

TelePresence Architecture and Design



Cisco TelePresence Cisco on Cisco

Julie Nordquist:

Hello, and welcome to the Cisco on Cisco seminar. I'm Julie Nordquist, program manager IT, and part of the Cisco on Cisco team. Today's presentation is the second of a four-part program on how Cisco supports virtual meetings with customers using Cisco TelePresence. In our last segment, we looked at TelePresence technology. In this segment, we'll explore TelePresence architecture and design, and you'll learn how Cisco IT has integrated this technology into its own network. In subsequent segments, we'll take a look at Cisco IT's solution deployment, as well as support and management. I'd like to introduce our Cisco IT technical expert for today's show on TelePresence architecture and design, Suresha Bhat. Suresha is a manager in emerging technologies IT. He's been with Cisco for more than eight years, and he currently manages a global deployment of more than 110 TelePresence solutions to Cisco offices. Welcome, Suresha.

Suresha Bhat:

Thank you, Julie. It's great to be here. Our key objective with the architecture and design was to use all parts of the TelePresence product, and then integrate that with the existing infrastructure components, and the Cisco network.

Julie Nordquist:

So sit back and enjoy as we explore Cisco IT's TelePresence architecture and design. Let's take a look at today's segment.

Agenda

Suresha Bhat:

Thank you. Today we're going to cover the architecture components of the TelePresence solution, and also the impact and the design of network.

TelePresence Technology Overview

Suresha Bhat:

As you can see, TelePresence is basically an IP telephony solution. It's an application that plugs into the network. So the key component here is really to work with the IP-based phones, and the call processing system, such as Call Manager in the networking segment. And also, we will look at how simple it is to integrate the solution on the Call Manager network, how to look at your call manager infrastructure, and enable this as an IP telephony solution.

TelePresence Technology Overview

Suresha Bhat: Key component of designing a solution for Cisco is that when we started doing the design in the beginning, it was a brand new product, we were still understanding the product, and we were also trying to understand how we would integrate this into the existing infrastructure we have. One of the key components is the group ware. Today, in Cisco, we use Microsoft Exchange. So we worked with the TSBU to integrate the whole Cisco TelePresence manager solution with Exchange and Active Directory to enable the single click call operation. The ultimate value is really how simple it is to use the system that we saw in the last segment, and also to create the whole service around TelePresence. We are talking about multiple solutions, multiple technologies, and also about 80 or so TelePresence end-points deployed at this time. So managing the day-to-day activities, call management, the support, the service, it's very critical. So architecture and design should take all of this into consideration as well.

Julie Nordquist: How many different types of solutions are there?

Suresha Bhat: Essentially, there are two types of endpoints, the 3000 and 1000, and then we have the Cisco TelePresence manager that sits in the data center for your network to manage the endpoints, then there is the Outlook integrations piece, where Active Directory as the room repository, and then the whole integration between them, and of course the network is the main component as well. So these are all the different components of the solution that we need to put together to build the Cisco solution.

Julie Nordquist: Great, thank you.

TelePresence Technology Overview

Suresha Bhat: Welcome. So let's look at what are the significant approaches we took to enable the solution? First of all, for architecting this kind of solution, we need to have a program mentality, and it's not just a patch the components together and put on your network, because you're going to load that network very heavily, you're going to have quality of service. And also, we need to target this as an enterprise solution for enterprise customers. We have so many other enterprise applications such as our ordering tools, our voice communications on the network, so we do not want to impact any of those, as well. TelePresence usually a priority three solution, from in terms of business criticality, but a meeting on it is a business critical solution, business critical meeting. So it has to have a severity one on that kind of a call, so you need to be able to solve that problem or get the solution, get the answer, immediately. So implementing that kind of architecture is very important from a program perspective, and also, we wanted to create AutoEye based on this deployment. So whatever we did, we took the program approach where we said, "Okay, if we did this architecture, is it going to generate the AutoEye we have earmarked for this at the end of deployment?" So a very critical thing to consider when you do a design or architecture is to look at what is the return on investment on this kind of solution? That also forces us to think in terms of efficiency and optimization from day one, and also the quality instead of, actually, "Okay, let's put the solution together and then look at tweaking it." We didn't take that approach. We didn't take that approach. We started thinking about how do we enable this from the day one itself.

Cisco TelePresence Architecture

Suresha Bhat: So let me now introduce you to the whole – all our architecture components. I mentioned them at a high level, now I will drill down into each of these major components and dependencies that needs to be ready. So if you're trying to launch TelePresence solution within your

environment and architecture, you got to look at all of these components.

Cisco TelePresence Architecture

Suresha Bhat: The first and foremost one is the Call Manager. As we discussed earlier, TelePresence solutions act like an IP telephone, and they look like an IP telephone to the Call Manager. The basic specification for the Call Manager selection is it should be at least 5.04 and above. We have – Cisco has released the 5.1 release, and it's going to release the 5.2 very soon. All those versions work perfectly with the TelePresence environment.

