

IP Telephony – Best Practices Part 2

Hello, and welcome back to Part 2 of the “Cisco on Cisco” Seminar on IP Telephony Operations Best Practices. I'm Rich Gore, Manager of the Cisco At Work Team. And of course the theme of our show today is Part 2 of the IP Telephony Best Practices at Cisco, where we're going to be continuing our overview of how Cisco IT manages its global deployment of IP Telephony. Now with us, joining us for Part 2 is Kevin O'Healey, one of our IT Engineers responsible for IP Telephony Operations at Cisco. Now Kevin has been with Cisco for three and a half years, plus about four or five days. And has been part of the IP Telephony Operations Support Team all of that time. Kevin, thank you very much for coming back for Part 2. You're welcome, Rich. I'm glad to be here today and I'm looking forward to picking up where we left off in session one and continuing on with an additional discussion on IP Telephony best practices. There's a lot more to cover, thank you. You're very welcome.

AGENDA

In session one, we discussed IP Telephony best practices as it relates to Cisco CallManager, voice quality and security. Today we're going to pick up where we left off and discuss monitoring an IP Telephony environment, supporting an IP Telephony environment, as well as considerations for a DialPlan. And then lastly, we'll discuss IP Telephony in small offices, home offices and lab environments, as well as final considerations to take into account when deploying IP Telephony.

MONITORING

Our first topic of discussion today is monitoring the IP Telephony environment. From a very basic level, you could utilize ICMP or ping test to monitor the individual components within your environment. This could include the CallManager servers or any routers or switches that you have deployed, as well as your Voice Gateways. Using ICMP or ping is going to verify that the box is powered on and is alive on the network. But it doesn't necessarily provide information as to the services that might be running on the box and if they're functioning properly. Going a little deeper, you can also utilize SNMP monitoring for the individual components on each device. For example, memory and CPU utilization on a server can be monitored. And then on a router or a switch, interface utilization could also be verified, as well as any errors that might be incrementing. And then on all of the components, both routers, switches and servers, SNMP could be used to monitor the power supplies and disk drives to notify you of any outages.

AUDIENCE QUESTION: **Q.** So is SNMP something that you would use on a regular basis? Or just when troubleshooting? Or when do you collect this information?

A. We use it on a regular basis. So we have it provisioned multiple levels of monitoring. So we do monitor all of our devices using ICMP just to notify us if the box is up or down on the network. That's the first level of reporting information. SNMP gives us the ability to go a little deeper, look at individual components on each of the devices.

AUDIENCE QUESTION: **Q.** And you're gathering this on a continuous basis, the SNMP information? **A.** Yes, yes, we do.

AUDIENCE QUESTION: **Q.** Okay, and you also mentioned something about interface utilization. What's the purpose of that? Why do you want to know what's going on, on the interface, in terms of utilization?

A. What we're looking for is to see if any of our given links are oversubscribed. So if we're passing more data than we anticipated over a given link. But then we also look for interface errors. For example, collisions or CRC errors that might indicate a bad cable or speed duplex

mismatch or any other errors that might affect the functionality of the system.

AUDIENCE COMMENT: Interesting, thank you.

MONITORING

All of the Voice Gateways and trunks within the environment should be monitored. Within Cisco, we utilize Microsoft Performance Monitor with a CallManager object to monitor the registration of our MGCP Gateways. Each of the MGCP Gateways registers to a specific CallManager and we can use Performance Monitor to verify that the device is registered. We can also use Performance Monitor to look at the utilization of the individual gateways to determine how many channels of a given gateway are currently in use. Something that's often overlooked is Voicemail port utilization. How many ports are necessary to interface from the CallManager environment into the Voicemail system? Both inbound ports, outbound ports for message notification and ports required to ... or message notification. Something to keep in mind is that if your users are utilizing Cisco Unity, where they have the option of accessing their Voicemail over the Web, this may reduce the number of inbound Voicemail ports that you need for your solution. Using the Cisco Unity tool set in the tools depot, you do have the ability to run reports and determine what the port utilization is on all of your Voicemail ports. Another monitoring suggestion or recommendation is to perform dial tone and TFTP availability testing. There are several third party applications that can do this for you. Dial tone testing is typically done by a simulated end-station that simulates an IP phone. It generates an off-hook event sent to CallManager and waits and checks to see what the delay to dial tone is, or how long it takes to receive a response from CallManager for the phone to provide dial tone. It verifies that CallManager has the necessary CPU cycles to provide dial tone or the dial tone indication. And it also checks for any latency on the network or delay that might affect a delay to dial tone. TFTP testing works in a similar fashion with a simulated end client; for instance, a simulated IP phone that generates a TFTP request to the TFTP server. Generally, this would be a phone configuration file or a phone device load. The goal is to download the file from TFTP, which is going to verify that the TFTP service is started and the TFTP application is functioning properly. Additional testing tools would be simulated calls between end-stations on the network, either simulated phones or real phones. And the goal is to look at packet loss or any jitter that might be affecting voice quality, as well as to ensure that the call setup and the call path are completed successfully. The goal is to set up automated testing tools that run on a regular basis. Depending on how large your environment this is, you might choose to run a test call, for instance, every 15 minutes or every hour. One thing to consider is, replicate the user experience. The simulated end clients are typically deployed on a PC or a laptop attached to an access layer switch. Be sure that it's attached to the same access layer switch that the end phones are also attached to. This is going to replicate any network delay that might exist. Also look at the VLAN provisioning. A laptop attached to a switch would generally be in the data VLAN which would not have the same QoS configuration as a phone on the voice VLAN.

