Avaya IP Telephony Implementation Guide

Communication Manager 3.1 Avaya Labs

ABSTRACT

This document gives implementation guidelines for the Avaya MultiVantage™ Communications Applications. Configurations and recommendations are given for various Avaya Media Servers and Gateways, as well as Avaya 4600 Series IP Telephones. This document also provides information on virtual local area networks (VLANs), and guidelines for configuring Avaya and Cisco networking equipment in VoIP applications.

The intent of this document is to provide training on IP telephony, and to give guidelines for implementing Avaya solutions. It is intended to supplement the product documentation, not replace it. This document covers the Avaya Communication Manager 2.2 through 3.1, and the Avaya 4600 Series IP Telephone 1.8 and later, with limited information regarding previous and future versions.

External posting: www.avaya.com.

May 2006
COMPAS ID 95180

Avaya IP Telephony Implementation Guide

Table of Contents

Trunks

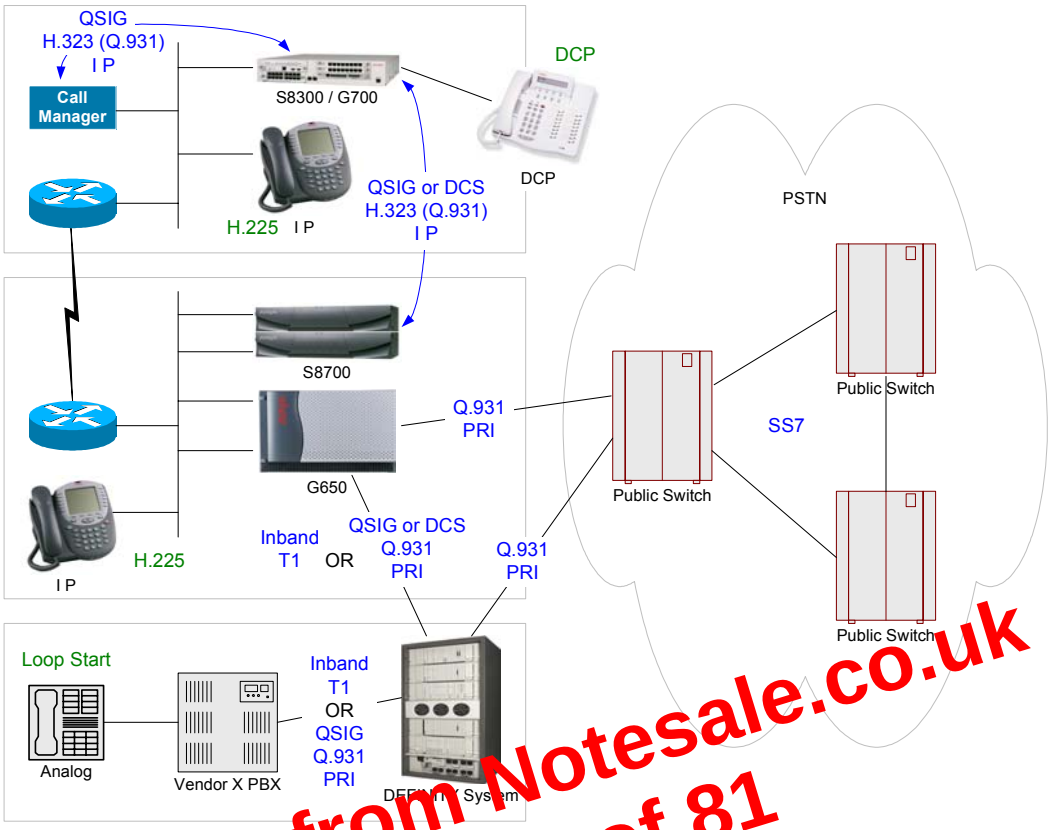

Figure 9: Trunks

This figure illustrates a broad picture to put trunks into context.