Julie Nordquist: So there's no downward compatibility?

Suresha Bhat: Not below the 5.04 version. You definitely need several enhancements we have in that product to be able to take advantage of the video solutions.

Julie Nordquist: How many call managers did you use in your deployment?

Suresha Bhat: Interesting question. Today we have used only one call manager with redundancy in our deployment. There is a particular reason for that. TelePresence, at this time, only supports one call manager cluster from our single click scheduling perspective. What was really important was one of our architectural and design considerations was we should have the single click in operation at Cisco. You can, actually, make calls without the TelePresence manager or the single click option, but it would be direct calling another site. But we wanted to take the whole scheduling, and also trying to figure out how to call another site out of the picture, and make it you walk into the room, you have the meeting on your phone, click and you're in the call. For that reason, we needed to connect all of our endpoints into one call manager cluster. So what we did was we chose one in our San Jose location, which just happens to be the corporate headquarters, and we dedicated one of the clusters, which was in production, and we upgraded that cluster to 5.04. Now, there's a very important point there. That is, you don't have to upgrade all of your infrastructure to the 5.04 release, because you just have selecting one and then upgrading it. If you have any other voice devices on that, my recommendation is to just dedicate that one to TelePresence and take the rest of the devices off it. It's not a problem, but we just want to maintain consistency, and it also helps in separation of – other IP device issues with TelePresence.

Julie Nordquist: So using just one cluster, is it still scalable and resilient?

Suresha Bhat: Absolutely. We now have close to 80+ solutions, 80+ TelePresence endpoints on that cluster. It scales all the way up to 150 and 300, about that is fine, and also it helps us a lot because now we can look at one cluster and manage it for TelePresence, and very soon BU has the road map to release multi-cluster support for TelePresence. One of the reasons that we did also help us in terms of adoption, that is whether it is in U.S. or Russia or India or in any other place, all TelePresence equipments have one standard numbering scheme. So, for Cisco, it is like 8424, and then like the last four digits. So every knows the 424 is now TelePresence. So maybe not a bad thing in the beginning when we are looking for a high level of adoption, and people do now remember the numbers to call, and they walk in and make a lot of ad hoc calls. It's not a great thing from an entitlement perspective, but adoption, it can be better, and from a support standpoint, we have a great support team that manages device services that includes the Call Manager. So it's very easy to support that because it's in one place, and also –

Julie Nordquist: I'm not going to ask you about entitlement right now because I know

we're going to cover it a little bit later.

Suresha Bhat:

Yeah.

Julie Nordquist:

I'll keep that in mind to talk to you a little bit more about that later.

Suresha Bhat:

Yep, we will drill down on that one as well. So Call Manager is a significant component of the whole solution, as it connects and it enables the calling.

Julie Nordquist:

Terrific.

Cisco TelePresence Architecture

Suresha Bhat:

Next one is the actual infrastructure, the whole LAN readiness. As you can see from the slide, I have got couple of key points in there. LAN readiness, and also quality of service and call admission control. Now, we talked about at a high level on quality of service, and impact of that of that on TelePresence. I will drill down on that in my later slides, but readiness of LAN, and LAN is really critical. It's not just about having the bandwidth between your sites where you're going to deploy TelePresence, it is also what kind of infrastructure you have, what kind of switches, routers you have in your edge and core to enable the TelePresence to go through without any network drop or call drop. So it's very important to design and look at all those considerations from an architecture standpoint. We will be drilling down on this particular aspect in our later slides, so I'll highlight one key point here. You see call admission control. Let me explain it a little bit better.

Call admission control is about how many calls to allow in a given bandwidth or a given quality of service. Let's say there are three TelePresence calls going on and we have enough bandwidth to cover that, and the fourth call comes in. Now, we can allow that call and have all the quality – the quality of all the calls drop, or we can restrict that caller, and let the quality of the other calls sustain. Today at Cisco, we do not have that implemented even though some of the solutions are available.

Julie Nordquist:

Call admission control?

Suresha Bhat:

Absolutely. We do have sufficient bandwidth to support multiple calling, and by placing the location, by choosing the locations for TelePresence deployment wisely and connecting them to through this San Jose hub, and also our Cisco network backbone has been upgraded to support the different calling and number of deployments we have. It's really a Cisco unique thing, because we have so many TelePresence deployed, and by the time customers get ready to deploy that much, definitely there will be call admission control. That we can talk about in a later seminar.

Julie Nordquist:

Great.