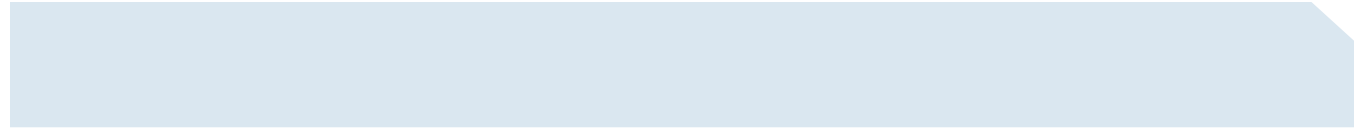
AUDIENCE QUESTION: **Q.** So Cisco is actually doing all these things? That we have simulated phones that are sending simulated calls and actually getting real dial tone messages on the network?

A. Yes, we have a series of different tools that we utilize, a combination of homegrown, as well as third party utilities that we use. The benefit is, that it gives us a real-time indication of how the systems operate. So if we have a simulated end client, it can generate an off-hook indication to CallManager, and it checks to see how long it takes to receive a play dial tone message back from CallManager.

AUDIENCE QUESTION: **Q.** How much call volume does all this stuff generate for Cisco?

A. For Cisco, not very much. But for other clients, we've had requests, how do they provision it? We know of some clients that place test calls every 15 minutes, for instance. So it's really dependent upon the individual client and what their goals are. Some clients like to place calls over their WAN to make sure that call routing on net is functioning properly. Others use it as a mechanism to test their PSTN trunks.

AUDIENCE COMMENT: Makes sense. Monitoring configuration backups also plays an important piece into IP Telephony best practices.



All of the IP Telephony components, including CallManager servers, Voice Gateways, IOS routers and LAN switches, need to have a configuration backup performed on a regular basis. For instance, a CallManager Publisher should be backed up every night. This will make it very easy to restore from backup if you were to experience a hardware failure or, for instance, a multiple drive failure. All of your Voice Gateways, routers and switches should also have a regular backup. These generally don't change as often, so perhaps once a week a backup should be performed. It will be much easier to restore a system from backup, than it will be to rebuild that system from memory.

AUDIENCE QUESTION: **Q.** So what would happen if you actually did have to rebuild a system from memory, if for some reason we hadn't been doing backups? How difficult is that and what would we have lost?

A. Sure. In the case of a CallManager Publisher, for instance, and that wasn't backed up, you would need to restore that system from memory, from scratch. That means that all of the devices, gateways, phones, conference bridges, transcoders, route patterns, everything would have to be done manually from memory. In the case of a router or a switch, this is easier to restore from memory, but it's still very time consuming.

AUDIENCE QUESTION: **Q.** And when you say from memory, you mean from the person's memory? **A.** From memory.

AUDIENCE COMMENT: Oh, god. Okay, so I understand. Now I understand the importance of backups, because you'd never get it out of my memory. It's very important that all of the systems, Voice Gateways, routers and switches, have a configuration backup. Even an older backup is helpful, but ideally, a recent backup makes it very easy to simply deploy new hardware and then restore that backup.

AUDIENCE COMMENT: I'm sold. And then as part of this, automated daily reporting should be in place. The benefit is, if a system is not configured to be backed up or if the backup were to fail, the support staff is notified and can perform a manual backup.

SUPPORT

Next, let's turn our attention to supporting the IP Telephony deployment. First, is going to be patches and upgrades. Within the environment, establish an upgrade policy. And every company is different as far as how they address the upgrade policy. We have some customers who will deploy a specific release of CallManager and utilize that release for an extended period of time, a year or longer. Other customers always want to have the latest and greatest feature sets. And so they deploy the latest version of CallManager once it becomes available.

AUDIENCE COMMENT: It sounds like Cisco. Yes. So one of the key things to do is to have an upgrade policy for when you're going to apply patches. If there is a critical patch that is released by Microsoft and is deemed to be critical by Cisco and released on cisco.com, you should have a policy within your organization for how often you're going to apply those critical patches, whether it's every 48 hours or 72 hours of a release, for instance. There are several different options when performing a remote upgrade or patch. One of them is going to be BNC or the integrated lights out board as a second option. This is going to enable you to remotely upgrade a system over the network. Of course you still have the option of utilizing the keyboard and mouse directly attached to the server. Terminal services are not supported for remote upgrades and patches on CallManager servers.

AUDIENCE QUESTION: **Q.** Don't we have IT people at every place where we have a CallManager cluster? So why would we ever do this remotely?