- PSTN trunks use the Signaling System 7 (SS7) signaling protocol. This protocol is not relevant to private, enterprise telephony systems.
- Private systems, such as the IP-Connect and DEFINITY servers in this illustration, commonly connect to public switches using ISDN PRI trunks with Q.931 signaling.
- Two private systems commonly connect to one another using T1 trunks with inband signaling, or ISDN PRI trunks with Q.931 signaling. This is illustrated in the trunks connecting the DEFINITY server to the IP-Connect, and to the Vendor X PBX.
- QSIG is a standard, feature-rich signaling protocol for private systems, and it "rides on top of" Q.931 as illustrated between the DEFINITY server and Vendor X PBX. DCS is the Avaya proprietary equivalent to QSIG, which also rides on top of Q.931 as illustrated between the IP-Connect and DEFINITY server.
- Gatekeepers, such as the S8700, S8300 and S8500, and Cisco Call Manager in this illustration, can connect to one another using IP trunks. The medium is IP and the signaling protocol is H.323, but Q.931 is still the fundamental component of H.323 that does the call signaling. And, as with ISDN PRI trunks, QSIG or DCS can be overlaid on top of Q.931.

QSIG is the standard signaling protocol that provides the feature-richness expected in enterprises. Generally speaking, traditional telephony systems support a broad range of QSIG features, while pure IP telephony systems support a very limited range. Due to the history and leadership of Avaya in traditional telephony, all Avaya call servers – whether traditional, IP-enabled, or pure IP – support virtually the same broad range of QSIG features.

DSCP

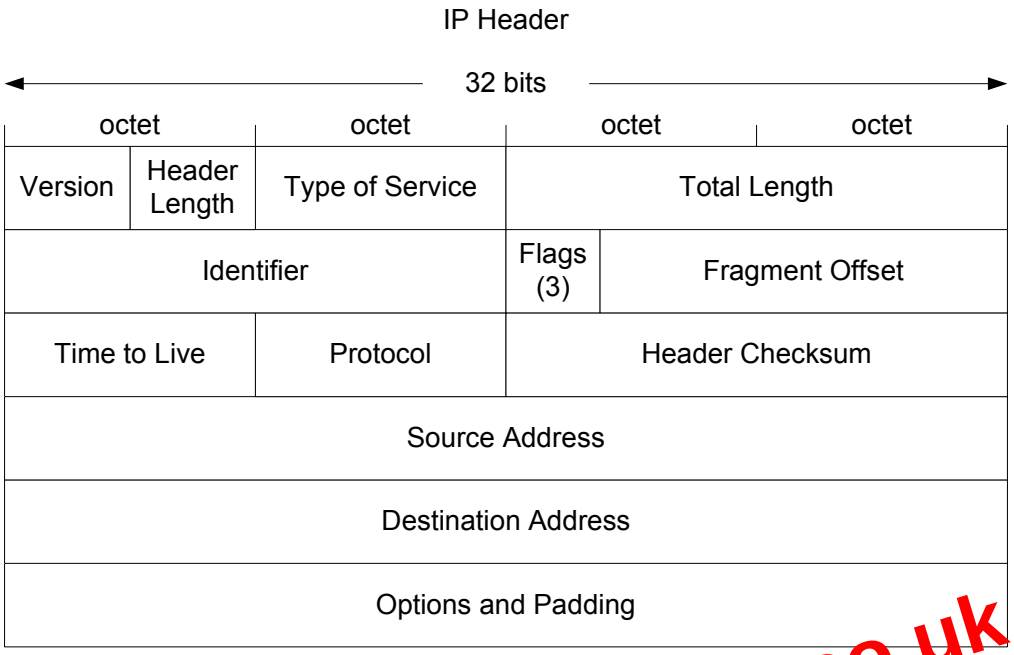

Figure 14: IP header

The figure above shows the IP header with its 8-bit Type of Service (ToS) field. The ToS field contains three IP Precedence bits and four Type of Service bits as follows:

Bits 0-2 IP Precedence	000 001 010 011 100 101 110 111	Routine Priority Immediate Flash Flash Override CRITIC/ECP Internetwork Control Network Control
Bit 3 Delay	0 1	Normal Low
Bit 4 Throughput	0 1	Normal High
Bit 5 Reliability	0 1	Normal High
Bit 6 Monetary Cost	0 1	Normal Low
Bit 7 Reserved		Always set to 0

Figure 15: Original scheme for IP ToS field

This original scheme was not widely used, and the IETF came up with a new marking method for IP called Differentiated Services Code Points (DSCP, RFC 2474/2475). DSCP utilizes the first six bits of the ToS field and ranges in value from 0 to 63. The following figure shows the original ToS scheme and DSCP in relation to the eight bits of the ToS field.