Cisco TelePresence Architecture

Suresha Bhat:

Okay. Then comes the important aspect of TelePresence, the whole room, the ambiance. We are a technology company, and focusing on the network, Call Manager was not a hard thing for us. But getting the room ready was a difficult challenge, because usually we have left this to the experts, facilities teams, and then a company would tell you they know everything about how to construct a room wisely. But with TelePresence, you've got to look at the whole overall impact on the equipment and the call as well. So it's very important that we choose a project team from facilities that coordination – that actually cross-functionally engages with the product, with the project, and also specifying the design requirements. The length, the width, the height, the lighting, the sound, acoustics, is very critical. Cisco publishes a

CTX level, what we call a Cisco TelePresence Experience level, room readiness assessment document, which tell you how to access the room, and how to get that room prepared. It's usually left with the customers to select a nice room, but a lot of recommendations are made from TSBU for that kind of selection.

- Julie Nordquist:* What does TSBU stand for?
- Suresha Bhat:* TelePresence Systems BU. Thank you for clarifying that. In addition, if you look at our web page, Cisco.com, there is tons of material available. Also are advanced services, which engage us on TelePresence deployment, will be able to work with you to specify those downloads and help you validate your readiness, along with CTX services.
- Julie Nordquist:* So Cisco IT has really been key to finding those standards and helping to have the product do what we intended it to do?
- Suresha Bhat:* Yes, you could say that. In fact, it was a very joint collaboration. We worked with the CTX team and TSBU, and we not only experimented with different settings, because we were developing those rooms, but also we were able to contribute to, okay, what is permissible and what is realistic. What is the ideal standard to strive for, what kind of positioning of a vacuum, the HVAC systems, the A/C systems would result in a better sound parameters. Those are the design considerations, really, and it was new, but it was very interesting to learn to learn all those and contribute to that part.
- Julie Nordquist:* So there can be some variation, it's just a matter of what you're – you have a certain parameters that you're trying to hit, and then you'll have a good experience with the room readiness.
- Suresha Bhat:* You're right. It depends on the environment itself. Let's say your office, your location is close to airport, then you have to have strict measures to soundproof the room. If you're in a secluded place, maybe a remote office, there is not much of ambience noise, then definitely not an issues to get that sound level. And also the positioning of corridors, sometimes some offices have kitchens next door, so that's a lot of clanking sound that you will not see during a normal day, but sometimes when there's heavy activity. So those are the things to watch out for, and we have contributed all that information to the whole room readiness assessment from TSBU.
- Julie Nordquist:* Great.

Cisco TelePresence Architecture

- Suresha Bhat:* Thank you. The next one is really the logistics. You may ask, "Why are we talking about logistics in architecture and design?" It's very key that when we plan on TelePresence deployment, logistics not only stands for shipping, equipment ordering, but it's also readiness to support these endpoints during the deployment. Architecture and design should take care of support as well. We have a _____ design process where we look at end-to-end, how we're going to get the quality, which influences day two operations as well. So we have actually, when we order product for a particular location, we order sufficient spares for right – support if any issues happen. Also, we look at all kinds of things. For example, if you're shipping a product to Europe, you need to have the European PAL adapters, and you also have to look at UPS for backup. Let's say you're having a very important call, power goes away – most companies have this kind of backup, but TelePresence we will put at least some kind of 15 minute UPS backup, because we have some critical locations like New York, London city, where we have a significant amount of financial customers, where we can't let the call go down. So we have so much of focus on logistics as well. My recommendation is that if you're developing TelePresence, look at logistics as a very key

part of the design element itself.

Julie Nordquist: In combination with this as well as the room readiness, how long does it typically take to get a room ready?

Suresha Bhat: About 30-60 days, depending on how quickly our facilities team can respond to a request. The essential elements of room readiness are basically the painting of the room in a unified – uniform color, so when it appears on screen, it is in the right color. TSBU has about six to eight recommendations for different types of color, and at Cisco we have chosen to pain them all standards, so when we are talking to each location, there is the standard experience from a quality perspective, but it is not mandatory. In addition, you need to also, sometimes – the room may be ready from an acoustics perspective, but making sure that you measure that. There is a CTX kit that actually has in some instruments for lighting and acoustic measurements, and then of course getting the right lighting installed in there. We have actually had great experience where we have brought up some rooms in about two weeks, and some times the permits for a particular place, with the government, takes a long time. So you have to plan for 30-60 days, and get the permits needed, and then put all the infrastructure in place. In some cases, you need to probably move a wall here and there, so that takes some time, but ideally, you should be able to find a corner room, seal off the walls, no windows of course, no glass, hopefully no glass there, because there's the vibrations, and then use the standard procedure that's in our document.

Julie Nordquist: Great. Thank you.

