in EMEA. We were able to take that staff and integrate it with the rest of our network operations staff. And the overall effort put into the Wide Area Network was reduced by about 20% to 30%. So we reduced our overhead 20% to 30% because we offloaded a lot of that planning into the Service Provider cloud. So the Service Provider would have to pick up the cost of doing that. How did that affect pricing? Well, not necessarily planning, but actually the operational element of it. Because we are talking about an IP network now and the Service Provider is operating at the level of an IP

network, as opposed to individual circuits. Ah, true. Their OSS systems, their operational support systems, are built in such a way that they are monitoring the IP network. So in a traditional network, the Enterprise monitors the IP level. The Service Provider may monitor the individual links underneath. In this particular case, the Service Provider takes and plays a part in monitoring the IP portion of the network, and many times resolves issues before we even see them in a traditional model. Mm-hm, okay. Our capacity planning efforts were reduced, as I mentioned, because we no longer had to look at per circuit, but rather per site bandwidth. Our network engineering efforts were reduced as Cisco leveraged the Service Provider's planning process. And the troubleshooting efforts were reduced as Cisco leverages that Service Provider process, as I mentioned a little bit earlier. Now the impact to pricing, which you asked about a little bit earlier, was that originally that was one of the reasons that you saw a price increase versus a traditional network. But at the same time, you have to balance that, and that MPLS VPN is a shared network, like traditional Frame Relay. It's not a time division multiplexed network. It's a statistical multiplexed network. And so in general, what that results in is that because the bandwidth is shared, you compare all the factors and at the end of the day it ends up being less expensive for Enterprises to use this versus the traditional TDM oriented networks. Impressive. Also, our network availability improves and I think this is another important element. How'd that happen? Well, what happened was the average daily downtime went down from six to two minutes for our sites through the proactive co-management of the service. This is what I was talking about earlier with OSS. In a traditional model, many of the Service Providers wait for you to call them to report a circuit. Now they are proactively monitoring the IP layer. They're able to go in and do that. Now I should note, this depends upon not just the MPLS VPN element, but the MPLS VPN element plus the investment in your Service Provider's operational and support systems. That is the primary requirement for these types of services, that they really begin monitoring that more and providing instant reaction to those types of network failures, as opposed to waiting for the customer to call. So it is the proactive element in addition. Network reliability has improved because the number of our priority one cases, or the critical cases for our infrastructure, had decreased. Very nice.

CISCO'S INTELLIGENT NETWORK INFRASTRUCTURE MPLS IP/VPN LESSONS LEARNED – EMEA

Given nearly equivalent bandwidth, and Cisco's bandwidth requirements, we saw that raw bandwidth costs between IP VPN and the circuit costs, at the time we did this were within 10% of each other. So that really shows how we had worked from the early days where it was a full mesh pricing model to a model where it was nearly equivalent. And nowadays it's actually a savings to go with MPLS VPN. We were the first major Enterprise implementation of IP VPN on this scale; I should say one of them. There were several others that took place, as well. And so for that, the engineering effort that we put into it as an early adopter was slightly larger than you would find in Enterprise using today to attach to the service. That's an important lesson or result that we have to put forward. We noted that IP VPN troubleshooting can be slightly more complex. And this was also in the earlier days of the service where having the visibility into the network and having a quote unquote green Service Provider, as well as our experience in it, we found it helped if the Enterprise understood the foundations of MPLS VPN. But it should be noted, the Enterprise doesn't have to know anything about MPLS to use these services. Why not? Because they connect to it at an IP level only and the MPLS is only found within the Service Provider network. So it's transparent, the MPLS level of the services is transparent. It is transparent to the Enterprise. Interesting. And finally, the migration plan is important to ensure that you have continued service throughout the transition. And where this is important is your site selection and how you bring sites onto the network, in order to minimize your overlapped costs of having both networks in place at once, as well as getting sites cut over to it. That type of planning is critical and for us it was let's establish the hubs onto this network first. And we understood those hubs would be overlapping throughout the transition. And then begin cutting over the edge offices and doing them in such a way that we could begin to peel back our other infrastructure. As we implemented, turn that off and generate the, close down that overlap. Makes sense. Going back just a second, you mentioned IP VPN troubleshooting can be slightly more complex. Is that still true or is that from the early days? I think that one has to admit that it is slightly more complex because you have an interaction at the IP level between Service Providers, and between the Enterprise IT teams. That said, the gap is closing on the knowledge that's necessary, right. Service Providers are more in tune. Their support systems are much better than they were when we started because they weren't very manual. And so we have had a significant change in the gap that was present. We still find that it is a good thing for our engineers to understand the basics of the technology underneath. But they don't have to be a CCIE in MPLS in order to do that. Very good.