8-bit Type of Service field							
IP Precedence bits			Type of Service bits				0
0	1	2	3	4	5	6	7
DSCP bits						0	0

WAN Link Speed	L3 Packet Size					
	64 bytes	128 bytes	256 bytes	512 bytes	1024 bytes	1500 bytes
56 kbps	9 ms	18 ms	36 ms	72 ms	144 ms	214 ms
64 kbps	8 ms	16 ms	32 ms	64 ms	128 ms	187 ms
128 kbps	4 ms	8 ms	16 ms	32 ms	64 ms	93 ms
256 kbps	2 ms	4 ms	8 ms	16 ms	32 ms	46 ms
512 kbps	1 ms	2 ms	4 ms	8 ms	16 ms	23 ms
768 kbps	640 *u*s	1.28 ms	2.56 ms	5.12 ms	10.24 ms	15 ms

Table 5: Cisco seralization delay matrix

Consult Cisco's documentation for detailed information regarding traffic shaping and LFI, and be especially careful with LFI. On one hand it reduces the serialization delay, but on the other it increases the amount of L2 overhead. This is because a single L3 packet that was once transported in a single L2 frame, is now fragmented and transported in multiple L2 frames. Configure the fragment size to be as large as possible while still allowing for acceptable voice quality.

Instead of implementing LFI, some choose to simply lower the MTU size to reduce serialization delay. Two possible reasons for this are that LFI may not be supported on a given interface, or that lowering the MTU is easier to configure. As explained in section 2.2 under the heading "L3 Fragmentation (MTU)," lowering the MTU (L3 fragmentation) is much less efficient than LFI (L2 fragmentation) because it incurs additional L3 overhead as well as additional L2 overhead. Lowering the MTU is generally not advisable and may not provide any added value, because it adds more traffic to the WAN link than LFI. The added congestion resulting from the increase in traffic may effectively negate any benefit gained from reducing serialization delay. One should have a thorough understanding of the traffic traversing the WAN link before changing the MTU.

Because of all these configuration variables, properly implementing QoS on a router is no trivial task. However, it is on the router where QoS is needed most, because most WAN circuits terminate on routers. Appendix F contains examples of implementing QoS on Cisco routers. This appendix does not contain configurations for all the topics discussed in this document, but it gives the reader a place to start.

QoS Guidelines

There is no all-inclusive set of rules regarding the implementation of QoS, because all networks and their traffic characteristics are unique. It is good practice to baseline the VoIP response (ie, voice quality) on a network without QoS, and then apply QoS as necessary. Conversely, it is very bad practice to enable multiple QoS features simultaneously, not knowing what effects, if any, each feature is introducing. If voice quality is acceptable without QoS, then the simplest design may be a wise choice. If voice quality is not acceptable, or if QoS is desired for contingencies such as unexpected traffic storms, the best place to begin implementing QoS is on the WAN link(s). Then QoS can be implemented on the LAN segments as necessary.

One caution to keep in mind about QoS is regarding the processor load on network devices. Simple routing and switching technologies have been around for many years and have advanced significantly. Packet forwarding at L2 and L3 is commonly done in hardware (Cisco calls this *fast switching* [2 p.5-18], "switching" being used as a generic term here), without heavy processor intervention. When selection criteria such as QoS and other policies are added to the routing and switching function, it inherently requires more processing resources from the network device. Many of the newer devices can handle this additional processing in hardware, resulting in maintained speed without a significant processor burden. However, to implement QoS, some devices must take a hardware function and move it to software (Cisco calls this *process switching* [2 p.5-18]). Process switching not only reduces the speed of packet forwarding, but it also adds a processor penalty that can be significant. This can result in an overall performance degradation from the network device, and even device failure.

Each network device must be examined individually to determine if enabling QoS will reduce its overall effectiveness by moving a hardware function to software, or for any other reason. Since most QoS

ip-interface

Options are **change**, **display**, and **list**. This form is used to configure individual IP boards. The first step is to associate a board Type and Slot # to a previously defined Node Name, and to give the board a Subnet Mask and default Gateway and assign it to a Network Region. For example, the board type C-LAN in slot 01A05 can be associated with the node name "c-lan_80" defined earlier. This assigns the IP address 192.168.80.10 to the C-LAN board in slot 01A05. Then the board can be given the mask 255.255.255.0 with default gateway 192.168.80.254. The board can also be assigned to network region 1.