Cisco TelePresence Architecture

Suresha Bhat: Thanks. Another key aspect of this that we need to consider from an architecture and design, which I will be drilling down in a bit in the later slides, is the whole integration with Exchange or Outlook or Lotus Notes for one click scheduling. Now, that is not out only benefit. The benefit is also entitlement. The entitlement is really about identifying how many meetings you're having, how to prioritize those meetings, because when you deploy TelePresence, they're going to be a few of them, maybe eventually a lot of them, but they are limited availability resources. You have to focus on providing these resources to defesions that have a significant value proposition to you. So you have to have a mechanism where you can identify a meeting request as higher value than maybe a one on one call, or just saying hi to some friends somewhere in a remote location. So you need to have a process. So integration with Outlook is not only for single click dialing, but is also to enable measuring the value of TelePresence. That actually leads to my next slide.

Cisco TelePresence Architecture

Suresha Bhat: It's about value matrix and support infrastructure. So we got to think about all these four – all these six dependencies to really enable TelePresence from an architecture/design standpoint. So technology, ambience, and then value matrix. These are the key components of the architecture.

One-Way Latency: Phase I and II

Suresha Bhat: So let me talk about the – another key artifact of network and quality of service. You may be thinking, "What about latency?" I'm sure when I talk about networks, people immediately start thinking about latency and jitter and packet-loss. So what we did at Cisco was to assess all of our phase one and phase two rollout sites for network latency, jitter, and packet-loss. It was not an easy operation, because we had a lot of sites, but we were able to do it very, very well. Since we are Cisco, we

definitely know our network, so we were able to measure to the latency. Yes, Julie?

Julie Nordquist:

I was just – this slide is kind of small, so can you tell us what we're actually seeing here?

Suresha Bhat:

Absolutely. As you can see, it's really about one site to another site network latency in terms of milliseconds, and it's also one way. So if you look at, for example, Bangalore to – let me highlight that particular part.

One-Way Latency: Phase I and II

Suresha Bhat:

So if you look at Bangalore to San Jose or another site, there – I cannot read it very clearly, either – 182 millisecond, one-way. Now, very interestingly, TelePresence provides you a great flexibility in terms of what should be your network bandwidth to support a good call. From the TelePresence glass to the network, we have actually minimized that by about 100 milliseconds in one location, so at two endpoints, you have up to 200 milliseconds. So with overall solution capability, you can go up to 400 milliseconds roundtrip, or 200 milliseconds one-way, without any loss of quality, and that's at 1080p, highest resolution. And let's say you have some high-level latency or jitter, TelePresence solution is intelligent enough to drop down to the next available better in resolution, for example 720p, and then up to 480p. After that, it drops the call. So that is very significant part of design, is to know what your site's latency, jitter and packet-losses are, and this also helps us do another thing. It helps us to identify a critical site from a business perspective, that needs to have TelePresence, and then we can go and fix the network. May be a switch replacement, may be a line card replacement, then we can get the network latency within the permissible limits. This is significant component of network design.

Julie Nordquist:

So make sure that you have the latency there before you go and deploy it?

Suresha Bhat:

Yes, making sure that we are under the permissible latency, and also the jitter and packet-loss. The reason I talked about jitter a lot is TelePresence is all about lip-synch, and we don't want to hear the words after the person has spoken. It can be like a badly dubbed Chinese movie.

Julie Nordquist:

Right.

Suresha Bhat:

So we want to make sure that that is as realistic as possible. TelePresence endpoints, the actual codecs, do a very good job of synching up these operations, but we want to ensure that network is ready to go. Cisco Advanced Services offers a lot of services in this area, so when we engage with a customer, they pull out a lot of the insight, how to get the whole network ready as well.

Julie Nordquist:

Also, I'm looking at this slide, and I'm not seeing anything red. I don't know if it's color-coded or not, but does this mean that most of the sites on here are ready and waiting for TelePresence?

Suresha Bhat:

Excellent observation. Actually, most of these sites – actually, all of these sites have TelePresence already deployed. You're right, we did have red in the initial phases when we started out in October for sites like Bangalore to, let's say Calgary or New York, where it was like, in this case, it's _____ about 182. We had them red because our permissible limit at that time from high level observations was 150 milliseconds, but ever since that time we have improved our TelePresence products a lot, much better as well, so we now can go up to 200-250 milliseconds, one way, without any red zone, and that's significant improvement of the process. But it's a significant

improvement from an IT perspective as well, because we've been constantly going at sites which are pretty long from a network latency perspective, and have tested them.

Julie Nordquist: Great.

Types of Switches and Circuits

Suresha Bhat: I spoke about couple of the components that go into the network earlier. So these are recommended products that you put in access and core – in the edge as well as access. So I'll highlight these, and I'll talk about that from a network diagram perspective. So the recommendation from Cisco is to use 6,500 Ks of ultrar with each chaise for the routing purposes, and then also 3845 switches. They key here is that these are developed for high bandwidth, high definition video traffic. These support that kind of traffic very well. Let's say you do not have these products in your network, but Advanced Services will help you assess the existing products, existing solution you have on your network, as a customer, to say whether that will actually meet the TelePresence requirements. So we should now resume that because we have a video enabled router or a switch on a network, it will support the TelePresence traffic. As you know, high-definition video and at that quality has a different way of bursting, and most of our audience are technical, so I can talk about bursting for a couple of seconds here.