A. Today, we do. Today, we consolidated and collapsed our CallManager clusters back to primary locations. In the past, we had clusters in remote locations where we didn't have onsite support staff.

AUDIENCE COMMENT: Thank you. But even today, it is much easier, from a support perspective, to upgrade a system in the data center from a cube or a desktop, rather than going down to the data center and going through the additional security provisions, as far as gaining access.

AUDIENCE QUESTION: **Q.** Okay, so remotely, it could be just a building away? **A.** Yes.

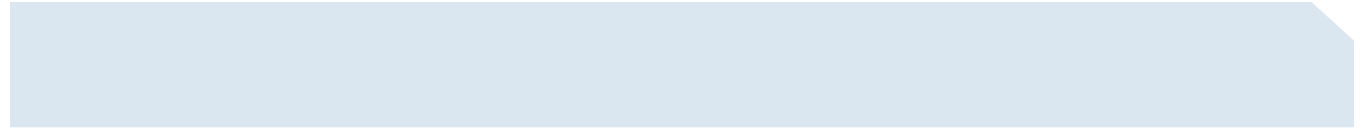
AUDIENCE COMMENT: Okay, that makes sense, thanks. It also enables us to apply our upgrades across multiple servers at the same time, because you're able to log into multiple servers, different clusters, for instance. Within your environment, depending on your hardware platform, you may choose to pull a redundant hard drive before performing a major upgrade for CallManager. This is something that within Cisco we have done in the past. And it's dependent upon the hardware platform that you have for your IP Telephony servers. What you can do is power down the server, pull one of the existing drives, power the server back on and then mirror the existing data onto an additional drive. The benefit is, if you were to encounter an upgrade failure, for instance, you can very easily restore the system to its previous configuration by simply reverting to the pulled drive. If you choose to use this process, however, test it in the lab first, so you're familiar with how to pull the drive, as well as how to restore from that pulled drive. Our recommendation, however, is going to be to always make sure that you have a recent backup of all of your systems before you perform an upgrade.

AUDIENCE COMMENT: So this pulled drive thing, that seems like it's manual intervention. You can't do that remotely. You cannot do that remotely. And generally we would do this only on a major upgrade, but having a current backup is still our first line of, in the event we were to have an upgrade failure, this is how we recover from the system.

AUDIENCE QUESTION: That makes sense. **Q.** So how often do we do the manual hard drive -- oh, so you just -- so just during very major upgrades?

A. Very major upgrades when we're going from one major release to another. But also, we utilize this in lab environments, where we can power down a production server, pull the redundant drive and then take that redundant drive to a lab environment and work through the upgrade process with our production database.

AUDIENCE COMMENT: Oh, interesting, okay. One final note, as far as the patch and upgrade process, is to utilize the change management procedures for all updates. The benefit of the change management process is that all teams are aware of changes to the environment. So, for example, if the network team makes a change to the LAN infrastructure, the voice team is made aware of that, so that we can determine if there's any impact to the voice services. There are a few different things you can do to minimize the impact during the upgrade. And these are things that we utilize within Cisco as well. The first is going to be to stagger the Cisco CallManager reboots. We utilize a specific order for when we perform upgrades on the servers. For example, we first upgrade the Publisher server, then the TFTP servers and then the primary CallManager servers. During our primary CallManager upgrades, we register all of the phones to their secondary subscribers in an alternate data center. This minimizes the impact to the client. Their phones will only register once from the primary CallManager to the secondary CallManager. Once the primary servers are upgraded and rebooted, the phones re-register back to the primary CallManagers. And then we can upgrade the secondary CallManagers without any further impact. During CallManager upgrades, you may wish to monitor the TFTP process. Internally, we utilize Microsoft Performance Monitor with the TFTP object. For major CallManager upgrades that may include a new device load for phones or gateways, it's important to monitor the TFTP performance to ensure that the TFTP service is keeping up with all of the TFTP requests. That there isn't an overload of network utilization, for instance, to the TFTP server. If you find that too many devices are requesting the new device load, you can stagger the CallManager reboots to minimize the number of devices that will request that file at the same time. Before upgrading the system, be sure to disable any alerting that you have. If your alerting is configured properly and you were to shut down a CallManager service, you should start to receive a combination of e-mails and pages notifying you of device registration and un-registration. So be sure to disable alerting before the upgrade,



but re-enable it after the upgrade. And then finally, perform post upgrade DialPlan and Voicemail testing. After every upgrade, place a series of test calls. Ideally, you will have a standardized DialPlan test that includes local calling, long distance, international, as well as inter-office calling, as well as the ability to dial other offices on net and off net. Something else that we found works very well, is before performing an upgrade, is to take a snapshot of the environment. Take a snapshot of all the registered phones, gateways, conference bridges and transcoders and this way you have something to compare to after the upgrade.

AUDIENCE QUESTION: **Q.** Take a snapshot in what fashion?

A. So, for example, using Performance Monitor with the CallManager object, we can look at all of the registered components on a given CallManager server.

AUDIENCE QUESTION: **Q.** So you, what, you printed out or save it to a file someplace? And then do the same thing afterward and compare the two?

