Avaya Application Solutions

IP Telephony Deployment Guide

555-245-600
Issue 2
November 2003

Notice

Every effort was made to ensure that the information in this document was complete and accurate at the time of printing. However, information is subject to change.

Warranty

Avaya Inc. provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language as well as information regarding support for this product, while under warranty, is available through the following Web site: http://www.avaya.com/support.

Preventing Toll Fraud

"Toll fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there may be a risk of toll fraud associated with your system and that, if toll fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Fraud Intervention

If you suspect that you are being victimized by toll fraud and you need technical assistance or support, in the United States and Canada, call the Technical Service Center's Toll Fraud Intervention Hotline at 1-800-643-2353.

How to Get Help

For additional support telephone numbers, go to the Avaya support Web site: http://www.avaya.com/support. If you are:

- Within the United States, click the *Escalation Management* link. Then click the appropriate link for the type of support you need.
- Outside the United States, click the *Escalation Management* link. Then click the *International Services* link that includes telephone numbers for the international Centers of Excellence.

Providing Telecommunications Security

Telecommunications security (of voice, data, and/or video communications) is the prevention of any type of intrusion to (that is, either unauthorized or malicious access to or use of) your company's telecommunications equipment by some party.

Your company's "telecommunications equipment" includes both this Avaya product and any other voice/data/video equipment that could be accessed via this Avaya product (that is, "networked equipment").

An "outside party" is anyone who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf. Whereas, a "malicious party" is anyone (including someone who may be otherwise authorized) who accesses your telecommunications equipment with either malicious or mischievous intent.

Such intrusions may be either to/through synchronous (time-multiplexed and/or circuit-based) or asynchronous (character-, message-, or packet-based) equipment or interfaces for reasons of:

- Utilization (of capabilities special to the accessed equipment)
- Theft (such as, of intellectual property, financial assets, or toll facility access)
- Eavesdropping (privacy invasions to humans)
- Mischief (troubling, but apparently innocuous, tampering)
- Harm (such as harmful tampering, data loss or alteration, regardless of motive or intent)

Be aware that there may be a risk of unauthorized intrusions associated with your system and/or its networked equipment. Also realize that, if such an intrusion should occur, it could result in a variety of losses to your company (including but not limited to, human/data privacy, intellectual property, material assets, financial resources, labor costs, and/or legal costs).

Responsibility for Your Company's Telecommunications Security

The final responsibility for securing both this system and its networked equipment rests with you - Avaya's customer system administrator, your telecommunications peers, and your managers. Base the fulfillment of your responsibility on acquired knowledge and resources from a variety of sources including but not limited to:

- Installation documents
- System administration documents
- Security documents
- Hardware-/software-based security tools
- Shared information between you and your peers
- Telecommunications security experts

To prevent intrusions to your telecommunications equipment, you and your peers should carefully program and configure:

- Your Avaya-provided telecommunications systems and their interfaces
- Your Avaya-provided software applications, as well as their underlying hardware/software platforms and interfaces
- Any other equipment networked to your Avaya products

TCP/IP Facilities

Customers may experience differences in product performance, reliability and security depending upon network configurations/design and topologies, even when the product performs as warranted.

Standards Compliance

Avaya Inc. is not responsible for any radio or television interference caused by unauthorized modifications of this equipment or the substitution or attachment of connecting cables and equipment other than those specified by Avaya Inc. The correction of interference caused by such unauthorized modifications, substitution or attachment will be the responsibility of the user. Pursuant to Part 15 of the Federal Communications Commission (FCC) Rules, the user is cautioned that changes or modifications not expressly approved by Avaya Inc. could void the user's authority to operate this equipment.

Product Safety Standards

This product complies with and conforms to the following international Product Safety standards as applicable:

Safety of Information Technology Equipment, IEC 60950, 3rd Edition including all relevant national deviations as listed in Compliance with IEC for Electrical Equipment (IECEE) CB-96A.

Safety of Information Technology Equipment, CAN/CSA-C22.2 No. 60950-00 / UL 60950, 3rd Edition

Safety Requirements for Customer Equipment, ACA Technical Standard (TS) 001 - 1997

One or more of the following Mexican national standards, as applicable: NOM 001 SCFI 1993, NOM SCFI 016 1993, NOM 019 SCFI 1998

The equipment described in this document may contain Class 1 LASER Device(s). These devices comply with the following standards:

- EN 60825-1, Edition 1.1, 1998-01
- 21 CFR 1040.10 and CFR 1040.11.

The LASER devices operate within the following parameters:

- Maximum power output: -5 dBm to -8 dBm
- Center Wavelength: 1310 nm to 1360 nm

Luokan 1 Laserlaite

Klass 1 Laser Apparat

Use of controls or adjustments or performance of procedures other than those specified herein may result in hazardous radiation exposures. Contact your Avaya representative for more laser product information.

Electromagnetic Compatibility (EMC) Standards

This product complies with and conforms to the following international EMC standards and all relevant national deviations:

Limits and Methods of Measurement of Radio Interference of Information Technology Equipment, CISPR 22:1997 and EN55022:1998.

Information Technology Equipment – Immunity Characteristics – Limits and Methods of Measurement, CISPR 24:1997 and EN55024:1998, including:

- Electrostatic Discharge (ESD) IEC 61000-4-2
- Radiated Immunity IEC 61000-4-3
- Electrical Fast Transient IEC 61000-4-4
- Lightning Effects IEC 61000-4-5
- Conducted Immunity IEC 61000-4-6
- Mains Frequency Magnetic Field IEC 61000-4-8
- Voltage Dips and Variations IEC 61000-4-11
- Powerline Harmonics IEC 61000-3-2
- Voltage Fluctuations and Flicker IEC 61000-3-3

Federal Communications Commission Statement

Part 15:

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Part 68: Answer-Supervision Signaling

Allowing this equipment to be operated in a manner that does not provide proper answer-supervision signaling is in violation of Part 68 rules. This equipment returns answer-supervision signals to the public switched network when:

- answered by the called station,
- answered by the attendant, or
- routed to a recorded announcement that can be administered by the customer premises equipment (CPE) user.

This equipment returns answer-supervision signals on all direct inward dialed (DID) calls forwarded back to the public switched telephone network. Permissible exceptions are:

- A call is unanswered.
- A busy tone is received.
- A reorder tone is received.

Avaya attests that this registered equipment is capable of providing users access to interstate providers of operator services through the use of access codes. Modification of this equipment by call aggregators to block access dialing codes is a violation of the Telephone Operator Consumers Act of 1990.

REN Number

For MCC1, SCC1, CMC1, G600, and G650 Media Gateways:

This equipment complies with Part 68 of the FCC rules. On either the rear or inside the front cover of this equipment is a label that contains, among other information, the FCC registration number, and ringer equivalence number (REN) for this equipment. If requested, this information must be provided to the telephone company.

For G350 and G700 Media Gateways:

This equipment complies with Part 68 of the FCC rules and the requirements adopted by the ACTA. On the rear of this equipment is a label that contains, among other information, a product identifier in the format US:AAAEQ##TXXXX. The digits represented by ## are the ringer equivalence number (REN) without a decimal point (for example, 03 is a REN of 0.3). If requested, this number must be provided to the telephone company.

For all media gateways:

The REN is used to determine the quantity of devices that may be connected to the telephone line. Excessive RENs on the telephone line may result in devices not ringing in response to an incoming call. In most, but not all areas, the sum of RENs should not exceed 5.0. To be certain of the number of devices that may be connected to a line, as determined by the total RENs, contact the local telephone company.

REN is not required for some types of analog or digital facilities.

Means of Connection

Connection of this equipment to the telephone network is shown in the following tables.

For MCC1, SCC1, CMC1, G600, and G650 Media Gateways:

Manufacturer's Port Identifier	FIC Code	SOC/REN/ A.S. Code	Network Jacks
Off premises station	OL13C	9.0F	RJ2GX, RJ21X, RJ11C
DID trunk	02RV2-T	0.0B	RJ2GX, RJ21X
CO trunk	02GS2	0.3A	RJ21X
	02LS2	0.3A	RJ21X
Tie trunk	TL31M	9.0F	RJ2GX
Basic Rate Interface	02IS5	6.0F, 6.0Y	RJ49C
1.544 digital interface	04DU9-BN	6.0F	RJ48C, RJ48M
	04DU9-IKN	6.0F	RJ48C, RJ48M
	04DU9-ISN	6.0F	RJ48C, RJ48M
120A4 channel service unit	04DU9-DN	6.0Y	RJ48C

For G350 and G700 Media Gateways:

Manufacturer's Port Identifier	FIC Code	SOC/REN/ A.S. Code	Network Jacks
Ground Start CO trunk	02GS2	1.0A	RJ11C
DID trunk	02RV2-T	AS.0	RJ11C
Loop Start CO trunk	02LS2	0.5A	RJ11C
1.544 digital interface	04DU9-BN	6.0Y	RJ48C
	04DU9-DN	6.0Y	RJ48C
	04DU9-IKN	6.0Y	RJ48C
	04DU9-ISN	6.0Y	RJ48C
Basic Rate Interface	02IS5	6.0F	RJ49C

For all media gateways:

If the terminal equipment (for example, the media server or media gateway) causes harm to the telephone network, the telephone company will notify you in advance that temporary discontinuance of service may be required. But if advance notice is not practical, the telephone company will notify the customer as soon as possible. Also, you will be advised of your right to file a complaint with the FCC if you believe it is necessary.

The telephone company may make changes in its facilities, equipment, operations or procedures that could affect the operation of the equipment. If this happens, the telephone company will provide advance notice in order for you to make necessary modifications to maintain uninterrupted service.

If trouble is experienced with this equipment, for repair or warranty information, please contact the Technical Service Center at 1-800-242- 2121 or contact your local Avaya representative. If the equipment is causing harm to the telephone network, the telephone company may request that you disconnect the equipment until the problem is resolved.

A plug and jack used to connect this equipment to the premises wiring and telephone network must comply with the applicable FCC Part 68 rules and requirements adopted by the ACTA. A compliant telephone cord and modular plug is provided with this product. It is designed to be connected to a compatible modular jack that is also compliant. It is recommended that repairs be performed by Avaya certified technicians.

The equipment cannot be used on public coin phone service provided by the telephone company. Connection to party line service is subject to state tariffs. Contact the state public utility commission, public service commission or corporation commission for information.

This equipment, if it uses a telephone receiver, is hearing aid compatible.

Canadian Department of Communications (DOC) Interference Information

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

This equipment meets the applicable Industry Canada Terminal Equipment Technical Specifications. This is confirmed by the registration number. The abbreviation, IC, before the registration number signifies that registration was performed based on a Declaration of Conformity indicating that Industry Canada technical specifications were met. It does not imply that Industry Canada approved the equipment.

Declarations of Conformity

United States FCC Part 68 Supplier's Declaration of Conformity (SDoC)

Avaya Inc. in the United States of America hereby certifies that the equipment described in this document and bearing a TIA TSB-168 label identification number complies with the FCC's Rules and Regulations 47 CFR Part 68, and the Administrative Council on Terminal Attachments (ACTA) adopted technical criteria.

Avaya further asserts that Avaya handset-equipped terminal equipment described in this document complies with Paragraph 68.316 of the FCC Rules and Regulations defining Hearing Aid Compatibility and is deemed compatible with hearing aids.

Copies of SDoCs signed by the Responsible Party in the U. S. can be obtained by contacting your local sales representative and are available on the following Web site: http://www.avaya.com/support.

All Avaya media servers and media gateways are compliant with FCC Part 68, but many have been registered with the FCC before the SDoC process was available. A list of all Avaya registered products may be found at: http://www.part68.org by conducting a search using "Avaya" as manufacturer.

European Union Declarations of Conformity

CE

Avaya Inc. declares that the equipment specified in this document bearing the "CE" (*Conformité Europeénne*) mark conforms to the European Union Radio and Telecommunications Terminal Equipment Directive (1999/5/EC), including the Electromagnetic Compatibility Directive (89/336/EEC) and Low Voltage Directive (73/23/EEC). This equipment has been certified to meet CTR3 Basic Rate Interface (BRI) and CTR4 Primary Rate Interface (PRI) and subsets thereof in CTR12 and CTR13, as applicable.

Copies of these Declarations of Conformity (DoCs) can be obtained by contacting your local sales representative and are available on the following Web site: http://www.avaya.com/support.

Japan

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may occur, in which case, the user may be required to take corrective actions.

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

To order copies of this and other documents:

Call: Avaya Publications Center
Voice 1.800.457.1235 or 1.207.866.6701
FAX 1.800.457.1764 or 1.207.626.7269

Write: Globalware Solutions
200 Ward Hill Avenue
Haverhill, MA 01835 USA
Attention: Avaya Account Management

E-mail: totalware@gwsmail.com

For the most current versions of documentation, go to the Avaya support Web site: http://www.avaya.com/support.

Contents

Section 1. Avaya Application Solutions product guide

Avaya Application Solutions

This chapter contains general discussions of the Avaya Application Solutions product line:

- Avaya Communication Manager
- Avaya Media Servers
- Avaya DEFINITY Servers
- Avaya Media Gateways
- Avaya Integrated Management
- Avaya communication devices
- Avaya Communication Manager applications

The next-generation Avaya Application Solutions portfolio powered by Avaya Communication Manager delivers on the promise of IP by offering a no-compromise approach to convergence in terms of reliability and functionality. “No compromise” means that Avaya allows customers to migrate to IP Telephony without compromising on features (all features are maintained or expanded), interfaces (all existing telephones and lines are supported, along with new IP Telephones, softphones, and IP trunks), or reliability. Avaya Communication Manager is the centerpiece of Avaya Application Solutions.

Communication Manager runs on a variety of Avaya Media Servers, provides control to Avaya Media Gateways and Avaya Communications Devices, and can operate in a distributed or networked call processing environment. Figure 1, Avaya Application Solutions, on page 20 and Figure 2, Communication Manager traffic flow, on page 21 summarize the Avaya Application Solutions.

Figure 1: Avaya Application Solutions

Figure 2: Communication Manager traffic flow

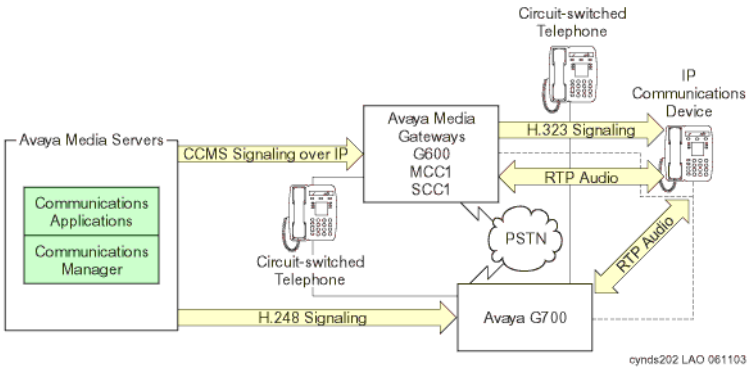

Communication Manager is the next generation of Avaya call processing software. Communication Manager is an open, scalable, highly reliable, and secure telephony application. Communication Manager operates on Avaya Media Servers, and on the existing family of DEFINITY and Avaya IP600 communication servers.

Communication Manager carries forward all the current DEFINITY capabilities, plus all the enhancements that enable enterprises to take advantage of new, distributed technologies, increased scalability, and redundancy. Communication Manager is evolved from DEFINITY software and delivers no-compromise, enterprise IP Telephony.

Avaya Media Gateways support voice traffic and signaling traffic that is routed between circuit-switched networks and packet-switched networks. The Gateways support all the applications and adjuncts that can be used with the Avaya DEFINITY Enterprise Communications Servers (DEFINITY ECS). These Gateways work with standards- based data networks and easily connect with the Public Switched Telephone Network (PSTN).

Communication Manager is extensible to IP, digital and analog telephones, and wireless business solutions. Avaya Communication Devices work with the full feature set of Communication Manager to help enterprises be more productive by providing anytime, anywhere access to maximize business continuity.

Avaya Communication Manager

Avaya Communication Manager provides user and system management functionality, intelligent call routing, application integration and extensibility, and enterprise communications networking. Communication Manager operates on Avaya Media Servers, and on the existing family of DEFINITY and Avaya IP600 communication servers. For more information on the Avaya Application Solutions related features of Communication Manager, see Call processing.

The following additional resource provides even more details on Communication Manager:

www.avaya.com/

Avaya Media Servers

An Avaya Media Server provides centralized, enterprise-class call processing. This call processing can be distributed across a multi-protocol network (including IP) to support a highly diversified network architecture that consists of headquarters, branch, remote, small, and home offices.

Linux-based servers

The Avaya S8700 Media Server and the Avaya S8300 Media Server are Linux-based Media Servers. These servers support:

- Distributed IP Networking and centralized call processing across multi-service networks
- Dual server design with hot fail-over
- Redundant LAN Interfaces and remote survivable call processing

Windows-based servers

The Avaya S8100 Media Server (formerly known as Avaya DEFINITY ONE and Avaya IP600) is based on a Windows 2000 operating system. The S8100 offers an integrated, all-in-one office solution that includes:

- Messaging
- Call Center
- Management

For more information on the architecture and the functionality of the Media Servers, see:

- *Hardware Solutions Guide*
- For US BusinessPartners:
  https://www.avaya.com/doc/gpp/public/pss/category/cs/eclips/media_servers_gateways/index.html

Avaya DEFINITY Servers

Avaya Communication Manager also runs on the following DEFINITY Servers, which can be IP-enabled:

- Avaya DEFINTY Server R
- Avaya DEFINITY Server SI
- Avaya DEFINITY Server CSI

These servers run on the Oryx/Pecos proprietary operating system, and function in the same way as the Media Servers in Figure 2, Communication Manager traffic flow, on page 21. These servers fit into Avaya CMC1, SCC1, and MCC1 Media Gateways.

The focus of this document is network design incorporating the newer Communication Manager platforms. Therefore, the DEFINITY Servers are only discussed briefly here. For detailed discussions on Communication Manager supported by these DEFINITY Servers, see:

- For US BusinessPartners:
  https://www.avaya.com/doc/gpp/public/pss/category/cs/eclips/definity_servers/index.html

Avaya Media Gateways

An Avaya Media Gateway supports both bearer traffic and signaling traffic that is routed between packet-switched networks and circuit-switched networks. Communication Manager running on Avaya Media Servers controls voice and signaling over a variety of stackable and modular Media Gateways:

- Avaya G600 Media Gateway
- Avaya G700 Media Gateway
- Avaya CMC1 Media Gateway
- Avaya SCC1 Media Gateway
- Avaya MCC1 Media Gateway
- MultiTech MultiVoIP Gateway

The Media Gateways contain the network and the endpoint interfaces, as well as call classification, announcement boards, and so on. Through these interfaces, Communication Manager performs gateway/gatekeeper functions. For more information on the Media Gateways, see Call processing.

The following additional resources provide even more details on the Avaya Media Gateways:

For US Business Partners:
https://www.avaya.com/doc/gpp/public/pss/category/cs/eclips/media_servers_gateways/index.html

Avaya Integrated Management

Avaya Integrated Management is systems-management software for managing converged voice and data networks. The applications include network management, fault management, performance management, configuration management, directory management, and policy management functionality.

- Avaya Site Administration
- Avaya Terminal Emulator
- Avaya Communication Manager Configuration Manager
- Avaya Communication Manager Fault and Performance Manager
- Avaya Communication Manager Proxy Agent
- Avaya VoIP Monitoring Manager
- Avaya Directory Enabled Management
- Avaya Terminal Configuration

For more information on Avaya Integrated Management, see:

- Avaya network management products (Integrated Management) on page 162
- For US BusinessPartners:
  https://www.avaya.com/doc/gpp/public/pss/category/system_network_management.html

Avaya communication devices

Avaya Communication Manager provides intelligent control for these smart devices:

- Avaya IP Telephones: 4600 Series (4602, 4606, 4612, 4620, 4624, 4630)
- Avaya digital telephones: 6400 Series, 2402, and 2420
- Avaya analog telephone (6200 Series, 2500, and 2554)
- Avaya IP Softphone
- Avaya IP Softphone for Pocket PC
- Avaya IP Agent
- Extension to Cellular Application
- DEFINITY Wireless DECT System
- Avaya Wireless Telephone Solutions

For more information about Avaya smart devices, see:

- Wireless on page 102
- For US BusinessPartners: https://www.avaya.com/doc/gpp/public/pss/category/telephone_end_user_devices.html

Avaya Communication Manager applications

Avaya Communication Manager has embedded capabilities for:

- Call Center
- Compact Call Center
- Computer Telephony Integration (CTI)
- Messaging
- Conferencing systems
- Unified Communication Center

These applications are briefly discussed in Communication applications on page 79.

Call Center

The Avaya Call Center solution is built on proven and innovative automatic call distribution (ACD) technology. This technology offers a suite of call routing capabilities that help agents handle calls more effectively. Customers can select from a powerful assortment of features, capabilities, and applications that are specially designed to enhance call center operations:

- Agent Access
- Avaya Call Management System
- Avaya Call Management System Supervisor
- Avaya Basic Call Management System
- Avaya Business Advocate

- Call Center
  - Avaya Call Center Basic
  - Avaya Call Center Deluxe
  - Avaya Call Center Elite
- Call Recording
- CallMaster® series digital telephones
- Computer Telephony (ASAI)
- Avaya Visual Vectors
- Avaya IP Agent
- Avaya Network Reporting
- Avaya Virtual Routing

Compact Call Center

- Basic Call Management
- Reporting Desktop
- Computer Telephony

For more information on these subjects, see:

US BusinessPartners: https://www.avaya.com/doc/gpp/public/pss/category/call_center_crm.html

Computer Telephony Integration (CTI)

CTI opens up Application Programmer Interfaces, which can be used to control the server from an external application. For more information on CTI, see:

US BusinessPartners: https://www.avaya.com/doc/gpp/public/pss/category/call_center_crm.html

Messaging

The following messaging systems are supported by Avaya Communication Manager:

- INTUITY™ Messaging Systems
- Aria® Messaging Systems
- Serenade® Messaging Systems
- Unified Messenger®

For more information on Avaya messaging products, see:

- For US BusinessPartners: https://www.avaya.com/doc/gpp/public/pss/category/messaging.html

Conferencing systems

For more information on Conferencing, see:

- For US BusinessPartners: https://www.avaya.com/doc/gpp/public/pss/category/conferencing.html

Unified Communication Center

For more information on the Unified Communication Center (UCC), see:

- For US Business Partners:
  https://www.avaya.com/doc/gpp/public/pss/category/unified_communication.html

Avaya Application Solutions platforms

The Avaya Communication Manager portfolio covers small, medium, and large enterprises with advanced communications needs between 40 and 36,000 stations per system. This chapter provides an overview of the Avaya Communication Manager platforms architecture that supports Avaya Application Solutions components and features.

Figure 3: Avaya Application Solutions platforms

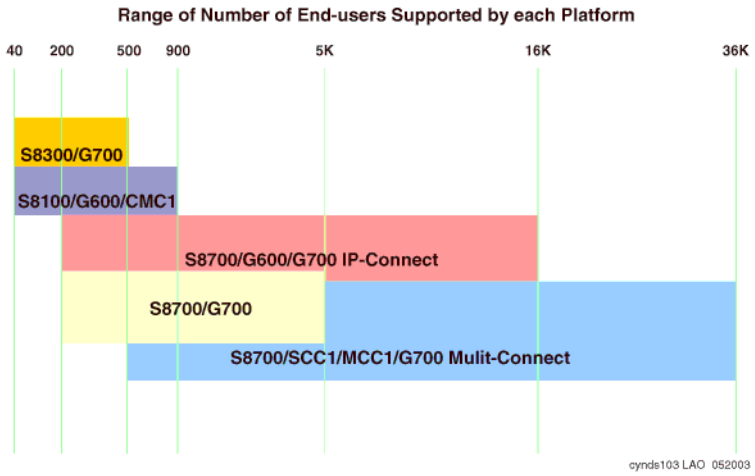

Table 1, Avaya Application Solutions comparison matrix, on page 28 contains an overview of the different Avaya Media Servers that are discussed in this chapter.

Table 1: Avaya Application Solutions comparison matrix

	Avaya S8100 Media Server	Avaya S8300 Media Server	Avaya S8700 Media Server IP-Connect Solution	Avaya S8700 Media Server Multi-Connect Solution
Operating system	Windows 2000	Pre-2.0 — Linux (Redhat v6.2) 2.0 — Linux (Redhat v8.0)	Pre-2.0 — Linux (Redhat v6.2) 2.0 — Linux (Redhat v8.0)	Pre-2.0 — Linux (Redhat v6.2) 2.0 — Linux (Redhat v8.0)
Processor	Intel Pentium III 500 MHz 20-GB hard disk drive 256 MB of RAM	Intel Pentium class server 20-GB hard disk drive 256 MB of RAM	Intel Pentium Class Server 20-GB hard disk drive 256 MB of RAM Removable Flash card backup	Intel Pentium Class Server 20-GB hard disk drive 256 MB of RAM Removable Flash card backup
General Business analog equivalent BHCC rate	Up to 5,000	Up to 10,000	Up to 100,000	Up to 300,000
Media Gateways	G600-up to 3 MCM1-up to 3	G700-up to 50	G600-up to 64 G700-up to 250	MCC1-up to 64 SCC1-up to 64 G700-up to 250
Number of IP endpoints supported	Up to 450 with G600 Up to 240 with MCM1	Up to 450	Up to 12,000	Up to 12,000
Reliability / survivability	No options	Local Spare Processor	Duplicated Processor Local Spare Processor backup for G700	Duplicated Processor Duplicated control network connection Duplicated bearer connectivity Local Spare Processor backup for G700

Small to mid-size enterprise

S8100 Media Server and Avaya G600 Media Gateway or CMC1 Gateways

The Avaya S8100 Media Server running the Avaya Communication Manager functions as the call controller, or gatekeeper, for the Communications Solution that supports up to 900 stations. This Communication Manager server is an “all in one” solution with telephony, multimedia messaging, and system management collocated on the same processor.

The S8100 with either the G600 or the CMC1 Media Gateway brings together Avaya products to create a new standard in multi-service IP Telephony:

- Call processing that runs on a customized surround-supported Windows 2000 server operating system, co-resident with sophisticated and integrated voice/fax mail, call centers, announcements, Call Detail Recording (CDR), and Web-based system administration. SNMP agent is also supported as an option to the standard Expert systems interface.
- Complete IP gateway and IP gatekeeper functions that support 100% TCP/IP transport of merged voice and data.

This solution is a stand-alone platform that does not require any additional external resources, and supports both traditional telephony and IP Telephony. This platform does not support remote gateways. It is scalable to supporting up to three G600 or three CMC1 Media Gateways, with a capacity of up to 900 stations, and has call center capacity of 2,000 BHCC. The platform can be networked to other PBX systems and Communication Manager platforms through both IP and circuit-switched links (Figure 6, S8100 Networked to other PBX and Communication Manager platforms, on page 32).

Table 2: Capacity limits and scalability for S8100 / CMC1 / G600

Feature	Capacity
General Business Busy Hour Call Completion	Up to 5,000
Number of IP stations	Up to 450 with G600 240 with CMC1
Number of stations	900
Number of trunks (independent of the number of users)	Up to 400
Total number of ports	Up to 1,300
Number of G600\CMC1 Media Gateways	Up to 3

S8100 with the G600 Media Gateway (Figure 4, Avaya S8100 Media Server with G600, on page 30) is designed for customers in the 25 to 60 line size or smaller, with a growth potential to 450 IP endpoints and 300 trunks. G600 is designed for communication environments that emphasize IP data and IP Telephony.

Figure 4: Avaya S8100 Media Server with G600

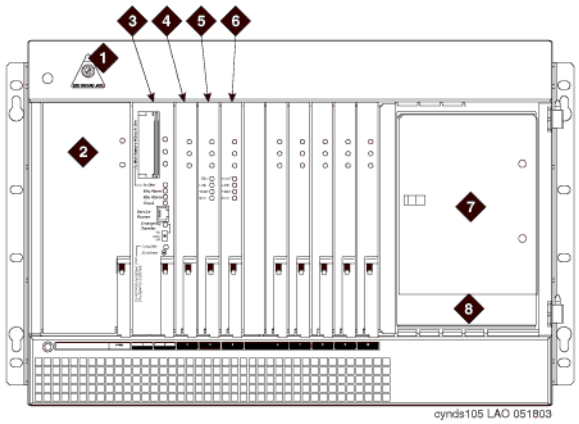

Figure notes

1	ESD ground jack	**5**	TN2302 Media Processor
2	650A power supply	**6**	TN799DP Control-LAN (C-LAN)
3	TN2314 Processor	**7**	Accessory compartment
4	TN744E Call Classifier/Detector	**8**	Fiber-optic pass-through

S8100 with the CMC1 Media Gateway (Figure 5, Avaya S8100 Media Server with CMC1, on page 31) is designed for customers in the 25 to 40 line size or smaller, with a growth potential of 240 stations and 300 trunks. The CMC1 is a compact cabinet with 10 universal port slots (standard reliability only). The CMC1 is designed for wall mounting, or can be floor-mounted or table-mounted where required.

Figure 5: Avaya S8100 Media Server with CMC1

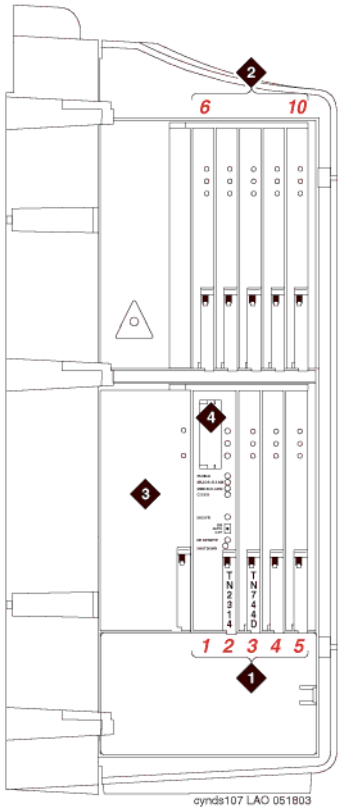

Figure notes

1	Slots 1-5	**3**	650 A power unit
2	Slots 6-10	**4**	PCMCIA flash disk

NOTE:
The S8100 Media Server must use either all CMC1 Media Gateways or all G600 Media Gateways. A single system cannot have a mix of CMC1 Media Gateways and G600 Media Gateways.

The S8100 solution uses Communication Manager to offer:

- Expert system diagnostic capability
- Virtual enterprise networking
- Remote worker and IP agent capability
- Full-featured telephony services
- Call Center

Two of the major applications that run on the S8100 are:

- Co-resident INTUITY AUDIX (R5.01) voice messaging
- Co-resident Avaya Site Administration (Release 1.9)

Figure 6: S8100 Networked to other PBX and Communication Manager platforms

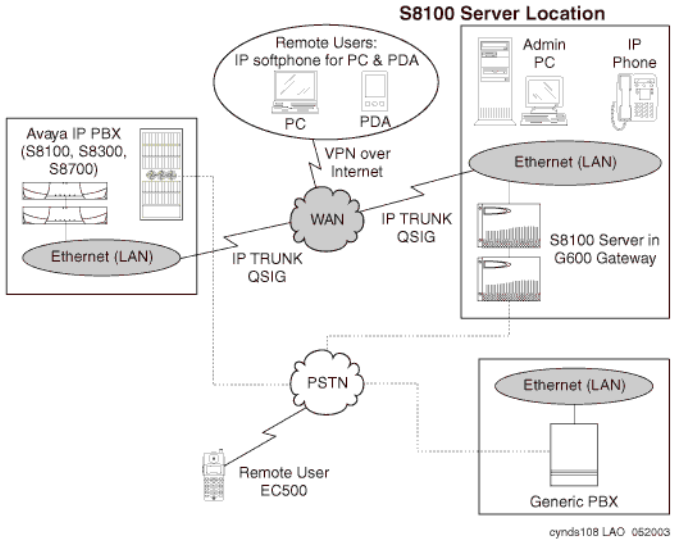

Avaya S8300 Media Server and Avaya G700 Media Gateway

The S8300 Media Server and G700 Media Gateway solution (Figure 7, Avaya G700 Media Gateway with the S8300 Media Server, on page 32) seamlessly delivers voice, fax, and messaging capabilities over an IP network. This unique solution converges the power of the Avaya Communication Manager feature set with the power of distributed Ethernet switching from the P330 Stackable Switching System.