TelePresence actually has high burst in the beginning, when you initiate the call, and then as you know, it's H.264 standards are all about getting most content in the beginning, and then stabilizing that network variation, language variation. So we have a high level of video burst, and then it slows down, smoothes – and TelePresence endpoint codecs have a great way of smoothing the traffic as well, but to accommodate that high level bursting, you need buffers. You need memory buffers in the products. So these products provide that capability very well.

Julie Nordquist: So that's why you've chosen these particular products?

Suresha Bhat: Yes, and –

Julie Nordquist: What are the reasons?

Suresha Bhat: Absolutely. These are next generation products, and they support video traffic very well. They are actually avid products, they have audio/video/data integration capability, and provides you that end-to-end quality of service that you can enable as well. That's why we chose these two products. It was recommended by John Defford, by IT in Advanced Services, along with the TSBU.

Cisco TelePresence High Level Architecture

Suresha Bhat: So now let's look at the whole infrastructure, how we connect these components from a solutions perspective. As you can see in this diagram – I'm just showing you a high level topology of this. The TelePresence endpoints are at the bottom. They are represented by the three codecs, the plasmas, and the two IP phones. As I had mentioned earlier, you're really looking at an application like a telephone connected to your network, and once the call – of course there is the Call Manager infrastructure, and there is the Exchange and Active Directory servers with TelePresence manager. On the top you see the users' interaction with booking process, and also the Call Manager that actually manages the call pattern. And all of this is connected to your switches and routers in your network, this is access and edge, and essentially you're plugging this in to your existing network.

Julie Nordquist: I meant to ask this question the last slide, when you did pick those products, the 6500 and the 3845, what impact did that have on Cisco

IT?

- Suresha Bhat:* Good question. At many sites we had this technology already available to enable our network. At remote sites, some old sites where we had, for example, 3745 switches, we needed to upgrade them, or we needed to use clustering of 3745s to enable enough buffer capabilities. That's exactly what we're telling our customers, also. This is, really, you don't have to do too much, but you have to take the minimum steps needed to enable your network to manage the high bandwidth video fluctuations, but other than that, we didn't do too much on the network. We did have to replace some of the switches, and a great thing that happened was we were actually looking at, across the board, replacing some of these switches for our internal network as we have now new products, new capabilities, so we were actually embarking on that already, so it was just aligning with that. For TelePresence we didn't even have to do – actually, for example, many of our tier one sites, we call them, because of large presence, already had these capabilities. So we just had to enable quality of service. That takes about a couple of hours, and then testing it, end-to-end, and then we are ready to go.
- Julie Nordquist:* Because they're supporting other advanced technologies as well.
- Suresha Bhat:* Absolutely. Walrus, other VTA meeting place technologies. We have multi-point enabled video meeting place technology. So all these videos are going on the same network, so that was a good foundation to start with.
- Julie Nordquist:* Sounds good.
- Suresha Bhat:* Yeah. Many of our customers definitely have this kind of technology as well, so they can really look at TelePresence as taking that whole video service to the next level from a technology perspective.
- Julie Nordquist:* Makes sense.

Cisco TelePresence High Level Architecture

- Suresha Bhat:* Absolutely. Now, let me speak about the overall architecture from different tiers of perspective as well. We talked about the basic and underlying architecture and design where we connect our endpoints, now we are also connecting across globally, and as you can see, between our U.S. to – in the U.S., we have, pretty much, leased lines from the circuit providers, so we can definitely enable quality of service, but in Europe, we have MPLS. So we actually rent the bandwidth, and we needed to get quality of service enabled through the quality to the service provider, so we actually reached out to them and we got this quality of service enabled. For example, our London city office is connected by MPLS. Our Paris, Munich office is all MPLS network. So we were able to test TelePresence on that network as well. Within U.S., we have different flavors, such as DS3, which is a 45-megabit connection, or OC3, which is a 150-megabit connection, and we have E3s in Europe, which are basically 34 megabits. So when we were talking about quality of service earlier, we need to have a basic measure, how much of the bandwidth we should reserve for TelePresence traffic. So we reserved about 30% on an OC3, and on DS3 and E3 we enabled about 47%. Now we have actually about 30-40% as an average across different networks.
- Julie Nordquist:* How did you come up with that specific number?
- Suresha Bhat:* Excellent question. As you know, TelePresence 3000 takes about 15 megabits per second, and that is about 28, 25% of the bandwidth, and we give some buffer as about 5 megabits. This allows us to put another 1000 on the same location.