A. And the benefit is, we now have a before and after snapshot of what our environment looks like. We can determine if an individual gateway didn't re-register, for example. Or if phones are still registered to their secondary CallManager and did not fail back to the primary CallManager. It's a final check before we turn the system back over to our production users; let's make sure everything is as it was before we started.

AUDIENCE COMMENT: Good precaution, yes. Continuing on with the ... of IP Telephony is case management.

SUPPORT

Provide an escalation path for all of the support teams. Typically, this will be in the format of a Tier 1/Tier 2/Tier 3 support path where all cases are first opened with Tier 1 and if necessary, escalated to Tier 2 or Tier 3. It's important to identify subject matter experts in all of the technologies with which IP Telephony is associated. This includes voice technologies as well as LAN and WAN. A subject matter expert in voice should be aware of the necessary QoS configurations required for high voice quality, as well as all the mechanisms for interfacing to the PSTN. By having the subject matter experts available ahead of time, it makes it very easy to escalate a critical case which requires more in depth knowledge. And then also, enforce escalations through the proper channels. All cases should be opened at Tier 1 and then, if necessary, escalate it to Tier 2 or Tier 3. The benefit of having all cases open at a certain level is that cases that have a similar symptom or a similar route cause, can be aggregated and you can start to look at what the big picture is. Identify any trends that might be helpful in the troubleshooting process. As part of the support process, there should be sufficient documentation for the environment. The only thing worse than no documentation is incorrect documentation, because then you don't know what to trust. Develop implementation and support documentation and provide it to all of your support staff. It should be in a centralized location that all of the support staff can access. An Internet Web Server works well, for example. FAQs for frequent problems should also be identified and documented. And they should be provided to both the Tier 1 support team, as well as an Internet Website where the users can access it directly before opening up a case. And then lastly, assign the correct user privileges based on job requirements. Not all members of the support staff need privilege or enable access to the IOS routers, for example. So access rights should be based on the job requirements. In the case of CallManager, you can utilize multi-level access, MLA, to provide granular access to the individual CallManager pages. For example, we can restrict Tier 1 to have access to global directory where they have the ability to reset a user password and view phone configuration pages, but not make any changes. Tier 2 would have the ability to make changes to phone configurations, but not change any of the DialPlan components, route patterns or translation patterns. And Tier 3 would have full access to all of the systems. In the case of an IOS router or Catalyst switch, make use of the user and privilege access on the individual devices to restrict what access rights individual users have.

DIAL PLAN CONSIDERATIONS

Next, let's discuss a few different DialPlan considerations. The first thing is to understand the existing DialPlan when you're considering a migration from an existing telephony environment to an IP Telephony environment. A few of the things that you should consider are current call volume and current call flow. How are calls being routed today? Are calls being routed over the PSTN? Are they being routed on net? And what is the call flow? For instance, are most of the users dialing out to the PSTN or other users within the same building?

AUDIENCE QUESTION: **Q.** So is the general goal to take your current DialPlan pre-IP Telephony and then duplicate it in an IP Telephony environment? Or are there some changes you need to make?

A. There are a couple of changes that ideally you would make. The first is to establish what your current grade of service is. What level of service is being provided based on the existing infrastructure?

AUDIENCE QUESTION: **Q.** You mean level of blocking, for instance?

A. For instance, a 1% blocking rate with a certain percentage of a retry rate. So establish a grade of service. The next is, where possible, try to aggregate your traffic. Having a more efficient utilization of long distance trunks, for example. So one method of doing this in a centralized call processing model is to have a centralized site with your long distance trunks. And then several remote sites that route calls to the PSTN first on net to the central site and then perform tail-end hop-off out to the PSTN. Are we actually doing that? In many instances, we are. The benefit is that we can negotiate lower per minute rate charges, because we have higher call volume in certain locations.

AUDIENCE COMMENT: Interesting, I didn't know that. So, there are many different things to look at. And also what you're trying to identify is, what is the busy hour and how much traffic do you have on the system? So once you've identified the busy hour, generally done by looking at the number of calls, as well as the duration of the calls, you can determine how much actual voice traffic you have running in your system. There are various models, such as the ... models, where you can determine the actual number of channels or trunks that you need.

AUDIENCE QUESTION: **Q.** You know, something I didn't hear when we migrated to IP Telephony, I did not hear our IP network engineers saying, oh, my god, we have to add, we have to double our WAN capacity because we're putting IP Telephony on or because we just put IP Telephony on. So is telephony a very large component of WAN traffic right now?

A. I will say it's a very large component and it is a significant piece.

AUDIENCE COMMENT: Important, yes. What we do, is we allocate strict priority for voice up to a certain percentage of our WAN links. And what that requires is knowledge from our perspective on the voice side as well as on the data side, how much traffic does each WAN circuit currently have and how much of that can we dedicate for voice? The goal is a more efficient utilization of the existing WAN traffic. If we have available bandwidth, we'll route the call on net. If the bandwidth isn't available, we'll route the call to the PSTN.

AUDIENCE COMMENT: Ah, that makes sense. So the real goal is to more efficiently use our existing WAN traffic and, where possible, utilize a lower cost method of accessing the PSTN, utilizing .