Figure 7: Avaya G700 Media Gateway with the S8300 Media Server

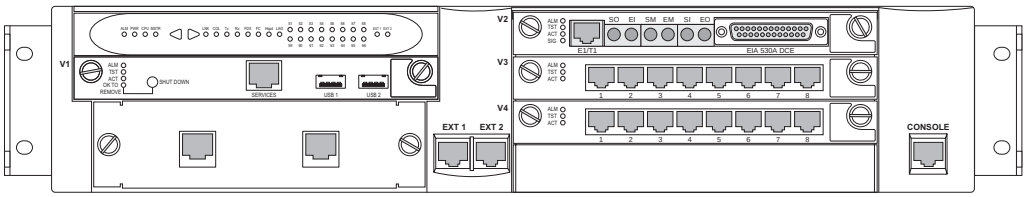

The S8300 / G700 is an stand-alone solution with a capacity of up to 450 IP and non-IP endpoints. A Linux-based S8300 Media Server can support up to 50 G700 Media Gateways. An S8300 Media Server and G700 Media Gateway solution is comprised of:

- A G700 Media Gateway. A G700 Media Gateway is always required. The G700 can host an S8300 Media Server, or various media modules depending on the telephony needs at a particular location.
- The S8300 Media Server. The S8300 Media Server is inserted into a media module slot. If present, the S8300 supports the Communication Manager that provides call-processing capabilities and features for the system. The S8300 can be configured as the primary call controller or as a Local Survivable Processor (LSP) standby server for another S8300 Media Server in the configuration.

Multiple G700 Media Gateways can be connected to each other through an Octaplane 8-Gbps stacking fabric, and Avaya P330 Expansion Modules, which allows adding additional Ethernet ports, fiber interfaces, ATM access or WAN access modules without additional switches. The system can be networked to other PBXs and Communication Manager platforms through an IP network.

Some of the key characteristics of this platform are

- Expert System Diagnostic Capability
- Hot-swappable Media Modules
- Co-resident INTUITY AUDIX messaging.

The platform is scalable, and has survivability and redundancy capability through a Local Survivable Processor (LSP), which supports all of the features of Communication Manager.

Figure 8: Avaya S8300/G700 in a stand-alone configuration

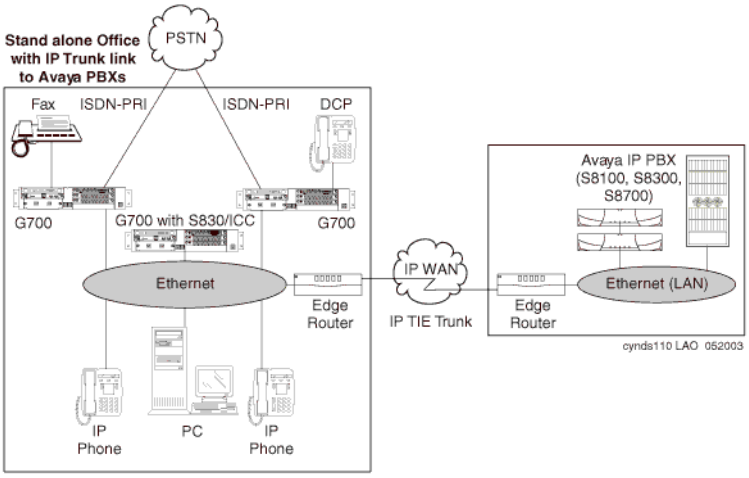

Table 3: Capacity limits for S8300 / G700

Feature	Capacity
General Business Busy Hour Call Completion	Up to 10,000
Number of IP users (cannot be combined with non-IP Users)	Up to 450
Number of non-IP users (cannot be combined with IP Users)	Up to 450
Number of trunks (independent of the number of users)	Up to 450
Total number of ports	Up to 900

Table 4: Scalability for S8300 / G700

Feature	Capacity
Number of G700 Media Gateways	Up to 50
Number of LSPs per server	Up to 10
Minimum ratio of LSPs to server	1:50

G700 hardware architecture

The design of the Media Gateway motherboard hardware brings together a multitude of hardware functions into a single 2U 19-inch rack-mountable enclosure. Integrated on the motherboard are:

- A gateway function that bridges the IP and telephony domains
- An Ethernet switching function and associated management features through an integrated Layer 2 switch architecture
- Processing elements that are necessary to support traditional telephony interfaces, such as trunks and analog/DCP lines

These processing elements are controlled by Communication Manager, thus offering the complete set of Communication Manager call features to both IP users and traditional telephony users.

From a hardware perspective, the G700 Media Gateway is an enclosure with an internal power supply and a motherboard. This design that provides the hardware resources for the Gateway functions, and electrical connectivity for four media modules, one Cascade module, and one Expansion module. Designed to resemble the P330, the enclosure houses the power supply and the motherboard, and provides the physical support to allow the insertion of the various modules. Figure 9, Avaya G700 Media Gateway (front view), on page 34 shows the Media Gateway enclosure.

Figure 9: Avaya G700 Media Gateway (front view)

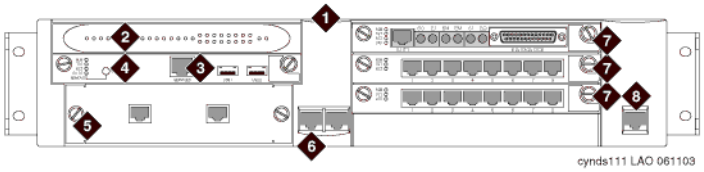

Figure notes

1	S8300 Media Server	**5**	Avaya P330 Expansion Module
2	LED board	**6**	10/100 Base T Ethernet ports
3	USB port (modem connection)	**7**	Media Modules
4	Network Interface Card (NIC) for Services laptop	**8**	Serial Command Line Interface (CLI) connector

The four media module slots can be populated with any combination of media module types:

- T1/E1 with integrated CSU/DSU (MM710)
- 8-port analog line/trunk (MM711)
- 8-port DCP line (MM712)
- 8-port BRI trunk (MM720)
- VoIP Engine (MM760)
- Internal Communications Controller (ICC-only 1 per gateway; must be in the first slot)

The Cascade module comes from the Converged Infrastructure LAN Switches product line, and provides the Octaplane interface:

- One full-duplex 4-Gbps Ethernet port (8 Gbps bandwidth) for high-speed interconnection of up to 10 MGs and P330 data switches in a stack arrangement
- Expansion module interface allows the use of expansion modules in the gateway. These expansion modules also allow WAN access routing.

The G700 motherboard hardware design involves three major blocks:

- A DSP engine and associated packet processor complex. This complex performs IP/UDP/RTP processing, echo cancellation, G.711 A/μ, G.729 (with or without silence suppression), fax relay, silence suppression, jitter buffer management, packet loss concealment, and so on.
- A Gateway Processor complex. This complex is the master controller of the Gateway, and controls all resources inside the Gateway under the direction of the Gateway Controller. Examples of the functions implemented here include the Media Module Manager, Tone/Clock, PKTINT, Announcements (record/playback), and H.248 signaling to the Gateway Controller.
- An Intel 960 processor complex. This complex is based on the architecture of the P330 data switch. This complex provides an eight-port Layer 2 switch function, and the i960 manages the Expansion and Cascade modules.

These major blocks are interconnected through two major communication paths: an Ethernet link and the Time Division Multiplexed (TDM) bus similar to that in a port network. In addition, the motherboard provides electrical and physical connectivity for four media modules.

VoIP Engine complex

The internal VoIP Engine block is where PCM voice samples are encoded and put into IP packets, and vice-versa. This block implements all the functions that are normally associated with a Gateway. Such functions include packet loss concealment, jitter buffer management, transcoding, and so on. The VoIP Engine also performs a T.38 fax relay function, whereby the VoIP Engine sends fax tones out-of-band over an IP network.

The VoIP Engine of the G700 motherboard has three major components: two Digital Signal Processors (DSPs), and a Motorola MPC8260 processor. The DSPs together provide the same VoIP channel capacities as the TN2302 Prowler circuit pack: 64 G.711 channels or 32 G.729 channels.

Each additional VoIP Media Module (MM760) increases the VoIP channel capacity of a G700 media gateway by the equivalent of a TN2302 circuit pack.

The Media Gateway Processor

The Media Gateway Processor (MGP) is the master controller of the Media Gateway. The Motorola 860T processor in this complex implements the H.248 protocol to communicate with the Gateway Controller. Under the direction of the Gateway Controller, the 860T Gateway Processor controls the flow of data through the Gateway. The 860T processor communicates with other processors in the system – the VoIP Engine processor, the i960 processor, and any processors on media modules – through either the control channel of the TDM bus, or an Ethernet link (the i960 processor connects only through Ethernet).

Functions implemented within the MGP complex include:

- Management of the media modules (reset control, board and interface insertion, and so on)
- Termination of the LAPD protocol running on the D-channel of E1/T1 trunks and BRI lines and trunks (32 channels capacity).
- Recorded announcement playback (15 playback channels, 1 record channel)
- Tone detection and generation (15 ports of tone detection)
- System clock generation and synchronization to an external network timing reference
- Download agent for the media modules
- License/translation storage
- System maintenance
- H.248 signaling
- Connection management

S8300 MG Controller architecture

The S8300 Internal Communications Controller (ICC) complex is implemented as a Media Module to provide Communication Manager call control and INTUITY AUDIX messaging. The controller targets the small-line-size, cost-conscious portion of the market, and as such, must be cost competitive with other solutions. The controller is based on standard Intel IA32 architecture, and runs the industry-standard Linux operating system.

The ICC runs the following co-resident applications:

- H.248 Media Gateway Controller
- H.323 GateKeeper
- Communication Manager Feature Server
- INTUITY AUDIX Messaging system

The S8300 ICC can be ordered both with and without INTUITY AUDIX support.

The ICC is a custom form-factor “media module” (MM) that plugs into the first slot in a G700 media gateway. The faceplate provides connectivity for two USB devices, and an Ethernet port for technician access. The faceplate also has operational LEDs and a shutdown switch. The media module backplane connector provides the interfaces for the internal 10/100 Ethernet bus and the TDM Bus.

An ICC can control up to 50 G700 media gateways, 450 stations, and 450 trunks.

Avaya IA770 INTUITY AUDIX Messaging Option for S8300

The Avaya IA770 INTUITY AUDIX Messaging Application, (IA770), optionally embedded on the Avaya S8300 Server, delivers voice, fax, and e-mail to enhance and simplify the communications and the exchange of information within both small enterprises, and the smaller locations of large enterprises. The IA770 uses the Linux operating system, which is consistent with the operating system of the Media Gateway.

The IA770 supports INTUITY digital (TCP/IP) and AMIS networking protocols. More extensive networking can be provided with the Avaya Interchange.

The IA770 consists of license-file-activated software that resides on the S8300 Server, and an ICC daughter card, which is field-installable and upgradeable.

S8300 as the Local Spare Processor (LSP)

The LSP is an S8300 server media module that is located in a G700 Media Gateway. The LSP provides survivability when the primary controller, either an S8300 ICC or an S8700 External Communications Controller (ECC), is inaccessible. Each system can have multiple LSPs. Each LSP has a copy of the primary controller's translations. The translations are updated regularly from the primary controller b way of a virtual link through an IP network. Typically, all LSPs are in idle mode, where the LSPs are not processing any calls. When the Media Gateway's Processor (MGP) or individual IP endpoints perceive the primary controller to be unreachable, the MGP or the IP endpoints attempt to register with an LSP. The LSP does not actively take over when the primary controller becomes unreachable, but waits for MGPs and IP endpoints to register with it. Each LSP runs in license-normal mode until IP Telephones or MGPs register with it, which triggers the LSP to move into a license-error mode. Each LSP can run in active mode for a maximum of 10 days per outage before it must be reset manually. A "reset 3" command on the LSP forces devices registered with it to return to their primary controller. When the customer resets the LSP, the 10-day license timer is reset after it contacts the primary controller.

Voice Announcement over the LAN

Voice Announcement over the LAN (VAL) capabilities are co-resident on the Avaya G700 Media Gateway. This G700 VAL announcement capability allows backup and restore of announcements to an external PC or a file server on the customer's local area network (LAN), in addition to internal backup in Flash memory. The announcements are stored as industry-standard waveform (.wav) files. This enables customers to create high-quality, studio announcements, save the announcements to their PC or server, and then share the same announcements with multiple Avaya Application Solutions. Other features of the G700 VAL announcement offer include:

- A G700 VAL announcement source functions the same as the TN2501AP for administration, recording, file handling using FTP, playback, and measurements.
- Each G700 VAL announcement source used is counted as a VAL board towards the Maximum VAL boards on the customer-options screen. The S8300 Media Server now comes with a license entitlement for using up to 50 VAL circuit packs. The S8700 Media Server comes with a license entitlement for one, and requires the purchase of additional licenses to enable the maximum of 50 which applies to both TN2501AP and G700 VAL sources.
- Voice quality is impacted when played over IP. However, quality is acceptable even with 2 hops and 10-msec delay.

- The use of G700 VAL sourced announcements impacts that gateway's overall occupancy, and IP Telephony resources (for example, high use global announcements such as the main greeting and some VOAs) should be handled by TN2501 circuit packs if the agents are not homed to that G700.
- FTP access for the G700 announcements use the same IP address as the address that is assigned to the G700 when installed (this address is displayed on the Media-Gateway form)

Avaya DEFINITY Server CSI and Avaya CMC1 Media Gateway

The Avaya Communication Manager with an Avaya DEFINITY® Server CSI offer is targeted at small to mid-size customers, or the satellite offices of large corporations who require sophisticated applications. This solution supports 50 to 500 telephones. The small to mid-size customer segment tends to be multi-site with 2 to10 locations. These customers are interested in applications such as voice messaging, networking, shared voice mail, and small call center applications. Satellite offices of large corporations tend to require the same applications that are used at their headquarters location.

Avaya DEFINITY Server SI and Avaya SCC1 Media Gateway or Avaya MCC1 Media Gateway

This solution is ideal for locations with 250 to 1,000 employees. Tis solution can support up to 2,400 stations with a seamless migration to larger Avaya Media Servers. For more information on these two platforms, see:

http://www.avaya.com/

Mid-market to large enterprise

Avaya S8700 Media Server Multi-Connect configuration

S8700 Media Server with an MCC1 or SCC1 Media Gateway is targeted at Avaya's largest customers. These customers are typically experiencing rapid growth, and looking for ways to consolidate their network. These are customers who require high-end applications such as DEFINITY Call Center Solutions, CTI applications, Unified Messaging, multimedia conferencing, and voice/data network integration, and are evolving to an IP-intensive environment. This solution supports up to 36,000 telephones.

This solution is also targeted at smaller customers who made an investment in DEFINITY, and are looking for a smooth transition into industry-standard processors that will enable expanded communications capabilities.

This is a large-office solution with the Media Server in the headquarter locations, with optional servers/gateways in the branch offices. There is also the option of duplicated headquarters with branch and remote offices. The Linux-based Media Servers support up to 36,000 stations, 12,000 of which can be IP, and 8,000 trunks, with the capacity limit of 300,000 General Business call mix BHCC.

Table 5: Capacities for S8700 / MCC1 / SCC1 / G700 Multi-Connect

Feature	Capacity
General Business Busy Hour Call Completion	Up to 300,000
Number of stations	Up to 36,000
Number of IP users (IP trunks + IP stations)	Up to 12,000
Number of trunks (independent of the number of users)	Up to 8,000
Total number of ports	Up to 44,000
Number of IP users G700/Total IP users (cannot be combined with non-IP users)	Up to 5,000/12,000
Number of non-IP users G700/total users (cannot be combined with IP users)	Up to 5,000/36,000

Because calls are processed through both IP and traditional circuit switch links, this platform is called the Multi-Connect Solution. The Media Server and the Gateways call control traffic is over a private dedicated Ethernet network that is provided by Avaya. The inter-Port Network TDM traffic flow is supported by a TDM-based Center Stage Switch (CSS), or an Asynchronous Transfer Mode (ATM) switch.

This solution is scalable to up to 44 Port Networks (PNs) through CSS configuration, and up to 64 PNs in an ATM configuration. The Multi-Connect solution has three reliability options:

- **Standard**. S8700 Media Server, with memory shadowing, two uninterruptible power supplies (UPS), one switch, and one IPSI in each IPSI-connected PN
- **High**. Standard reliability, plus a second switch and a second IPSI in each IPSI-connected PN. This design provides for a second redundant call control network
- **Critical**. High reliability plus duplication of the bearer network

Table 6: Scalability for S8700 / MCC1 / SCC1 / G700 Multi-Connect[1]

Feature	Capacity
Number of G700 Media Gateways	Up to 250
Number of MCC1 / SCC1 Port Networks	Up to 44 with CSS 64 with ATM
Number of LSPs per server	Up to 50
Number of G700 Gateways per LSP	Up to 50

1 The information in this table represents the maximum number of calls the S8700 Media Server can execute when unconstrained by other factors such as TDM bus limitations, call duration, or a small number of telephones. We assume processor occupancy of 0.90, at which various delay criteria such as cut through is preserved, and is just at the threshold where call shedding would commence.

S8700 Media Server

The Avaya S8700 Media Server always consists of two Intel-based servers running on a Linux operating system. In S8700 Multi-Connect and IP-Connect configurations, the S8700 Media Server provides the main feature and management processing capabilities of the system. The Media server is connected to other system and external components primarily through IP networks.

S8700 external features

- Six 10/100 Ethernet NICs per server, which are used as follows:
  - — Dual control network connections
  - — A heartbeat link to the duplicated server
  - — Administrative access from the corporate network
  - — Technician access
  - — One unused
- A CD-ROM for on-site software loading
- A PCMCIA Flash disk for translations backup
- USB ports for remote access connections (modems and other auxiliary devices)
- A reset button
- Support for global power
- A fiber-channel interface to support server duplication

Figure 10: Avaya S8700 external features

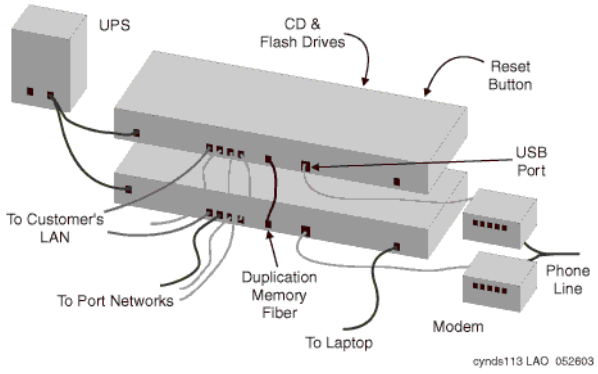

UPS or power backup The S8700 Media Servers always require power backup to avoid power problems, and to ensure graceful shutdown of the system processes if the power fails. The AS1 700-VA UPS provides approximately 30 minutes of power backup. Combinations of battery extension modules and a 1500-VA UPS provide up to 8 hours of power backup.

The AS1 UPS units use SNMP traps to send an alarm when power fails. This action initiates a graceful shutdown process of the Linux server, including the call processing software.

USB modem Each S8700 Media Server supports a Universal Serial Bus (USB) modem. For customers with an Avaya service contract, the modem is used to send alarms to the Avaya Services organization, and to facilitate maintenance by Avaya Services personnel.

Internal hardware elements

The server has the following specifications:

- 256 MB of main memory
- IDE hard disk for booting Linux and Communication Manager
- IDE CD-ROM for software installations and upgrades
- One USB port for modem support
- One PCMCIA slot for removable Flash card support

Figure 11: Avaya S8700 Media Server schematic

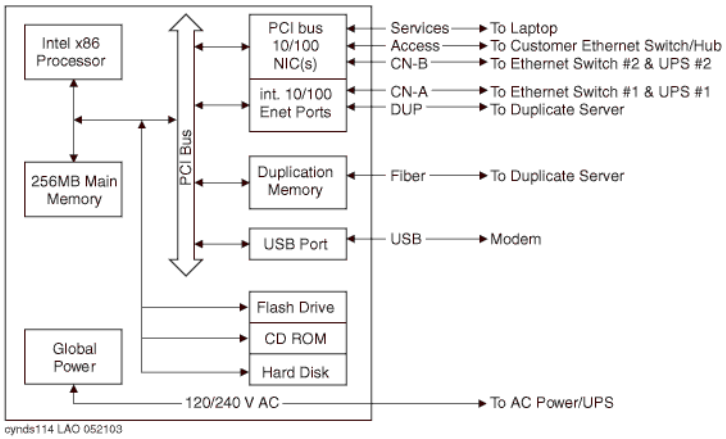

Other components

The S8700 in a Multi-Connect solution also includes the following components:

- An Avaya P133, P134, P333, or P334 Ethernet switch with duplication option
- One or more IP Server Interface (IPSI) circuit packs (TN2312AP)
- A Center Stage Switch (CSS) or an ATM Switch for bearer connectivity
- One or more MCC1 or SCC1 Media Gateways, also known as port networks (PNs)

Figure 12: Avaya S8700/MCC1 Multi-Connect major components

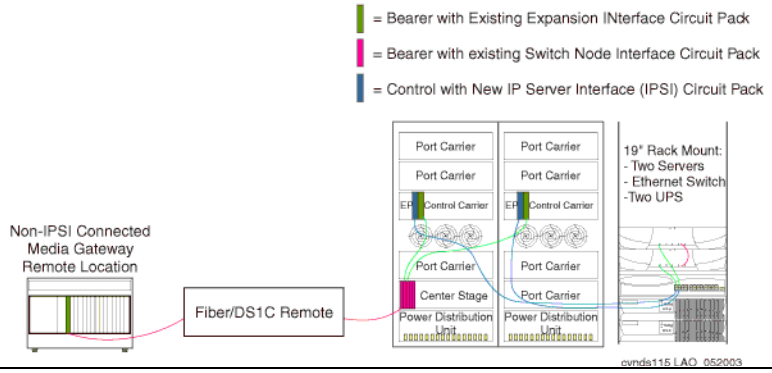

Control network through an Avaya Ethernet switch

When designing S8700 Multi-Connect systems, a control network connects the servers to the IPSIs through a private 10/100 BaseT Ethernet. It consists of two separate Ethernet networks made from Avaya Ethernet switches. Control network A connects to the primary IPSIs, and control network B connects to the secondary IPSIs (Figure 13, S8700 Multi-Connect control network, on page 42).

Figure 13: S8700 Multi-Connect control network

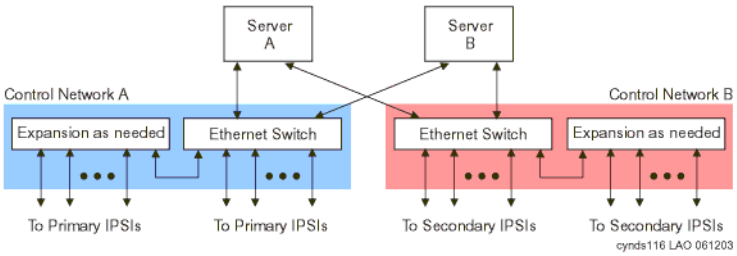

For detailed hardware feature and functionality of the switches, use this link:

http://www1.avaya.com/enterprise/solutions/multiservice/

Circuit packs that support IP signaling and media traffic

Figure 14: S8700 / MCC1 signaling path

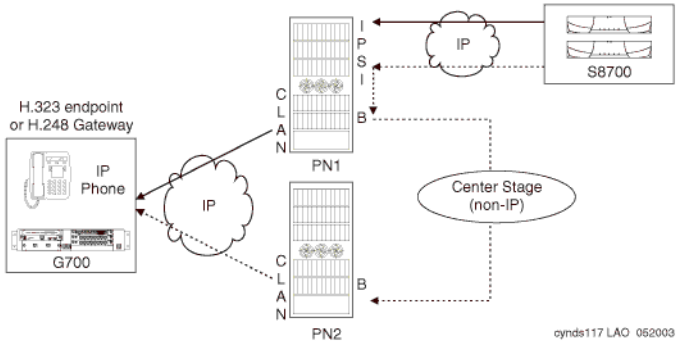

Figure 15: S8700 Multi-Connect - a basic phone call

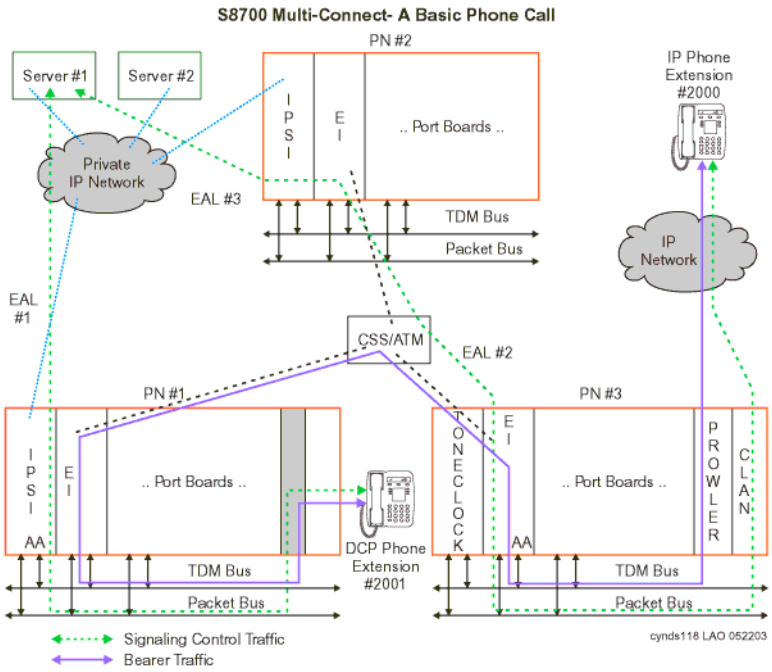

IP Server Interface (TN2312AP) The IP Server Interface (IPSI) is the communication interface between the server and the Media Gateways (port networks). The IPSI is responsible for gateway control, and for tunneling call control messages back to the S8700.

One IPSI circuit pack is required per IPSI-connected Media Gateway for standard reliability. Duplicated IPSI circuit packs are required per IPSI-connected Media Gateway for high reliability and critical reliability.

The IPSI is located in the tone/clock slots, and provides the following functions:

- PKTINT packet bus interface
- Archangel TDM bus interface
- Tone/Clock functionality found on the TN2182B Tone/Clock circuit pack
- Ethernet interface for technician access
- Ethernet interface for connectivity to Services laptop computers
- Maintenance board interface for communication with the EPN maintenance board

Each IPSI typically controls up to five gateways by tunneling control messages over the center stage (TDM) network to the PNs that do not have IPSIs. For locations with high IP Telephone traffic, Avaya recommends a greater number of IPSI circuit packs.

An IPSI cannot be placed in:

- A PN that has a Stratum-3 clock interface
- A remote PN that uses a DS1 converter
- A Survivable Remote Expansion Port Network (SREPN)

The IPSI supports the following functions:

- Supports eight global Call Classification ports
- Supports network diagnostic capabilities
- Provides PN clock generation and synchronization for Stratum-4 type II only
- Provides PN tone generation
- Provides distributed PN packet interface
- Supports the download of IPSI firmware
- Provides serial number support for License File feature activation

Control LAN (TN799DP) The TN799DP Control LAN (C-LAN) circuit pack acts as front-end processor and concentrator and provides the gateway between the public IP Telephony network and the S8700 system All H.323 signaling messages between IP Telephony endpoints and the S8700 servers must pass through the C-LAN. The connectivity path between the IP endpoint and the server is as follows:

> Endpoint ⇔ IP Network ⇔ C-LAN ⇔ PN backplane ⇔ IPSI ⇔ IP network ⇔ S8700 Media Server

The C-LAN circuit pack is used for all IP call signaling for both IP trunks and stations. This circuit pack also provides TCP/IP connectivity to such adjuncts and synchronous applications as Call Management System (CMS) and INTUITY AUDIX.

This circuit pack also supports firmware download capability for all firmware-downloadable circuit packs in a PN, which allows administrators to remotely update the firmware or application code of circuit packs such as the TN799DP (C-LAN) or TN2302AP Media Processor.

The S8700 platforms support a maximum of 64 C-LAN circuit packs per system. The number of C-LAN circuit packs that are required depends on the number of IP endpoints that are connected, and the options that the endpoints use. For example, it might be advantageous to segregate IP voice control traffic from device control traffic.

IP Media Processor (TN2302AP) (The TN2302 (MedPro) circuit pack is a media processor that provides a gateway between the TDM bus and the Ethernet network for the audio streams.

The S8700 Multi-Connect requires resources on a TN2302 circuit pack for IP Telephony bearer communications. The MedPro includes a 10/100 BaseT Ethernet interface to support IP trunks and H.323 endpoints. The MedPro can perform echo cancellation, silence suppression, dual-tone multi-frequency (DTMF) detection, and conferencing.

As shown in Figure 16, TN2302AP Media Processor (MedPro) operation, on page 45, the Media Processor converts circuit-switched audio streams to packet-switched streams. The Media Processor supports multiple codecs, so it can compress audio samples during packetization. When needed for conference calls, it can also take multiple audio streams, sum them together, and send the resulting audio stream to multiple recipients on the network.

Figure 16: TN2302AP Media Processor (MedPro) operation

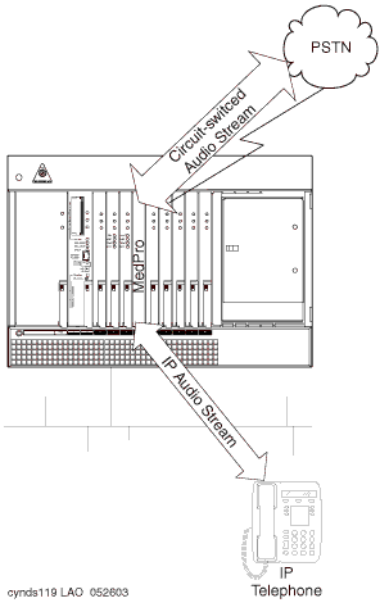

To do the job, the circuit pack has a set of DSP resources. These resources are deployed dynamically and flexibly to any of a number of tasks, including:

- Originating and terminating IP-based packet-switched audio streams
- Establishing and maintaining an RTCP control channel for each IP audio channel
- Compressing and decompressing audio (for example, G.729 to G.711)
- Terminating TCP for an incoming T.120 data stream, and transcoding it to H.221- compliant format for transmission onto the TDM bus and vice-versa
- Summing multiple audio channels into a composite signal for audio conferencing

The S8700 Media Server is responsible for sending messages to the circuit pack to allocate and to configure the DSP resources to the required task and connecting multiple resources into a chain that performs the desired media processing function. In addition, the Media Server sends the information to the destination of these audio streams.

Since H.323 allows any of several different codecs to be used for encoding an audio stream on the IP network, the Prowler board is able to use any of the following codecs:

- G.711
- G.723.1
- G.729 (A, B)

In the same way that a MedPro interfaces with IP Telephony endpoints, it can connect to another MedPro to interconnect two or more Avaya switches in an IP network over an IP trunk.

Media Gateways

The MCC1 Media Gateway, the SCC1 Media Gateway, and G700 Media Gateway are supported in a Multi-Connect configuration. S8700 Multi-Connect can have a mixture of MCC1 and SCC1 cabinets in a system. However, the type of cabinet cannot be split within a Port Network.

Multi-Carrier Cabinet (MCC1) Media Gateway The MCC1 Media Gateway can contain up to five of the following carriers:

- A Port Carrier that contains one or more of the following:
  - Port circuit packs
  - VOIP conversion resources
  - Service circuit packs
  - Tone clocks
  - Expansion Interface (EI) circuit packs
- A Switch Node Carrier that contains Switch Node Interface circuit packs that compose the Center Stage Switch (CSS).
- An Expansion Control Carrier that contains service slots and port slots.

The MCC1 Media Gateways can support a maximum of 98 trunk or line port circuit packs.

Figure 17: MCC1 Media Gateway

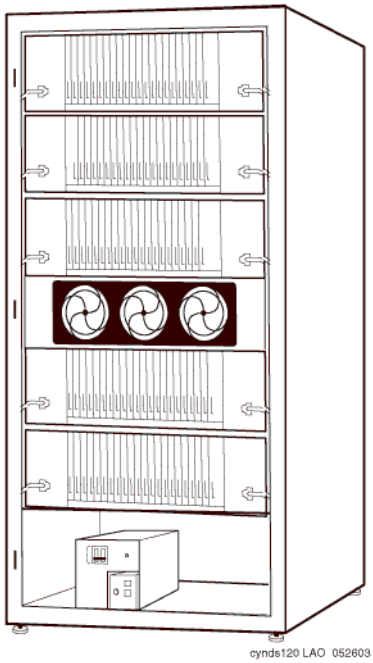

Single-Carrier Cabinet (SCC1) Media Gateway The SCC1 Media Gateway consists of a single carrier. Up to four SCC1 Media Gateways can be connected together in one location to form one port network. There are two types of SCC1 Media Gateways:

- An Expansion Control Cabinet that contains service slots and port slots.
- A Port Cabinet that contains ports and interfaces to an Expansion Control Cabinet.