802.1p/Q tagging for an IP board is also enabled or disabled on this form. A number (including 0) in the VLAN field indicates the VID, and it means that tagging is enabled on the board with that VID. Although most implementations where tagging is enabled should use VID 0, other VIDs are permitted as well. The letter 'n' in this column means that tagging is disabled on the board, and a blank means that tagging is not supported on the board. To properly enable L2 tagging on the C-LAN and MedPro/MR320 boards, follow the instructions in section 2.3 under the heading "Rules for 802.1p/Q Tagging."

The speed and duplex settings for an IP board are configured on this form under the Ethernet Options heading.

The TN2602AP MR320 board has a VOIP Channels parameter to indicate how many channels are active on the board. While this parameter is configurable, it is restricted by licensing. The initial licensing options are to purchase a number of boards with 80 channels each, and a number of boards with 320 channels each.

The 2602AP MR320 board can be administered for Critical Reliable Bearer.

The Shared Virtual Address is the virtual ip-address owned by the currently active MR320 and must be administered in the same subnet as the "real" ip-addresses. The duplicated MR320s also share a virtual MAC address that is automatically assigned by one of four virtual MAC tables. Each Virtual MAC table contains 64 cached AVAYA owned MAC addresses and each table can be displayed with display virtual-mac-table SAT command.

The C-LAN board parameter Number of CLAN Sockets Before Warning. This is related to the information in section 3.4, heading "C-LAN Capacity and Recommendations." This parameter only dictates when a warning is triggered and does not affect the total number of TCP sockets supported by the C-LAN. Although the recommended number of sockets on a C-LAN may be less than 400, it is advisable in many cases to wait until 400 (default value) to trigger an alarm.

The parameter Receive Buffer TCP Window Size should be left at the default value of 8320. The default value should only be changed by AVAYA Services. The Allow H.323 Endpoints and the Allow H.248 Endpoints fields are administered to allow or disallow registration of endpoints and gateways on the C-LAN. The Gatekeeper Priority parameter is used for Alternate Gatekeeper lists and is available when H.323 endpoints are allowed to register. The lower the number the greater the priority.

data-module

Options are **change**, **display**, and **list**. This form is used to assign an extension (required for call processing) to a C-LAN board, and to specify other parameters. The Extension can be any valid extension in the dial plan, and does not have to be a DID extension. The Type is Ethernet. The Port is the board slot # appended with the number 17 (ie, 01A0517). The Link number can be any available number from the output of the **display communication-interface links** command. The Name is the previously defined node name (ie, "c-lan_80").

There are network address translation (NAT) options for direct IP-IP audio. Since Avaya Communication Manager 1.3, Avaya has permitted shuffling between endpoints that are separated by NAT. NAT has been a hurdle for VoIP due to the fact that the address in the IP header is translated, but embedded IP addresses in the H.323 messages are not translated. This hurdle has been overcome to some extent with the "NAT shuffling" feature in Communication Manager, without the need for H.323-aware NAT devices. See "NAT Tutorial and Avaya Communication Manager 1.3 NAT Shuffling Feature" at www.avaya.com.

Note: In addition to the **ip-network-region** form, shuffling and hairpinning must be enabled on two other forms: the **system-parameters features** form, page 16; and the **station** form, page 2, for each station.

The RTCP monitoring feature is used with the Avaya VoIP Monitoring Manager (VMM). Enabling this feature causes the audio endpoints in this region to send periodic RTCP reports to VMM. VMM uses these reports to keep a history of audio quality for all reporting endpoints. The default server parameters are configured on the **system-parameters ip-options** form. If the default settings are not desired in any given network region, specific settings can be applied on a per region basis.

The RSVP feature requires careful integration with the IP network and must not be enabled without the supporting IP network configurations. These configurations can be cumbersome and require a significant amount of network overhead. A better call admission control (CAC) mechanism is native to Communication Manager as of 2.0 and is explained in detail in the "Avaya Communication Manager Network Region Configuration Guide" at www.avaya.com.

The H.323 Link Bounce Recovery parameters, the RSVP list on page 2 of this form, and the inter-region connectivity matrix beginning on page 3 of this form are covered in detail in a separate document. See the "Avaya Communication Manager Network Region Configuration Guide" at www.avaya.com.