AUDIENCE QUESTION: **Q.** Do we do a lot of PSTN overflow within our network?

A. It depends on the individual site of how much bandwidth they have. So we do have mechanisms in place where if there's insufficient bandwidth, CallManager will automatically, without any involvement from the user, reroute the call on the PSTN. And in the event that our long distance trunks are unavailable, we'll utilize local trunks for that long distance call.

AUDIENCE COMMENT: I'm just thinking, the overflow would get kind of expensive after a while and we'd want to start putting more of that onto the network. We do, and we do look at what our total utilization for our WAN links are. And as sites change, we do go back and revisit, do we need to add additional bandwidth? Or in some instances, what we found is that, we're routing more calls on net, fewer calls off net. And so we're able to actually reduce the number of trunks to the PSTN.

AUDIENCE COMMENT: Beautiful. And as Rich and I just discussed, be sure to engineer the solution to aggregate traffic and trunking together to take advantage of more efficient and cost-effective call routing, for example, utilizing tail-end and hop-off.

DIAL PLAN CONSIDERATIONS

Hello, and welcome back to Part 2 of the Cisco on Cisco -- There we go. Kevin. ..., there we go. Kevin of the ... Yes. Ah, sure, and I can change my name to Rich --

IPT IN SMALL OFFICE, HOME OFFICE, AND LABS

Next, let's discuss IP Telephony in small offices, home offices and labs and some of the best practices. For small offices, you have the option of utilizing centralized call processing where a single CallManager cluster is located in a campus location or a central location. And then several remote offices are distributed and the phones registered to CallManager in the central site over the WAN. You can utilize the survivable remote site telephony gateways to provide redundancy for the phones. In the event of a WAN failure, where the phone loses registration with CallManager, the phone will automatically re-register to the SRST Gateway. And there are a few different devices which can provide SRST Gateway capabilities, either IOS routers or the Communications Media module for the Catalyst switches. As part of SRST deployment, be sure to verify SRST fallback. For instance, implement an access control list to prevent the phone keep-alives to CallManager. The phone will lose connectivity to CallManager and will register to the SRST Gateway. Verify this process. Then also, while in fallback mode, test the DialPlan, local, long distance and international dialing, Voicemail functionality and then also 9-1-1 dialing. Make sure that the users have access to emergency services.

AUDIENCE QUESTION: So we use SRST in all of our remote sites, since we have our CallManagers clustered in very large hubs. **Q.** Do we, on a routine basis, actually test, actually failover the WAN link or I guess block the calls to the CallManager and use -- how often do we do that? Is that once a year?

A. When we first deploy SRST in our remote sites, we take that opportunity to make sure everything is functioning properly. Then anytime that we do changes to the remote network we take that opportunity to test again. And so if we're performing standardization across all of our sites, for example, we utilize that opportunity, test the DialPlan, verify SRST fallback.

AUDIENCE QUESTION: **Q.** Has it ever actually taken place in real life where we've needed SRST fallback where the link to the CallManager has gone down?

A. We've had a few instances. We're very fortunate and that for the most part we have stable WAN links. But there have been a few instances where we have transitioned into fallback mode, we can monitor the system, we can go in and look at the log, see the phone registrations and also verify dialing in that regard as well.

AUDIENCE QUESTION: **Q.** How did it work? **A.** Very well, in fact.

AUDIENCE COMMENT: Oh, good, okay. The way that we have our SRST configured, is we have local DSP resources for conferencing and transcoding. And then we also have our Voice Gateways, our connectivity to the PSTN on the SRST Gateways. So in the event that we

lose connectivity to CallManager, the phones at the remote sites still have full functionality inbound and outbound to the PSTN.

AUDIENCE COMMENT: Lovely. As part of the SRST configuration, you should also distribute your Voice Gateways across multiple devices. So ideally, multiple WAN Gateways to which the PSTN Gateways are attached. In this way, if there were a power failure or a hardware failure affecting one device, the alternate device still has connectivity to the PSTN for inbound and outbound connectivity. For smaller sites where you only have a single PRI, you can implement FXO or INB backup on the alternate device. Essentially dedicated analog lines, so you still have the ability to place outbound calls if the PRI were to go down or if the device to which the PRI is terminated were to go down. Also, provide shared resources in addition to local resources for redundancy. So a combination of DSPs for conference bridges, as well as transcoders, can be deployed locally on the SRST Gateway or on an alternate gateway local to the site. And then additional DSP resources can be available in the centralized location. If the DSPs at the local site are fully consumed, the user's phone can take advantage of the DSPs in the central site to complete that conference bridge. And also, pay careful attention when developing the DialPlan for the remote sites. Be aware of sites specific dialing requirements, for instance, seven digits versus 10 digit local dialing. This will be dependent upon where in the country you operate and which service providers you utilize for accessing the PSTN. The key is to be aware of the differences in dialing patterns before deploying a site. And then for a centralized call processing deployment, you also have the option of implementing automated alternate routing, AAR, to protect against an out of bandwidth condition. If there's insufficient bandwidth from a call from one office on a CallManager cluster to another office on the same CallManager cluster, CallManager will automatically reroute that call through the PSTN without any involvement from the user.