Figure 18: SCC1 Media Gateway

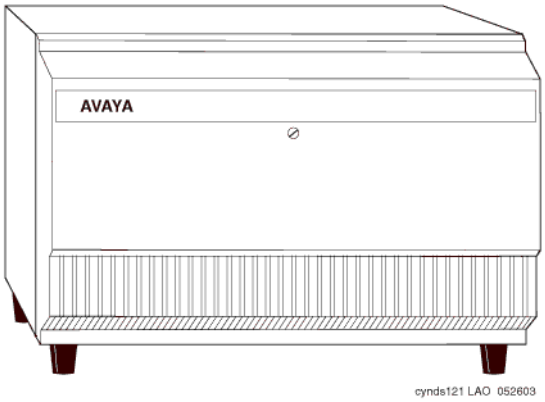

Non-IPSI connected Media Gateway Typically, one of every five Port Networks (PNs) contains one or two IPSI circuit packs. The remaining PNs are referred to as *non-IPSI connected*. Non-IPSI connected PNs get their control information from the servers through one of the PNs that does contain an IPSI. Such control messages are “tunneled” through the circuit-switched network. The system software controls this communication and allocation. The software automatically routes the control messages through an appropriate IPSI. There is no need to administer which IPSI controls the non-IPSI connected PNs. The system automatically allocates those resources, and also compensates for any component failure.

Remote MCC1\SCC1 Media Gateways The dedicated control network for an S8700 with MCC1 or SCC1 Media Gateway can be extended to an IPSI in a remote media gateway. But for cost effectiveness and straightforward installation, Avaya recommends that all of the IPSI-connected media gateways be collocated with the S8700 and Ethernet switches. The circuit-switched network dictates the available options.

Non-IPSI connected media gateways’ circuit-switched network can be extended through all the options available with DEFINITY G3r. Center Stage Switch configurations can use fiber extenders or DS1-Converter (DS1-C) facilities, allowing the media gateway separation to be essentially limitless. When ATM-PNC is used, the media gateway separation is also essentially limitless (see ATM network on page 48).

Remote G700 Media Gateway The S8700 Media Server can provide the call processing features for a remote G700 media gateway over an H.248 link. In this configuration, S8700 can support up to 250 G700 Media Gateways. An S8300 media module that is located in a G700 Media Gateway in a remote location provides survivability when the primary controller is inaccessible. For more information, see S8300 as the Local Spare Processor (LSP) on page 37.

Center Stage Switch

The Center Stage Switch (CSS) is a connection hub that provides Port Network communication. A CSS can be used when the Multi-Connect Solution that is composed of more than three Port Networks. Often, the CSS is incorporated into smaller configurations to allow for growth. The CSS consists of from one to three switch nodes (SN), which reside in a Port Network carrier. SNs are composed of one or two switch node carriers, depending on whether the solution is being duplicated for critical reliability. Port Network expansion depends on internal SN-to-SN traffic, according to the following guidelines:

- 1 SN expands from 1 to up to 15 PNs.
- 2 SNs expands to up to 29 PNs.
- 3 SNs expands to up to 44 PNs.

ATM network

The Asynchronous Transfer Mode (ATM) switch is a replacement option for the CSS, or for the direct-connect switch. Several Avaya ATM switch types can provide Port Network connectivity. Non-Avaya ATM switches that comply with the ATM standards that are set by the European Union can also provide Port Network connectivity.

With S8700 Multi-Connect, ATM-Port Network Connectivity (ATM-PNC) allows any ATM switch or ATM network that complies with specified standards and capacities to serve as the means to connect to the PN. In this type of configuration, the ATM switch or network replaces the CSS. ATM-PNC is used to connect port networks within a single switch. The WAN Spare Processor (WSP) is not supported. One ATM supports up to 64 PNs.

S8700 Multi-Connect configuration for higher availability

When used with the MCC1, SCC1 or G700 Media Gateway, the S8700 Media Server has the following reliability options:

- Standard reliability configuration
- High reliability configuration
- Critical reliability configuration

Standard reliability configuration

The standard reliability option is the most basic option that consists of the following components:

- Two S8700 Media Servers
- Server-to-IPSI control is not duplicated
- One UPS unit for each S8700 Media Server. Using two UPS units ensures that a single UPS failure or repair operation does not disable the system.
- One IPSI in each IPSI-connected port network
- Circuit-switched traffic between port networks is carried on a simplex network that is made up of one Expansion Interface (EI) in each port network. The EIs are cabled with lightguide fiber to either the Center Stage Switch (CSS) or an Asynchronous Transfer Mode (ATM) switch.

Figure 19, S8700 Multi-Connect in a standard reliability configuration, on page 49 shows an example of a standard reliability configuration.

Figure 19: S8700 Multi-Connect in a standard reliability configuration

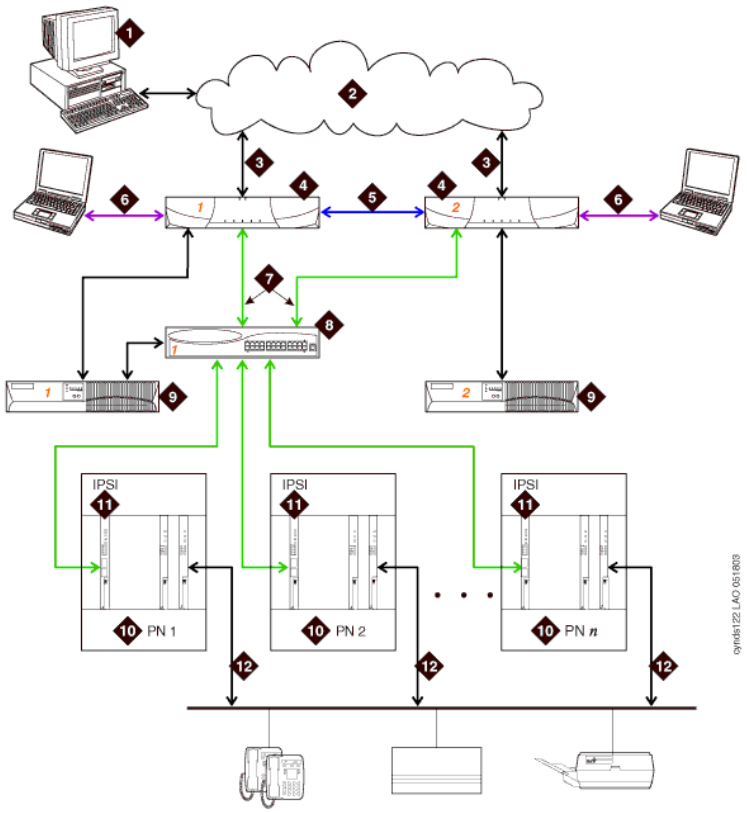

Figure notes

1 The Administration PC accesses the S8700 Media Server over the corporate data network.

2 Corporate IP network.

3 Corporate IP network interface. The Ethernet 4 link from the S8700 Media Server to the data network.[1]

4 Two S8700 Media Servers are always present. One server is in active mode, and the other server is on standby.

5 Duplication interface, default Ethernet 2. The dedicated Ethernet connection between the S8700 Media Servers.

6 Services interface, default Ethernet 1. The server's dedicated Ethernet connection from the S8700 Media Server to a laptop computer (active only during on-site administration or on-site maintenance).

7 Network control A interface, default Ethernet 0. The server's Ethernet connection to one or two Ethernet switches. This Ethernet link carries the control signals for the S8700 Multi-Connect PNs.

8 Ethernet switch. At least one Ethernet switch is required to support the S8700 Multi-Connect control network. If many PNs are present, two Ethernet switches can be daisy-chained together to provide sufficient Ethernet connections to the IPSI boards in the PNs.

9 UPS. Keeps the S8700 Media Servers and the Ethernet switches functional during brief power outages.

10 PN. Provides the telecommunications functions of the S8700 Multi-Connect Media Server.

11 IPSI. The IPSI circuit pack carries the control network signals to the PNs, and provides tone clock functionality.

12 Bearer connectivity over Center Stage Switch or ATM.

1 The Ethernet connection to the corporate network in this figure is a nondedicated network. IP addresses for the various components of the S8700 Multi-Connect Media Server must be administered to prevent conflicts with other equipment that shares the network. In the default S8700 Multi-Connect configuration, all other Ethernet connections operate on their own closed LANs.

High reliability configuration

The high reliability configuration option builds on the standard reliability option. The high reliability option duplicates components, so that no single point of failure exists in the control network. The high reliability configuration consists of the following components:

- Two S8700 Media Servers
- Two IPSI circuit packs in each IPSI-connected port network
- Two Ethernet switches
- Two UPS units

Circuit-switched traffic between port networks is carried on a simplex network that is made up of one Expansion Interface (EI) in each port network. The EIs are cabled with lightguide fiber to either the Center Stage Switch (CSS) or an Asynchronous Transfer Mode (ATM) switch.

Figure 20, S8700 Multi-Connect in a high reliability configuration, on page 51 shows an example of a high reliability configuration.

Figure 20: S8700 Multi-Connect in a high reliability configuration

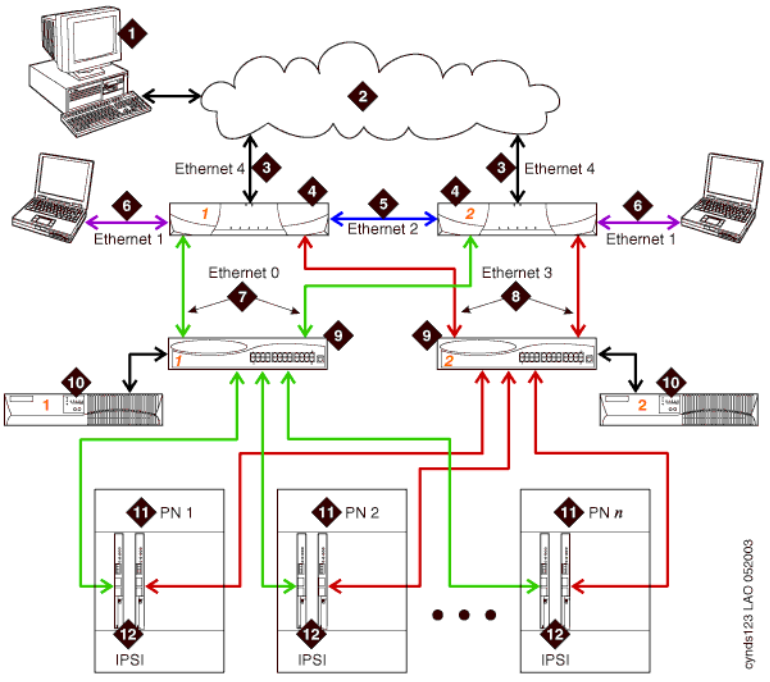

Figure notes

1 The Administration PC is used to access the S8700 Media Server over the corporate data network.

2 Corporate IP network.

3 Corporate IP network interface. The Ethernet 4 link from the S8700 Media Server to the data network.[1]

4 Two S8700 Media Servers are always present. One server is in active mode, and the other server is on standby.

5 Duplication interface, default Ethernet 2. The dedicated Ethernet connection between the S8700 Media Servers.

6 Services interface, default Ethernet 1. The server's dedicated Ethernet connection from the S8700 Media Server to a laptop computer (active only during on-site administration or on-site maintenance).

7 Network control A interface, default Ethernet 0. The server's Ethernet connection to one or two Ethernet switches. This Ethernet link carries the control signals for the S8700 Multi-Connect PNs.

8 Network control B interface, default Ethernet 3. The server's Ethernet connection to one or two Ethernet switches. This Ethernet link carries the control signals for the S8700 Multi-Connect PNs.

9 Ethernet switches. If many PNs are present, two Ethernet switches can be daisy-chained together to provide sufficient Ethernet connections to the IPSI boards in the PNs.

10 Duplicated UPSs. Keeps the S8700 Media Servers and the Ethernet switches functional during brief power outages.

11 Port Networks. Provides the telecommunications functions of the S8700 Multi-Connect Media Server.

12 Duplicated IPSI circuit packs

1 The Ethernet connection to the corporate network in this figure is a nondedicated network. IP addresses for the various components of the S8700 Multi-Connect Media Server must be administered to prevent conflicts with other equipment that shares the network. In the default S8700 Multi-Connect configuration, all other Ethernet connections operate on their own closed LANs.

Critical reliability configuration

The critical reliability configuration option is built upon the high reliability configuration. In the critical reliability configuration, the bearer network has duplicated components so that there is no single point of failure. The critical reliability configuration consists of the following components:

- Two S8700 Media Servers
- Two IPSI circuit packs in each IPSI-connected port network
- Two Ethernet switches
- Two UPS units
- Two CSS/ATM EI (Expansion Interface) in every port network

S8700 Multi-Connect survivability

In addition to the high reliability of the duplicated S8700 Multi-Connect Media Servers, the S8300 Media Server in a Local Survivable Processor (LSP) configuration can be used to provide survivability. Additional recovery capability is embedded in the Communication Manager that resides on the S8700 Media Server.

Avaya S8700 Media Server IP-Connect configuration

The S8700 IP-Connect configuration is an all-IP solution that is built on open IP network connection. This solution is designed for medium to large enterprises. The main difference between the IP-Connect Solution and Multi-Connect solution is that IP-Connect uses the IP network for all inter-gateway communication, where Multi-Connect also uses CSS or ATM networking.

As Table 7, Capacity limits for S8700 / G600 / G700 IP-Connect solution, on page 54 shows, the system supports up to 12,000 IP stations, 8,000 non-IP stations, and 8,000 IP and non-IP trunks, with the capacity of 100,000 General Business call mix BHCC. The IP-Connect platform is scalable to 64 Port Networks, each of which can house up to four G600s and up to 250 G700 Media Gateways. The Media Server complex still consists of duplicated S8700 servers. One server is active, and the other server is on standby. Table 8, Scalability for S8700 / G600 / G700 IP-Connect solution, on page 54 shows the IP-Connect G700 capacities.

Figure 21, Avaya S8700 with remote G600 / G700 Media Gateways, on page 53 shows an example of an S8700 with remote G600/G700 Media Gateways.

Figure 21: Avaya S8700 with remote G600 / G700 Media Gateways

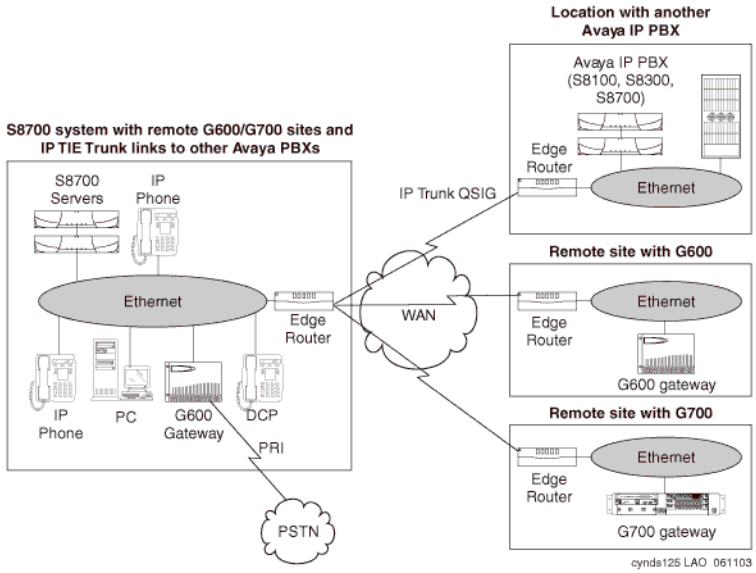

Figure 22, S8700 IP-Connect - a basic phone call, on page 53 shows a call through an S8700 IP-Connect system.

Figure 22: S8700 IP-Connect - a basic phone call

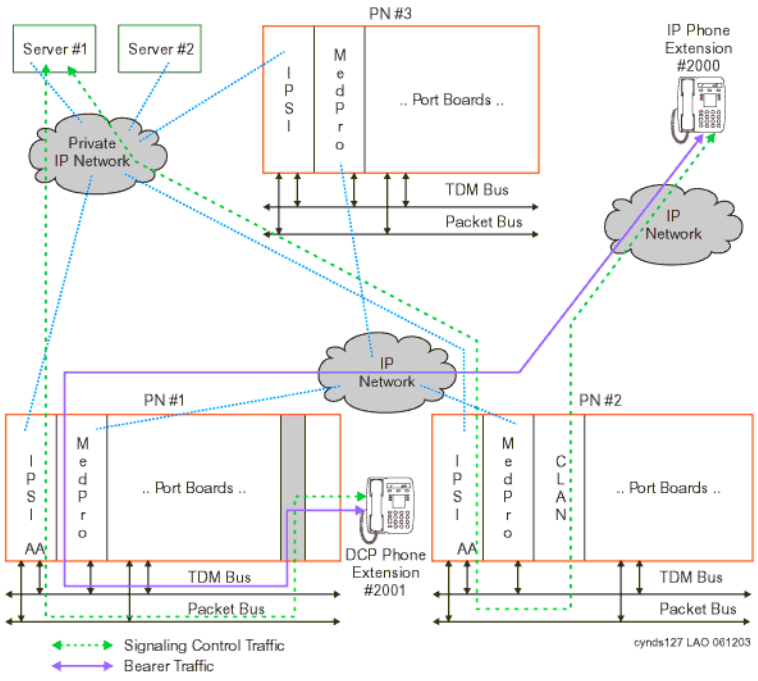

Table 7: Capacity limits for S8700 / G600 / G700 IP-Connect solution

Feature	Capacity
General Business Busy Hour Call Completion	Up to 100,000 analog equivalent
Number of stations	Up to 12,000
Number of IP users (IP trunks + IP stations)	Up to 12,000
Number of non-IP users	Up to 8,000
Number of trunks (independent of the number of users)	Up to 4,000
Total number of ports	Up to 16,000
Number of IP users G700/total IP users (cannot be combined with non-IP users)	Up to 5,000/12,000
Number of non-IP users G700/total users (cannot be combined with IP users)	Up to 5,000/12,000

Table 8: Scalability for S8700 / G600 / G700 IP-Connect solution

Feature	Capacity
Number of G700 Media Gateways	Up to 250
Number of G600 Media Gateways (port networks)	Up to 64 PNs, each houses up to four G600s
Number of LSPs per server	Up to 50
Number of G700 gateways per LSP	Up to 50

Main components

The S8700 IP-Connect consists of the following main components:

- Duplicated S8700 Media Servers
- Two UPS units, one for each server
- Two Abstract Control Modem (ACM) compliant Universal Serial Bus (USB) modem
- At least one IPSI
- TN799DP C-LAN (for IP endpoint signaling)
- At least one TN2302AP IP Media Processor to support inter-port and intra-port network media connectivity
- The G600 Media Gateway
- Avaya Communication Manager

Figure 23, S8700 Media Server IP-Connect major components, on page 55 shows the main S8700 IP-Connect components mounted in an open EIA-310-D- compliant, 19-inch data rack.

Figure 23: S8700 Media Server IP-Connect major components

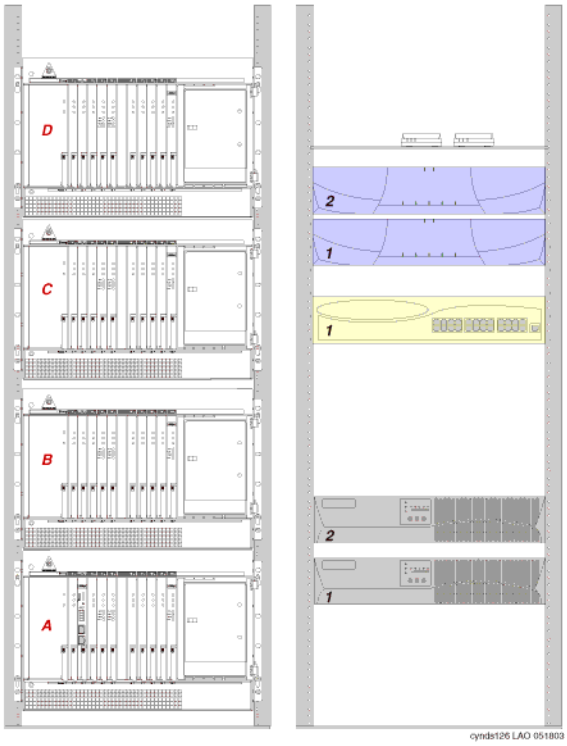

The left data rack contains a stack of four G600 Media Gateways that are labeled A through D.

The right data rack contains the following (from top to bottom):

- Two USB-compliant modems
- Two S8700 Media Servers
- One Avaya Ethernet switch
- Two AS1 UPS units

G600 Media Gateways / port networks

The S8700 Media Server for IP-Connect supports a maximum of 64 G600 Media Gateways, each of which can house up to four rack-mounted cabinets. The S8700 IP-Connect solution does not support the CMC1 Media Gateway, the SCC1 Media Gateway, or the MCC1 Media Gateway.

The G600 Media Gateway has the following characteristics:

- Connects to the S8700 Media Server through the IPSI circuit pack across an IP network.
- 19 inches (48.26 cm) wide, 13 inches (33.02 cm) high, and 21 (53.34 cm) inches deep.
- 10 universal slots plus 1 power supply slot in each Media Gateway.
- A maximum of 64 port networks.
- A maximum of four rack-mountable cabinets (RMCs). The RMC for each individual G600 Media Gateway must be collocated in the same 19-inch data rack due to cable length restrictions.
- A G600 Media Gateway consists of a control RMC that is designated A, and second, third, and fourth optional RMCs that are designated as B, C, and D, respectively.
- The circuit packs are inserted and removed from the front of the cabinet. Cabinet I/O is through the back, and through a front cable pass-through slot on the right.
- The G600 Media Gateway is AC powered only. The G600 Media Gateway has no internal batteries, and internal DC power is not an option. Avaya recommends an external UPS.
- An RJ45 patch panel is recommended for cross-connect to the data network or the wall field.

S8700 IP-Connect reliability configurations

The S8700 Media Servers are duplicated. The control and bearer links are not duplicated. The clock functionality is provided by the IPSI circuit pack in each port network. As an all-IP solution, the S8700 IP-Connect only supports IP media gateways. The S8700 IP0-Connect does not support traditional CCS- or ATM-connected media gateways.

S8700 IP-Connect configuration

Figure 24: S8700 / G600 IP-Connect configuration

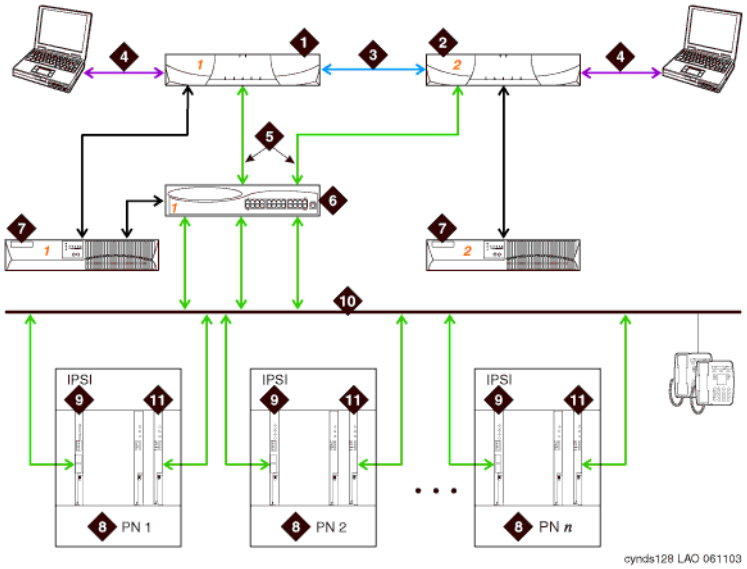

Figure notes

1 Two S8700 Media Servers. One server is in a active mode, and the other server is on standby.

2 Duplication Interface. The Ethernet connection between the two S8700 Media Servers.

3 A dedicated Ethernet connection to a laptop computer. This connection is active only during on-site administration or maintenance, and the Services interface can link to the non-active server through a telnet session.

4 Connection from the servers to the Ethernet switch.

5 Ethernet switch. A device that provides port multiplication on a LAN by creating more than one network segment.

6 UPS units. Two UPS units are required

7 Port Network. An optional configuration of Media Gateways that provides increased port capacity.

8 IPSI. A circuit pack that transports control messages over IP. The IPSI circuit pack is used so that the S8700 Media Server can communicate with the Port Networks.

9 Customer LAN.

10 TN799 Control-LAN (C-LAN)

Avaya DEFINITY Servers R and SI

Avaya Communication Manager with an Avaya DEFINITY® Server R or Avaya DEFINITY® Server SI offers are targeted at the application-intensive market segment. This market segment is characterized by multiple locations that use many applications today, with the anticipation of adding more. Circuit-switched voice communication is critical for these customers. However, the strategic direction is IP Telephony, including voice/data network integration, unified messaging, and multimedia conferencing and collaboration. Most customers who use these solutions also use some form of centralized decision-making for all their sites. DEFINITY offers a flexible solution, scaling up to 25,000 endpoints. Thus, DEFINITY offers the ability to grow without the need to change an entire communications system.

For more information about these two platforms, see the following Web sites:

Customers and Business Partners:

http://www.avaya.com/

Other Avaya IP Telephony servers

Avaya IP Office

Avaya IP Office is another standalone Avaya platform that supports IP Telephony for the small to mid-size market.

Avaya IP Office is an IP PBX for 10 to180 stations. Avaya IP Office is not part of the Avaya Application Solutions offer, and thus is not covered extensively in this document. For more information about the IP Office, see:

http://www.avaya.com/ac/common/index.jhtml?location=M1H1005G1002F2012P3035N4240

Greenfield deployment

This chapter explains how to implement Avaya Application Solutions components in a Greenfield site. A Greenfield site is a business or an organization that does not have an existing communication system. Most Greenfield systems are deployed into new businesses and organizations, and these systems tend to be smaller in size. Occasionally, an established large organization may completely remove its existing system and install a new system. In these cases, the incumbent system is usually a leased service, such as a centrex service from a telephony service provider.

In general, most organizations want to protect their investment in their PBX communications system. Avaya provides ways for our circuit switched PBX customers to evolve from circuit switched systems to IP-enabled systems. This solution provides most of the advantages of IP Telephony with minimal equipment upgrades to an enterprise's existing PBX. The evolution approach is described in Evolution from circuit-switched to IP on page 67.

Components needed for Greenfield deployment

In a Greenfield deployment, the primary connection medium is IP. To provide the greatest flexibility and the lowest costs for a converged solution, most endpoints should be IP Telephones or IP softphones. A mixture of IP endpoints and circuit-switched endpoints places increased demand on Media Processor resources, and thus increases the cost of the deployment. Intersite communications should also be IP based. This can be done either through direct connections between IP Telephones or through IP trunks. Circuit-switched or TDM-based communications should be kept to a minimum. The primary TDM connections should be for PSTN access, where necessary, and connections to any analog telephones, modems, or fax machines that exist (Figure 25, A Greenfield IP Telephony deployment, on page 60).

In a Greenfield deployment, the emphasis is on IP Telephony. Multi-Connect systems that emphasize TDM connections are not generally recommended, except in special circumstances. Those circumstances include when there is a need for:

- Critical reliability
- Significant analog or DCP endpoints

Figure 25: A Greenfield IP Telephony deployment

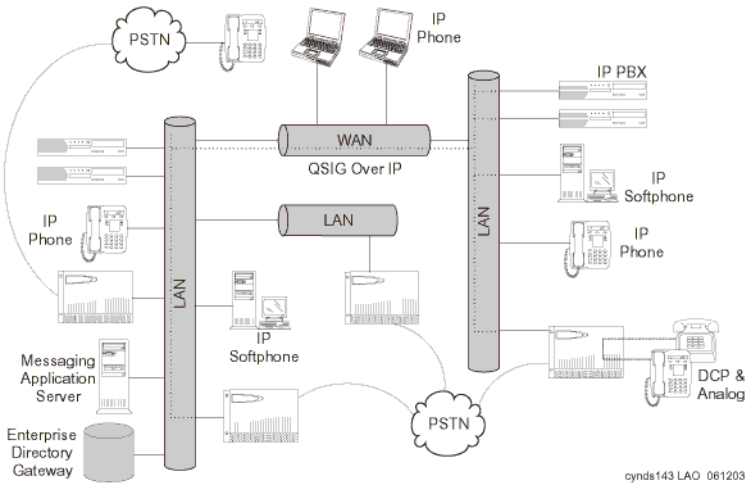

Media Server (H.323 Gatekeeper)

The Media Servers are responsible for running Avaya Communication Manager and controlling the Media Gateways and endpoints. The Media Servers control the dial plan translations and call routing, call setup and teardown, Call Detail Record (CDR) generation, traffic management. The Media Servers also offer H.323 gatekeeper functionality, and provide the extensive telephony features that are included with Avaya Communication Manager.

Avaya's Linux-based servers include:

- S8300 Media Server (Media Module resides in the Avaya G700 Media Gateway)
- S8700 Media Server

Avaya Communication Manager

Communication Manager IP capabilities and applications support voice over an IP network, and ensure that remote workers have access to communication system features from their PCs. Communication Manager also provides standards-based control between Media Servers and Media Gateways, which allows the communications infrastructure to be distributed to the edge of the network. The Communication Manager IP engine offers features that enable users to increase the quality of voice communications. Quality of Service (QoS) features enable users to optimize voice quality by assisting some routers in prioritizing audio traffic. Communication Manager Media Processors allow for hairpinning and shuffling. These features make voice communications more efficient by reducing both per-port costs and IP bandwidth usage. Avaya IP Telephony Solutions support trunks, IP communications devices, IP Port Networks, and IP control for Media Gateways. Avaya IP Telephony Solutions are implemented using various IP Media Processor circuit packs inside the Avaya Media Gateways. The IP Media Processors provide H.323 trunk connections, and H.323 voice processing for IP Telephones. H.323 signaling is handled by a C-LAN circuit pack or native processor Ethernet connectivity. The IP network can be extended across geographically disparate locations. With Communication Manager ISDN, Distributed Communication Services (DCS+), or QSIG services, Communication Manager can extend feature transparency, centralized voice mail, centralized attendant service, call center applications, and enhanced call routing across IP trunks.

For more information on Communication Manager architecture, see the Call processing chapter.

Media Gateways and Port Networks

Avaya Media Gateways support voice and signaling traffic that is routed between circuit-switched networks and packet-switched networks. Avaya Media Gateways support all the applications and the adjuncts that are supported by the Avaya DEFINITY® Enterprise Communications Servers, accommodating Call Center and Customer Relationship Management applications, messaging, remote workers, and remote offices. Avaya Media Gateways work with standards-based IP networks, and connect easily with the Public Switched Telephone Network (PSTN). The IP network infrastructure provides support for the communication between the Media Servers and the Media Gateways.

In a Greenfield installation, the recommended gateways are the G600 and the G700. The G600 houses traditional circuit switch boards and boards that support IP Telephony. The G700 houses Avaya Media Modules that provide ports for non-IP endpoints, including analog and DCP telephones.

Greenfield configurations

S8300 / G700 standalone (small-to-midsize enterprise)

An S8300 Server with a G700 gateway is designed for a small to mid-size office. Each G700 supports between 40 and 450 stations. The S8300 server can support up to 50 G700 Gateways. An S8300 server does not support G600 gateways or traditional port networks. An S8300 server can, however, be networked with other systems through IP trunks or PSTN tie trunks.

The S8300 Server fits into a Media Module slot in the G700 Media Gateway. As shown in Figure 26, A G700 system, on page 62, the G700 is a 2U 19-inch rack-mountable chassis. The S8300 Server contains a built-in Ethernet switch, an IP expansion module slot, four Media Module slots, and an Octaplane stacking module slot. The built-in IP Telephony module has the same functionality as the TN2302AP Media Processor (MedPro) circuit pack. An extra VoIP Media Module can be inserted in the G700 for extra media-processing resources. Other Media Modules support traditional endpoints.

Figure 26: A G700 system

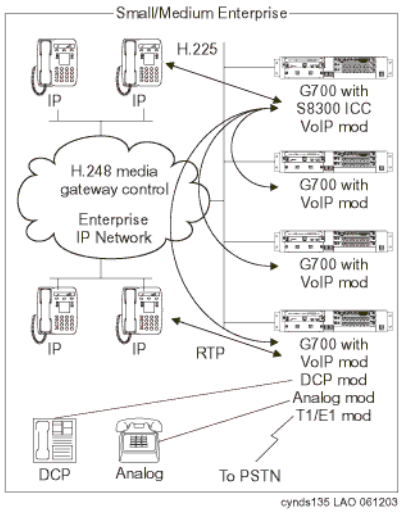

Medium-to-large enterprise solutions

S8700 / G600 IP-Connect

The S8700 IP-Connect system (Figure 27, S8700 IP-Connect system, on page 63) is a scalable solution that supports up to 64 G600 Media Gateways, each of which house up to four rack-mounted cabinets. This solution can support up to 12,000 IP stations. The S8700 Media Servers can be networked with other systems through IP or circuit-switched trunks to provide for significantly larger telephony networks. The control link between the Media Servers and the Media Gateways traverses the enterprise IP network. All G600 Media Gateways require IPSI circuit packs to provide the Gateway's control link. There is no traditional circuit switch (Center Stage Switch), and the media traffic flow is entirely through the enterprise data network. Each G600 has at least one Media Processor circuit pack, which provides the gateway between the TDM bus and the circuit pack's Ethernet connection for the audio streams. Each G600 also has at least one C-LAN, which provides H.323 signaling to IP endpoints.