- Julie Nordquist:* Oh, great.
- Suresha Bhat:* That's an interesting observation from a design perspective because once you have a unit in there, when people started opting, they want one more very quickly. So it's very important to assess the existing bandwidth, also assess the traffic on that existing bandwidth, and look at providing what's not being used and prioritizing that for TelePresence, and that's what we have done as well. In sites where you have, probably, let's say, a DS3, and you are utilizing close to 80% of that, you might want to think about upgrading that. Let me share a tidbit of information. When we went to service providers to upgrade a DS3 line to OC3, we found out that it cost us less to get more bandwidth, so that was very interesting observation. So we went back and asked for all of our circuits to be replaced so we could get better bandwidth at lower cost. It's not just necessary to use only a limited available bandwidth, but you may be surprised to find that you get a lot of actual bandwidth for less money, and then you can dedicate that to TelePresence traffic. From a quality of service perspective, you can use it if the TelePresence call is not going through. So this is the majority of our network topology, and as I said, there is no magic to it. It's really enabling that quality of service across the different type of networks. So one thing I wanted to highlight here is that if you're thinking, "How is TelePresence going to work on E3, or a DS3 or a MPLS?" this slide explains that: It works on all these different flavors. We have already proved it.

Cisco QoS Strategy Overview

- Suresha Bhat:* So let me talk about quality of service strategy because a lot of people are already thinking, "Wait, quality of service is great. What type of quality of service is it? Is it PQ, class-based and what have you?" I can talk about Cisco's perspective here, because we have different type of recommendations depending on your traffic, depending on your network topology and design that can vary. So just from the Cisco perspective I'll speak about it today. We enabled class-based cuing for TelePresence. We created a separate TelePresence class in our quality of service pipe, and we gave IP precedence of five for audio and four for video. Here's the philosophy: Audio is always very important. You might lose video on a TelePresence quality service when you don't have enough bandwidth available for some reason, but your audio is at least guaranteed, and you can still continue the conversation. That's really keeping the service level up concept that I talked about. After that, you can see how Cisco divides different type of precedence levels. Of course, the control systems are the highest precedence, they get six or seven, and then we get audio which is a five, and then we have video, TelePresence video at four, and then we have classified rest of the audio, video and regular traffic at lower priority. Here is the thing; you might have a lot of video going through your traffic already. There could be – if you downloads, or it could be corporate communications, and it could be VT Advantage traffic, and what have you. When you put a high bandwidth, high value application like TelePresence, it's very essential to separate those video traffics out. In Cisco, what we did was we took all other traffic such as those downloads and etc., and we put on a scavenger class, so that it's never going to impact the TelePresence or other important corporate communication video quality. So we wanted to ensure that for TelePresence.
- Julie Nordquist:* Oh, that's great.
- Suresha Bhat:* So, essentially, from this slide our summary here is that this is how we did it, but you might have better ways to do that. So work with our Advanced Services to identify what's best for you to implement as a quality of service strategy for your company.

Generic Field Sales Office Topology

- Suresha Bhat:* So, I spoke about a couple of the edge, core, and where these things

are connected. So usually, this – if you're from networking background, you'll understand the circular things are the routers and the square ones are the switch. So we put the 3845 switch in there, and it actually has very good compatibility with the TelePresence type of bandwidth replacing application, and in addition you can also do 3745 stacks. We have certified that solution as well, and also enterprise ESC solution and engineering team recommends PQ for quality of service as well. There is another options there, depending on the quality and type of network you have, and of course the 65KE routers for the networking. That really supports TelePresence type of traffic very well.

Cisco TelePresence Global Process Flow

- Suresha Bhat:* So let me now talk about the important aspect of TelePresence, the whole entitlement and schedule management. This is the hard part, and a lot of network experts are there. They can design the network, they can get the network to a readiness state very quickly, but getting an entitlement implemented is a business challenge. So I would like to take a moment explain this business architecture and how we implemented that at Cisco. If you have actually attended any of Cisco EBCs or if you have seen the information on our website, this has been one of the significant contributions Cisco IT has done for the TelePresence deployment. As I said earlier, TelePresence is all about simply booking the conference rooms in Outlook. That simplicity comes from implementing the solution underneath it. So let me explain the process, and then I'll drill down to how we enabled it from a technology perspective
- Julie Nordquist:* Sounds good.
- Suresha Bhat:* First of all, the user actually goes to Outlook, and in Outlook we have enabled, if you are familiar with Outlook, and this could be Lotus Notes or other calendaring solutions like Google, Yahoo, which will be coming up very shortly. So if you got to Outlook, you will see conference rooms available. What we have done in our global address list is that we have created a filter for TelePresence rooms, globally, and we have theatres like Europe, Americas, and then of course emerging market, and Asia-Pack. In there, you will see all the rooms that are marked as TelePresence listed. So it is very easy now for people. They don't have to hunt down what is a TelePresence room. They select two rooms into their Outlook resources section. They add attendees to the meeting, locally and remotely, it doesn't help to – you don't have to differentiate, and then they click a button. That's all the users need to do, and they're already doing that, because they are booking conference rooms, and attending the meetings. So they book two rooms. Within five minutes – by the way, the key important factor here is the rooms should be available, and Outlook takes care of the availability, free, busy, look-up and what have you. So once the rooms are booked, the users get confirmation from TelePresence manager that those two rooms are booked and their TelePresence meeting has been confirmed.
- Julie Nordquist:* So they can't book a room if it's not available?
- Suresha Bhat:* Absolutely.
- Julie Nordquist:* That there wouldn't be duplicate, yeah –
- Suresha Bhat:* Yes, and even if they go ahead and book it, like many people do sometimes, TelePresence manager checks that availability and comes back to you and tells, "You don't have two rooms available free that can be scheduled. Looks like one of the rooms has been booked." So it gives you that messaging as well. What we have also done is that in that messaging, there is a link to our web form. Now, this is a very simple web form, you can do it by our regular web application. What we did was we took an existing service solution we were already using for