AUDIENCE COMMENT: Very nice. It works very well.

IPT IN SMALL OFFICES, HOME OFFICES, AND LABS

When deploying IP Telephony for the home user, there are a few different things to consider. First, is going to be QoS on the home router. Within the user's home network, prioritize voice traffic over data traffic. In this way, the user has the ability to use their phone, while still accessing the Web or e-mail and their voice traffic has priority. Keep in mind, that there is not priority or QoS over the Internet for voice traffic. So this is really going to be helpful on the user's home network, but it is still possible that they might experience packet loss or delay within the Internet itself. Broadband Internet access for home users works best, ideally, with a minimum of 256K uplink and 768K downlink. The benefit is going to be, the user can have one or more phone conversations at the same time accessing the Web and e-mail. A hardware VPN device with encryption works well. Another option is going to be a software VPN. So, for instance, many of your employees who travel may choose to utilize a software VPN on their laptop with Cisco IP Communicator. This also works well for users who are working from a hotel or a remote location. And then implement a distributed VPN infrastructure with multiple points of access. And, ideally, less than 200 milliseconds of roundtrip delay from the user's home office back to the corporate backbone. This is going to ensure that they have sufficient delay to provide high voice quality. And then lastly, standardize on a codec selection for all remote users. Typically this is going to be a low bandwidth codec, such as G729. This is going to ensure that they have sufficient bandwidth on their lower speed home connection to the Internet, while still providing voice quality.

IPT IN SMALL OFFICES, HOME OFFICES, AND LABS

When deploying IP Telephony in the lab, a few of the things to consider would be IP addressing. All of the IP addressing in labs should utilize RFC 1918 address space. This is going to be private address space that is not publicly routable. First, this will prevent external access from your lab environment out into the Internet. But another benefit is that it's going to preserve your public address space for your production devices. And lastly, it makes it very easy for you to identify private address space, utilize the access control lists to restrict where that lab address space has access within your environment. All devices within the lab environment should utilize a proxy server to go to the Internet. This is going to simplify the process of identifying where lab users access the Internet, what sites they access. As well as simplifying the process of maintaining security. Access control lists should be implemented to restrict which traffic can be sourced from

the lab. And all traffic within a lab or all address space should be registered address space.

AUDIENCE QUESTION: **Q.** What do you mean? So you're using ACLs to restrict certain calls from getting out of the lab? What types of calls would you not want to leave the lab space?

A. So, for example, all of our lab space is registered address space. And so if I, as a user, register a certain address space for a lab, but then I go and add additional private address space within my lab that should not be routable on the production environment. It should only be registered address space that's permitted to access production.

AUDIENCE COMMENT: Okay, that makes sense, thanks. And the benefit is that it makes it very easy for the lab administrators to restrict access or shut down my lab address space if I were infected with a virus, for instance. And then for interfacing from a lab environment to the production environment, implement only static routing. The goal is to eliminate any problems in a production environment as a result of misconfiguration for a dynamic routing protocol in the lab. And then, as I just mentioned, all antivirus software on lab PCs needs to be up to date, just as it is in a production environment. And there needs to be a mechanism in place to cut off an infected device in the lab from the production network.

ADDITIONAL CONSIDERATIONS

Next is a discussion of additional considerations that we believe are important.

ADDITIONAL CONSIDERATIONS

First is going to be a QoS policy consistently applied within your environment. Create and document a QoS policy for your deployment of IP Telephony. This should include prioritized traffic, dedicated bandwidth for voice, as well as a certain amount of bandwidth dedicated to each of your data applications. Be sure to test and then implement the policy consistently. The QoS policy needs to be applied to all of the devices within the environment. It needs to be performed end-to-end to ensure high voice quality. Next is standardized configurations. Develop a standard for all of the configurations within your environment. All CallManager servers should be deployed utilizing the same standard configuration. All of the DialPlan considerations should be deployed consistently. Likewise, all of the routers, switches, as well as Voice Gateways. The key is to develop a standard and then to deploy it consistently on all of the devices. Management metrics are going to help you define a successful IP Telephony deployment. Before you can determine whether or not IP Telephony was successful, you first need to determine what measurements or statistics you're going to use to make that determination. For example, is it going to be based on device uptime or availability? Or is it going to be based on voice quality measurements? The goal is to identify and prioritize what will be measured to define success. And then develop automated systems that are going to help you determine if you're able to meet what those measurements are. And then lastly is documentation. Create an IT Operations Internet Website which can store and serve documents related to the IP Telephony environment. As I mentioned previously, the only thing worse than no documentation, is incorrect documentation. So identify an owner of the site and for each individual document, because as technologies change and as the infrastructure changes, the related documentation will also need to be updated.

Q AND A

AUDIENCE COMMENT: Well, Kevin, you've given this presentation with customers in several different times and places and we've accumulated some questions that have been important to customers. So since we have a little bit of extra time, what I'd like to do is to ask some of those questions and since we have you here, ask you to answer them. Let's go ahead and give it a shot.