Figure 27: S8700 IP-Connect system

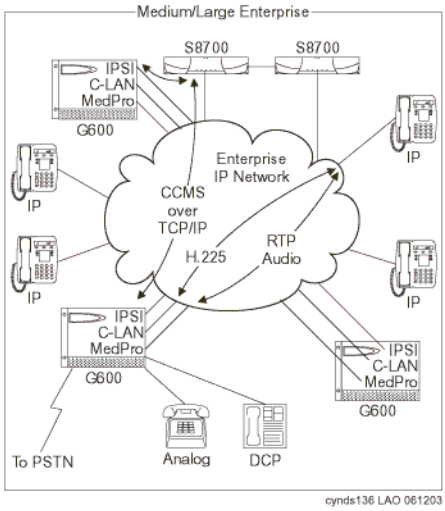

S8700 IP-Connect with remote G700s

The IP-Connect solution can be expanded to support a remote office with G700 Media Gateways in addition to G600 Gateways (Figure 28, S8700 IP-Connect with remote G700s, on page 63). This solution is designed for enterprises that require a high number of IP stations, but a low number of PSTN or traditional circuit-switched connections. The S8700 Server is the call controller that communicates with the G700 Gateways through the C-LAN. In this configuration, the C-LAN circuit pack acts as the front-end processor for both the G700 Media Gateways and IP endpoints.

Figure 28: S8700 IP-Connect with remote G700s

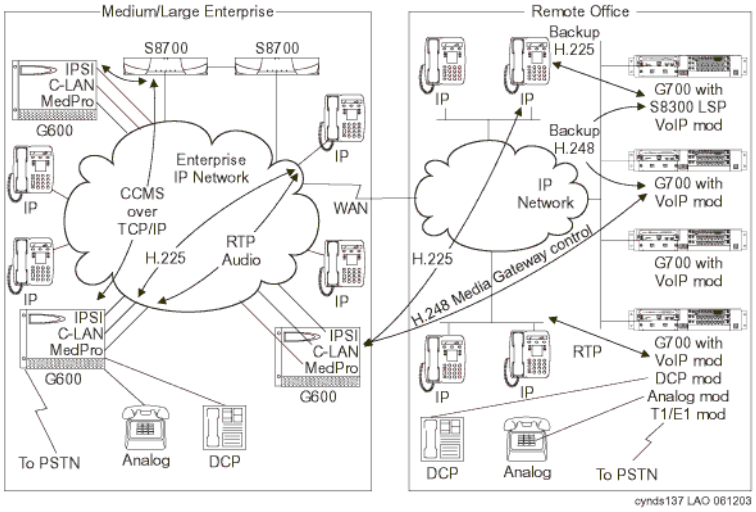

Required circuit packs for S8700 configuration

The circuit packs that are required for IP Telephony in a Communication Manager system include:

- TN2312 IP Server Interface (IPSI) for Port Network control
- TN799 Control LAN (C-LAN) for signaling and TCP/IP socket termination
- TN2302 Media Processor (MedPro) for the media flow

These circuit packs can reside in the G600 Media Gateways in widely-distributed locations.

The signaling connectivity path between the endpoint and the servers in the S8700 configuration is shown in Figure 31, Signaling flow, on page 65.

As shown in Figure 29, Signaling path (S8700 / G600 configuration), on page 64, an IP Telephone sends all IP Telephony signaling traffic to the C-LAN. The C-LAN multiplexes IP Telephone signaling messages, and sends them to the S8700 server through the IPSI.

In an S8300 configuration, a C-LAN circuit pack is not needed. All signaling traffic is sent directly to the S8300 Ethernet interface. The S8300 Server performs all C-LAN functions natively. The connectivity between the endpoint and the server is:

Endpoint⇔IP network⇔S8300 Media Server

Figure 29: Signaling path (S8700 / G600 configuration)

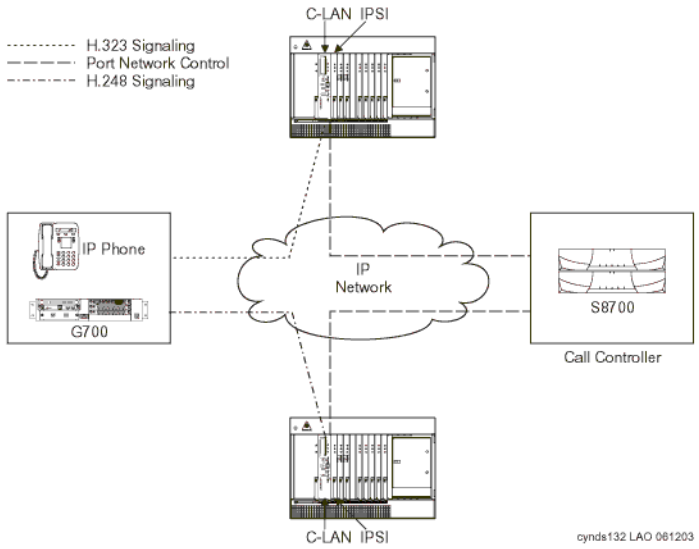

NOTE:
In the IP-Connect S8700 / G600 configuration each Port Network has an IPSI circuit pack.

As Figure 30, Media flow path (S8700 IP-Connect configuration), on page 65 shows, an IP Telephone sends all media streams to the MedPro. Once a call is established, if the remote endpoint is another IP Telephone, the media stream might shuffle (be redirected to the other endpoint) without requiring MedPro resources. Media Processors are also used to transport media streams in IP tie trunks.

Figure 30: Media flow path (S8700 IP-Connect configuration)

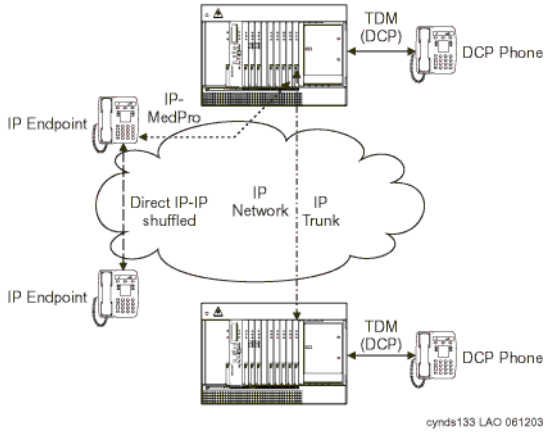

For detailed characteristics of the IPSI, C-LAN, and Media Processor circuit packs, see the Avaya Application Solutions platforms chapter.

Figure 31: Signaling flow

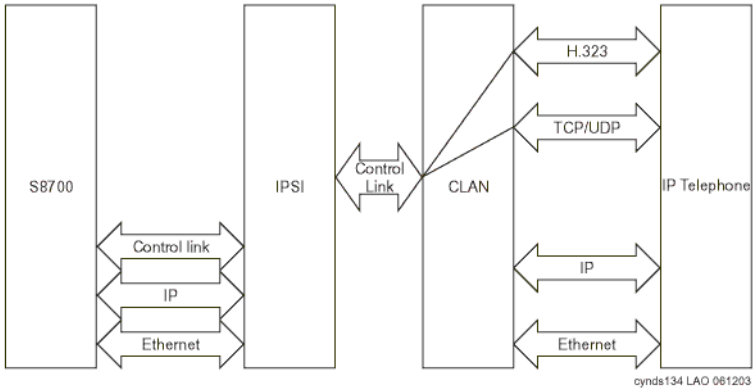

Communication devices

Avaya stations include IP Telephones and IP softphones. Avaya also supports IP trunks. In addition, all the Media Gateways support traditional terminals, such as analog, BRI, and DCP telephones. For detailed descriptions of Avaya IP Telephony endpoints, see the Terminals chapter.

Evolution from circuit-switched to IP

The Avaya DEFINITY® Enterprise Communications Server G3r has been the flagship product in the DEFINITY family of communications servers. As technology changed, Avaya was able to leverage the rapid advances in microprocessor technology to increase the capacity and processing power of the traditional DEFINITY platform (Figure 32, Traditional DEFINITY configuration, on page 67) to benefit our customers.

Figure 32: Traditional DEFINITY configuration

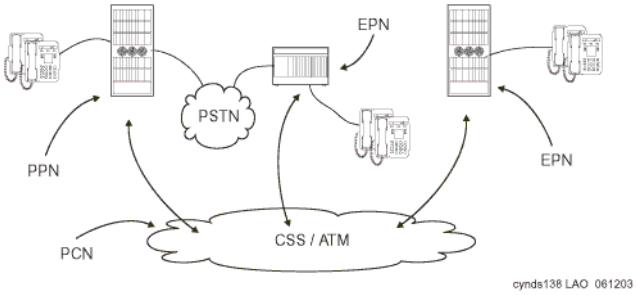

Avaya also has the objective to protect our customers' communications investment in the Avaya DEFINITY platform by helping our customers leverage their existing investments in Avaya solutions. Upgrading allows a customer to make a smooth transition to IP Telephony technology without sacrificing the features or reliability of their current DEFINITY. A customer can make small incremental investments to move from the circuit-switched world to a full IP PBX while retaining their investment in TDM-based equipment and connections. On the endpoints, moving to IP Telephony allows simplified moves, adds, and changes. It also simplifies the building wiring plan by sharing one Ethernet connection with both the IP Telephone and the desktop PC. It also adds IP mobility while retaining the rich set of DEFINITY features. For both IP Telephone and traditional circuit-switched telephone users, migrating to IP Telephony offers the opportunity to bypass tolls, and route traditionally metered long-distance calls across an unmetered IP network instead, saving operational costs.

With the S8700 Multi-Connect solution, Avaya is delivering a high-capacity server and a migration path from DEFINITY. The S8700 Media Server uses an industry-standard Linux operating system on an industry-standard server, which enables all endpoints to use Communication Manager. This solution allows customers to migrate to IP Telephony and to a higher performance processor without sacrificing the reliability of the G3r platform.

There are three stages to upgrading from a DEFINITY G3r to an Avaya Communication Manager IP PBX:

1. Replace the G3r processor with industry-standard S8700 servers.
2. Add IP circuit packs (C-LAN and MedPro) to support IP endpoints.
3. Consolidate multiple systems into a single system to simplify administration. Support network or processor failure conditions with LSPs deployed at remote sites.

Steps 1 and 2 can be reversed. The next five diagrams show the migration from circuit-switched DEFINITY to an IP-enabled S8700 Multi-Connect system with server consolidation and LSP survivability at a remote site.

Migration from DEFINITY Server R to S8700 Multi-Connect

Phase 1: Processor replacement

This section explains how an existing non-IP Avaya Communication Manager PBX can evolve to an IP Telephony-based solution. We will examine the case of an existing system that is based on the traditional PPN/EPN architecture, which will be applicable to all the G3 platforms.

When designing S8700 Multi-Connect systems, the call control network must be a nonrouted private network that is disconnected from the internal IP network of the enterprise. The equipment for this private network is provided as part of the S8700 Multi-Connect solution. See Voice quality network requirements for more information on setting up an IP network that can support IP Telephony.

The S8700 Media Server controls Media Gateways in a different manner than DEFINITY controlled Port Networks. With DEFINITY, the control path signaling shared the same transport media as the bearer channels. In a Multi-Connect system, call control signaling is established from the S8700 Media Server over an Ethernet connection to the TN 2312 IP Server Interface (IPSI) in the IPSI-connected Media Gateway. Non-IPSI connected Port Networks get their control information from the servers through the Center Stage Switch to one of the Port Networks that does contain an IPSI. An IPSI can support up to five Port Networks.

The full migration from a G3r to an S8700 server Multi-Connect system, with traditional or ATM center stage, involves the following simplified steps:

1. Decide which EPNs are to be IPSI connected, and replace processor complexes with IPSIs.
2. Install servers.
3. Install Ethernet switches.
4. Install UPS units.
5. For IPSI-connected port networks, upgrade each EPN.
6. Connect the duplication links.
7. Connect the servers and the IPSIs to the control LAN.
8. Sequentially bring up the duplicated servers.

Figure 33, S8700 Media Servers (Multi-Connect configuration), on page 69 shows the completion of Phase 1, an S8700 Multi-Connect system that supports only traditional endpoints.

Figure 33: S8700 Media Servers (Multi-Connect configuration)

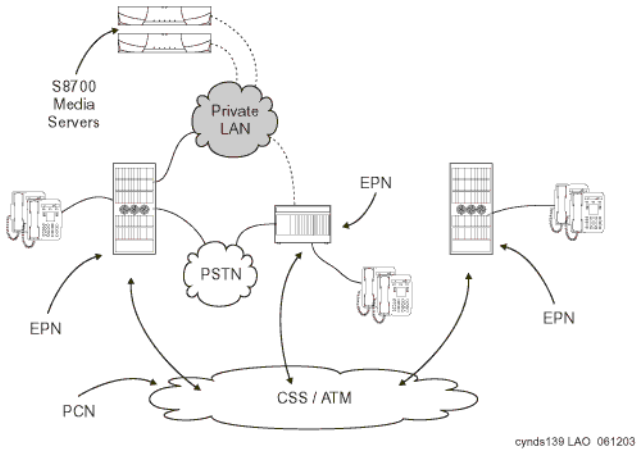

NOTE:
In the traditional PBX system, signaling and bearer traffic for all calls connects through the TDM buses within Port Networks and the ATM or traditional center stage

Phase 2: IP-enable the Port Networks to support IP endpoints

Port Networks, with the addition of IP enabling circuit packs, are able to serve as Media Gateways, representing the integration of IP and TDM telephony.

IP-enabling the existing system incrementally for IP endpoint support requires only two circuit packs:

- TN799DP Control-LAN (C-LAN) for IP call signaling
- TN2302 Media Processor (MedPro) for IP audio media processing, including media streams that are intended for IP softphones and IP Telephones

Signaling and bearer communication can connect through both the traditional TDM/center stage route and the IP network infrastructure. This gentle migration to IP Telephony (Figure 34, IP-enabled DEFINITY configuration, on page 70) might have minimal impact on an existing, non-IP system, while simultaneously enabling all new IP endpoints to fully access all the Communication Manager features.

NOTE:
If Phase 2 is implemented before Phase 1, the system will resemble Figure 34, IP-enabled DEFINITY configuration, on page 70. Also, a system that implements Phase 2 but not Phase 1 cannot support as many IP endpoints as a system that implements both Phase 1 and Phase 2 (see Avaya DEFINITY Servers R and SI on page 58).

Figure 34: IP-enabled DEFINITY configuration

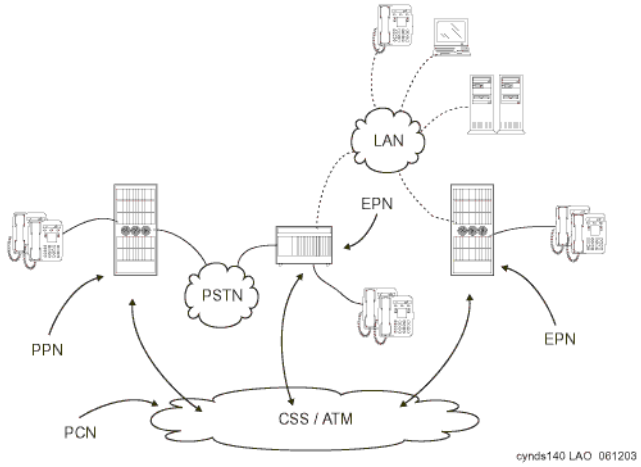

At this stage, the media flow between two IP endpoints can be “shuffled.” That is, the media flow proceeds directly between both endpoints without requiring Media Processor resources. Shuffling may be used across multiple sites or multiple Avaya switches. Likewise, calls between an IP endpoint at one site and a circuit-switched endpoint at another site can be shuffled so that the media stream flows between the IP Telephone and the Media Processor circuit pack in the Port Network that is connected to the circuit-switched endpoint. By using the IP network to the greatest extent possible, enterprises can minimize the use of expensive circuit-switched trunks.

Figure 35: IP-enabling the S8700 Multi-Connect configuration

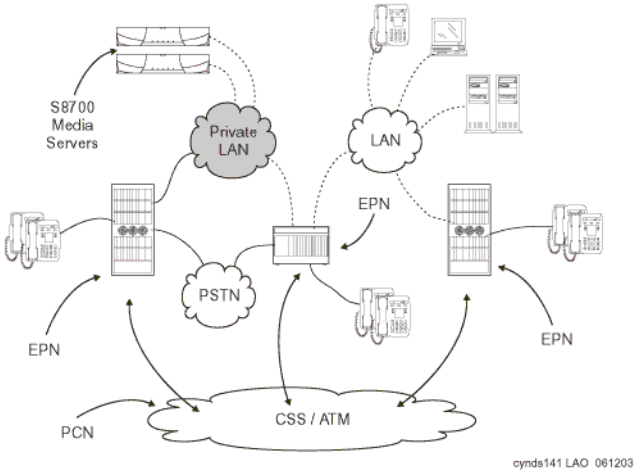

Phase 3: server consolidation

Traditionally, some enterprises have elected to use multiple DEFINITY systems at remote sites to protect against a circuit failure on the center stage network bringing down an entire remote site. With the decision to run multiple servers came the need for additional administrative resources and a more complex dial plan. Today, through the use of IP Telephony technology and the enhanced processing capabilities of the S8700 Media Servers, Avaya has a solution to consolidate smaller remote DEFINITY servers, such as ProLogix or DEFINITY ONE into an S8700 Multi-Connect system, while maintaining remote site survivability in the event of a network or processor failure. By consolidating multiple DEFINITY servers into one S8700 system, an enterprise can realize cost savings in simplified administration and a simplified dial plan. With support for up to 36,000 endpoints, the S8700 system has the scalability to support a remote site's server consolidation.

Consolidating remote site servers into an S8700 system requires a G700 Media Gateway with S8300 Local Spare Processor (Figure 36, S8700 / G700 system with S8300 as Local Spare Processor (LSP), on page 71). In the event of a network or processor failure, the S8300 server takes over active call processing and gateway and endpoint management for the remote site, allowing continued operation with no loss of features until the outage is repaired.

Figure 36: S8700 / G700 system with S8300 as Local Spare Processor (LSP)

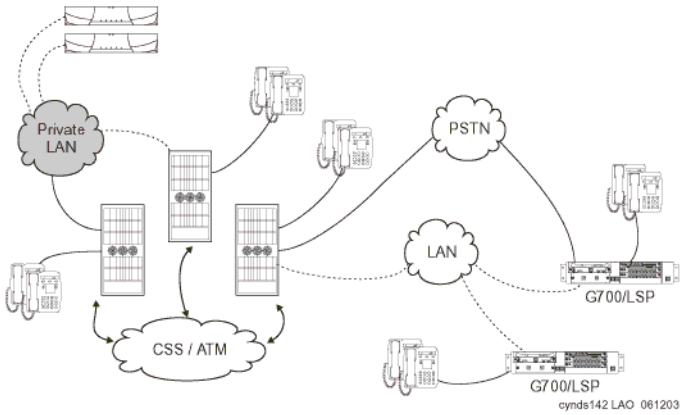

Because the G700 relies on IP Telephony technology, this option is especially attractive to customers who decide to use a majority of IP endpoints at the remote site. This solution will, however, continue to support analog endpoints and DCP endpoints. Analog trunks and ISDN trunks are also supported. To decrease operational expenses, the circuit-switched trunks back to the main site can be replaced with IP trunks.

Call processing

This chapter explains the features, the strengths, and the architecture of Communication Manager call processing.

Communication Manager capabilities

This section lists and explains the major features of Avaya Communication Manager that are over and above traditional PBX features:

- Terminal mobility
- User login
- IP softphone
- Analog terminal

This chapter emphasizes the call processing components of Communication Manager and its architecture, and briefly discusses IP-related applications in the areas of telephony, convergence, networking and call routing, mobility, telecommuting, and remote office. This chapter is not an exhaustive resource for Communication Manager features.

Communication Manager operates on the Avaya Media Servers, and on the existing family of DEFINITY and Avaya IP600 communication servers. Communication Manager seeks to solve business challenges by powering voice communications and integrating with value-added applications. Communication Manager provides user and system management functionality, intelligent call routing, application integration and extensibility, and Enterprise Communications networking.

For more information on Communication Manager, see:

- *Avaya Communication Manager Feature List* contains details about the full capabilities of Communication Manager (more than 700 features) by release, application area, or both:
  - Avaya U.S. BusinessPartners

    https://www.avaya.com/doc/gpp/public/pss/category/cs/eclips/multivantage_software/index.html
- *Overview for Avaya Communication Manager*, 555-233-767, contains descriptions of each feature.
- *Highlights of Avaya Communication Manager*, 555-233-778, provides a delta view of new features in Communication Manager Release 1.1 and Release 1.2.

  Both of these documents are available at:

  http://support.avaya.com/japple/css/japple?PAGE=avaya.css.ProductDetail&temp.groupID=107529&temp.selectedFamily=152705&temp.selectedProduct=136527

Voice and multimedia networking

Intelligent networking and call routing

With Avaya Communication Manager, servers can use IP trunks across an IP network to communicate between switches without the need for dedicated leased lines. With Communication Manager, IP trunks can use Distributed Communication Services (DCS+) or QSIG Services to extend feature transparency, centralized voice mail, centralized attendant service, call center applications, and enhanced call routing across IP trunks.

IP Port Network / Media Gateway connectivity

IP PNC allows S8700 Media Servers and G600 Media Gateways to be connected over IP networks. Avaya Communication Manager uses a proprietary method to package voice and signaling messages over IP. This method allows deployment of communications systems throughout a customer's data network.

H.248 Media Gateway control

Communication Manager uses the standards-based H.248 media gateway control protocol to perform call control of Avaya G700 Media Gateways. H.248 defines a framework of call control signaling between the intelligent Media Servers and multiple Media Gateways. H.248 controls both IP (H.323) and non-IP connections into a media gateway. H.248 has been extended by Avaya to also tunnel proprietary CCMS messages, to allow for enhanced call handling.

Call Processing

Communication Manager gatekeepers

A gatekeeper is an H.323 entity on the network that provides address translation and controls access to the network for H.323 endpoints. For Communication Manager platforms, these are the Avaya S8300 and S8700 Media Servers through the C-LAN circuit packs. H.323 RAS (Registration, Admission, and Status) Protocol messages are exchanged between them and the IP endpoints for the endpoint registration.

All IP endpoints (IP softphones, IP agents, and IP Telephones) H.323 voice applications should register with an Avaya gatekeeper before any calls are attempted. Communication Manager enforces call signaling (Q.931) and call control (H.245) channels from endpoints to terminate on the gatekeeper. This allows Communication Manager to provide many of its calling features to H.323 calls.

Registration and alternate gatekeeper list

The RAS protocol is used by the IP endpoint to discover and register with the Communication Manager gatekeeper. The discovery mechanism uses unicast IP facilities.

When registration with the original gatekeeper (C-LAN or S8300) IP address is successful, the switch sends back the IP addresses of all the C-LANs (or LSPs) in the IP Telephone's network region. These addresses are used if the call signaling to the original C-LAN circuit pack fails. Figure 37, Discovery and registration process to the gatekeeper, on page 75 shows the registration process.

Figure 37: Discovery and registration process to the gatekeeper

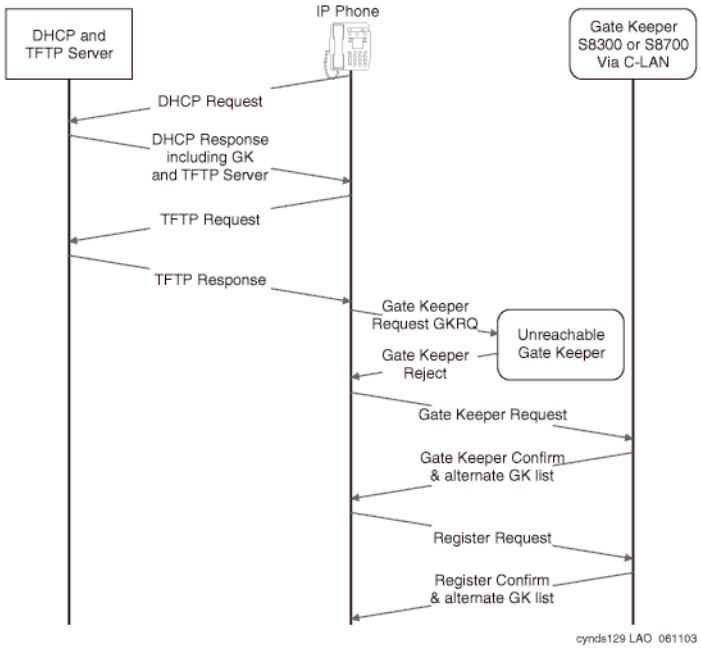

Call signaling

Communication Manager implements the gatekeeper routed call model of H.323. The registration process that is described above allows the endpoint and the Communication Manager gatekeeper to exchange addresses to establish a TCP connection for a "call signaling" channel (the H.323/H.225 channel). Once the TCP connection is established for call signaling, the H.225.0/Q.931 signaling protocol is used over that connection to route the call and exchange addresses necessary to establish a second TCP connection. This second TCP connection is used for "media control" (the H.245 channel).

When Communication Manager chooses to route the media flow streams through the switch, it selects and allocates available media processor resources, and sets the corresponding circuit packs up to receive and send the media stream or streams from/to the endpoints using the negotiated capabilities for each terminal. Each terminal is told to send its media stream or streams to the appropriate Media Processor circuit pack. The switch connects the two media streams, and thus completes the bearer path between the terminals.

Media stream handling

Media processing

The basic functions of the TN2302 Media Processing (MedPro) circuit pack include:

- Taking media streams off the IP network, terminating RTP/UDP (adjusting for variable delay in arrival rate), and converting them into PCM audio for transmission on the TDM bus.
- Taking media streams from the TDM bus, encoding them with the proper codec, and transmitting them as RTP packets to an IP endpoint.
- Originating and terminating an RTCP control channel for each media stream.

The particulars of the media conversion that is to be performed on each media stream are controlled by Communication Manager. The Quality of Service (QoS) information obtained from the RTCP channel is passed from the circuit pack to Communication Manager.

DTMF tone handling

The Media Processor circuit pack listens for and detects DTMF tones coming from the TDM bus, strips them out of the audio stream, and sends a message to the Media Server indicating that it has done so. The Media Server in turn generates and sends the appropriate H.245 tone message to the endpoint that is receiving the audio stream. The receiving endpoint then plays the specified tone. Compressed codecs, such as G.729, generally do a poor job of passing DTMF tones. By sending tones out of band, fidelity is maintained. This method is useful when connecting to a voice mail or an integrated voice response (IVR) system, where DTMF digits are used to navigate through prompts.

When this capability is used on an H.323 tie trunk between Communication Manager switches, the switch that receiving the H.245 tone message plays the required tone onto all the ports receiving the audio stream.

Media stream for audio conferencing

When calls between IP endpoints are conferenced, the media streams must be routed through the MedPro circuit pack.

Communication Manager allows the audio streams from different parties to come into different Media Processor circuit packs. Each Media Processor sends its received signal to the TDM bus in Pulse Code Modulation (PCM) format. All the other processors serving endpoints on the call can then receive and sum the audio signals coming from all parties, and send the resultant composite audio stream to the IP parties that it supports.

Figure 38, MedPro support of a three-party audio conference, on page 77 provides an example to show how the MedPro circuit pack is configured for a three-party H.323 audio conference using G.729. This conference is conventional in that it uses TDM bus timeslots to allow each party to listen to all of the other parties. However, a more efficient form of conferencing is possible when all the parties are IP endpoints, where the audio streams are multiplexed directly on the Media Processor circuit pack, and no TDM timeslots are used. To establish the configuration shown in Figure 38, MedPro support of a three-party audio conference, on page 77, the DSP resources on the MedPro are allocated as needed to audio conferencing. Communication Manager balances the available media processing resources, effectively sharing load among multiple MedPros.

Figure 38: MedPro support of a three-party audio conference

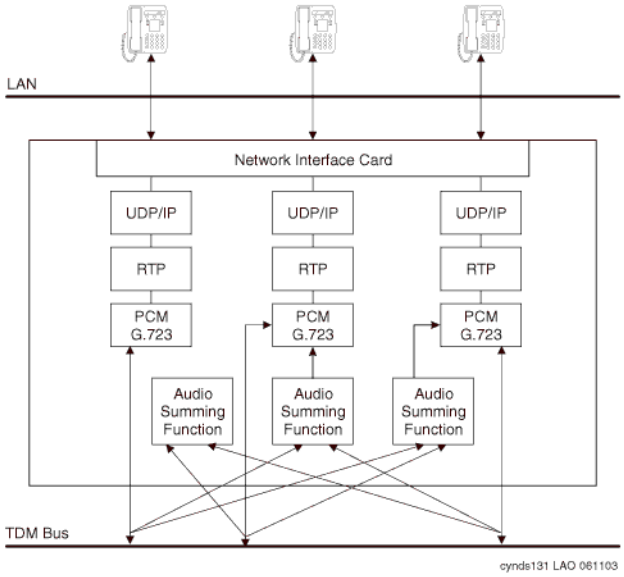

Separation of Bearer and Signaling (SBS)

In an Avaya IP Telephony system, call signaling and bearer traffic may be routed over separate paths. This is useful for a remote branch office with only limited WAN bandwidth back to headquarters. Call signaling traffic can be routed across the WAN, while bearer traffic is sent over the PSTN.

IP-based (H.323) trunks

In circuit switched networks, trunks provide the means to interconnect PBXs with each other and to the PSTN. Connection to the public network allows PBX station users to call and be called by terminals that are not part of the PBX private network of the PBX. An analogous arrangement exists in packet-switched IP networks.

H.323 trunks connect H.323 systems or gateways over IP networks, similar to circuit-switched tie trunks.

A set of Communication Manager switches can each be attached to an IP network, and voice and fax calls can flow between them in the usual manner except that the call signaling and audio/fax streams are carried over the IP network. The signaling is carried on a TCP connection through the C-LAN circuit packs, and the audio and fax streams are carried between switches through the Media Processor circuit packs.

The benefits of using IP trunks include:

- Reducing long distance voice and fax expenses
- Facilitating global communications
- Providing a fully functional network with data and voice
- Converging and optimizing networks by using the available network resources

IP trunk calls can be compressed to save network bandwidth. Repeated compression and decompression (transcoding) results in a loss of data at each stage and degrades the final quality of the signal. The maximum recommended number of compression cycles on a call is three. Normal corporate voice calls or fax calls typically go through fewer than three compression cycles.

H.323 trunks can also connect to other vendors' compliant PBXs that serve as an H.323 gateway or gatekeeper.

IP tie trunks

IP tie trunks are used to connect switches to one another. When an H.323 trunk is used to interconnect two switches, the trunk can also carry standard (QSIG) and proprietary (DCS+) signaling for interswitch feature transparency. The location of each other node (switch) in the network is administered, and node selection is based on the dial plan and call routing features such as AAR/ARS.

H.323 tie trunks are administered as a new type of trunk group. Instead of administering ports as members of the trunk group, only the number of channels must be specified. Each channel is analogous to a member trunk. In addition, an H.323 tie trunk can be made a member of a signaling group so that a virtual D-channel can be administered and used to carry feature transparency information.

Trunk signaling

Several variations of H.323 signaling must be accommodated for the variety of trunks supported by Communication Manager. These are specified as options in the trunk group administration. When the H.323 trunk is used as a tie trunk to another vendor's switch, gateway, or gatekeeper, Communication Manager sets up a separate TCP connection for each call.

For more on trunk signaling, see:

http://support.avaya.com/elmodocs2/comm_mgr/r1_3/cd823_3/233767_4/233767_4.pdf

Mobility

IP Telephones or IP Softphones

IP Telephones allow access to the features of Communication Manager without having to be tied to one location. One of the major benefits of IP Telephones is that you can move the telephones around on an IP network just by unplugging them and plugging them in somewhere else. One of the main benefits of IP Softphones is that you can load them on a laptop computer, and connect them to the Communication Manager switch from almost anywhere. Users can place calls, and handle multiple calls on their PCs. For detail description and features of IP Telephones and softphones see the Terminals chapter.