Inter-Gateway Alternate Routing (IGAR) on page 2 of this form is a new feature for Communication Manager 3.0. This feature is covered in detail in the "Avaya Communication Manager Network Region Configuration Guide" at www.avaya.com. Related to IGAR is a new parameter on the **cabinet** form to assign the cabinet to a network region. The assignment of a cabinet to a network region, which is a concept new to Communication Manager 3.0, applies primarily to IGAR. It has no relation to IP boards in that cabinet, and it does not assign traditional resources attached to that cabinet, such as non-IP stations and trunks, to a network region.

ip-network-map

Options are **change** and **display**. This form is used to assign stations to Communication Manager network regions by IP address range or subnet. If a station's IP address does not fall into any of the ranges configured on this form, the station is assigned to the same network region as the gatekeeper it registers with. Whether by assignment on this form or by inheritance, it is very important to assign IP stations to the proper network region. To understand how these methods of network region assignment affect the station, see the "Avaya Communication Manager Network Region Configuration Guide" at www.avaya.com.

The VLAN column is used to send a VID to IP phones. This field should only be used if DHCP option 176 is not available. If such is the case, then two rows are required on this form: one row for the data VLAN through which the phone passes, and another row for the voice VLAN on which the phone finally resides, with both rows containing the voice VID. The resulting functionality would be as follows.

- IP phone boots and obtains address on data VLAN.

- Media Encryption: New to Avaya Communication Manager 2.1. This parameter permits media encryption between the two Avaya systems joined by this IP trunk. Selecting 'y' invokes a passphrase, and both ends of the IP trunk must have the identical passphrase. This facilitates a key exchange between the systems, which makes media encryption possible between endpoints on the two systems, as long as the **ip-codec-set** forms on both systems are configured with matching encryption options. In other words, enabling encryption on the **ip-codec-set** form permits encryption within a system. Media encryption between two systems is possible when they have compatible codec sets and encryption options, and are connected by an IP trunk with this feature enabled.
- DTMF over IP: See the section below for the **system-parameters ip-options** form.
- Calls Share IP Signaling Connection: 'y' if the far-end is an Avaya device, 'n' if it is another vendor's device. 'y' means that a single H.225 signaling connection is used for all trunk members (all calls), and 'n' means that each trunk member (each call) uses a separate signaling connection. The G150, R300, and MultiVOIP gateway require this to be set to 'y'.
- Bypass if IP Threshold Exceeded: Part of a feature commonly referred to as "TDM fallback" or "IP trunk bypass." This parameter has to do with whether or not a TDM fallback trunk is utilized when the IP network fails or performs poorly between the near-end and far-end gatekeepers. The thresholds for this fail-over are configured in the **system-parameters ip-options** form, as described in Appendix G. Appendix G is a Q&A discussion on the IP trunk bypass feature and associated issues related to IP trunks.
- Direct IP-IP Audio Connections: 'y' typically, same as with endpoints.
- IP Audio Hairpinning: 'y', unless G150s, R300s, or MultiVOIP gateways can talk across the trunk.

The LRQ Required parameter allows IP trunk availability to be determined on a per call basis. When this option is enabled a RAS-Location Request (LRQ) message is sent to the far-end gatekeeper prior to each call over the IP trunk. The far-end gatekeeper responds with a RAS-Location Confirm (LCF) message and the call proceeds. The absence of an LCF from the far-end gatekeeper indicates that the call cannot proceed. If this occurs and the near-end gatekeeper is configured with the necessary route pattern, the next preferred trunk in the route pattern is used for the call as follows.

- Send LRQ.
- Wait 2sec for LCF (1sec as of Communication Manager 3.0).
- Send LRQ.
- Wait 2sec for LCF (1sec as of Communication Manager 3.0).
- Go to next preferred trunk in route pattern (4sec total per call for Communication Manager pre-3.0; 2sec total per call as of 3.0).

The LRQ feature affects individual calls, whereas the IP trunk bypass feature affects entire IP trunks. The IP trunk bypass feature takes some time to detect a problem in the IP network and put the signaling-group into bypass state. When this happens, with the appropriate route pattern in place, it results in all calls being routed onto the next preferred trunk. The LRQ feature speeds up per call re-routes until IP trunk bypass is established, so the two features can work in conjunction.
When LRQ is enabled the near-end listen port must be 1719. This means that the far-end gatekeeper must have its far-end listen port set to 1719. If the far-end gatekeeper is an Avaya call server and also has LRQ enabled (near-end listen port is 1719), then the near-end gatekeeper must have its far-end listen port set to 1719. Also, when LRQ is enabled calls cannot share the IP signaling connection, so this parameter must be set to 'n'. Each call establishes signaling across the IP trunk after a successful LRQ/LCF exchange.
For information about IP trunking with the Cisco Call Manager, see "Avaya S8300 Media Server and Avaya S8700 Media Server Networked with Cisco Call Manager using H.323 Signaling and IP Trunk Groups" at www.avaya.com.

which is to have only two VLANs configured on a trunk port connected to an IP phone, so that broadcasts from non-essential VLANs are not permitted to bog down the link to the IP phone.