AUDIENCE QUESTION: **Q.** Okay, let's start with question number one. Okay, the first question is, we're interested in putting IP phones in public areas, like conference rooms and lobbies where non-employees have access to them. What phone configuration options can we use to help secure this type of deployment?

A. This is a fairly common question. Users/customers want to deploy phones, but they're concerned about the security ramifications of putting a phone in a public area. There are a couple of different things you can do. First, is to prevent against toll fraud. Calling search spaces should be applied that restrict where that phone has the ability to call. For example, a phone in a lobby might be able to dial local and 9-1-1 and perhaps within the office, but does not have the ability to dial long distance or international.

AUDIENCE COMMENT: It makes sense. Something else to consider is the directories. Do you want users to be able to go to the directory and perform a search on all of the corporate users? Do you want someone to be able to walk into the lobby and determine the CEO's phone number, for instance? So that should be restricted within the phone configuration page. Couple of additional items would be to disable the PC port on the back of the phone.

AUDIENCE COMMENT: Ah, very good. This prevents someone from coming into the environment with a laptop, plug in their PC into the back of the phone, obtaining an IP address and accessing the corporate network. Likewise, video should also be disabled on the phone. Gratuitous ARP should be disabled to prevent a man-in-the-middle attack. And then also, PC access to the voice VLAN should be disabled. This prevents someone from attaching a packet capture to the back of the phone and being able to capture the packets of the actual voice traffic.

AUDIENCE QUESTION: Makes sense. Okay, so let's go to question number two. **Q.** What tools are available to help me to plan and size a deployment of IP Telephony? Also, how do I determine how many trunks I will need to the PSTN to support all of my users?

A. So a two part question. As far as the first question, which is what tools are available to plan and size an IP Telephony deployment? The first thing I would suggest is looking at the IP Telephony readiness assessment at cisco.com.

AUDIENCE COMMENT: Ah, okay. This is a good resource for somebody who's looking at IP Telephony and isn't sure where to they stand in so far as their network, are they ready for voice? It's a questionnaire that the administrator can provide the answers, as far as, what their environment looks like today. Do they have QoS in place? How stable is their environment? How large of a deployment are you looking at? A few other things to look at or consider, would be the voice bandwidth calculator. It let's you start to look at, what is the bandwidth requirement in order to provide voice over my WAN, for instance? The voice bandwidth calculator is very user friendly. Enables a user to select a codec, G711, for instance. And then determine a Voice over Technology, Voice over Frame, Voice over ATM, Voice over IP. And then ultimately and transport protocol, PPP, Frame Relay, etcetera. And it will determine how much bandwidth is required on a per call basis. So that's very good for determining how much WAN bandwidth you need. To address the second part of the question, which is, how do you determine how many trunks you need to the PSTN? The best place to look is going to be your call detail records on your existing environment. Determine what the busy hour is, how many calls you have in that hour, what the duration of the calls are, to determine how much actual call traffic you have. Identify a grade of service that needs to be provided. And then utilize the appropriate calculator. For instance, a ... calculator. And this will tell you exactly how many trunks are required, based on your call traffic and the grade of service you're looking for.

AUDIENCE COMMENT: It makes sense. Okay, question number three. **Q.** What performance monitor alerts do you use to monitor CallManager? And is it best to monitor everything you can or only a handful of items?

A. So I'm going to go back to one of my primary focuses and that's on standardization. So, first, every server needs to have the same base performance monitor objects monitored. For example, drive utilization. So this needs to be performed on all servers. Memory utilization. But then other servers, such as the TFTP server or a CallManager might have different objects that are being monitored.

AUDIENCE COMMENT: Ah, interesting. So in the case of TFTP, we would monitor the TFTP heartbeat. This is an indication of how long that service has been started. The heartbeat drops below a certain level, it's an indication to us that that service has restarted or the

server has restarted. CallManager will monitor the CallManager heartbeat, as well as registered devices, registered gateways, conference bridges, transcoders.

AUDIENCE QUESTION: **Q.** And the monitoring tools to do this, are they part of CallManager?

A. And CallManager heartbeat is a new one for me, so I'm assuming. Microsoft Performance Monitor ships with the operating system itself. And then when you install the Cisco IP Telephony server, the appropriate objects for that server or that service, are added to Microsoft Performance Monitor.

AUDIENCE COMMENT: How nice. So, for instance, when you install TFTP service on a server, the TFTP object is added to Microsoft Performance Monitor.

AUDIENCE COMMENT: Oh, interesting, okay, thank you. And then as far as what to monitor. Monitor everything or only a subset. What we have found, is that if you monitor everything, you generate a lot more alarms and you start to ignore them. So it's better to only monitor those components that are critical. Those things that you're very concerned about. And the lesser items, it's not as important to monitor those.

AUDIENCE COMMENT: Okay, makes sense or at least not alarm on them. Yes.

AUDIENCE QUESTION: Okay, the fourth question. **Q.** When voice and data become merged, is it best to combine the voice and the data teams as well?