getting feedback from users, and we repurposed that to just ask some basic question about meeting. So as a user gets a confirmation about the meeting, in that link they can click and go to this entitlement form. All we ask, really, is, what is this meeting about, who are you meeting with, and is there any sales implications to this, is it travel reduction. It's a list of about eight or nine drop-down values that asks what is the meeting all about. This tells us – this gives us the value information to prioritize the meetings if there is any conflict.

So how do we resolve conflict? We have an alias within our organization, and help this process. So if somebody wants to and has to get a particular room, and there is already a meeting booked for that session, they can escalate, they can open the case, and our entitlement help desk literally looks at the value proposition for both the meetings, and then bases – and basically based on this process, Cisco executives meeting with customers to help them from resolution of an issue or closing a deal gets the highest priority, then, Cisco customer demonstrations, and then comes internal partner meetings and resolver meetings and what have you. Our internal meeting at the lowest priority doesn't mean that we don't allow our internal audience to use. I said this is only to resolve conflict, but all our rooms, all our conference rooms with TelePresence are available to all the audience in Cisco to book. So we do not actually lock it down for any particular group of people. We enable everybody to have access to it.

- Julie Nordquist:* But you do want to get the most high value uses out of it, so therefore, if there is conflict, you want to pick the highest value use for it?
- Suresha Bhat:* Yes, and that's very well summarized. That's really all there is about entitlement. Once –
- Julie Nordquist:* I imaging that's where the matrix come in, where you're actually able to take those, that information you're getting from your survey, and putting that back into understanding are you getting the value that you want to from the product.
- Suresha Bhat:* Perfect. Perfect. That's exactly what we're doing, and we do that on a weekly basis. So we summarize this matrix across all tiers, across all rooms, and across all type of segments, such as business value, and then also from group perspective, which organization. This is a matrix that people really love, and they can see which rooms are utilized how, so it also tells us something. We can put other TelePresence equipment in there. We can enable other mechanisms for increasing the value. Maybe we open up the rooms for 12 hours to 24 hours. We do that today, and that also forces you to think, "How do I support the solutions at this high utilization?" So those are the things we learn, and then we implement the solution.
- Julie Nordquist:* You just mentioned something. I'm just curious, are the rooms – how long are the rooms open?
- Suresha Bhat:* Excellent question. We open them from probably morning 6:00 in the local time zone to evening 10:00. That's official hours, but they're not locked. They're accessible by just Cisco badge access, so any employee will be, or anyone who has Cisco badges can go in, book the rooms, and utilize the rooms any time of the day, to be honest with you. But we use the normal business hours, because –
- Julie Nordquist:* Can you get support?
- Suresha Bhat:* Absolutely, you can –
- Julie Nordquist:* Past-hours, though?
- Suresha Bhat:* Yes, you can get support for the equipment and the call and what have

you, but there is a very important factor to think. Many of the locations in buildings sometimes could be leased, and they may be under control by somebody else. For example, for the ASE in the room, and also for access to the building. For example, some buildings may be secure for access at nights, so that needs to be – each company, each customer need to work out the details, but we at Cisco open them – keep them open any time because TelePresence changes the whole global work culture. People are coming at night and what have you. So, on average, if we take about ten-hour business day as an average hour for our matrix calculation, so if a room is booked for eight hours, you've got 80% utilization. That's how we calculate it.

Julie Nordquist:

Okay, interesting.

Suresha Bhat:

So once the meetings are scheduled, you are actually seeing that meeting on the phone, and then people click on it and they're in the meeting. What we also do is we actually send out a satisfaction survey to most of these meeting owners and schedulers to talk about their experience in the meeting, and also how the process worked. So we have collected so much information, so it helps us provide that matrix back to the organization.

Julie Nordquist:

And back to the BU to help make their product better.

Suresha Bhat:

Absolutely, and that's the whole intention.