Extension to Cellular

Extension to Cellular is an integrated mobility solution that offers users the freedom to work anywhere, anytime, using any type of cellular or wireless telephone. With Extension to Cellular, calls to an office number are extended to a cellular telephone, allowing users to receive work-related calls wherever they are and whenever they need. Additionally, the cellular telephone can be administered so that when a user calls into the office, the user's name and office telephone number appear in the caller ID display of the telephone being called. When the Extension to Cellular cell phone is administered to send office caller ID, the user also has the option of picking up an ongoing Extension to Cellular cell phone call on the office telephone when the user enters the office.

Extension to Cellular works over PRI as well as an IP trunk interface. The cell phone user receives the same features and capabilities for incoming calls as a caller ID-enabled analog telephone that is connected directly to the Avaya Communications Server. Extension to Cellular provides this capability regardless of the cell phone's cellular service provider or the cellular standard in use.

Communication applications

Call Center

The Avaya Call Center provides a total solution for a customer's sales and service needs. Building on the performance and flexibility of the Avaya Communication Manager, customers can select from a powerful assortment of features, capabilities, and applications that are specially designed to enhance call center operations.

The objective of this offer, which involves new and existing versions of Avaya Media Servers and Communication Manager, as well as a host of attached call center peripherals, is to improve Avaya's Call Center offers by supporting increased capacities. These capabilities include 6-digit and 7-digit extensions, LAN backup of Call Management System for the High Availability offer, and customer requested enhancements to be made available in a single global release.

Avaya Call Center applications are designed to efficiently connect each caller with the representative who is best suited to serve that caller. Avaya Communication Manager begins the process by capturing information about the caller even before the call is routed. That information is integrated with existing databases, and the combined data is used to match caller to agent.

Avaya Communication Manager integrates with a variety of Call Center applications like the Avaya Call Management System S) for real-time reporting and performance statistics, and with Avaya Business Advocate for expert predictive routing according to incoming calls, not just historical data.

Avaya Call Management System (CMS)

The Avaya Call Management System collects call traffic data, formats management reports, and provides an administration interface for Automatic Call Distribution (ACD) on your Communication Manager. CMS helps enterprises manage the people, traffic load, and equipment in an ACD environment by answering such questions as:

- How many calls are we handling?
- How many callers abandon their calls before talking with an agent?
- Are all agents handling a fair share of the calling load?
- Are our lines busy often enough to warrant adding additional ones?
- How has traffic changed in a given ACD hunt group over the past year?

Site Statistics for Remote Port Networks forwards location IDs to CMS to provide call center site-specific reports.

CMS reliability and redundancy

Dual Links to CMS provides an additional TCP/IP link to a separate CMS for full, duplicated CMS data collection functionality and high availability CMS configuration. The same data are sent to both servers, and the administration can be done from either server. The ACD data is delivered over different network routes to prevent any data loss from such conditions as ACD link failures, CMS hardware or software failures, CMS maintenance, or CMS upgrades.

Computer Telephony Integration (CTI)

Computer Telephony Integration (CTI) enables Communication Manager to be controlled by external applications, and allows integration of customer databases of information with call control features. CTI is a LAN-based solution that consists of server software that runs in a client/server configuration.

CTI opens up Application Programmer Interfaces like ASAI, Telephony Services Application Programming Interface (TSAPI), and Java Telephony Application Programming Interface (JTAPI), which can be used to control the server from an external application.

Application Programming Interfaces (APIs)

Communication Manager supports the following APIs to interface with other applications:

- Adjunct Switch Application Interface (ASAI) allows adjunct applications to access a collection of Communication Manager features and services. Integration with adjuncts occurs through APIs. ASAI is part of Avaya Computer Telephony.
- DEFINITY Application Programming Interface (DAPI) for accessing control and data paths within Communication Manager.
- Java Telephony Application Programming Interface (JTAPI) is an open API supported by Avaya Computer Telephony that enables integration to Communication Manager ASAI.
- Telephony Application Programming Interface (TAPI).
- Telephony Services Application Programming Interface (TSAPI) is an open API supported by Avaya Computer Telephony that allows integration to Communication Manager ASAI.

Best Services Routing (BSR) polling

Best Service Routing (BSR) polling over QSIG Call Independent Signaling Connections (CISCs) and Temporary Signaling Connections (TSCs) provides the ability to do BSR polling between multiple site over H.323 IP trunks without requiring an ISDN PRI B-channel. QSIG CISC/TSCs are used by BSR polling software to reduce the need for the IP Media Processor circuit pack, thereby making BSR a cost-effective, multi-site solution for an enterprise-wide contact center.

Meet-me conferencing

Meet-me conferencing provides conferencing of up to six parties from any communication device that is internal or external to the business network. This feature does not require any special hardware. Meet-me conferencing uses a software approach that is based on Vector Directory Number (VDN) vectors and announcements. An announcement source is necessary to use meet-me conferencing. Supported announcement sources include:

- Voice Announcements over the LAN (VAL) circuit pack
- G700 local announcement
- Integrated Announcement circuit pack
- An external source

Avaya LAN switching products

This chapter discusses how Avaya products add value to an IP Telephony deployment.

Converged infrastructure LAN switches

P330

The Avaya P330 Ethernet stackable switching system (Figure 39, Avaya P330 Ethernet stackable switching system, on page 83) features modular functionality, high port density, carrier-class reliability, and building-block simplicity, allowing you to phase in exactly the features and functions for today's application-driven networks. In combination with the Avaya IP Telephony Media Gateways and Avaya IP Telephones, the Avaya P330 stack becomes an important component of a complete, integrated data and voice solution for an enterprise.

Figure 39: Avaya P330 Ethernet stackable switching system

The Avaya P330 family of affordable, stackable Ethernet switches can operate as part of a total workgroup solution. You can stack the P330 switches to allow for pay-as-you-grow scalability, from a few to hundreds of ports and N+1 redundancy. Adding a single P330 series routing switch brings multilayer routing to the entire switch stack. Everything comes together with building-block simplicity, and often with no configuration.

The P330 family includes:

- Avaya Octaplane® stacking fabric
- Up to 640 10/100BASE-Tx ports, and up to 100 100/1000BASE-T ports in a stack
- Up to 120 Gigabit Interface Converter (GBIC) Small Form-factor Pluggable (SFP) ports in a stack
- Uplink modules with 10/100/1000, Asynchronous Transfer Mode (ATM) and wide area network (WAN) connection
- Layer 2 switching with the P333T, P333T-PWR, P334T, and P332MF
- Multilayer functionality (Layer 2 switching and Layer 3 routing) with the multilayer P332G-ML, multilayer P332GT-ML, P333R and P333R-LB
- Load Balancing with P333R-LB
- Power over Ethernet with the P333T-PWR (complies with IEEE 802.3af draft standard)
- Avaya Switch Architecture for Extreme Resiliency (SAFER) Technology™ with port, switch, and stack redundancy
- New Equipment Building Standards (NEBS) Level 3 certified
- Multiple management options, including command line interface, Web-based manager, and Avaya MultiService Network Manager management software
- Multiple monitoring options including four groups of remote monitoring (RMON) and AnyLayer SMON (Switch Monitoring) application
- Every member of the Avaya P330 family of products is available in both AC and DC versions. A stacking slot enables you to assemble a stack of up to 10 switches with superb resilience and redundancy, allowing all ports to be dedicated to networking.

C460

For enterprises looking to deploy Avaya Communication Manager Communications Applications, the Avaya™ C460 converged multi-layer switch is a highly resilient network platform designed to provide high-availability support for mixed data and IP Telephony deployments.

The Avaya C460 features a compact modular six-slot chassis with the following main characteristics:

- Four I/O slots and two Supervisor slots
- Fully redundant architecture (including switching fabric, supervisor modules and PSUs)
- Power over Ethernet (PoE) support with the FE ports
- High density - up to 192 FE PoE ports and 48 GE ports
- Fabric switching throughput of 64Gbps at Layer 2 and 48Mpps routing at Layer 3
- Policy and QoS mechanisms
- Full router functionality
- Wire-speed Layer 3 forwarding on all ports
- Optimal use of physical chassis size (10U)
- 300W or 1000W (for PoE support) power supplies

The C460 full redundancy (supervisor and fabric, power supply, link and port interfaces, router processor, and fans), high port density and powerful Layer 2 and Layer 3 wire-speed switching engine make it suitable for robust network infrastructure. The C460 offers advanced management and monitoring capabilities using complete GUI tools, including the SMON and Any-layer SMON applications in the Avaya Information Management software.

The C460's available I/O modules include:

- 48 10/100 PoE port Inline Power module
- 48 10/100 PoE port Inline Power + 2 GBIC (SFP) Gigabit Ethernet port module
- 12 GBIC (SFP) ports Gigabit Ethernet module
- 48 10/100 port Ethernet module
- 48 10/100 port Ethernet + 2 GBIC (SFP) Gigabit Ethernet port module

The C460 extends Avaya Convergence solutions to the network edge by providing advanced network capabilities, including Quality of Service (QoS), high performance, advanced power management, security and manageability. Designing a converged network infrastructure using this highly-resilient, modular, high-performance solution ensures a lifespan of the network, which reduces the cost of ownership and improves return on investment.

With its flexible configuration options and high-capacity performance, the C460 can also be deployed as a distribution layer switch or as the network backbone for small to medium enterprises looking for a reliable modular solution.

For enterprises deploying Avaya Communications Manager for mission-critical call center and large-scale campus environments, the C460 offers an ideal IP Telephony platform that combines fault tolerance, network responsiveness for business continuity, and integrated management and monitoring for converged networks.

P580

The Avaya P580 MultiService switch (Figure 40, Avaya P580 and P882 MultiService switches, on page 86) brings powerful switching and routing capabilities, improved virtual local area network (VLAN) performance, and Quality of Service to your communications network.

Figure 40: Avaya P580 and P882 MultiService switches

Equally at home with Layer 2 switching and Layer 3 routing, the P580 brings increased flexibility to your LAN, and can support your next-generation data network. The P580 is also ready for voice and video, as well as data. The P580 delivers the bandwidth, the network interfaces, and the Quality of Service (QoS) that you need. It can be configured for resilience and first-class management capability. Simply stated, the P580 is the right choice for both bandwidth-starved campus backbones and high-performance workgroups.

The P580 MultiService Switch offers:

- Backplane capacity of 55 gigabits per second (Gbps)
- Up to 40,000,000 packets per second (pps) Layer 2 switching
- Up to 40,000,000 pps Layer 3 routing
- Layer 2 and Multilayer (Layer 2/Layer 3) modules
- Up to 288 10/100 Ethernet ports
- Up to 144 fiber Fast Ethernet ports
- Up to 48 full-duplex Gigabit Ethernet ports
- Up to three 10-Gigabit ports per switch
- ATM uplinks at OC-3c/STM-1 and OC-12c/STM-4c speeds for multiservice networking
- Fault tolerant power, switch, links, management
- Unique Open Trunk™ VLAN interoperability
- Class of Service/Quality of Service
- Eight priority queues in hardware
- Interoperability with existing Avaya 50-Series media modules or 80-Series modules

The P580 takes the form of a compact 10.5-inch high rack-mount chassis with six payload slots. The slots accept both Avaya 50-Series and 80-Series media modules. Spare crossbar switch elements, crossbar controllers, and power supplies can be mounted for superb N+1 redundancy. The maximum port density is 48 ports per module.

P882

The Avaya P882 MultiService switch (see Figure 40, Avaya P580 and P882 MultiService switches, on page 86) delivers the QoS that you need for next-generation applications and delay-sensitive, time-critical traffic.

Heading the Avaya family of network switches, the P882 combines powerful backbone switching, high-port density, performance, and capacity. This high-density, highly scalable backbone switch supports data, voice, and video, with up to 768 ports and up to 139 Gbps backplane bandwidth. The P882 is enabled for voice/data/video applications without the need for expensive forklift upgrades.

The Avaya P882 provides:

- Up to 139 Gbps switching/routing performance
- Up to 768 10/100 ports per switch
- Up to 128 Gigabit Ethernet ports
- Up to 384 100BASE-FX ports
- Up to eight 10-Gigabit ports per switch
- Data, voice, and video application capabilities
- ATM uplinks at OC-3c/STM-1 and OC-12c/STM-4c speeds for multiservice networking
- QoS for guaranteed delivery of time-critical, delay-sensitive traffic
- Interoperability with existing Avaya 50-Series media modules or 80-Series modules for Investment protection
- Avaya Switch Architecture For Extreme Resiliency (SAFER) Technology™ for no single point of failure

The P882 incorporates a rack-mount chassis with 16 payload slots. The slots accept both Avaya 50-Series and 80-Series media modules. Spare crossbar switch elements, crossbar controllers, power supplies, and fans can all be mounted for superb N+1 redundancy. The maximum port density is 48 ports per module.

Avaya Power over Ethernet switches

Available PoE Switches options

In a Power over Ethernet (PoE) switch, power and data are combined in the Ethernet switch and sent on a single cable incorporating Power over Ethernet technology, thus simplifying the power management and cabling plant and saving rack space.

Avaya offers the P333T-PWR PoE switch.

P333T-PWR

The P333T-PWR has all the features of the Avaya P333T:

- Same Layer 2 & stack features like P333T
- 24x10/100BASE-T ports
- Expansion slot for all optional P330 Ethernet ATM & WAN expansion modules

- Slot for the Octaplane Stacking Module
- Up to 10 units or a mixed with other P330 family switches in a single P330 stack

Besides the added ability to apply power to PDs, the P333T-PWR has the same L2 features and QoS capabilities as the other P330 family products. Ports with no PD connection have all the L2 features and meaning as before. All Layer 2 features can be configured on the P333T-PWR. However, LAG, Port Redundancy and Inter-module are not applicable for powered ports, even though they can be configured.

All of the high-end management for the Layer 2 and for the newly-added power features are available as before through the CLI, Embedded Web Manager, & MSNM.

All ports on the switch provide class 0 (0.5-16.5 Watt) in-line power to any PD over the signal pairs (1/2 & 3/6).

NOTE:
The expansion module ports are not powered.

P333T-PWR power is carried over the signal leads providing remote -48V power feeds on all 24 ports. This allows the PD to be up to 100m away from the switch. Each port performs a standard compatibility detection process before power is supplied to the Ethernet lines. If the PD is removed or the link is interrupted, the port polling mechanism notices and power is cut off to the inline port while the detection process is applied again.

The P333T-PWR applies power to the port only after it detects that a PD is actually connected to the port. Each PD has a resistance range known as a "signature." According to the "signature," the P333T-PWR knows what power has to be supplied to the device.

Load detection is performed every 240 ms. All ports are checked for the resistance "signature" on a port-by-port basis. Only non-powered ports participate in the periodic load detection. Disconnection of ports that were previously powered is detected by sensing a power drop. Disconnected ports then automatically join the periodic load detection cycle.

Each port of the P333T-PWR is protected against channel overload, short circuit, and reversed polarity that might be caused by faulty connection between two feeding channels or by a crossed cable connection.

Power priority mechanism

Since the internal power supply of the P333T-PWR may not be capable of driving all the ports simultaneously, a priority mechanism is implemented. This priority mechanism determines the order in which ports will be powered on after boot, and powered off if the power resources of the module are exhausted. Three user-configurable port power priority levels are available: low, high & critical. Within each priority level the lower the port number, the higher the priority (by default all the ports have low priority).

Disconnected power will be automatically reconnected to the PDs based on their priority, whenever there is an available power budget. Immediately after the P333T-PWR has booted-up, it starts to supply power to the ports where a load is detected. Ports are powered up one after another, based on the port priority, until the 182W limit is reached. Power calculation is based on the actual power consumption of the PD. After this, no more ports are powered up until the total power consumption drops lower than 182W. The remaining 18W are spare for change in the power draw of PDs.

Backup power supply

Please check the Avaya support site for the recommended power supply solution:

support.avaya.com

Midspan Power Unit

Description

The official name for this device is the 1152A1 Power Unit, but the Midspan Power Unit can also be called a powered data unit (PDU) or a power over Ethernet (POE) device. The Midspan Power Unit is 1U in height (1.75 inches or 4.44 cm) and has 24 RJ45 data input jacks on the bottom row, and 24 data and power output RJ45 jacks. Data flow is unaffected if power is disrupted and if the endpoint does not require power. An example is a laptop computer that is connected to the 1152A1. The computer does not receive power from the 1152A1. If the 120-volt power is disrupted to the 1152A1, the computer data stream would not be affected. The 1152A1 unit provides a maximum of 200 watts or a peak of 16.8 watts per port. This unit powers any device that conforms to the 25-K Ohm resistive signature defined in the IEEE 802.3-2003 af standard. This unit also powers devices that use the nonstandard capacitive signature, such as Cisco IP telephones. The 1152A1 provides positive voltage on pins 4/5 and negative voltage on pins 7/8, which is one of the three methods as described by the IEEE 802.3af standard.

Designed usage

The Midspan Power Unit is designed to mount in a 19-inch data rack, or can be stacked up to four units high using the optional rubber feet. Its niche is to provide power to only those IP endpoints that need power. The alternative is to have a switch that incorporates power. However, any nonpowered device that uses that switch is not using the power capabilities of the switch, and does not justify the higher price per port of that switch. The Midspan Power Unit solves this problem by providing power without altering the network topology.

The 1152A1 can be collocated with the data equipment or closer to the endpoints. In all cases, IEEE 802.3af capable IP devices must connect directly to this PDU. The PDU cannot power any device if a hub or a switch is between itself and the endpoint because it will not sense the resistive signature needed to authorize the release of power.

Power modes (Avaya IP Telephones)

The Avaya IP Telephone has four different power modes:

- Ethernet spare pairs (4/5 and 7/8)
- Ethernet signaling pairs (1/2 and 3/6)
- Traditional telephony (7/8)
- (4630 model only) External transformer with a barrel connector

The 1152A1 power unit powers only through pairs 4/5 (+) and 7/8 (-).

Barrel connector through brick transformer

This brick type transformer provides 5 watts of power to the telephone. The Avaya telephone treats this brick as the primary power source, and will *not* accept power from the Ethernet cable if the barrel is seated into the telephone, with or without the brick attached to AC power.

Ethernet cable through 1152A1 PDU

Adequate power from the 1152A1 is supplied to the generation 2 telephones over the Ethernet cable. Category 5 or better cable is required for Fast Ethernet to function from the IP Telephone.

Power using adapters

Generation 1 telephones can receive power from the 1152A1 through an in-line adapter. This adapter provides the resistive signature so that the 1152A1 allows power to flow to the telephone. The generation-2 telephone does not need an adapter, but it might mistakenly be used on a generation-2 telephone. Both generation phones work as designed through all tests performed in Avaya labs.

Interoperability with Converged Infrastructure LAN and Cisco switches

Interoperability was tested with 24 generation 1 telephones and 2 IP Telephones attached to the Mid-Span Power unit. Data was supplied through the following switches:

- P550
- P333R
- P333T
- Cisco Catalyst 3524 (Powered Ethernet Blade)
- Cisco Catalyst 4003
- Cisco Catalyst 6509 (Powered Ethernet Blade)

All data switches pass data correctly through the 1152A1 unit, and the telephones work equally well with all of them. This is expected because Ethernet is an open standard.

Interoperability with Wireless Access Point products

The 1152A1 unit can also power Avaya's Wireless Access Point systems. The AP1, AP2, or AP3 act as a bridge between the wireless and the wired LAN. This system requires a 5-volt power supply that can be replaced by a splitter, which fits in the same cavity as the original power converter and allows power over the Ethernet, eliminating the need to find a power source close to the unit.

Converged infrastructure security gateways

VSUs

Avaya's line of VPN concentrators, called VPN Service Units (VSUs, see Figure 41, Avaya VSU 1000, on page 91), enable your business to securely connect remote users, branch offices, business partners, and customers, and take full advantage of the cost savings and productivity-enhancing benefits of virtual private networks (VPNs).

Figure 41: Avaya VSU 1000

VSU gateways are dedicated, hardware-based VPN gateways that overlay remote access and site-to-site VPN and firewall services on an existing enterprise network. VSUs provide all the benefits of VPNs without creating a performance bottleneck that slows down the network. Since all VPN services integrate easily and transparently into an existing enterprise network, customers enjoy easy access to network resources, which translates to increased productivity, and improved business efficiency. VSUs employ the strong two-factor authentication, data integrity and confidentiality services offered by IPSec, using either digital certificates or pre-shared secrets. VSUs also offer extended authentication mechanisms for remote users.

Unlike firewall or router-based VPN solutions, VSU gateways are designed specifically to handle high-bandwidth VPN traffic, using a dedicated, high-performance IPSec packet-processing engine and real-time data compression. They offer wire-speed performance that ranges from 16 Mbps to 100 Mbps for 3DES-encrypted IPSec traffic, and they can bridge at even higher speeds with non-VPN traffic. VSUs operating by default in bridge mode, and seamlessly layer into the network behind an access router, or in parallel with an existing firewall.

The VSU Series of VPN Gateways consists of 6 models:

- The VSU 5 product family (VSU 5 Gateway and VSU 5X Gateway) for small branch offices and home telecommuter offices
- The VSU 100 Gateway for small and medium businesses
- The VSU 2000 Gateway for branch offices
- The VSU 5000 Gateway, the VSU 7500 Gateway, and the VSU 10000 Gateway for large enterprises and managed data service providers

VSUs are centrally administered and configured using the VPNmanager® policy-based management application. They support remote access services with the VPNremote® desktop VPN client software.

VSU Gateways provide powerful levels of performance, manageability, and scalability. Whether managing thousands of remote access users, or delivering Voice over IP with service level agreements, networks can perform better, faster, and more reliably.

VPN Client

VPNremote® Client offers cost-effective, easy-to-install remote VPN connectivity that helps increase the productivity of telecommuters and mobile workers by providing secure, simple-to-use access to your enterprise network from any Internet access point.

VPNremote Client is a compatible with Microsoft Windows software, and provides secure, authenticated access to enterprise network resources and applications over the Internet. This application leverages the benefits of global access and cost-effective public network features to support a remote or a mobile work force. VPNremote Client not only provides support for data applications, but also delivers voice-over-VPN that enables you to use the Avaya IP softphone for secure, convenient telephony from your laptop computer. To protect the integrity and confidentiality of data that travels outside of an enterprise network, VPNremote Client uses standards-based IPSec technology to provide strong two-factor authentication, robust 3DES encryption, and data compression.

VPNremote Client overcomes the complexities that are typical of deploying a remote access solution. Easy installation and dynamic configuration dramatically reduces the burden for both end users and administrators. The intuitive graphical user interface-based Connection Manager helps you easily log on to your VPN by selecting a preconfigured user profile and entering your password. User profiles can also be exported to enable users to connect to the VPN from any computer using VPNremote. VPNremote also supports mobility by allowing users to securely connect to many wireless LAN systems.

Terminals

Avaya offers a wide range of communications devices to meet any company's unique needs. Since Avaya Communication Manager is extensible to IP, digital and analog telephones, and wireless business solutions, the spectrum is covered, regardless of your environment. IP Telephones and IP softphones allow access to the features of Communication Manager from more than one location. One of the major benefits of IP Telephones is that you can move the telephones around on a LAN just by unplugging them and plugging them in somewhere else. One of the main benefits of IP softphones is that you can load them on a laptop computer, and then use the modem on the computer to connect the softphones to the switch from almost anywhere.

Avaya IP Softphone

Figure 42: Avaya IP Softphone

Avaya IP Softphone is for employees who work remotely, on the road or at home. Accessed through a simple graphical interface on the screen of a PC or laptop computer, the IP softphone gives mobile workers the full suite of Avaya Communication Manager features and functions, whenever and wherever they need them.

- Patented technology for high-quality IP Telephony
- Full access to your personalized desktop telephone
- Windows PC feature set
- Microsoft Outlook integration (autodials from your Contacts list)
- Multiple call appearances
- Single or dual connect options
- Directory Access (LDAP)

Softphone operating modes

Road Warrior mode

The Road Warrior mode is used when there is only a single telephone line available to access the IP network and Communication Manager, or when broadband Internet access is available, for example, in homes or hotel rooms.

In Road Warrior mode the voice (audio) travels across the IP network along with call signaling traffic for an IP Telephony configuration that uses the industry-standard H.323 protocol, and offers a great amount of flexibility due to the widespread availability of IP networking.

Telecommuter configuration

The Telecommuter (Agent at Home) configuration is ideally suited for users who work from a remote office and have two lines for remote access, or in situations where only low-speed Internet connections are available. With this option, feature/access control and signaling is maintained and delivered across the IP network (using H.323), but the voice is delivered across a second line to either a public switched telephone network (PSTN), or digital line to help ensure toll-quality voice. This capability can be extended to a cellular, PCS, or GSM telephone. In the Telecommuter configuration, the Avaya communications server “binds” the two connections as a single transaction or session.

IP Telephone mode

The Avaya IP Telephone mode enables users to log into and control their Avaya IP Telephone from the Avaya IP softphone. Users can speak and listen through their telephone, but unlike the Telecommuter configuration, users can make and handle calls from both the Avaya IP softphone interface and the IP Telephone. This feature can help improve productivity by integrating the IP softphone capabilities and Personal Information Managers such as Microsoft Outlook with the IP Telephone. The Avaya IP Telephone configuration is supported on the Avaya 4606, 4612, 4620, 4624, and 4630 IP Telephones with Release 1.7 or later.

Avaya IP Agent

Avaya IP Agent is a Windows-based softphone application that is specifically designed to accommodate contact center agents who work remotely or in an office location. It runs on Windows 98, Windows 2000, Windows XP, or Windows NT® 4.0 PCs, enabling agents to work from their PC, anywhere, through remote connectivity to their corporate network. Agents have access to the full range of Avaya agent capabilities using a graphical user interface with standard drag-and-drop conventions.

- Screen pops are based on dialed number identification service (DNIS), automated number identification (ANI), and prompted digits.
- The integrated call history feature provides agents with a detailed view of calls made and received.
- “Road Warrior” and “Telecommuter” modes

Avaya Softconsole

Avaya Softconsole™ is a software attendant console that builds on the features of the popular Avaya 302 Attendant Console.

- Searches internal and external directories
- Displays detailed caller information on up to six calls simultaneously
- New interface
- Comprehensive setup wizards
- E-mail integration
- Enhanced directory capabilities
- Choice of two IP connections or DCP connection
- Voice over IP configuration (telecommuter)
- Dual connection (road warrior) for toll-quality audio
- DCP connection using the CallMaster VI
- Integrated iClarity for IP audio
- Directory lookup and dialing
- Integrated with directory management to support up to 100 directory databases
- Permanent and per-call notes

Avaya IP Softphone for Pocket PC

The Avaya IP Softphone for Pocket PC (Figure 43, Avaya IP Softphone for Pocket PC, on page 95) brings the full capabilities of Avaya IP softphone to your Windows CE handheld device, such as the Compaq iPAQ Pocket PC and gives you the full list of features and functions of Avaya Communication Manager on your pocket PC.

Figure 43: Avaya IP Softphone for Pocket PC

The Avaya IP Softphone for Pocket PC is a downloadable application for customers who own an Avaya IP softphone license. It delivers the full set of Communication Manager call features through a graphical display of your Avaya multiline telephone, with its identical extension number, speed dial buttons, and personal feature settings. Mobile workers can receive calls virtually anywhere, and remote workers can connect to your enterprise with wireless local area networks (LANs) and virtual private networks (VPNs).

The latest release of the Avaya IP Softphone for Pocket PC provides new user productivity features and global support with multiple languages.

Features

- **CTI control of IP Telephones**. Improve user productivity with Avaya IP Softphone for Pocket PC's ability to control the following Avaya IP Telephones

2420	4602	6408D+	8405D
	4606	6416D+	8401D
	4612	6424D+	8411D
	4620		8434
	4624		
	4630		

- **Globalization**. Supports multiple languages through language packs.
- Emergency Call Handling 911.
- Ability to modify E 911 station feature settings.
- Ability to modify the look and feel of the graphical user interface with a swap skin capability.
- Call log history.

Avaya 4600 Series IP Telephones

Figure 44: Avaya 4602 IP Telephone

Figure 45: Avaya 4606 IP Telephone

Figure 46: Avaya 4612 IP Telephone

Figure 47: Avaya 4620 IP Telephone

Figure 48: Avaya 4624 IP Telephone

Figure 49: Avaya 4630 IP screenphone

The IP Telephone is a physical voice terminal that provides IP Telephony. Avaya IP Telephones bring the rich features and functions of Avaya Communication Manager directly to the desktop. They are an essential part of converged voice and data networks that are built with the Avaya Application Solutions components. These telephones deliver an extensive set of features, high audio quality, and have an attractive streamlined design.

The 4600 Series IP Telephone sets (or terminals) are a product platform of terminals that support Avaya Application Solutions. The sets operate and function similar to the 6400 series Digital Communications Protocol (DCP) terminals when connected to the Avaya servers with Avaya Communication Manager, but provide added functions that are not possible within the digital terminal product line.

Networking coordination

The IP terminals use the Internet Protocol to communicate with the systems to which they are attached. The protocol is H.323 with proprietary signaling added to provide access to the full functionality that is available in the Avaya servers running Avaya Communication Manager.

IP Telephones are intended to connect to the customer's data network. These networks inherently contain products from many different vendors, and thus are less controlled than circuit switched networks. Therefore a Network Assessment is recommended as outlined in Network assessment offer.

The 4600 series voice terminals cannot be connected directly to the Public Switched Telephone Network (PSTN). They can only be connected to an Ethernet-based IP network. Therefore, network connection issues related to direct connection to the PSTN do not apply.

The IP Telephones provide connectivity to multiple external devices through a single Ethernet connection. The telephones contain an integrated switch or hub to connect the user's PC to the network through the telephone, thus requiring only a single Ethernet connection for both devices.

The 4600 series IP Telephones require the Avaya Application Solutions Circuit packs TN799C or higher C-LAN board for registration to the gatekeeper and signaling stream, and the TN2302AP (MedPro) for call setup and media stream. In the case of the G700 Media Gateway, the VoIP Engine or VoIP Media Module provides the support for media stream.

IP Telephone installations require a Trivial File Transfer Protocol (TFTP) server on the network for software transfer to the IP Telephones. Avaya strongly recommends DHCP servers. IP Telephone investment protection is supported by the software download capability. Future software releases for the telephones that offer feature enhancements or protocol changes are possible without changing the hardware platform. Software updates are provided as required to correct bugs, implement changes, and add new features and capabilities.

For customers who want an all switched network, the 4620 telephone contains an integrated Layer 2 switch. In addition, the 30A Ethernet Switch Base adds fully switched capability to the 4612 and the 4624 telephones.

Features and applications

Table 9: Avaya 4600 series IP Telephone features and applications ***1 of 2***

Feature	Application
Speakerphone	High-quality, built-in speakerphone with echo cancellation, directional microphone, and a tuned speaker cavity provides the highest audio quality.
Infrared capabilities	An infrared data association (IrDA) port is provided on the front of each IP Telephone for Personal Digital Assistant (PDA) and PC application integration (not available on 4602 sets). With the built-in IrDA port, users can communicate with and command the telephone using a personal digital assistant (PDA) or other infrared-equipped device.
Speed Dialing	This feature allows the user to store telephone numbers that are dialed at the touch of a feature button. Users can program up to 120 speed-dial buttons.
Call Log Application	*(4620 and 4630 sets only)* This feature stores and displays information, such as identifying the call as Outgoing, Incoming Answered, and Incoming Unanswered. It presents information about all calls in a given category. • Up to 90 entries on 4620 sets • Up to 100 entries on 4630 sets

1 of 2

Table 9: Avaya 4600 series IP Telephone features and applications ***2 of 2***

Feature	Application
Web Browser Application	*(4620 and 4630 sets only)* This feature provides Web Access to HTML Web-based information. The 4620 Web Access application is analogous to the application on the 4630. However, different display capabilities cause the 4620 telephones to have a simpler, less capable Web interface than 4630 sets.
Features that are common to the 4600 Series	• G.711, G.729A/B Voice Coders.
	• QoS options for UDP Port selection, Diffserv, 802.1p/Q.
	• Support for Simple Network Management Protocol (SNMP), Version 2.
	• DHCP client and Statically (Manual) Configurable IP Addressing.
	• NetMeeting™ compatible.
	• Multiple power options, including support for power over Ethernet LAN technology.
	• 10/100 Base T Ethernet connections.
	• Integrated Ethernet Hub - optional connection (PC to telephone).
	• Infrared (IrDA) port.
	• Built-in headset jack.
	• Full-duplex speakerphone with echo cancellation.
	• Feature buttons for Conference, Transfer, Drop, Hold, Redial, Mute, Speaker, Voice Mail, and so on.
	• Set angle of 15^{o} for display visibility with optional wall/desk stand.
	• Special handset supports AB styles.
	• Message Waiting Indicator.
	• Hearing aid compatibility.

2 of 2

Avaya 4620 IP Telephone

The 4620 is an Avaya IP Screenphone providing a large-screen, graphic display (168x132 dots), a Web Access application, and an integrated switch to connect a PC without requiring an additional wall jack.