Q AND A

So that brings us to the questions and answers I know that have been collected and I'll turn it over to you. So, yes, you presented this kind of information to customers and to IT people several times. And in the process we've collected some of the questions that have been asked in previous sessions. We do have a little bit of time left, so I'd like to take this time, since we have Craig here, to answer a few of them. So the first question was --what are some of the major benefits of our MPLS VPN network or using MPLS VPN for our network in Europe? Why is it better than using a SONET and leased line network, which is what we had before? Well, in general, it was the two categories that I spoke of. It was the reduction in the amount of effort we have to spend in managing the core of the network, plus the enhanced services that can go on top of that, the IP PSTN network. Okay, very good. Somebody had pointed out that while we have an MPLS VPN in Europe, we use a SONET hierarchical SONET infrastructure in the United States. Why doesn't Cisco build a similar MPLS VPN network in the US? Well, it's actually a matter of managing contracts, managing expiration dates and managing priorities. We have MPLS VPNs as our direction. We have them throughout EMEA as well as in Australia and New Zealand. In the United States, we have just one year prior to the MPLS VPN network that we built in Europe. We had gone through a Wide Area Network redesign for the US. And in the US we had looked at the possibility of going to IP VPNs, but the coverage just wasn't there at that time. I see. So we established a nine hub network based on SONET, interconnections between those hubs. And then brought the satellite offices, connected them into each of those hubs, using internet Data Centers as collocation facilities for the equipment behind that network. And that has worked very, very well. The management of it is much like that of a traditional network. We're still looking at fiber paths. We're still doing the capacity planning of that core. But it's worked very well. And so as that contract expires as we get the opportunity to go back and that will transition towards an IP VPN network. Oh, I see, okay. Okay, thank you. Well, that's actually about all the questions we have, in fact, all the time we have to answer them. So thank you very much. And for the rest of you there is more information available for you.

MORE ROUTING AND SWITCHING RESOURCES

You can get more information about Cisco IT deployments by going to the Cisco IT at Work website, which is the URL at the top of the page you're looking at now. You can read those Cisco IT case studies to find out about what sorts of technologies we've deployed, what benefits we've gained from those deployments, what lessons we've learned form them. And there are also some operational practices and some presentations about the technologies and about the business to help you learn more. You can also find more information on some of our other IP networks, some design guides, some operational practices. There's also several other documents and white papers and presentations on Cisco.com. Below that you'll see a toll-free number where you can call for more information or to place an order. And you can also order Cisco resources on the web from the URL at the bottom of that page.

THANK YOU FOR WATCHING

So I'd like to thank all of you for watching and for spending this time with us, and for being interested in what the Global Technology Seminar series is all about. We hope that you've enjoyed this show and that it has helped answer some of your questions about MPLS VPN, and how we're using it in the Wide Area Network. And I'd like to thank Craig very much for spending this time with us and for sharing with us your expertise, your experience, your enthusiasm for this technology. Thank you. Well, it was a pleasure to be here, Rich, and I'm very happy to have been able to address the topic and a lot of the concerns and questions that were presented. We hope you've enjoyed this show. We'll see you soon. And thanks for watching.

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

Printed in the USA C78-337379-00 02/06