Julie Nordquist:

Great.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat:

Let me take a moment and explain how this whole thing works. So if you look at the red zone, you have Exchange solution, we have the Cisco TelePresence manager, and then the Call Manager in the mix, and we'll talk about how this whole interaction works to enable that single click from a technology and design perspective.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat:

So we talked about the basic connectivity, so we have one single call manager cluster, we have multiple clusters of Exchange servers that this is a key, and then we have also Active Directory in the mix that actually manages the room information. So let me give you the simple low-down of this by couple of slides here.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat:

First of all is the TelePresence equipment, and you have – that has, the phone has the schedule display capabilities, but the codec itself can store the schedule up to 30 days. So once the schedule is pushed to the endpoint, it can store that there, it can display that schedule, and five minutes before the meeting, the meeting is activated. So if you see a meeting for 10:00 and you are there at 9:00 you cannot activate that meeting. That's very important thing. Also, if you have a private meeting, you do not want to show the details, we can do that through the scheduling process as well.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat:

Now we'll talk about the whole Call Manager and Cisco TelePresence manager infraction. As you know, the codec really makes a – when it makes a call, the call is diverted through the Call Manager, and when you click on that scheduled event, we need to know which number, what room we are calling, and how the Call Manager will route it. Routing is taken care by the Call Manager design, but we need to know what number and what room to call.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat: So now I am highlighting the Cisco TelePresence manager. This is the brains of the operation. This is the solution that takes the information from Exchange; it takes the information from Active Directory about the room, and validates it. Essentially, you have a lot of rooms in your organization, and each room has a TelePresence system. Somewhere we need to associate these two. The number, the telephone number from the endpoint, and the room name. That's a unique identifier in many cases. It could be a conference ID and what have you. Cisco uses conference IDs as the key to associate the room name. Now, these are also stored in our Active Directory, and these are also stored in our Exchange servers.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat: We have a very good implementation of Exchange for calendaring purposes, so we have multiple Exchange servers in different geographies, and they replicate across each other, so we use front-end servers as the point of contact for the CTM. That way we can actually now change – if we change the schedule, it gets replicated all over within five minutes. That's a very critical thing, because you change your mind, "Okay, I wanted to have a meeting at 9:00 and I'm going to have it at 9:30," within a minute this is replicated across and the schedule shows up on the phone. Now, you're talking about changing something in Outlook that shows up on the phone, probably in Paris or in Singapore, within a minute. That's really the key here.

Cisco TelePresence: Flow of Schedule Data

Suresha Bhat: Now the whole Active Directory interaction. At Cisco, we use Active Directory as the single source of information about our assets and resources. What Cisco TelePresence manager does is it actually takes information from the rooms. For example, let's say we booked room A and room B. It goes and checks the Active Directory to see if room A and room B valid rooms, or are they active, are they really the right TelePresence rooms, because they're identified as TelePresence rooms, and then validates the schedule. Once it validates that these are legally certified rooms, it takes the schedule, and then through the availability of the codec, if the codec is up and running at that time, it pushes the schedule to the codec at the endpoint, and then codec itself pushes that schedule on the phone. And codec also keeps the schedule for 30 days. Let's say, if for some reason, we lost the Cisco TelePresence manager for about two hours, we took it down for maintenance purposes, all the meeting that are scheduled will still happen, and then if there is any changes, once CTM comes back up, it will synchronize those schedule as well with the codec. At the same time, if you change Outlook calendaring and CTM is on but the codec is off, you're actually doing some maintenance work replacing the parts, CTM keeps that information. As soon as the codec comes back up, it pushes that schedule. So CTM is really, I would say from a networking infrastructure and design perspective, is the brain, so you want to design it well, and we have also got hard backup for CTM as well. So this is really how we enabled the whole entitlement operation and the matrix gathering.

Julie Nordquist: Was it difficult to coordinate all that, integrate it?

Suresha Bhat: No it was not, but it did take some time. As we were actually embarking on this journey it was new product if you'll remember, so we had to put some effort in there, but it was a cross-functional team that pulled it together.

Julie Nordquist: Thank you.

Suresha Bhat: That's all I have for today.

Julie Nordquist: Thanks, Suresha.

Suresha Bhat: Thank you, Julie.

To learn more about Cisco IT experience

Julie Nordquist: For additional information on Cisco TelePresence, or to learn more about other Cisco IT experiences with Cisco technologies and solutions, please visit the Cisco on Cisco website at the URL you see here now. This website gives you access to more than 100 case studies and operational best practices on a variety of Cisco IT deployments.

Cisco

Julie Nordquist: We'd like to thank our viewers for spending time with us and being interested in the Global Technology seminar series, and please be sure to view the other segments of our program on Cisco IT's experiences with TelePresence, including TelePresence technology development, as well as support and management.

We hope that you've enjoyed this program. See you soon.



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

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

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

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

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

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