Avaya 4630 IP Screenphone

This full color, touch screen, Web access IP Telephone includes six telephony-related applications, designed for ease of use with a menu-based interface. It can display a variety of information, including Web pages specially downsized for small-format displays. Sample applications are concierge desks at hotels, airline frequent travelers clubs, financial services kiosks, and as an executive desktop phone.

- Multi-button capabilities supported by Avaya Communication Manager; 3 to 5 call appearances plus 21 feature buttons
- Speed dial provides 120 speed dial “buttons” organized into groups for easier access; names, numbers, and group names are user-programmable
- Call log lists of up to 100 incoming and outgoing calls
- Access to corporate directory information on a Lightweight Directory Access Protocol (LDAP) server
- Web access provides “browsing” access to HTML Web-based information, including support for downloaded Java™ applets
- Access to multimedia messaging capabilities of the Avaya DEFINITY AUDIX® or Avaya INTUITY™ AUDIX systems using Avaya www.messenger

Communication Manager support for the 4600 IP Telephone Series

The 4606, 4612, and 4624 IP Telephones are supported beginning with DEFINITY Enterprise Communications Servers (ECS) Release 8.4, and Release 9. The IP Telephone sets will NOT operate on any previous software releases but will operate with all later releases. In addition to operating with all Communication Manager platforms, the endpoints also operate with Avaya DEFINITY ECS G3r or G3si, DEFINITY One, DEFINITY ProLogix, Avaya IP600 Communications Server, as well as Avaya IP Office.

The 4630 telephone has been developed as an IP endpoint for the Avaya Media Servers using Avaya Communication Manager and DEFINITY Servers. DEFINITY Release 10 and Communication Manager Release 1.0 are the first releases of software to provide native support for the 4630 Screenphone. The Screenphone will NOT operate with full functionality on DEFINITY Release 9.5, and it will NOT work at all on any previous releases of the targeted systems.

The Avaya 4602 and 4620 IP Telephone have been developed as IP endpoints for the Avaya Media Servers using Avaya Communication Manager Release 1.1. These IP Telephone sets will NOT operate on any previous releases, but will operate with all later releases. These phones must be natively administered in Avaya Communication Manager 1.1 in order to support the automatic button labeling and new features. Provision for labeling the buttons manually has not been made.

Wireless

Avaya's family of wireless business systems offers mobility solutions that help your employees stay connected and remain productive from wherever their work takes them—whether they are in the office, moving around campus, or around the country.

Avaya Extension to Cellular

This is software that, when combined with a cellular phone, offers one-number access for business connectivity anytime, anywhere with no missed calls. Avaya Communication Manager enables the Avaya Extension to Cellular to transparently bridge calls from the Avaya server to any digital cellular phone, regardless of service provider or cellular standard.

- One-number portability allows for a high level of accessibility because your office number is bridged to your digital cell telephone.
- Simultaneous ringing keeps you and your associates in touch, so you can respond quickly to urgent enterprise matters without delay.
- Software only solution does not require the expense of a wireless office service. It can utilize your existing cellular telephone and service coupled with Communication Manager.

Other digital wireless systems

In addition to Extension to Cellular, Avaya also has other TDM-based digital wireless systems available, including Avaya TransTalk 9000 Digital Wireless System and Avaya DEFINITY® Wireless Business System.

Section 2. Deploying IP Telephony

Traffic engineering

This chapter provides an introduction to traffic engineering. Specifically, this chapter discusses various traffic models, algorithms, and resource sizing.

This section includes the following topics:

Introduction

The process of configuring, engineering, and deploying a Communication Manager system, or a network of Communication Manager systems, begins with specifying the quantity and the nature of the endpoints to be accommodated. Principles of traffic engineering are then applied to determine the quantity and the placement of the various necessary resources. Once the designed configuration adheres to all specifications and system constraints, the process is finished.

This discussion of the configuration, engineering, and deployment processes is intended as an overview that is suitable for a fairly general audience. One example that is designed to exercise all aspects of these processes continues throughout the chapter to present the finer points of network design.

Design inputs

This section summarizes the essential design elements that the customer must specify.

Topology

An Avaya Communication Manager system consists of a server and all of the equipment under that server's control. Such equipment may be geographically dispersed among a variety of sites, and the equipment at each site may be segregated into distinct logical collections known as Network Regions. In cases where one server is insufficient for controlling all of the equipment, multiple Avaya systems can be networked together. So, a *Network Region* is a component of a *site*, which is a component of a *system*, which is a component of a *network*.

A single Avaya Communication Manager system is comprised of one or more *Network Regions*. Each Network Region is a logical grouping of endpoints, including stations, trunks, and Media Gateways. Customers can choose to establish various Network Regions on the basis of geography, business sectors, or any of a variety of other considerations. For example, a customer with facilities in both New York and Los Angeles might choose to use a single Communication Manager system, with one Network Region in each of the two cities. Another possibility is to assign two Network Regions to each city. In that case, each such geographical grouping of Network Regions is said to comprise a *site*.

Alternatively, that same customer might want to administer three Network Regions, where one region corresponds with Sales and Marketing, another with Customer Support and Services, and a third with Research and Development. In this case, the Network Regions are established independently of geographical considerations, because associates from each of the three distinct business sectors may be physically located in both cities. Yet another possibility is to construct Network Regions to correspond with IP subnets.

The various Network Regions within a Communication Manager system are interconnected by an IP network. An IP network can consist of local area networks (LANs), wide area networks (WANs), or a combination of both LANs and WANs. A common approach is to use a LAN at each site, and interconnect those LANs through a WAN. Because Network Regions are used to specify differences between the treatment of intrasite and intersite traffic, or to properly select localized media resources for optimal voice quality, Network Regions should not span multiple geographical locations.

A Communication Manager system can operate as an independent entity, or can be networked together with other Communication Manager systems. For networked systems, the various Communication Manager systems in the network are generally interconnected by IP tie trunks. If the two members of a given pair of Communication Manager systems in a network are not directly interconnected by tie trunks, calls between the two systems must be tandemed through other Communication Manager systems in the network.

When there is a need to accommodate endpoints in various geographic locations, the customer has the choice to either set up a single Communication Manager system with a site at each location, or use a network of multiple Communication Manager systems to span the locations. The choice of which one is more appropriate pertains to the issue of scalability. An extremely large number of endpoints might mandate the use of multiple systems.

While Communication Manager systems have been designed with an IP infrastructure, they also support circuit-switched endpoints, and the full complement of traditional DEFINITY features. However, customers usually realize a significant advantage when those customers implement an IP-oriented solution for systems that are geographically dispersed.

Each endpoint and Media Gateway is assigned to a Network Region when its IP address is administered. Also, each Network Region is administered with a codec preference list, which is a list of up to five codecs that are supported by that Network Region. Uncompressed G.711 and compressed G.729 are the most commonly used codecs in Communication Manager systems. Each Communication Manager system is administered with the Internetwork Region Connection Management (IRCM) matrix, which provides enough information to specify which codecs to use when completing a call between Network Regions.

Conversely, if the IRCM does not specify a codec set between two Network Regions, calls cannot be completed between those regions over an IP connection. For instance, the manager of an office building can use a single Communication Manager system to service all the individual lessees, with a separate Network Region for each company. Those Network Regions generally would not be connected by the IRCM because independent companies would be unwilling to share each others' resources. Subsequent sections of this chapter further explain sharing resources across connected Network Regions.

Multiple Communication Manager systems are often networked together by IP tie trunks, although circuit-switched tie trunks can also be used. If the two members of a given pair of Communication Manager systems in a network are not directly interconnected by tie trunks, calls between the two systems must be routed through other Communication Manager systems in the network, or through the public switched telephone network (PSTN).

Although Avaya products are IP enabled, the products must interface with circuit-switched endpoints and systems. For example, Communication Manager systems require circuit-switched trunks to access the PSTN because central offices today are not equipped for IP trunking. Some customers also prefer to continue to use their circuit-switched telephones in Communication Manager systems.

Circuit-switched endpoints interface to circuit packs that reside in either an H.248 Media Gateway or a traditional Port Network (PN). Although each Media Gateway belongs to one particular Network Region, no correlation exists between PNs and Network Regions. PNs are interconnected through a circuit-switched center stage or an ATM center stage (S8700 Multi-Connect systems) or an IP network (IP-Connect systems).

Endpoint specifications

Normally, a customer who submits a Request for Proposal (RFP) specifies the number of each type of station to place in each site, in each Communication Manager system in the network. Certain customers might want to specify station placement more precisely. For example, a customer might specify the exact population of circuit-switched stations on a Media Gateway.

The majority of customers know exactly how many of each station type are needed at each site, based on the population of their anticipated end-users. However, the issue of trunk sizing is not as straightforward. Trunk traffic is tightly coupled with station traffic because at least one party in every Communication Manager trunk call is a Communication Manager station (except in relatively rare cases in which a Communication Manager system is used to tandem calls between non-Communication Manager endpoints). That being the case, station traffic effectively induces trunk traffic. The given topology of trunk groups dictates which pairs of Communication Manager systems are directly connected by trunks, and which Communication Manager systems are directly connected to the PSTN (or other non-

Communication Manager systems). However, the size of each trunk group must be engineered with consideration of the amount of traffic that each such trunk group is anticipated to carry. A traffic engineer should either specify the number of trunks in each trunk group directly, or allow the configuration algorithm to size the trunk groups to a specified Grade of Service (GOS). This GOS is usually P01, which is 1% blocking. In some cases, customers might choose to over-engineer or under-engineer certain trunk groups based on nontraffic considerations, such as reliability, cost, security, and so on.

Endpoint traffic usage

Traffic usage is typically expressed in Erlangs, which represent the average number of busy servers in a given server group. For example, if a group of stations carries 100 Erlangs of call usage, that means the average number of those stations that are busy at any given time is 100. The usage of a single station, when expressed in Erlangs, represents the fraction of time that the station is in use. So, a station that carries 0.1 Erlang of usage is in use 10% of the time.

The most common way to specify total station usage is to multiply the usage per station by the total number of stations. A traffic engineer can either explicitly specify the per-station usage for each group of stations, or allow the configuration algorithm to specify per-station usages automatically, using default values. Common defaults for station traffic usage in general business scenarios are:

- **Light** traffic—0.056 Erlangs per station (stations average 5.6% usage)
- **Moderate** traffic—0.11 Erlangs per station (stations average 11% usage)
- **Heavy** traffic—0.17 Erlangs per station (stations average 17% usage)

The most commonly used default value for a general business system is 0.11 Erlangs per station. The most common way to determine trunk usage rates is to divide the total traffic load that is carried by each trunk group by the number of trunks in the group. It is difficult to assign a typical default value for usage per trunk. Such usage can vary greatly from system to system, and even from trunk group to trunk group within a particular system.

Traffic usage has two components:

- Average call duration (also known as call *hold time*)
- Average number of calls per hour

Systems are usually engineered to accommodate the busiest hour of a normal business day. The number of calls that are completed during that busiest hour is denoted by Busy Hour Calls Completed (BHCC). BHCC is not be confused with Busy Hour Calls Attempted (BHCA), which represents the total number of calls attempted during the busiest hour, regardless of how many of those calls are actually successfully completed. The general expression for the relationship between BHCC, average call duration, and usage is:

$$\text{Usage (Erlangs)} = \frac{\text{BHCC} \times \text{seconds per call}}{3600}$$

A commonly used default value for average call duration in a general business system is 200 seconds per call. Example 1 shows how to calculate the station usages using the data given.

Example 1

Assume that the customer has sites in Atlanta, Boston, and Cleveland that the customer wants to populate with the following endpoints (Table 10, Example 1 configuration data, on page 108).

Table 10: Example 1 configuration data

Endpoints	Atlanta	Boston	Cleveland
DCP Telephones	540	180	
IP Telephones	1,080	450	270
Analog stations	108	18	
Road Warriors	27		
Other			Two G350 Media Gateways, each of which supports 18 analog stations, and a suitable number of circuit-switched PSTN trunks

Additional design criteria

- Each site is to have a suitable number of PSTN trunks (which terminate on PNs in Atlanta and Boston, and on the G350 Media Gateways in Cleveland).
- This is a general business application (for example, no Call Center agents), where the average usage per station is assumed to be 0.11 Erlangs, and the average call duration is assumed to be 200 seconds.
- Each site consists of a single Network Region, and all three Network Regions are interconnected in the sense of the IRCM matrix.
- One-third of all calls are intercom calls (that is, calls between two stations), one-third are inbound PSTN trunk calls, and one-third are outbound PSTN trunk calls.

Preliminary calculations

Based on the assumption of 0.11 Erlangs per station, Table 11, Example 1 station usage by endpoint type, on page 108 shows the total station usage for each station category in the system.

Table 11: Example 1 station usage by endpoint type

Endpoints	Atlanta (Erlangs)	Boston (Erlangs)	Cleveland (Erlangs)
DCP Telephones	60	20	
IP Telephones	120	50	30
Analog stations	12	2	
Road Warriors	3		
Analog stations administered to G350 Media Gateways			4

Call usage rates

In the previous section, station usages and overall endpoint usages, including both stations and trunks, were discussed. The overall endpoint usage is sometimes referred to as port usage rate (PUR). The term station usage rate (SUR) applies when referring only to the stations. In general, a traffic usage rate, when expressed in Erlangs, represents the average number of busy servers in a given server group. So, SUR represents the average number of stations in a particular group that are simultaneously in use, while PUR represents the average number of endpoints, including stations and trunks, in a particular group that are simultaneously in use.

Similarly, the term call usage rate (CUR) represents the average number of simultaneous calls that are carried by a particular facility. In an environment where essentially every call is either inbound or outbound (such as a call center), CUR and SUR are equal, because there is exactly one Communication Manager station used in each call. However, in an environment such as a general business scenario in which some calls are intercom, some calls are inbound, and some calls are outbound (such as a General Business scenario), CUR and SUR are not equal, because some calls (the intercom calls) use two Communication Manager stations, and others (inbound and outbound calls) use only one Communication Manager station.

The next step in the configuration process is to determine the amount of traffic flow between Communication Manager systems in a network, and between the sites in each individual CM system. Those traffic flows can be further refined to identify the traffic flows between the various categories of endpoints within each site. All such traffic flows can be represented in tabular form.

Communities of interest

The various sites within a particular Communication Manager system comprise *communities of interest* (COI), in the sense that the endpoints in each particular site share some common trait or interest, usually geographical proximity. A COI matrix offers a convenient representation of the traffic flows between the various sites. For example, consider the COI matrix in Table 12, 3-site standalone community of interest (COI) matrix, on page 110 for a three-site, stand-alone Communication Manager system.

In practice, a COI matrix that is associated with a given system is populated with actual traffic values. In Table 12, 3-site standalone community of interest (COI) matrix, on page 110, each diagonal matrix entry represents intrasite call usage, and all other entries represent intersite call usage. The call usages used to populate the table can be determined empirically or through theoretical means. In some cases, actual call usage data can be obtained through polling an existing system. In other cases, it might be appropriate to apply a mathematical model to estimate the call usages.

Table 12: 3-site standalone community of interest (COI) matrix

		To endpoints in site __		
CUR		**1**	**2**	**3**
From endpoints in Site	**1**	Call usage generated by Site 1 endpoints, terminating at Site 1 endpoints	Call usage generated by Site 1 endpoints, terminating at Site 2 endpoints	Call usage generated by Site 1 endpoints, terminating at Site 3 endpoints
	2	Call usage generated by Site 2 endpoints, terminating at Site 1 endpoints	Call usage generated by Site 2 endpoints, terminating at Site 2 endpoints	Call usage generated by Site 2 endpoints, terminating at Site 3 endpoints
	3	Call usage generated by Site 3 endpoints, terminating at Site 1 endpoints	Call usage generated by Site 3 endpoints, terminating at Site 2 endpoints	Call usage generated by Site 3 stations, terminating at Site 3 endpoints

One of the first steps in the process is to distinguish between intercom call usage, inbound PSTN call usage, and outbound PSTN call usage. Inbound and outbound tie trunk usage must also be considered when working with multiple Communication Manager systems that networked together. However, that discussion is presented in a later section.

Although Avaya systems can be used as tandem switches for PSTN traffic, that possibility is not considered here. Traffic between two other Avaya systems in a network is the only traffic that can be routed through Communication Manager. So, in the case of a single stand-alone system, there is typically no tandem traffic. Therefore, because every call involves at least one station, one must be careful to reconcile the station usage with the call usage.

For example, suppose that the total station usage is 100 Erlangs, which could hypothetically correspond to 20 Erlangs of intercom call usage, 30 Erlangs of inbound PSTN usage, and 30 Erlangs of outbound PSTN usage:

- **Intercom** station usage = 40 Erlangs (2 Avaya stations per call x 20 Erlangs of intercom call usage)
- **Inbound** station usage = 30 Erlangs (1 Avaya station per call x 30 Erlangs of inbound call usage)
- **Outbound** station usage = 30 Erlangs (1 Avaya station per call x 30 Erlangs of outbound call usage)

The 40 Erlangs that are associated with intercom calls, plus the 30 Erlangs that are associated with inbound calls, plus the 30 Erlangs that are associated with outbound calls total 100 Erlangs of station usage.

Alternatively, 100 Erlangs of total station usage could also hypothetically correspond to 35 Erlangs of intercom call usage, 10 Erlangs of inbound PSTN usage, and 20 Erlangs of outbound PSTN usage. Using the procedure from the preceding example to verify this:

- **Intercom** station usage = 70 Erlangs (2 Avaya stations per call x 35 Erlangs of intercom call usage)
- **Inbound** station usage = 10 Erlangs (1 Avaya station per call x 10 Erlangs of inbound call usage)
- **Outbound** station usage = 20 Erlangs (1 Avaya station per call x 20 Erlangs of outbound call usage)

The 70 Erlangs that are associated with intercom calls, plus the 10 Erlangs that are associated with inbound calls, plus the 20 Erlangs that are associated with outbound calls total 100 Erlangs of station usage.

However, suppose that once again the station usage is 100 Erlangs. Assuming that there is no tandem traffic, this cannot correspond to 10 Erlangs of intercom call usage, 20 Erlangs of inbound PSTN usage, and 30 Erlangs of outbound PSTN usage.

- **Intercom** station usage = 20 Erlangs (2 Avaya stations per call x 10 Erlangs of intercom call usage)
- **Inbound** station usage = 20 Erlangs (1 Avaya station per call x 20 Erlangs of inbound call usage)
- **Outbound** station usage = 30 Erlangs (1 Avaya station per call x 30 Erlangs of outbound call usage)

The 20 Erlangs that are associated with intercom calls, plus the 20 Erlangs that are associated with inbound calls, plus the 30 Erlangs that are associated with outbound calls total 70 Erlangs, leaving 30 Erlangs of unaccounted station usage. This is a sign that the parsing of call traffic into intercom, inbound, and outbound might have been done erroneously. One possible explanation for the extra 30 Erlangs of station usage is adjunct traffic, such as stations that are connected to voice mail, providing that 30 Erlangs of voice mail calls makes sense in the model.

The bottom line is, regardless of what model is used to parse call traffic into its various components, one must be able to reconcile the overall station usage with the overall call usage. Specifically, if there is no tandem traffic, the following relationship must hold:

$$\text{SUR} = (2 \times \text{intercom CUR}) + \text{inbound PSTN CUR} + \text{outbound PSTN CUR}$$

Having established that point, several examples describe some methods for populating the COI matrix. For the sake of continuity, all of the examples are built upon Example 1.

Example 2

Uniform Distribution model

In the case of a stand-alone Avaya system, the Uniform Distribution model works on the assumption that when a given station places an intercom call, the call is equally likely to terminate at any of the other stations in the entire system. Analogous statements regarding this model can also be made for inbound trunk calls and outbound trunk calls. Specifically, any inbound call is equally likely to terminate at any of the stations in the system, and any outbound call is equally likely to have been originated by any of the stations in the system. The fundamental concept underlying the Uniform Distribution model is that stations are essentially indistinguishable from one another from a traffic engineering point of view. This model is usually the most appropriate option when engineering a system for which little or no information about the nature of the various stations exists. This model will now be applied to the system that is described in Example 1.

The design criteria for Example 1 was one-third of all calls being intercom, one-third being inbound PSTN, and one-third being outbound PSTN. From the station usages that are listed in Example 1, it follows that the total station usage in Atlanta is 195 Erlangs, the total in Boston is 72 Erlangs, and the total in Cleveland is 34 Erlangs, for a system-wide total of 301 Erlangs of station usage. Under the “one-third intercom, one-third inbound, one-third outbound” assumption, this corresponds to a system-wide total of 75 Erlangs of intercom call usage, 75 Erlangs of inbound call usage, and 75 Erlangs of outbound call usage (rounding to the nearest Erlang in each case). To verify this, first consider the fact that all three components are equal (each is 75 Erlangs) satisfies the “one-third, one-third, one-third” requirement.

Furthermore, since 75 Erlangs of intercom call usage corresponds to 150 Erlangs of station usage, 75 Erlangs of inbound call usage corresponds to 75 Erlangs of station usage, and 75 Erlangs of outbound call usage corresponds to 75 Erlangs of station usage, there is a total of 150 + 75 + 75 = 300 Erlangs of station usage. This agrees with the specified 301 Erlangs if one ignores error due to rounding off.

One could assume in this example that each PSTN trunk is capable of carrying both inbound calls and outbound calls. Trunks are normally engineered to a desired Grade of Service (GOS), or blocking level. A commonly used GOS for trunks is P01, which represents a nominal blocking rate of 1 out of every 100 call attempts. To determine how many trunks are needed to attain P01, one must know the call traffic load to be carried by those trunks. Both inbound call usage and outbound call usage are included in that load.

> **NOTE:**
> If IP softphone telecommuters were used in this example, they would have also contributed toward trunk load. Although the signaling link between a telecommuter and the Communication Manager system to which the telecommuter is registered is carried over IP, the media flow between the two uses a PSTN trunk.

Example 1 indicates that the total load to be carried by the trunks is 75 + 75 = 150 Erlangs, which accounts for both inbound and outbound PSTN call usage. Use of the standard Erlang blocking model indicates that 171 trunks (DS0s) would be required to carry the 150 Erlangs of trunk call usage at P01. However, one must consider the trunk selection process for PSTN calls.

Communication Manager uses a first-site-preference algorithm for outbound trunk calls. This algorithm specifies that all outbound calls first attempt to seize a trunk within the originating station's site, and tries to use a trunk in a different site if and only if it is blocked at its local trunks. For inbound PSTN trunk calls, the CO selects the trunk. Therefore, Communication Manager cannot use an analogous first-site-preference algorithm for inbound calls. However, such an algorithm can be effectively imposed by assigning different calling numbers for the three sites, which is typical in this example since the sites are in different area codes.

The goal of a first-site preference algorithm is to minimize intersite traffic. When this algorithm is used, there is intersite traffic if and only if it overflows to a trunk on another site after having been blocked at the trunks in its own site. Under the assumption that a first-site preference algorithm is used in this example, the trunks at the three individual sites must be sized independently, as opposed to all together. Initially, the overflow traffic is ignored, but that topic is discussed later in this example.

Since overflow traffic is ignored for the time being, intersite trunk traffic is zero, which implies that the off-diagonal entries of the inbound and outbound COI matrices will all be zero. To determine the values of the diagonal entries, which correspond to intrasite trunk usage, the Uniform Distribution model is applied. In particular, 65% (that is, 1755/2709) of the stations are in Atlanta, 24% (that is, 648/2709) of the stations are in Boston, and 11% (that is, 306/2709) of the stations are in Cleveland. Therefore, the Uniform Distribution model implies that 65% of the 75 Erlangs of inbound CUR (that is, 49 Erlangs) is assumed to terminate in Site 1 (Atlanta), 24% (that is, 18 Erlangs) is assumed to terminate in Site 2 (Boston), and 11% (that is, 8 Erlangs) is assumed to terminate in Site 3 (Cleveland). Similarly, 49 Erlangs of outbound CUR is assumed to originate in Site 1, 18 Erlangs is assumed to originate in Site 2, and 8 Erlangs is assumed to originate in Site 3.

It is instructive for this example to construct three different COI matrices rather than just one. Specifically, it is useful to construct one for intercom CUR, one for inbound CUR, and one for outbound CUR. The information from the previous paragraph can be used to populate the following inbound and outbound COI matrices (Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113):

Table 13: Inbound COI matrix for the Uniform Distribution model in Example 2

		To stations in Site __		
Inbound CUR		**1**	**2**	**3**
	1	49 Erlangs	0	0
	2	0	18 Erlangs	0
From trunks in Site ____	**3**	0	0	8 Erlangs

Table 14: Outbound COI matrix for Uniform Distribution Model in Example 2

		To trunks in Site __		
Outbound CUR		**1**	**2**	**3**
	1	49 Erlangs	0	0
	2	0	18 Erlangs	0
From stations in Site ____	**3**	0	0	8 Erlangs

Again, Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113 and Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113 are constructed without considering overflow traffic. Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113 and Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113 imply that the Site 1 PSTN trunks carry 98 Erlangs (49 inbound and 49 outbound) of traffic, the Site 2 trunks carry 36 Erlangs, and the Site 3 trunks carry 16 Erlangs. Applying the standard Erlang loss model with a P01 GOS to each of the three sites implies that at least 116 trunks are needed in Site 1, at least 49 trunks are needed in Site 2, and at least 26 trunks are needed in Site 3. Note that this constitutes a total of 191 trunks, as opposed to the estimate of 171 trunks that was obtained without sizing the three trunk groups separately. A total of 171 could be used to attain an overall grade of service of P01, but that would induce a large amount of intersite traffic. The use of 191 total trunks, distributed between the three sites as specified above, ensures that at least 99% of the calls are guaranteed to be *intrasite*.

In some cases, there might be factors that justify overengineering the trunk groups. For example, a customer who is based in North America most likely leases T1 trunk facilities between each of its sites and the appropriate COs. In this example, it might be reasonable to use five T1 facilities (that is, 120 DS0 channels) for Atlanta, three T1 facilities (that is, 72 DS0 channels) for Boston, and two T1 facilities (that is, 48 DS0 channels) for Cleveland. This yields an overall GOS much better than P01, and at the same time, the use of standardized equipment reduces costs. In fact, the use of Erlang's loss formula implies a blocking probability of 0.004 in Atlanta, and negligible blocking probabilities (that is, several orders of magnitude better than P01) for the other two sites. These extremely low-blocking probabilities justify the assumption that intersite trunk traffic (overflow traffic) is negligible in this example.

Finally, the entries for the intercom COI matrix must be determined. Of the 195 Erlangs of station usage in that site, 49 Erlangs are associated with inbound calls, and 49 Erlangs are associated with outbound calls. That leaves 195 - 49 - 49 = 97 Erlangs of station usage in the Atlanta site for intercom calls. Similarly, there are 72 - 18 - 18 = 36 Erlangs of station usage in the Boston site for intercom calls, and 34 - 8 - 8 = 18 Erlangs of station usage in the Cleveland site for intercom calls.

It is assumed that half of each individual station's usage is associated with calls that the station generates, and the other half is associated with calls that the station receives. Therefore, half of the 97 Erlangs of station usage (that is, 49 Erlangs) in the Atlanta site corresponds to intercom calls originated in the Atlanta site. Similarly, half of the 36 Erlangs of station usage (that is, 18 Erlangs) in the Boston site corresponds to intercom calls originated in Boston, and half of the 18 Erlangs of station usage (that is, 9 Erlangs) in the Cleveland site corresponds to intercom calls originated in Cleveland.

Using the percentages from earlier, the Uniform Distribution model implies that 65% of the intercom traffic originated by each station in Atlanta is terminated in Atlanta, 24% is terminated in Boston, and 11% is terminated in Cleveland. Applying those percentages to the 49 Erlangs of intercom traffic that is generated in Atlanta implies that 32 Erlangs of intercom call usage is generated in Atlanta for termination in Atlanta, 12 Erlangs of intercom call usage is generated in Atlanta for termination in Boston, and 5 Erlangs of intercom call usage is generated in Atlanta for termination in Cleveland. Analogous calculations can be made in relation to intercom traffic that is generated in Boston and in Cleveland. The results are tabulated in the intercom COI matrix that is associated with this example (Table 15, Intercom COI matrix for the Uniform Distribution model in Example 2, on page 114):

Table 15: Intercom COI matrix for the Uniform Distribution model in Example 2

		To stations in Site __ (all data in Erlangs)		
Intercom CUR		**1**	**2**	**3**
	1	32	12	5
	2	12	4	2
From stations in Site ____	**3**	6	2	1

The general formulas used to populate the COI matrix entries in Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113, Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113, and Table 15, Intercom COI matrix for the Uniform Distribution model in Example 2, on page 114, respectively, for the Uniform Distribution model are:

$$\text{Inbound CUR to Site } i = \left(\frac{\text{number of stations in Site } i}{\text{total number of stations}}\right) \times (\text{total inbound CUR})$$

$$\text{Outbound CUR from Site } i = \left(\frac{\text{number of stations in Site } i}{\text{total number of stations}}\right) \times (\text{total outbound CUR})$$

$$\text{Intercom CUR from Site } i \text{ to Site } j = \left(\frac{\text{number of stations in Site } j}{\text{total number of stations}}\right) \times \left(\begin{matrix}\text{total intercom CUR} \\ \text{originating in Site } i\end{matrix}\right)$$

Additional comments regarding Example 2

In the Uniform Distribution model introduced in Example 2, the relative weights that are associated with the various sites correspond to the distribution of stations throughout the sites. Alternatively, the weights could correspond to the relative overall station usages in the various sites. Such a model takes into account not only the number of stations, but also how busy the stations are. In Example 2, since every station is assumed to have the same usage (specifically, 0.11 Erlangs), the weights that are based on the number of stations per site are exactly the same as the weights that are based on the overall station usage per site. Such a model is not always appropriate. For example, consider a system with two sites, with 100 stations in each site. Suppose that the average usage per station in Site 1 is 0.1 Erlangs, and that the average usage per station in Site 2 is 0.2 Erlangs. In a Uniform Distribution model where the weights are based on station usage per endpoint, a caller in Site 1 is twice as likely to call a station in Site 2 than a station in Site 1 (because the total station usage in Site 2 is 20 Erlangs, and the total station usage in Site 1 is only 10 Erlangs). The general formulas used to populate the COI matrix entries in Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113, Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113, and Table 15, Intercom COI matrix for the Uniform Distribution model in Example 2, on page 114, respectively, for the Uniform Distribution model based on relative SUR are:

$$\text{Inbound CUR to Site } i = \left(\frac{\text{total station usage in Site } i}{\text{total station usage}}\right) \times (\text{total inbound CUR})$$

$$\text{Outbound CUR from Site } i = \left(\frac{\text{total station usage in Site } i}{\text{total station usage}}\right) \times (\text{total outbound CUR})$$

$$\text{Intercom CUR from Site } i \text{ to Site } j = \left(\frac{\text{total station usage in Site } j}{\text{total station usage}}\right) \times \left(\begin{matrix}\text{total intercom CUR} \\ \text{originating in Site } i\end{matrix}\right)$$

Example 3

Empirical approach for existing systems

Another possible means of populating the COI matrices exists for established systems. In such cases, the necessary information can be read from traffic reports. This method is particularly useful for customers who are considering an upgrade from their current equipment.

Expanded COI matrices

So far, all the discussion pertaining to COI matrices has focused on a macroscopic view of sites. In particular, all the COI matrices presented have dedicated one cell for each pair of sites. In preparation for Resource sizing on page 122, it is useful to partition each such cell into collections of smaller cells that describe the call flows between different communities of endpoint types within the sites.

One possible partitioning scheme for each site is to create the following three general endpoint categories:

- IP endpoints
- Circuit-switched endpoints
- PSTN trunks

Consider the COI matrix for a three-site, stand-alone Communication Manager system, as presented in Table 12, 3-site standalone community of interest (COI) matrix, on page 110. A suitable expansion of that matrix might take the form of the matrix in Expanded COI matrix for a three-site system on page 116 in which

- *I* represents IP endpoints
- *C* represents circuit-switched endpoints
- *P* represents PSTN trunks

This finer categorization of endpoints permits the use of a single COI matrix for intercom, inbound, and outbound call usage rates.

Table 16: Expanded COI matrix for a three-site system

			To endpoints in Site ___								
			1			2			3		
			I	C	P	I	C	P	I	C	P
From endpoints in Site ___	1	I									
		C									
		P									
	2	I									
		C									
		P									
	3	I									
		C									
		P									

Example 4

Expanded COI matrices

In this example, we revisit Example 2, which pertains to the Uniform Distribution model, in more detail. The various endpoints are grouped into the three categories that are referenced in Table 16, Expanded COI matrix for a three-site system, on page 116. The COI matrix in Table 15, Intercom COI matrix for the Uniform Distribution model in Example 2, on page 114 lists the intercom call usage rates between each pair of sites, including intrasite call usage. Those usage rates can be broken down into finer components. Table 17, Endpoints in a three-site system, on page 117 reviews the various endpoints in each site.