VLAN Binding Feature (P330/C360)

On the Avaya P330/C360, additional VLANs are added to a port using the VLAN binding feature. The port may be a trunk port (802.1Q tagging enabled) or an access port (no 802.1Q tagging). The port does not need to be a trunk to forward multiple VLANs, and for one application – connecting to an Avaya IP phone – it must not be a trunk (ie, do not issue the **set trunk** command). The following steps enable VLAN binding.

1. Verify that the port is configured with the desired port/native VLAN.
2. Add additional VLANs with one of the following **vlan-binding-mode** options.

Static option:	
set port vlan-binding-mode <mod/port> static	Put the port in bind-to-static mode.
set port static-vlan <mod/port> <vid>	Statically add another VLAN, in addition to the port/native VLAN.

----- OR -----

Configured option:	
set vlan <id>	Add a VLAN to the *configured* VLAN list. Type **show vlan** to see entire list.
set port vlan-binding-mode <mod/port> bind-to-configured	Apply the configured VLANs to the port and permit only those VLANs (**bind-to-all** permits all VLANs and not just the configured).

3. If the port is connected to a router or to another switch, trunking must be enabled with the command **set trunk <mod/port> dot1q**, which causes all egress frames to be tagged. However, if the port is connected to an Avaya IP phone with an attached PC, trunking must not be enabled so that none of the egress frames are tagged. This is necessary because most PCs do not understand tagged frames.

Setting the Priority without Trunking or VLAN binding (Single-VLAN Scenario)

With Avaya switches it is possible to set the L2 priority on the IP phone, even if the phone is not connected to a trunk or multi-VLAN port. That is, the Avaya switch does not need to be explicitly configured to accept priority-tagged Ethernet frames on a port with only the port/native VLAN configured. This is useful if the phone and the attached PC are on the same VLAN (same IP subnet), but the phone traffic requires higher priority. Simply enable 802.1Q tagging on the IP phone, set the priorities as desired, and set the VID to zero (0). Per the IEEE standard, a VID of zero assigns the Ethernet frame to the port/native VLAN.

Cisco switches behave differently in this scenario, depending on the hardware platforms and OS versions. Here are Avaya Labs test results with a sample of hardware platforms and OS versions.

Catalyst 6509 w/ CatOS 6.1(2)	Accepted VID zero for the native VLAN when 802.1Q trunking was enabled on the port. In this case, all but the native VLAN should be cleared off the trunk.
Catalyst 4000 w/ CatOS 6.3(3)	Would not accept VID zero for the native VLAN. Opened a case with Cisco TAC, and TAC engineer said it was a hardware problem in the 4000. Bug ID is CSCdr06231. Workaround is to enable 802.1Q trunking and tag with native VID instead of zero. Again, clear all but the native VLAN off the trunk.
Catalyst 3500XL w/ IOS 12.0(5)WC2	Accepted VID zero for the native VLAN when 802.1Q trunking was disabled on the port.
Conclusion	Note the hardware platform and OS version and consult Cisco's documentation, or call TAC.

Note that setting a L2 priority is only useful if QoS is enabled on the Eth-switch. Otherwise, the priority-tagged frames are treated no differently than clear frames.

Sample Multi-VLAN Scenario for Avaya P330 Code 3.2.8 and Cisco CatOS and IOS

Here is a sample multi-VLAN scenario. Suppose there is a Cisco router connected to a P330 switch that contains two VLANs, one for the VoIP devices and one for the PCs. To conserve ports and cabling, the PCs are connected to the phones and the phones are connected to the P330 switch.