A. Ah, interesting. Many of our customers ask this. Do we need to retain our Legacy TDM engineers as we migrate to IP Telephony? Or do we need to focus on data engineers? What we have found, is that you really need a combination of both skill sets. Personally, what I have seen is that our Legacy TDM engineers are very good at call routing. They understand how to route a call from one system to another. They understand trunking to the PSTN and grade of service.

AUDIENCE COMMENT: They also understand DialPlan, which is something that routing folks don't really know. It's very different. There is a significant difference between a DialPlan versus an IP routing table and how IP routes -- IP routing is done. A data engineer may have a better background in so far as QoS, as well as the network infrastructure. So what we have found works well, is a team of mixed skill sets. Over time, they learn each other skill sets. But in the short term, you have the necessary skill sets from both Legacy TDM, as well as data.

AUDIENCE QUESTION: **Q.** So combining the teams and then letting them learn from each other? **A.** Yes, works very well.

AUDIENCE QUESTION: Sounds good. So question five. **Q.** How does Cisco assign telephone numbers to employees? Do we do it by building and location or is it completely random? And how do you keep track of which numbers are available to each employee?

A. For our larger sites, it would be very nice if we could assign phone numbers based on building. Where this building has a particular range, another building has another range. From an administrative perspective, that's very easy. The down side is, is that by tomorrow, an employee from this building would have moved to another building. And so now our standard plan doesn't fit. So what we have found in a large campus environment is we deploy our telephone numbers across all of our buildings. We maintain several large blocks of numbers. And we monitor how many numbers out of each block are currently utilized. And so we try to keep it roughly equal. But any given building will have numbers from different ranges.

AUDIENCE COMMENT: Interesting, I didn't know that. The key though is to have a single source of record. One location where you have a master data base of all of the numbers and who owns them where they're in use.

AUDIENCE QUESTION: Makes sense, yes. Okay, and I'm not sure if we have any other questions. **Q.** Oh, there is another question, which is, what are some of the best practices for doing major CallManager upgrades? I remember we talked about hard drive swapping which sounded a little scary to me. So what are some of the best practices for doing major CallManager upgrades? And do you mirror drives or perform a CallManager backup prior to the upgrade?

A. Okay, so when it comes to you performing a major CallManager upgrade, there are a few different things we do. First, is we utilize the change management process. In this way, we notify everyone of the pending CallManager upgrade. And so anybody who is a stakeholder or will be affected by that change has advance notice. The next thing we do is take a current backup of our CallManager system. We'll also take a snapshot of the entire registered device on all of our CallManager servers, so we have a before and an after to compare it with. Depending on the upgrade, we may or may not choose to mirror drives. But since we do have a current backup, in the event that we did need to restore the system that is how we would do it.

AUDIENCE COMMENT: That does make sense, yes. And then following the actual upgrade, perform a DialPlan test. So DialPlan voice mail gets tested. And then we perform one final snapshot of all of the registered devices and compare it to the pre-upgrade snapshot.

AUDIENCE QUESTION: **Q.** You mentioned the DialPlan test before. Is there actually somebody who sits there after the backup or after the upgrade and then dials a local number and then dials a long distance number in?

A. Oh, okay. So what we have is a very standard script. So within our campus environment, for instance, we utilize the same numbers over and over again, so we have a consistent system test. We also utilize or have available to us, third party tools that can do that automated DialPlan testing for us. So some companies or organizations choose to go down that path. What we find, however, is that we can immediately identify if there is a problem without going any further into the test and address that problem.

AUDIENCE COMMENT: Makes sense, okay. Well that's about all the questions we have and actually all the time we have for questions today. So, Kevin, thank you very much for sharing your vast amount of experience in IP Telephony with us today. And thank you, all of you, for attending. I wanted to let you know that there are some more resources for you if you're interested in case studies or operational practices about how Cisco IT has actually

MORE IP COMMUNICATIONS RESOURCES

deployed and benefited from various IP Telephony and, in fact, other Cisco products within the Cisco IT environment. Then there's a URL you can go to take a look at a wide variety of case studies and other documents showing our best practices, our lessons learned. And some of the business benefits we've gained from these deployments. There's also, just below that on the screen, there's some URLs you can go to for more detailed operational practices and design guides and presentations on cisco.com on a variety of IP Communications topics. And if you're interested in calling and perhaps ordering some product or just getting some questions answered, there's an 800 number on the screen right below that. And if you're interested in looking for more collateral, more products that you can find from Cisco, there's a URL you can go to for Cisco marketplace. I know, for instance, that there's a CD with all the case studies and other IT material in that marketplace. So there are a lot of resources available to you. At any rate, thank you very much.

CISCO SYSTEMS

We really appreciate your attendance here. It shows your interest in IP Telephony and it matches our excitement to be able to share this information with you. So thank you very much for attending. And, Kevin, thank you very much for sharing your time with use. Rich, you're very welcome. I really enjoyed this.

AUDIENCE COMMENT: Great, okay, goodbye.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

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

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

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

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

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

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

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Use the Copyright Style for the Trademark block and the Job# info