Table 17: Endpoints in a three-site system

Endpoints	Atlanta	Boston	Cleveland
IP stations	1107 (1080 IP Telephones + 27 Road Warriors)	450 (450 IP Telephones)	270 (270 IP Telephones)
Circuit-switched stations	648 (540 DCP stations + 108 analog stations)	198 (180 DCP stations + 18 analog stations)	36 (36 analog stations)
PSTN trunks	120 (DS0) PSTN Trunks (5 T1 facilities)	48 (DS0) PSTN Trunks (2 T1 facilities)	24 (DS0) PSTN Trunks (1 T1 facility)

First consider the 32 Erlangs of intercom CUR between Site 1 stations (Table 15, Intercom COI matrix for the Uniform Distribution model in Example 2, on page 114). Site 1 (Atlanta) has a total of 1755 stations, 1107 of which are IP stations, and 648 of which are circuit-switched stations. So, 63% of the stations in Site 1 are IP, and 37% are circuit switched. Therefore, 63% of Site 1 intercom calls are generated by IP stations, and 63% of those calls are terminated by IP stations. Since 63% of 63% is 39.7%, 39.7% of Site 1 intercom calls are IP station to IP station. Similarly, 37% of the Site 1 intercom calls that are generated by IP stations are terminated by circuit-switched stations. Since 37% of 63% is 23.3%, 23.3% of Site 1 intercom calls are IP station to circuit-switched station.

Also, 37% of Site 1 intercom calls are generated by circuit-switched stations, and 63% of those calls are terminated by IP stations. Since 63% of 37% is 23.3%, 23.3% of Site 1 intercom calls are circuit-switched station to IP station. Finally, 37% of the Site 1 intercom calls that are generated by circuit-switched stations are terminated by circuit-switched stations. Since 37% of 37% is 13.7%, 13.7% of Site 1 intercom calls are circuit-switched station to circuit-switched station.

So, since 39.7% of Site 1 intercom calls are IP station to IP station, IP station to IP station call usage is 39.7% of the 32 Erlangs of overall Site 1 intercom CUR, or 12.7 Erlangs. Similarly, both the Site 1 IP station to circuit-switched station CUR and the Site 1 circuit-switched station to IP station CUR are equal to 23.3% of 32 Erlangs, or 7.5 Erlangs. Finally, the Site 1 circuit-switched station to circuit-switched station CUR is 13.7% of 32 Erlangs, or 4.4 Erlangs.

A similar process is used to break down the 12 Erlangs of intercom CUR into its components. There are a total of 648 stations in Site 2 (Atlanta), 450 of which are IP stations, and 198 of which are circuit-switched stations. So, 69% of the stations in Site 2 are IP, and 31% are circuit-switched. We have already determined that 63% of intercom calls that are generated in Site 1 are generated by IP stations. Similarly, 69% of intercom calls that are terminated in Site 2 are terminated by IP stations. Since 69% of 63% is 43.5%, 43.5% of Site 1 to Site 2 intercom calls are IP station to IP station. Also, 31% of intercom calls that are terminated in Site 2 are terminated by circuit-switched stations. Since 31% of 63% is 19.5%, 19.5% of Site 1 to Site 2 intercom calls are IP station to circuit-switched station.

In addition, 37% of intercom calls that are generated in Site 1 are generated by circuit-switched stations, and 69% of those calls are terminated by IP stations. Since 69% of 37% is 25.5%, 25.5% of Site 1 to Site 2 intercom calls are circuit-switched station to IP station. Finally, 37% of intercom calls that are generated in Site 1 are generated by circuit-switched stations, and 31% of those calls are terminated by circuit-switched stations. Since 31% of 37% is 11.5%, 11.5% of Site 1 to Site 2 intercom calls are circuit-switched station to circuit-switched station.

So, since 43.5% of Site 1 to Site 2 intercom calls are IP station to IP station, Site 1 IP station to Site 2 IP station CUR is 43.5% of the 12 Erlangs of overall Site 1 to Site 2 intercom CUR, or 5.2 Erlangs. Similarly, the Site 1 IP station to Site 2 circuit-switched station CUR is equal to 19.5% of 12 Erlangs, or 2.3 Erlangs, and the Site 1 circuit-switched station to Site 2 IP station CUR is equal to 25.5% of 12 Erlangs, or 3.1 Erlangs. Finally, the Site 1 circuit-switched station to Site 2 circuit-switched station CUR is 11.5% of 12 Erlangs, or 1.4 Erlangs.

The values for the remaining COI cells that correspond to intercom traffic for this example are calculated in a similar manner. Table 18, COI matrix for Example 4 (intercom CUR values only), on page 118 summarizes the results of that exercise:

Table 18: COI matrix for Example 4 (intercom CUR values only)

			To endpoints in site ___								
			1			2			3		
			I	C	P	I	C	P	I	C	P
From endpoints in site ___	1	I	12.7	7.5		5.2	2.3		2.8	0.37	
		C	7.5	4.4		3.1	1.4		1.6	0.22	
		P									
	2	I	5.2	3.1		1.9	0.85		1.2	0.16	
		C	2.3	1.4		0.85	0.37		0.54	0.07	
		P									
	3	I	2.8	1.6		1.2	0.54		0.78	0.10	
		C	0.37	0.22		0.16	0.07		0.10	0.01	
		P									

The general formula that is used to determine the expanded intercom CUR entries in Table 18, COI matrix for Example 4 (intercom CUR values only), on page 118 is:

CUR generated by stations of type t in Site i and terminated by stations of type t in Site j $= f_i^{\,t} \times f_j^{\,t} \times$ (intercom CUR from Site i to Site j)

where:

- "Type *t*" refers to IP or circuit-switched
- $f_i^t = \frac{\text{number of type } t \text{ stations in Site } i}{\text{total number of stations in Site } i}$
- $f_j^t = \frac{\text{number of type } t \text{ stations in Site } j}{\text{total number of stations in Site } j}$

Now that the intercom CURs have been determined, CURs that involve trunks will be addressed. First, because Communication Manager systems are rarely used to route PSTN traffic, all of the COI matrix entries that correspond to PSTN-PSTN traffic are zero. Next, Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113 and Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113 help us determine the entries that correspond to inbound and outbound PSTN traffic.

According to Table 13, Inbound COI matrix for the Uniform Distribution model in Example 2, on page 113, the inbound PSTN usage that arrives on Site 1 trunks and terminates at Site 1 stations is 49 Erlangs. We have already determined that 63% of the stations in Site 1 are IP and 37% are circuit switched. Therefore, the Uniform Distribution model implies that 63% of the 49 Erlangs (that is, 30.9 Erlangs) is inbound to Site 1 IP stations, and 37% of the 49 Erlangs (that is, 18.1 Erlangs) is inbound to Site 1 circuit-switched stations. Similarly, the Uniform Distribution model and Table 14, Outbound COI matrix for Uniform Distribution Model in Example 2, on page 113 together imply that 63% of the 49 Erlangs of Site 1 outbound PSTN usage (that is, 30.9 Erlangs) is outbound from Site 1 IP stations through Site 1 PTSN trunks, and 37% (that is, 18.1 Erlangs) is outbound from Site 1 circuit-switched stations through Site 1 PTSN trunks. Note that by an assumption in Example 2, Site 1 inbound and outbound traffic only terminates and originates at Site 1 stations. This completes the work for Site 1. Sites 2 and 3 are handled in a similar manner, and the resulting completed COI matrix for Example 4 is provided in Table 19, Completed COI matrix for Example 4, on page 119.

Table 19: Completed COI matrix for Example 4

			To endpoints in Site ___								
			1			2			3		
			I	C	P	I	C	P	I	C	P
From endpoints in Site ___	1	I	12.7	7.5	0	5.2	2.3	0	2.8	0.37	0
		C	7.5	4.4	0	3.1	1.4	0	1.6	0.22	0
		P	30.9	18.1	0	0	0	0	0	0	0
	2	I	5.2	3.1	0	1.9	0.85	0	1.2	0.16	0
		C	2.3	1.4	0	0.85	0.37	0	0.54	0.07	0
		P	0	0	0	12.5	5.5	0	0	0	0
	3	I	2.8	1.6	0	1.2	0.54	0	0.78	0.10	0
		C	0.37	0.22	0	0.16	0.07	0	0.10	0.01	0
		P	0	0	0	0	0	0	7.1	0.94	0

The general formula that is used to determine the expanded inbound and outbound CUR entries in Table 19, Completed COI matrix for Example 4, on page 119 is:

$$\begin{pmatrix}\text{Inbound CUR to stations of type } t \\ \text{in Site } j \text{ over PSTN trunks in Site } i\end{pmatrix} = f_j^{\,t} \times \begin{pmatrix}\text{inbound CUR from trunks in Site } i \\ \text{to stations in Site } j\end{pmatrix}$$

$$\begin{pmatrix}\text{Outbound CUR from stations of type } t \\ \text{in Site } i \text{ over PSTN trunks in Site } j\end{pmatrix} = f_i^{\,t} \times \begin{pmatrix}\text{outbound CUR from stations in Site } i \\ \text{to trunks in Site } j\end{pmatrix}$$

where:

- "Type *t*" refers to IP or circuit-switched
- $f_i^{\,t} = \dfrac{\text{number of type } t \text{ stations in Site } i}{\text{total number of stations in Site } i}$
- $f_j^{\,t} = \dfrac{\text{number of type } t \text{ stations in Site } j}{\text{total number of stations in Site } j}$

In general, one may choose to expand a COI matrix in any of several different possible ways, depending upon the needs of the problem. In the preceding example, separating the endpoints into IP, circuit-switched, and PSTN makes sense for the upcoming resource-sizing calculations, as will be seen later in this document. In other examples, other sets of categories may be more appropriate. Also, the number of categories per site is not limited to three.

COIs for multiple-site networks

The discussion of COIs up to this point has been limited to stand-alone Communication Manager systems. It is also possible to network several Communication Manager systems together. IP tie trunks serve as the most common mode of interconnectivity. However, circuit-switched tie trunks are also supported.

To engineer a network of multiple Communication Manager systems, one must to know the topology of sites within each of the individual systems, and the overall topology of the entire configuration. Consider the network of systems that Figure 50, Network of Avaya systems and system sites, on page 121 shows.

Figure 50: Network of Avaya systems and system sites

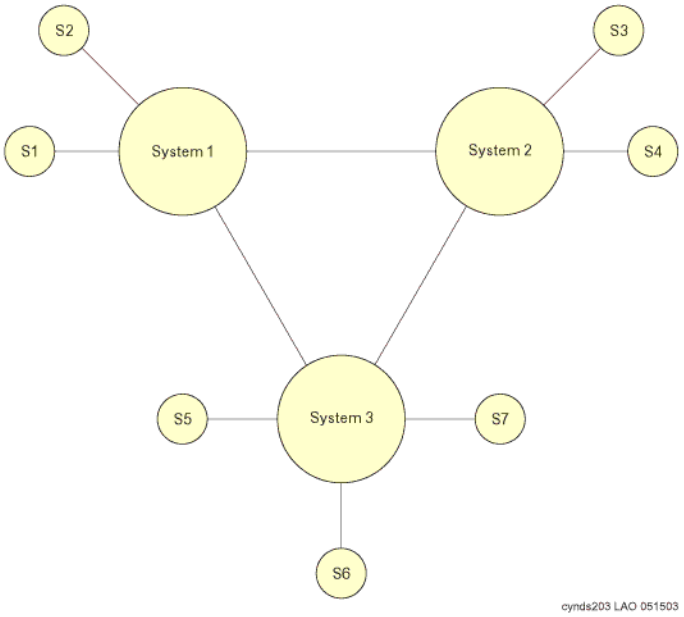

Figure 50, Network of Avaya systems and system sites, on page 121 shows three distinct Communication Manager systems, that are interconnected by IP trunk groups. This network has a total of seven sites, which are labeled “S1” through “S7” in the figure. Systems 1 and 2 each have two sites, and System 3 has three sites.

A seven-site COI matrix analogous to the three-site matrix in Table 12, 3-site standalone community of interest (COI) matrix, on page 110 can be constructed for the network shown in Figure 50, Network of Avaya systems and system sites, on page 121. A corresponding seven-site, expanded COI matrix, similar to the one in Table 16, Expanded COI matrix for a three-site system, on page 116, can also be constructed. However, when multiple systems are networked together, the additional step of engineering the tie trunk groups must be performed. To do this, the COI matrices are used to determine the traffic flow between each pair of Avaya systems.

In the network that is shown in Figure 50, Network of Avaya systems and system sites, on page 121, IP Trunk Group 1 carries calls between Sites 1 and 3, Sites 1 and 4, Sites 2 and 3, and Sites 2 and 4, in addition to a presumably small amount of overflow traffic that involves other sites. The traffic load that is associated with such calls is used to size that trunk group. Tie trunk groups are typically sized at either P01 (1% blocking) or P03 (3% blocking). In a system such as the one in Figure 50, Network of Avaya systems and system sites, on page 121, the traffic engineer must account for overflow traffic. The traditional Wilkinson model is an effective tool for doing so. However, for systems that have larger numbers of systems in the network, there can be many possible paths between a given pair of systems. In such cases, determining the hierarchy of paths to consider for calls between two systems is not always straightforward. The analysis involved in sizing the tie trunk groups in topologies such as those can be quite complex.

Resource sizing

Overview

The primary Communication Manager resources that have the potential to be bottlenecks are the TN799 C-LAN (Control LAN) circuit packs, the port network TDM bus pairs, the TN2302 Media Processing circuit packs, the TN2312 IPSI circuit packs, the server's processing capacity, and IP bandwidth. This section of the document provides a description of each of these resources, and a discussion about how to engineer them. This is the final stage of the design process.

Signaling resources

The TN799 C-LAN circuit pack provides the interface for a signaling channel between an IP endpoint and a packet bus (which ultimately interfaces with the Avaya server). When an IP endpoint, G350 MG, or G700 MG registers to a C-LAN circuit pack, a software object known as a C-LAN *socket* is allocated for that endpoint or gateway to use, for as long as it remains registered. C-LAN sockets are also required for the support of certain adjuncts.

Each C-LAN circuit pack can only support a finite number of C-LAN sockets. The total number of C-LAN circuit packs that are required to support a particular system depends on the total required number of C-LAN sockets, which in turn depends on the total number of IP endpoints, G350 MGs, G700 MGs, and adjuncts. Once the required number of C-LAN circuit packs is determined, their physical placement is not particularly important. An individual C-LAN circuit pack can support endpoints in different Network Regions, even those that are not administered to communicate with each other.

Media processing and TDM resources

The media processing resources on TN2302 Media Processing circuit packs on a PN, a G600 MG, or a G650 MG, provide the gateway for an audio channel between an IP endpoint and a circuit-switched TDM bus. On a G350 MG or G700 MG, the media processing resources reside on an on-board VoIP module. A G700 MG can accommodate an optional extra VoIP module as well. The media stream for a call between a circuit-switched endpoint and an IP endpoint on a PN or MG traverses the PN's or MG's TDM bus, a TN2302 Media Processing circuit pack or a VoIP module (as applicable), and an IP network. The media stream for a call between two circuit-switched endpoints on a single port network or Media Gateway uses that PN or Media Gateway's TDM bus, and does not require any media processing resources. However, the media stream for a call between two circuit-switched endpoints that reside on different circuit-switched facilities (that is, two different PNs, two different Media Gateways, or one PN and one Media Gateway) traverses each circuit-switched facility's TDM bus, a media processing resource on each circuit-switched facility (a Media Processing circuit pack or VoIP Media Module, as applicable), and an IP network. Figure 51, Examples of media streams between Avaya endpoints, on page 123 shows some examples of the various possible media streams.

Figure 51: Examples of media streams between Avaya endpoints

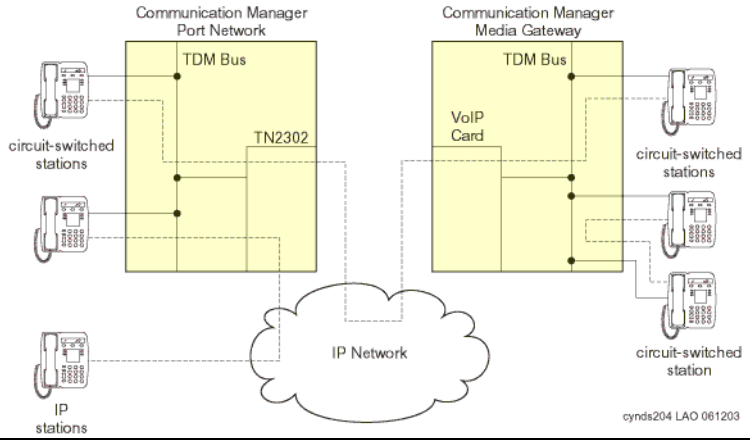

Although we stated that calls between two circuit-switched endpoints on different port networks use an IP connection, the use of a circuit-switched center stage between the two PNs is also supported. However, using circuit-switched facilities is not viable for interconnecting multiple Media Gateways, or for interconnecting PNs and Media Gateways.

Figure 51, Examples of media streams between Avaya endpoints, on page 123 provides some insight into how a call between an IP endpoint and a circuit-switched endpoint, as well as a call between two circuit-switched endpoints, utilizes media processing and TDM resources. Calls between IP endpoints are addressed first.

Communication Manager supports three general modes of connectivity between IP endpoints: *IP-TDM-IP* connectivity, *hairpinning*, and *shuffling*. Hairpinning can take one of two forms: *deep* or *shallow.* These various modes of connectivity are described in more detail below.

IP-TDM-IP connectivity

A call that uses IP-TDM-IP connectivity between two IP endpoints requires one bidirectional media processing “channel” for each IP endpoint involved, as well as a bidirectional TDM resource on every PN (or Media Gateway) that is involved in the call. This option most often applies in systems that use a circuit-switched center stage for interport network connectivity. In such a system, IP-TDM-IP is required in order for two IP endpoints in network regions not configured for connectivity (in the sense of the IRCM matrix) to talk to one another.

Hairpinning

Unlike the IP-TDM-IP connectivity option, hairpinning requires that all media processing resources for a given call reside on a single TN2302 Media Processing circuit pack or a single G350 or G700 Media Gateway VoIP Media Module. A hairpinned call is originally set up as an IP-TDM-IP call, but once the set-up process is complete, no TDM resources are required. However, resources on the Media Processing circuit pack or VoIP Media Module are required for the duration of the call. A Media Processing circuit pack and a VoIP Media Module each house an onboard Central Processing Unit (CPU) and Digital Signal Processors (DSPs).

Shuffling

A shuffled call relinquishes all TDM and media processing resources after call setup. Therefore, the media stream of a shuffled call traverses only an IP network. This is the most commonly used mode of connectivity between two IP endpoints in the same system.

Figure 52, Connectivity modes between two IP endpoints, on page 124 shows the various modes of connectivity between two IP endpoints.

Figure 52: Connectivity modes between two IP endpoints

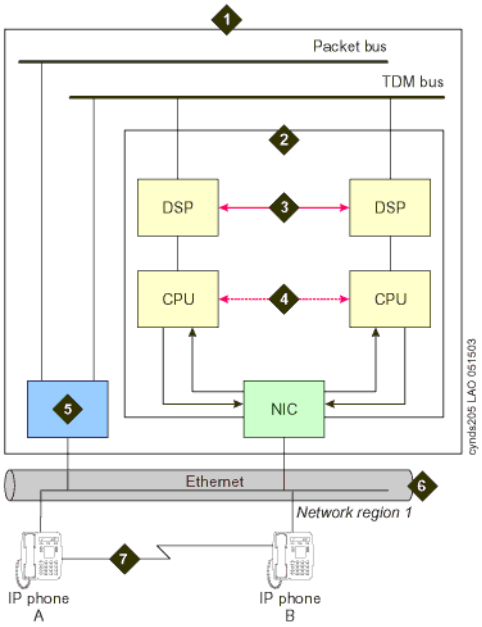

Figure notes

1	Avaya server	**5**	TN799DP Control LAN (C-LAN) circuit pack
2	TN2302AP IP Media Processor (MedPro) circuit pack	**6**	Customer LAN
3	Deep hairpinned audio connection	**7**	Shuffled audio connection
4	Shallow hairpinned audio connection		

At this point, we can quantify the TDM and media processing requirements for various call types. Throughout this discussion, calls between two IP stations are assumed to use shuffling. That being the case, an intrasite call between two IP endpoints requires neither TDM nor media processing resources, beyond the completion of the initial call set-up process. Each intrasite call between an IP endpoint and a circuit-switched endpoint (including PSTN trunks) requires one TDM resource and one media processing resource. Each of these resources is furnished by the PN or the Media Gateway to which the circuit-switched endpoint is administered. See Figure 51, Examples of media streams between Avaya endpoints, on page 123 for an example.

The TDM and media processing resources that are required for each intrasite call between two circuit-switched endpoints depends upon whether the call is intraport network or interport network. Specifically, each intraport network call requires one TDM resource (on the port network to which the two circuit-switched endpoints are administered), and no media processing resources. See Figure 51, Examples of media streams between Avaya endpoints, on page 123 for an example. Also, assuming that IP interport network connectivity is being used (as opposed to a center stage), each interport network call requires two TDM resources and two media processing resources. One of each of these resources is supplied by each of the PNs that is involved in the call. In the preceding discussion, everything that applies to a PN also applies to a Media Gateway.

In general, the TDM and media processing requirements for intersite calls are accounted for somewhat differently than the requirements for intrasite calls. Throughout this discussion, we assume that shuffling is implemented. When an IP endpoint is involved in an intersite call, it induces no TDM or media processing usage in its own site beyond the resources that are initially required for the call set-up process, regardless of the nature of the far-end party. On the other hand, when a circuit-switched endpoint (including PSTN trunks) is involved in an intersite call, one TDM resource and one media processing resource are required from the port network or Media Gateway to which it is administered, regardless of the nature of the far-end party.

The preceding discussion is summarized in Figure 53, Intra-site TDM and Media Processing resource requirements, on page 125 and Figure 54, Inter-site TDM and Media Processing resource requirements, on page 126.

Figure 53: Intra-site TDM and Media Processing resource requirements

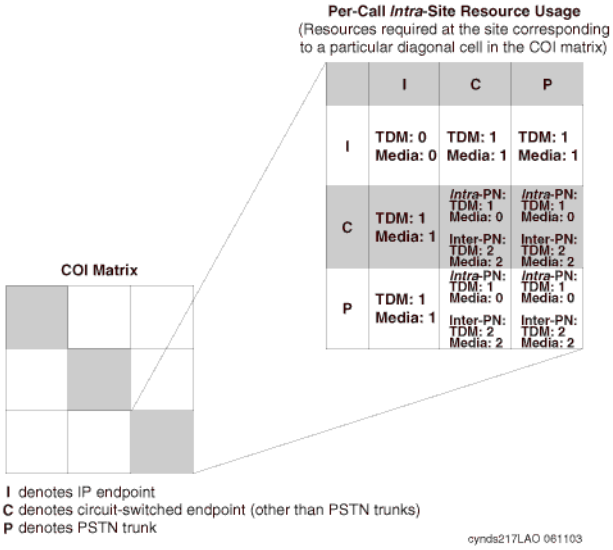

Figure 54: Inter-site TDM and Media Processing resource requirements

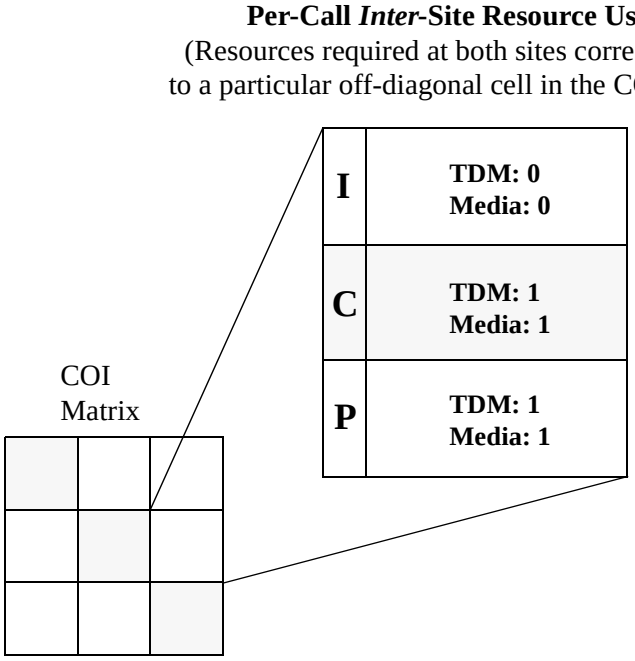

"I" denotes IP endpoint; "C" denotes circuit-switched endpoint (other than PSTN trunks);
"P" denotes PSTN trunk

In Figure 54, Inter-site TDM and Media Processing resource requirements, on page 126, the usages are presented on an endpoint-by-endpoint basis. For example, according to Figure 54, Inter-site TDM and Media Processing resource requirements, on page 126, an intersite call between an IP endpoint in Site 1 and a circuit-switched endpoint in Site 2 requires no TDM or media processing resources in Site 1, but does require one TDM resource and one media processing resource in Site 2.

The overall TDM usage and media processing usage for each site can be calculated from an expanded COI matrix, along with the information from Figure 53, Intra-site TDM and Media Processing resource requirements, on page 125 and Figure 54, Inter-site TDM and Media Processing resource requirements, on page 126. To illustrate, Example 4 will be further expanded.

Example 5

TDM and media processing usage

Consider the COI matrix in Table 19, Completed COI matrix for Example 4, on page 119 in Example 4. A set of nine cells corresponds to calls originated in Site 1 and terminated in Site 1 (that is, the upper left group of nine cells, arranged in a three-by-three submatrix). The uppermost and leftmost cell of those nine cells indicates that the IP-to-IP call usage for Site 1 intrasite calls is 12.7 Erlangs. The other four cells of those nine cells which fall in a row or column that is labeled "I" indicate that the total call usage between IP endpoints and circuit-switched endpoints (including PSTN trunks) within Site 1 is (7.5 + 30.9 + 7.5 + 30.9) = 76.8 Erlangs. The remaining four cells of those nine cells indicate that the total call usage between two circuit-switched endpoints (including PSTN trunks) within Site 1 is (4.4 + 18.1 + 18.1 + 0) = 40.6 Erlangs. Analogous numbers for intrasite usages that correspond to the other two sites are similarly derived.

Next, consider the three-by-three submatrix that corresponds to calls from Site 1 to Site 2. The total call usage from Site 1 to Site 2 which involves an IP endpoint in Site 1 can be determined by adding the three cell values of those nine cells that correspond to IP endpoints in Site 1. Specifically, the total is (5.2 + 2.3 + 0) = 7.5 Erlangs. The total call usage from Site 1 to Site 2 which involves a circuit-switched endpoint (including PSTN trunks) in Site 1 can be determined by adding the remaining six cell values of those nine. Specifically, that total is (3.1 + 1.4 + 0 + 0 + 0 + 0) = 4.5 Erlangs.

The total call usage from Site 1 to Site 2 which involves an IP endpoint in Site 2 can be determined by adding the three cell values of those nine that correspond to IP endpoints in Site 2. Specifically, the total is (5.2 + 3.1 + 0) = 8.3 Erlangs. And finally, the total call usage from Site 1 to Site 2 which involves a circuit-switched endpoint (including PSTN trunks) in Site 2 can be determined by adding the remaining six cell values of those nine. Specifically, that total is (2.3 + 1.4 + 0 + 0 + 0 + 0) = 3.7 Erlangs. Analogous numbers for the other five combinations of intersite usages are similarly derived. The results are shown in Table 20, Recategorization of CURs from Table 19, Completed COI matrix for Example 4, on page 119, on page 127.

Table 20: Recategorization of CURs from Table 19, Completed COI matrix for Example 4, on page 119

Endpoints		Site 1 (Atlanta)	Site 2 (Boston)	Site 3 (Cleveland)
Intrasite: I, I		12.7 E	1.9 E	0.78 E
Intrasite: I, C or P		76.8 E	26.7 E	14.4 E
Intrasite: C or P, C or P		40.6 E	11.4 E	1.9 E
Calls from Site 1 to Site 2	I	7.5 E	8.3 E	0
	C or P	4.5 E	3.7 E	0
Calls from Site 2 to Site 1	I	7.5 E	8.3 E	0
	C or P	4.5 E	3.7 E	0
Calls from Site 1 to Site 3	I	3.2 E	0	4.4 E
	C or P	1.8 E	0	0.59 E
Calls from Site 3 to Site 1	I	3.2 E	0	4.4 E
	C or P	1.8 E	0	0.59 E
Calls from Site 2 to Site 3	I	0	1.4 E	1.7 E
	C or P	0	0.61 E	0.23 E
Calls from Site 3 to Site 2	I	0	1.4 E	1.7 E
	C or P	0	0.61 E	0.23 E

Table 20, Recategorization of CURs from Table 19, Completed COI matrix for Example 4, on page 119, on page 127 provides a summary of call usage rates, which can be mapped to a table of TDM usage rates and media processing usage rates by using the information in Figure 53, Intra-site TDM and Media Processing resource requirements, on page 125 and Figure 54, Inter-site TDM and Media Processing resource requirements, on page 126. We assume that there is only one PN in Site 1, one in Site 2, and two G350 Media Gateways in Site 3. Under this assumption, which will be assessed shortly, all calls between circuit-switched endpoints in Sites 1 and 2 are assumed to be intra-Port Network. A minimum of two G350 Media Gateways is required to house the 36 analog telephones in Site 3. The results of this exercise are shown in Table 21, TDM and Media Processing usages (Erlangs) for Example 5, on page 128.

Table 21: TDM and Media Processing usages (Erlangs) for Example 5

Endpoints		Site 1 (Atlanta)	Site 2 (Boston)	Site 3 (Cleveland)
Intrasite: I, I		TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
Intrasite: I, C or P		TDM: 76.8 Media: 76.8	TDM: 26.7 Media: 26.7	TDM: 14.4 Media: 14.4
Intrasite: C or P, C or P		TDM: 40.6 Media: 0	TDM: 11.4 Media: 0	TDM: 1.9 Media: 0
Calls from Site 1 to Site 2	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 4.5 Media: 4.5	TDM: 3.7 Media: 3.7	TDM: 0 Media: 0
Calls from Site 2 to Site 1	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 4.5 Media: 4.5	TDM: 3.7 Media: 3.7	TDM: 0 Media: 0
Calls from Site 1 to Site 3	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 1.8 Media: 1.8	TDM: 0 Media: 0	TDM: 0.59 Media: 0.59
Calls from Site 3 to Site 1	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 1.8 Media: 1.8	TDM: 0 Media: 0	TDM: 0.59 Media: 0.59
Calls from Site 2 to Site 3	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 0 Media: 0	TDM: 0.61 Media: 0.61	TDM: 0.23 Media: 0.23
Calls from Site 3 to Site 2	I	TDM: 0 Media: 0	TDM: 0 Media: 0	TDM: 0 Media: 0
	C or P	TDM: 0 Media: 0	TDM: 0.61 Media: 0.61	TDM: 0.23 Media: 0.23
Totals		TDM: 130.0 Media: 89.4	TDM: 46.7 Media: 35.3	TDM: 17.9 Media: 16.0

The TDM usage rates of 130.0 Erlangs for Site 1 and 46.7 Erlangs for Site 2 can both be easily handled by the TDM facilities of a single PN, which is capable of carrying up to 200 Erlangs of TDM traffic at a P001 GOS. Therefore, the assumption that all calls between two circuit-switched endpoints is intra-Port Network is valid. If one PN was insufficient to support the TDM usage in one of the sites, the calculations would have been repeated under the assumption of two PNs. If a pair of PNs was still insufficient, the number would continually be incremented until there were enough port networks to handle the TDM usage in that particular site. Finally, the TDM resources on the two G350 Media Gateways are easily sufficient for supporting the 17.9 Erlangs of TDM traffic in Site 3.

> **NOTE:**
> The more PNs, the more inter-Port Network calls there are, and hence more TDM usage, since each interport network call requires resources in *each* PN that is involved in each call.

Each TN2302 Media Processing circuit pack (or Media Gateway VoIP Media Module) can support only a finite number of simultaneous calls. However, the exact number that can be supported varies according to the codecs of the calls to be supported. In general, compressed calls (for example, G.729 codec) require twice as many media processing resources as uncompressed calls (for example, G.711 codec). Also, calls utilizing AES media encryption require approximately 25% more media processing resources than unencrypted calls.

A TN2302 Media Processing circuit pack (or a MG VoIP Media Module) can support both compressed and uncompressed calls, as well as both encrypted and unencrypted calls, all simultaneously. Therefore, the general model for sizing the media processing resources is very complex. The model is a “batch arrival and service” model, and the details are beyond the scope of this document.