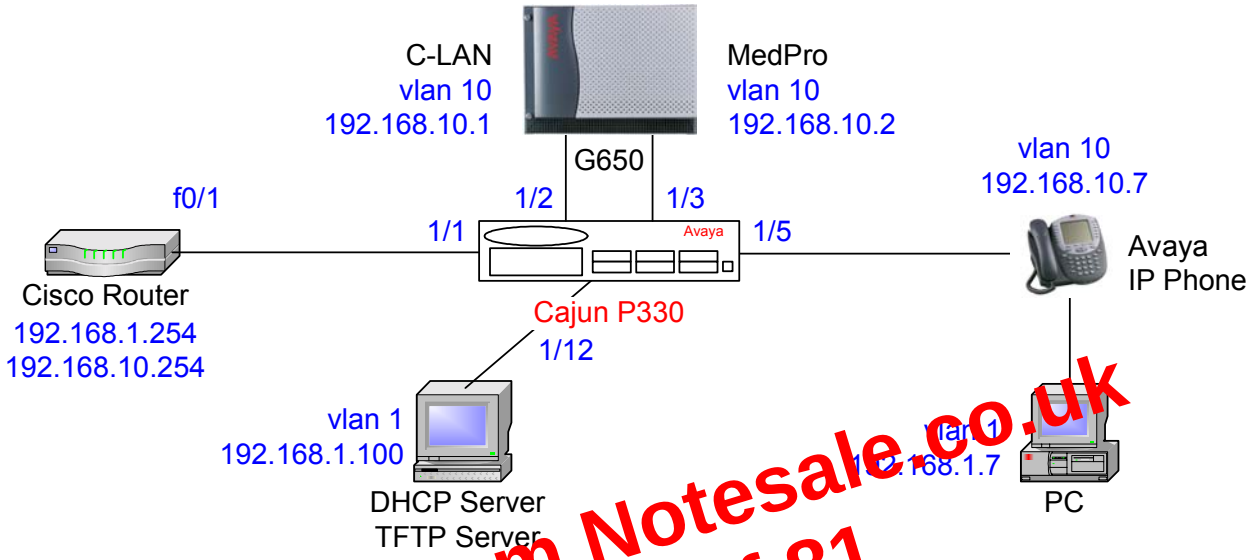

Cisco Router Configuration	
interface FastEthernet0/0 **description 802.1Q trunk interface** **!** **interface FastEthernet0/1.1** **encapsulation dot1q 1** **ip address 192.168.1.254 255.255.255.0** **!** **interface FastEthernet0/1.10** **encapsulation dot1q 10** **ip address 192.168.10.254 255.255.255.0** **ip helper-address 192.168.1.100**	To forward DHCP requests to the DHCP server.

P330/P360 configuration (bind-to-static option)	
	All ports have port/native VLAN 1 by default.
set port vlan-binding-mode 1/1 static	Port in static binding mode by default, but command shown.
set port static-vlan 1/1 10	In addition to v1, v10 statically bound to port.
set trunk 1/1 dot1q	Port connected to Cisco router is an 802.1Q trunk port.
set port spantree disable 1/1	Spanning Tree disabled at the port level.
set port vlan 10 ½	Port/native VLAN changed to 10 on this port.
set port spantree disable ½	Spanning Tree disabled at the port level.
set port level ½ 6	Port L2 (802.1p) priority set to 6.
set port vlan 10 1/3	
set port spantree disable 1/3	
set port level 1/3 6	

See section 3.5, heading "trunk-group and signaling-group" for details on the LRQ feature that applies to individual calls placed over an IP trunk. For the IP trunk as a whole the best method is the IP trunk bypass feature. In addition there is also a Maintenance Function. This function assesses the IP trunk every 15 minutes in a G3r or Linux platform, and every hour in a G3i platform. Without going into detail, the Maintenance Function determines whether the signaling group is in service or out of service. It can detect a network outage, but it does not assess network performance.

A third method was implemented as of Avaya Communication Manager 1.3. With this method a failure to set up a signaling link triggers the Maintenance Function to assess the IP trunk immediately. Assuming the failure to set up the signaling link is the result of a network outage, the Maintenance Function detects this and puts the signaling group out of service within one minute. For example, suppose there is an IP trunk between an S8700 system and an S8300/media-gateway. There is an outage in the IP network between the two systems and the S8700 discovers this after a measurement interval (IP trunk bypass feature). The S8700 puts the signaling group in bypass state and begins using the fallback TDM trunk. The S8300/media-gateway normally does not detect the outage until the next Maintenance Function cycle. However, if the S8300 attempts to place a call over the IP trunk and cannot establish a signaling link to the other end, this triggers the Maintenance Function immediately, which takes the signaling group out of service, causing the fallback TDM trunk to be used. So the S8300 detects the outage less than one minute after the first call attempt.

The scenario for severe congestion is different. In the case of severe congestion the S8700 detects the congestion and puts the signaling group in bypass state, the same as with a network outage. It then sends a message to the S8300 indicating this condition. (This message is also sent in the network outage case, but it doesn't reach the far end because of the outage.) The **status signaling-group** command at the S8700 shows the signaling group in bypass state. The same command at the S8300 shows the signaling group in far-end bypass state. In this condition both sides use the fallback TDM trunk until the S8700 puts the signaling group back into service.

Q3: As a follow-up to the previous question, what are the effects of the two sides not detecting the outage at exactly the same time?

Both sides accept incoming calls on TDM trunks, regardless of the state of IP trunks. So if side A detects an IP network outage and calls side B via the TDM trunk instead of the IP trunk, side B accepts the call. Side B continues to attempt using the IP trunk until it detects the outage, at which time it utilizes the TDM trunk for its outbound calls.

In the case of severe congestion, side A detects the congestion first, goes to bypass state, and starts using the TDM trunk. This causes side B to go to far-end bypass state and also use the TDM trunk. Eventually side B detects the congestion and goes to bypass state as well (unless the system is an S8300/media-gateway).

Q4: When the IP network recovers after an outage or severe congestion, do both sides discover this at the same time and start sending calls over the IP trunk at the same time? If not, what are the effects?

No, as with detecting the failure, detecting the recovery is also independent. But this is usually not a problem because both sides accept incoming calls on an IP trunk in bypass state. So if side A detects the IP network recovery first and calls side B while B is still in bypass state, side B accepts the call. However, the same is not true if side B is in out of service state.

The scenario for severe congestion is the same.

Q5: If the C-LAN or S8300 on one end of the IP trunk fails, does the IP trunk cover to a different C-LAN or S8300?

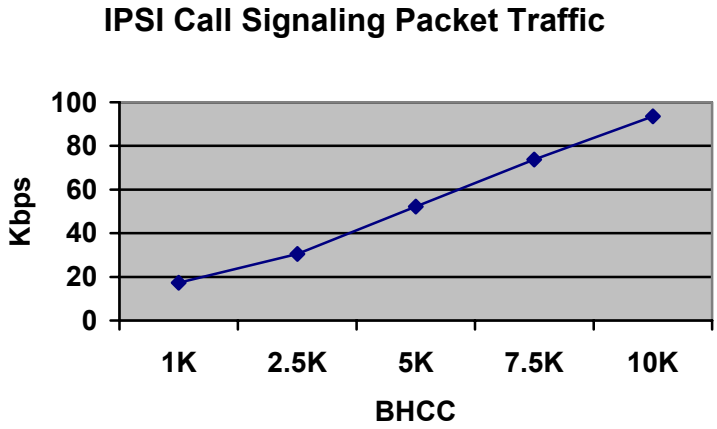

Provisioning for VoIP must include the Layer 2 overhead, which includes preambles, headers, flags, CRCs and ATM cell padding. The amount of overhead per VoIP call includes:

- Ethernet adds a 18 byte header, plus a 4 byte CRC plus an optional 4-byte 802.1Q Tag plus a 8 byte preamble for a total of up to 34 bytes per packet.
- Point-to-Point Protocol (PPP) adds 12 bytes of layer 2 overhead per packet.
- Multilink PPP adds 13 bytes per packet.
- Frame Relay adds 6 or 7 bytes per packet.
- ATM adds varying amounts of overhead depending on cell padding.
- IPSI encryption adds up-to 23 bytes (AES) for the encryption header and padding in addition to Layer 2 overhead.

IPSI bandwidth calculations should include the additional overhead on a per packet basis depending on the type of WAN link. For example, for a busy hour call completion rate of 5K calls (moderate general business traffic rate), the L2 overhead for a PPP link would be 61 PPS X 12 bytes/packet or 6.3 Kbps for additional PPP L2 overhead for a **minimum** of 58.1 Kbps. Encryption would add 23 additional bytes per packet or and additional 11.2 Kbps for a total of 69.3 Kbps.

A general rule of thumb for IPSI Control traffic bandwidth allocation is to add an additional 64Kbps of signaling bandwidth to the minimum required bandwidth in order to manage peak (burst) traffic loads and either round up or down to nearest DS0. Using the previous example of 5K busy hour calls using encrypted PPP links to control remote port networks you would guarantee 128Kbps (69.3Kbps + 64Kbps) for IPSI signaling bandwidth across the WAN link.

A standby IPSI consumes an additional 2.4 Kbps bandwidth on the standby link.