In practice, a fairly common strategy is to use an uncompressed codec for intrasite calls, and a compressed codec for intersite calls. This is due to the trade-off between bandwidth savings, increased media processing costs, and voice quality for compressed calls. If a private LAN is used for intrasite calls, bandwidth usage is of less concern than media processing cost and voice quality. However, for intersite calls, especially over a public WAN, the bandwidth savings offered by the use of compression outweighs the extra processing costs and slight degradation of voice quality.

Recall that any usage that is expressed in Erlangs represents the average number of busy servers at any given time. For the total media processing usages provided at the bottom of Table 21, TDM and Media Processing usages (Erlangs) for Example 5, on page 128, a “server” can be thought of as the set of media processing resources that is necessary to support a single bidirectional media stream through a media processing circuit pack. Consider the total of 89.4 Erlangs of media processing usage in Site 1. This usage consists of 76.8 Erlangs of intrasite usage, and 12.6 Erlangs of intersite usage. Assume that an uncompressed codec is used for the intrasite calls, and a compressed codec is used for the intersite calls. Since each compressed call requires twice as many media processing resources as each uncompressed call, the 12.6 Erlangs must be counted twice. Therefore, the media processing load is actually 76.8 + (2 x 12.6) = 102.0 Erlangs. Similarly, the total media processing loads in Sites 2 and 3 are 43.9 Erlangs and 17.6 Erlangs, respectively. Those numbers are also based on the assumption that media encryption was not used.

Table 22: Number of TN2302 Media Processors or G700 Media Gateway VoIP Modules required for a given carried load

Carried load (Erlangs)	Required number of TN2302 circuit packs	Carried load (Erlangs)	Required number of TN2302 circuit packs
43	1	634	11
98	2	695	12
155	3	756	13
213	4	817	14
272	5	879	15
332	6	940	16
392	7	1,001	17
452	8	1,063	18
512	9	1,125	19
573	10	1,187	20

Table 23: Number of G350 Media Gateway VoIP Modules required for a given carried load

Carried load (Erlangs)	Required number of G700 MGs	Carried load (Erlangs)	Required number of G700 MGs
18	1	155	6
43	2	184	7
70	3	213	8
98	4	243	9
126	5	272	10

Table 22, Number of TN2302 Media Processors or G700 Media Gateway VoIP Modules required for a given carried load, on page 130 implies that three TN2302 Media Processing circuit packs should be used in Site 1 (Atlanta), and two should be used in Site 2 (Boston). Table 23, Number of G350 Media Gateway VoIP Modules required for a given carried load, on page 130 implies that the media processing resources on the two G350 Media Gateways in Site 3 (Cleveland) are easily sufficient. The required number of port networks, MGs, and media processing resources for Example 5 is summarized in Table 24, TDM and Media Processing Requirements for Example 5, on page 131.

Table 24: TDM and Media Processing Requirements for Example 5

Site	TDM Requirement	Media Processing Requirement
1	1 PN	3 TN2302 boards
2	1 PN	2 TN2302 boards
3	The 2 G350 MGs are sufficient	The on-board VoIP resources on the 2 G350 MGs are sufficient

If more than one PN had been required in a particular site, intrasite calls between circuit-switched endpoints in that site would have contributed toward media processing usage because inter-Port Network calls between circuit-switched endpoints traverse an IP network. Since only one PN is required in each site in this example, the media-processing usage for calls between circuit-switched endpoints is zero in each site, as indicated in Table 21, TDM and Media Processing usages (Erlangs) for Example 5, on page 128.

TN2312 IPSI circuit packs

Sizing the TN2312 IP Server Interface (IPSI) circuit packs is a fairly straightforward process. The number of IPSI circuit packs that are required in the system depends on the total number of C-LAN sockets that are required, and the number of ISDN D-channels in the system. Specifically, each IPSI circuit pack supports up to a combined total of 2,480 C-LAN sockets and ISDN D-channels. This is a system-wide constraint, as opposed to a site-by-site constraint. For an IP-Connect system, each PN must house exactly one IPSI circuit pack, neglecting duplicated IPSI circuit packs for enhanced reliability. Therefore, if the C-LAN sockets and the ISDN D-channels indicate a need for more IPSI circuit packs than the required number of PNs to support the TDM usage, more PNs are needed (note that placing two active IPSI circuit packs in a single PN is not permitted). In other words, the number of PNs must be large enough to fulfill both the TDM and the IPSI requirements.

In a system utilizing a circuit-switched center stage an IPSI circuit pack is not required in each port network. However, there are restrictions pertaining to how many port networks can be supported by a single IPSI circuit pack.

If the number of port networks needs to be increased to satisfy the IPSI requirements, then the TDM and media processing engineering processes must be redone (since an increased number of port networks implies an increase in inter-port-network traffic). This is an iterative process.

Processing occupancy

The Busy Hour Call Attempt (BHCA) rate of a system is the total number of calls that are attempted within that system, during its busy hour. This is distinct from the Busy Hour Call Completion (BHCC) rate of a system, which counts only those calls that have actually been completed. The *call capacity* of a system refers to its BHCC rate.

In Communication Manager products, server occupancy (or processor occupancy, as applicable) can be broken down into three categories: static occupancy, Call Processing (CP) occupancy, and system management (SM) occupancy. The static component refers to keep-alive processes, the CP component refers to processes that are required to set up and tear down calls (as well as vectoring operations, in the case of call centers), and the SM component refers to background maintenance operations and system audits. In theory, static occupancy is a fixed overhead, and CP occupancy is directly proportional to the call rate. SM occupancy is allocated on an as-needed basis, such as for periodic maintenance functions. However, if the overall server occupancy exceeds a particular threshold, SM operations are postponed until a quieter traffic period.

Usually, the relationship between the sum of static and CP occupancy, as a function of BHCC, is linear, with a positive y-intercept, as illustrated in Figure 55, Relationship Between Processing Occupancy and BHCC Rate, on page 132. The slope of the line corresponds to the average processing cost per call, and the intercept corresponds to the idle (that is, no-load) occupancy. The average processing cost per call depends on the mix of calls that is being handled by the system, and how complex each type of call is. For general business calls, nearly all of the CP occupancy is associated with setting up and tearing down calls. The call processing that is required for maintaining the call once it has been established is negligible in comparison, regardless of how long the call lasts. In a call center, the additional cost of processing vectoring steps throughout the lifetime of a call must also be considered.

Figure 55: Relationship Between Processing Occupancy and BHCC Rate

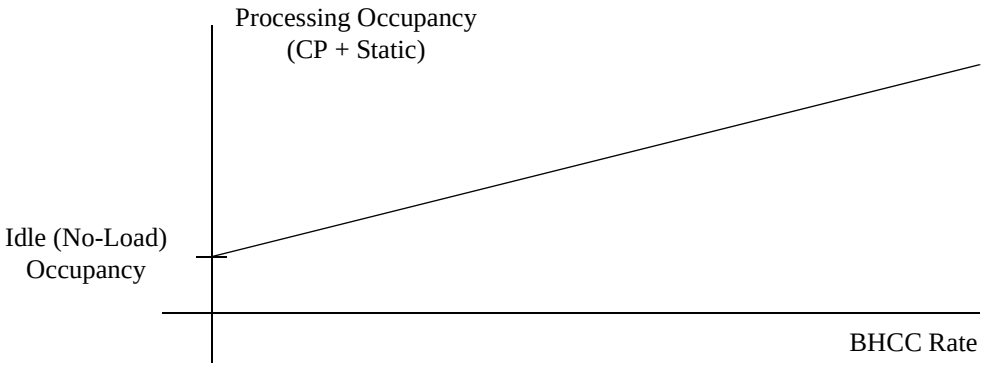

To determine the anticipated processor occupancy that is associated with a particular configuration, the average processing cost per call must be determined based on the anticipated volume of each type of call, and the complexity of the various call types. This average cost per call implies the slope of the line in relating static and CP occupancy to the BHCC rate. The intercept of that line, which corresponds to the no-load occupancy, depends on several factors, including which Communication Manager platform is being used, how many endpoints are administered, and so on.

Communication Manager systems are designed to keep the sum of static and CP occupancy below a particular threshold. This is done to allow a suitable amount of processing time for system management functions.

So for a given configuration, the various types of calls to be supported are identified, and the processing cost for each call type (based upon the complexity of the call) must be assessed. That information can then be used to determine the average processing cost per call, based on the anticipated relative frequencies of the various call types. The slope of the line relating the sum of static and CP occupancy can then be determined from the average processing cost per call. The intercept of that line is determined by information such as the Communication Manager platform used, the number of endpoints administered, and so on.

Therefore, for the given configuration, the specific linear model for the relationship between the sum of static and CP occupancy, as a function of BHCC, has been derived. Using the anticipated BHCC rate in that model yields the expected combined static and CP occupancy. If that value exceeds the preset threshold, the configuration is unacceptable for the anticipated call rate. In such a case, to support that call rate, either another platform must be considered, or multiple platforms must be networked together.

IP Bandwidth and Call Admission Control

IP bandwidth analysis for media streams begins with determining the number of bidirectional media streams that are associated with each type of call supported by the system. Throughout this discussion, calls between two IP stations are assumed to use shuffling. That being the case, Figure 52, Connectivity modes between two IP endpoints, on page 124 indicates that an intrasite call between two IP endpoints requires a single bidirectional media stream through the LAN at that site. Figure 51, Examples of media streams between Avaya endpoints, on page 123 indicates that each intrasite call between an IP endpoint and a circuit-switched endpoint (including PSTN trunks) also requires a single bidirectional media stream through the LAN at that site. In addition, Figure 51, Examples of media streams between Avaya endpoints, on page 123 indicates that each interport network intrasite call between two circuit-switched endpoints (including PSTN trunks) also requires a single bidirectional media stream through the LAN at that site (assuming that IP-Connect is used, as opposed to a circuit-switched center stage). In fact, the only intrasite call that does not require a single bidirectional media stream through the LAN at that site is an intraport network call between two circuit-switched endpoints which requires no IP resources because the call is completed solely across the circuit-switched TDM bus of the PN. Each intersite call requires exactly one bidirectional media stream through each participating site's LAN, as well as a single bidirectional media stream through the WAN that connects the two sites.

The preceding discussion is summarized in Figure 56, Required number of bidirectional IP media streams for intra-site calls, on page 133 and Figure 57, Required number of bidirectional IP media streams for inter-site calls, on page 134.

Figure 56: Required number of bidirectional IP media streams for intra-site calls

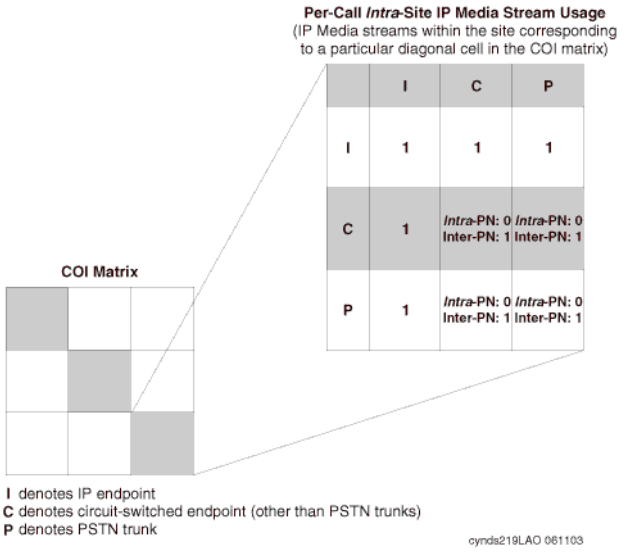

Figure 57: Required number of bidirectional IP media streams for inter-site calls

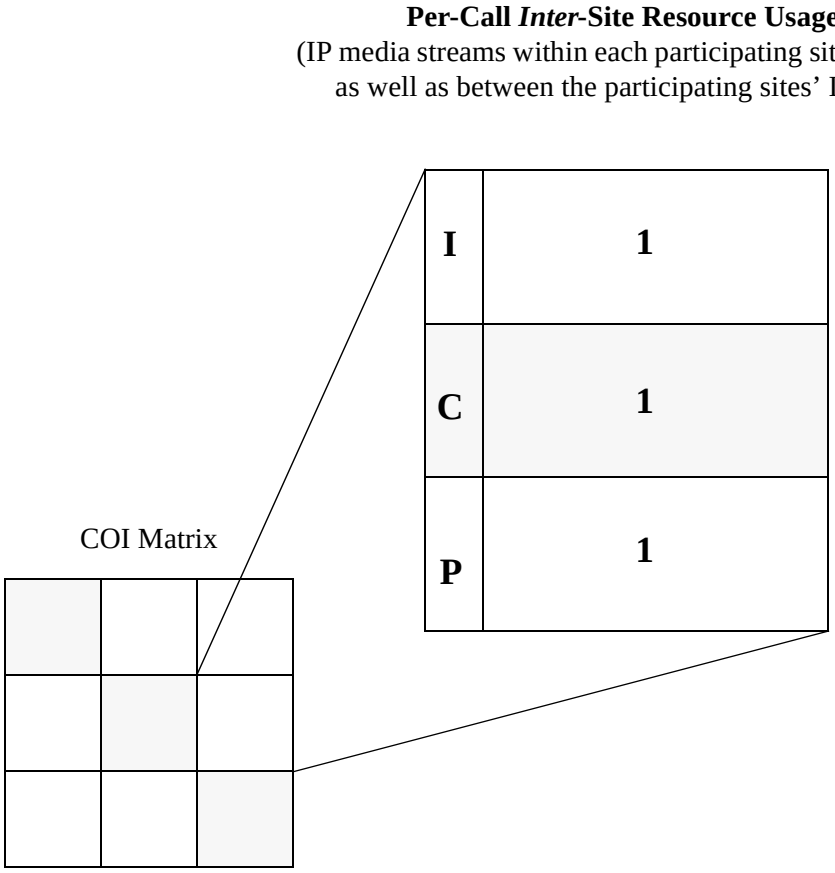

"I" denotes IP endpoint; "C" denotes circuit-switched endpoint (other than PSTN trunks); "P" denotes PSTN trunk

Figure 56, Required number of bidirectional IP media streams for intra-site calls, on page 133 and Figure 57, Required number of bidirectional IP media streams for inter-site calls, on page 134 provide information regarding the required number of bidirectional media streams per call. This information can be combined with call usage information to provide IP bandwidth usage estimates, as shown in Example 6.

Example 6

IP bandwidth considerations

The information in Figure 56, Required number of bidirectional IP media streams for intra-site calls, on page 133 and Figure 57, Required number of bidirectional IP media streams for inter-site calls, on page 134 along with the information in Table 20, Recategorization of CURs from Table 19, Completed COI matrix for Example 4, on page 119, on page 127 produces the following tables of bandwidth usages that are associated with the configuration in Example 4:

Table 25: IP LAN bandwidth usages (Erlangs) for Example 6

Endpoints	Site 1 (Atlanta)	Site 2 (Boston)	Site 3 (Cleveland)
Intrasite: I, I	12.7	1.9	0.78
Intrasite: I, C or P	76.8	26.7	14.4
Intrasite: C or P, C or P	0	0	0
Calls from site 1 to site2	12.0	12.0	0
Calls from site 2 to site 1	12.0	12.0	0
Calls from site 1 to site 3	5.0	0	5.0
Calls from site 3 to site 1	5.0	0	5.0
Calls from site 2 to site 3	0	2.0	2.0
Calls from site 3 to site 2	0	2.0	2.0
Totals	123.5	56.6	29.2

Table 26: IP WAN bandwidth usages (Erlangs) for Example 6

Endpoints	WAN bandwidth (Erlangs) between Sites 1 and 2	WAN bandwidth (Erlangs) between Sites 1 and 3	WAN bandwidth (Erlangs) between Sites 2 and 3
Calls from site 1 to site 2	12.0	0	0
Calls from site 2 to site 1	12.0	0	0
Calls from site 1 to site 3	0	5.0	0
Calls from site 3 to site 1	0	5.0	0
Calls from site 2 to site 3	0	0	2.0
Calls from site 3 to site 2	0	0	2.0
Totals	24.0	10.0	4.0

Table 25, IP LAN bandwidth usages (Erlangs) for Example 6, on page 135 and Table 26, IP WAN bandwidth usages (Erlangs) for Example 6, on page 135 express bandwidth usages in Erlangs, because each such usage actually represents the average number of simultaneous bidirectional media streams through the IP network in question. To convert those usages into bandwidth requirements in units of kilobits per second (kbps), one must know how many kbps each call requires. To answer that question, a closer look at IP packet structure is necessary.

An IP packet consists of a payload and some amount of overhead. The payload consists of actual sampled voice, and the overhead represents headers and trailers, which serve to navigate the packet to its proper destination. The overhead due to IP, UDP, and RTP is 40 bytes, while the Ethernet overhead is between 18 and 22 bytes (18 is assumed in this example). This represents a total overhead of 58 bytes (464 bits), regardless of the nature of the payload. For this example, Layer 2 (Ethernet) overhead is included in that total. At every router boundary, because Ethernet overhead is included in this example, our calculations are for bandwidth on a LAN. Because WAN protocol (for example, PPP) Layer 2 headers are generally smaller than Ethernet headers, WAN bandwidth is slightly less than LAN bandwidth.

The size of the payload depends on certain parameters that relate to the codec that is used. The two most common codecs that are used with Avaya products are uncompressed G.711 and compressed G.729. The transmission rates that are associated with those codecs are 64 kbps for G.711 (this is the Nyquist sampling rate for human voice) and 8 kbps for G.729.

The packet “size” is sometimes expressed in units of time (specifically, in milliseconds). The following formula yields the packet size, expressed in bits:

$$\text{Number of bits of payload per packet} = \left(\text{Transmission Rate (kbps)}\right) \times (\text{ms per Packet})$$

Table 27, Payload size per packet, on page 136 is populated using this formula, and provides the payload size per packet (expressed in bits) as a function of packet “size” (that is, ms per packet) and codec.

Table 27: Payload size per packet

Packet “size” (ms)	G.711 (bits)	G.729 (bits)
10	640	80
20	1280	160
30	1920	240
60	3840	480

Note that the number of bits of payload per packet depends on the packet “size,” but it is independent of the “sizes” of the individual frames that are contained in that packet. For example, a packet “size” of 60 ms could be referring to six 10-ms frames per packet, or three 20-ms frames per packet, or two 30-ms frames per packet, and so on. Presently, the most commonly used packet “sizes” are 20 ms. Both G.711 and G.729 codecs typically use two 10-ms frames per packet.

As stated earlier, there is an overhead of 464 bits per packet. So, the bandwidth (expressed in kbps) that is associated with a unidirectional media stream (assuming no Silence Suppression is used) is augmented from 64 kbps and 8 kbps (for G.711 and G.729, respectively) to account for this overhead. The results of this exercise are provided in Table 28, Bandwidth requirements for media streams, on page 137.

Table 28: Bandwidth requirements for media streams

Packet "size" (ms)	G.711 (kbps)	G.729 (kbps)
10	110.4	54.4
20	87.2	31.2
30	79.5	23.5
60	71.7	15.7

Note that the entries in Table 28, Bandwidth requirements for media streams, on page 137 correspond with a single (unidirectional) media stream. As we will see in the following example, the entries in Table 28, Bandwidth requirements for media streams, on page 137 are not multiplied by the *average* number of simultaneous streams, but rather by a much larger number that represents the 99.9th percentile for the simultaneous number of streams.

Example 7

LAN bandwidth

In Example 6, the total IP LAN bandwidth usage for each site was calculated, and expressed in Erlangs at the bottom of Table 25, IP LAN bandwidth usages (Erlangs) for Example 6, on page 135. Specifically, the total LAN bandwidth usage in Site 1 is 123.5 Erlangs, in Site 2 is 56.6 Erlangs, and in Site 3 is 29.2 Erlangs. This implies that the average number of bidirectional media streams that are simultaneously in use at any given time in Site 1 is 123.5. Analogous statements can also be made regarding Sites 2 and 3.

Every media stream across the IP LAN in any of the three sites is assumed to use the uncompressed G.711 codec, since bandwidth is relatively inexpensive within a private LAN, as opposed to a public WAN. Assume, for the sake of this example, a standard IP packet size of 20 ms. So for the G.711 codec, Table 28, Bandwidth requirements for media streams, on page 137 indicates that each media stream consumes 87.2 kbps of IP LAN bandwidth. It may be tempting at this point to simply multiply 87.2 kbps by 123.5 simultaneous bidirectional media streams, to arrive at the estimate for the overall LAN bandwidth needed for Site 1. However, 123.5 is merely the *average* number of simultaneous media streams, and approximately half of the time, there are at least 124 simultaneous media streams in use.

In this example, suppose that the goal is to supply enough bandwidth to adequately support the media streams at least 99.9% of the time. The standard infinite-server queueing model implies that less than 0.1% of the time there are at least 159 simultaneous media streams in the Site 1 LAN. So, it is sufficient to engineer the LAN bandwidth to support 158 simultaneous media streams. Therefore, the Site 1 LAN requires at least (158 simultaneous media streams) x (87.2 kbps per media stream) = 13.8 Mbps of bandwidth, in each direction. This result, along with the analogous results for Sites 2 and 3, are provided in Table 29, IP LAN bandwidth requirements in each direction, for Example 7, on page 138.

Table 29: IP LAN bandwidth requirements in each direction, for Example 7

Resource	Site 1 (Atlanta)	Site 2 (Boston)	Site 3 (Cleveland)
Simultaneous media streams for "P001"	158	81	47
LAN bandwidth (Mbps)	13.8	7.1	4.1

In Table 29, IP LAN bandwidth requirements in each direction, for Example 7, on page 138, the number of simultaneous media streams for "P001" represents the 99.9th percentile for the number of simultaneous unidirectional streams, as determined by applying the standard infinite-server queueing model.

A slight variation of the procedure that was used to determine LAN bandwidth in Example 7 can be used to determine WAN bandwidth. The use of compressed RTP (cRTP) is a means by which bandwidth is conserved. Specifically, the use of cRTP reduces the overhead due to IP, UDP, and RTP from 40 bytes to between 2 and 4 bytes (4 bytes are assumed for this example). Using the PPP overhead of 7 bytes (which would vary if ATM, HDLC, or Frame Relay were used) implies a total overhead of 11 bytes (88 bits) in this example. This implies the following table of WAN bandwidths (Table 30, IP WAN bandwidth requirements for media streams, on page 138), which assumes the use of cRTP:

Table 30: IP WAN bandwidth requirements for media streams

Packet "size" (ms)	G.711 (kbps)	G.729 (kbps)
10	72.8	16.8
20	68.4	12.4
30	66.9	10.9
60	65.5	9.5

This table can be used in the WAN bandwidth calculation for the system in Example 6.

Example 8

WAN bandwidth

In Example 6, the total IP WAN bandwidth usage between each pair of sites was calculated, and expressed in Erlangs at the bottom of Table 26, IP WAN bandwidth usages (Erlangs) for Example 6, on page 135. Specifically, the total WAN bandwidth usage between Sites 1 and 2 is 24.0 Erlangs, between Sites 1 and 3 is 10.0 Erlangs, and between Sites 2 and 3 is 4.0 Erlangs. This implies that the average number of media streams simultaneously in use at any given time between Sites 1 and 2 is 24. Analogous statements can also be made regarding WAN traffic between each of the other two pairs of sites.

Every media stream across the IP WAN, between any pair of sites, is assumed to use the compressed G.729 codec, since bandwidth is relatively inexpensive within a private LAN, as opposed to a public WAN. Assume, for the sake of this example, a standard IP packet size of 20 ms. For the G.729 codec, Table 30, IP WAN bandwidth requirements for media streams, on page 138 indicates that each (unidirectional) media stream consumes 12.4 kbps of IP WAN bandwidth. Similar to the case in Example 7, 24 is the *average* number of simultaneous bidirectional media streams. As in Example 7, the bandwidth is sized to a "GOS" of P001 ("GOS" in this context is actually a pseudo-GOS; true GOS is associated with a fixed number of channels, as is typical of circuit-switched systems). The standard infinite-server queueing model implies that less than 0.1% of the time there is at least 40 simultaneous media streams between Sites 1 and 2. So, it is sufficient to engineer the WAN bandwidth between those two sites to support 39 simultaneous media streams. Therefore, the WAN between Sites 1 and 2 requires at least (39 simultaneous media streams) x (12.4 kbps per media stream) = 484 kbps of bandwidth. This result, along with the analogous results for the WAN traffic between the other two pairs of sites, are provided in Table 31, IP WAN bandwidth requirements in each direction, for Example 8, on page 139.

Table 31: IP WAN bandwidth requirements in each direction, for Example 8

Requirement	Between sites 1 and 2	Between sites 1 and 3	Between sites 2 and 3
Simultaneous media streams for "P001"	39	20	10
LAN bandwidth (kbps)	484	248	124

In Table 31, IP WAN bandwidth requirements in each direction, for Example 8, on page 139, the number of simultaneous media streams for "P001" represents the 99.9th percentile for the number of simultaneous streams, as determined by applying the standard infinite-server queueing model.

To this point, all the discussion regarding bandwidth relates only to bearer traffic (media streams). Network packet traffic that is related to signaling is very different from the bearer traffic because it tends to occur in bursts. For example, while the bearer traffic that is associated with a particular call tends to involve a constant, steady stream of packets throughout the duration of that call, the signaling traffic for that same call tends to occur in bursts during call setup and teardown.

The bandwidth that is required for signaling is generally negligible in comparison to the bandwidth that is required for bearer traffic. However, since Avaya products use Separation of Bearer and Signaling (SBS), the bearer traffic and signaling traffic use distinct paths. Therefore, signaling bandwidth must be given its due consideration, despite the fact that it is negligible in comparison to bearer bandwidth.

Signaling traffic is more prone to bursts than bearer traffic because the former consists of messages that are associated with call set-ups and tear-downs, as opposed to traffic that is uniformly distributed throughout entire call durations. However, the "bursty" effect is somewhat assuaged for larger call volumes. Although the precise bandwidth requirement for a given configuration depends on the nature of the endpoints involved, a reasonable approach is to allocate an overhead of 50 bits per second (bps) for each IP endpoint in the network, as well as the following (as applicable) for every 1000 calls:

- 11 kbps for messaging between the S8700 server and an IPSI circuit pack on a G600 MG
- 8 kbps for the H.248 link between a C-LAN circuit pack and a G700 MG

NOTE:
The 11 kbps and 8 kbps associated with 1000 calls should not be amortized to produce estimates for systems with very light traffic. For example, 11 bits per second and 8 bits per second will not support an individual call.

Physical resource placement

As a default, resources should be balanced as uniformly as possible. For example, if 11 Media Processors are required in a Network Region that has three PNs, two of the PNs should house four Media Processors each, and the other PN should house the final three Media Processors. Advanced users should be able to manually override the resource-placement defaults. For example, there might be reasons beyond traffic engineering for specifying an unbalanced system or an overengineered resource pool, such as reliability, cost, security, physical constraints, and so on.

Final checks and adjustments

The final step in the design process is to verify that the final configuration proposal meets the following criteria:

- All endpoints and Media Gateways have been assigned to various Network Regions, sites, and/or Communication Manager systems, according to customer specifications.
- The placement of resources adheres to the physical capacities of the proposed platform.
- The number of PNs and/or Media Gateways is sufficient to handle the TDM traffic, the required number of IPSI circuit packs, and the required number of port circuit packs.
- The number of C-LAN circuit packs is sufficient to support the desired number of IP endpoints, Media Gateways, and certain adjuncts.
- The number of media processing circuit packs is sufficient to handle both calls involving IP endpoints, and interport network calls between circuit-switched endpoints, unless a circuit-switched center stage is used instead of IP-Connect.
- The anticipated call volume can be handled by the server.
- There is sufficient bandwidth in all IP networks to support the anticipated media traffic.

Security

This chapter discusses the security design and features for Avaya Communication Manager, and how to operate Avaya systems securely.

> **NOTE:**
> Because this information is valuable both to those who want to protect the system and to those who seek to “hack” into those systems, the information in this section is deliberately incomplete. For example, we discuss the use of one-time passwords for user authentication, but not the mechanism of how this feature works.

Earlier systems did not interface with the data network and were neither susceptible to the types of attacks that are prevalent on those networks, nor provided a gateway into such networks from which an attack might be launched. With the convergence of voice (IP Telephony) and data over corporate enterprise networks, this is no longer true.

The main topics included in this chapter are:

Your security policy

System security does not begin with the system itself, but with the people and the organizations that operate or use the system. One of the most important tools for securing a system is to have a written, published, and enforceable *security policy.* Your security policy should clearly address these questions:

What are you trying to protect?

The security policy usually attempts to protect information, whether the information is in the form of data (files) or conversations (digitized voice packets). Customers should assess the value of those assets that require protection, and compare the true costs of security to the value of those assets.

What are you protecting it from?

Most often, criminals, who are also called “hackers,” pose a significant threat to secure information. However, do not forget to look internally. A significant number of attacks come from within an enterprise. Your security policy should include rules about behavior, the consequences of bad behavior, a path of escalation, and a person to contact with regard to security issues.

How likely is a threat against these assets?

Security is always a trade-off. The more security, the more inconvenience and the more cost. To avoid the necessary inconvenience, some users are likely to subvert the security policy. For example, if you make passwords so complex so that the passwords are difficult to remember, people will write the passwords down. Users prefer easy access without security. Having to log on is inconvenient. However, everyone must endure some level of inconvenience if the system is going to be secure against attacks. The security policy must define this level of inconvenience to ensure that the security polity is not circumvented. In addition, management must support the policy, and establish clear rules for its enforcement, including the consequences for violating it. A security policy that does not establish consequences for violations quickly becomes irrelevant.

Recommendations for your security policy

Avaya recommends that you continuously review your security policy, and keep up with new threats and to make improvements each time a weakness is found. To effectively support your security policy, your company must allocate long-term resources to the development, implementation, and reassessment of the policy.

Avaya Communication Manager and Media Servers

This section discusses Avaya’s security designs:

- Built-in Linux security features
- One-time passwords
- Shell access
- Root access
- Remote access
- Secure access
- Monitoring and alarming
- Data encryption

Built-in Linux security features

Proprietary vs. open operating systems

Open operating systems such as Linux or a version of Microsoft Windows are often thought to be less secure environments compared to proprietary systems. To some extent this is true, but it is important to understand why Oryx-Pecos, Avaya's proprietary operating system for its legacy products, is more secure than an open operating system because it does not support the types of network connections that converged voice and data network configurations demand. So why not enhance Oryx-Pecos? Aside from the economic reasons, there is a security paradox: to make an operating system secure, reveal its inner most secrets. When the operating system software is publicly available and implemented in varying environments for a wide range of applications, there are many more eyes looking for security holes. The expertise of the entire technical community is brought to bear on the problem. Of the major operating systems (Unix, Linux, Windows), one is not inherently more secure than another. Each has inherent security flaws. All can be made secure through the application of a good security policy, which includes proper administration and configuration, and diligent application of vendor updates when security problems are discovered.

The Linux environment has a security advantage because

- Problems can be identified both by testing (hacking) and by reviewing the source code itself.
- Security "holes" tend to be fixed more quickly compared to proprietary operating systems.

Avaya capitalizes on Linux' security advantage

The Avaya S8700 and S8300 Media Servers run under the Linux operating system that has two important security features:

- Built-in protection against certain types of Denial of Service (DOS) attack, such as SYN floods, ping floods, malformed packets, oversized packets, sequence number spoofing, ping/finger of death, etc. Attacks are recognized at the lower levels of the software and their effect is blunted. (It is not possible for a target system to always provide service during a DOS attack. Rather, the protection is to automatically resume service as soon as the attack is removed.)
- The Linux kernel is compiled with a set of options to precisely tailor its operation to maximize security consistent with required operation of the system. These include a number of built-in firewall and filtering options. All file and directory permissions are set to minimize access as much as possible consistent with proper system operation. The disk drives of the S8700 and the S8300 servers contain multiple partitions, each of which is restricted according to the type of data that it contains. All unneeded services are disabled either permanently or through administration for those services. Disabled services and capabilities include NFS, SMB, X-windows, rcp, rlogin, and rexec. The system administrator has additional control of which services are visible from the multiple Ethernet interfaces that are connected to the enterprise LAN. Other Ethernet interfaces are permanently configured to restrict services.

One-time passwords

Standard login accounts use static passwords that can be used multiple times to log in to a system. Anyone who can monitor the login messages can also capture passwords, and use the passwords to gain access. You can administer the S8700 and the S8300 servers for one-time passwords that have a fixed-user name but not a fixed password. In this case, users must supply a unique, one-time password for each session, and even if the password is compromised, it cannot be reused. When a system is covered by an Avaya service contract, all logins that are accessed by Avaya Services technicians are protected by one-time passwords.

Shell access

Access to a "shell" from which arbitrary commands can be executed is not granted by default to a login on an S8700 or an S8300 server. When a login is created, the system administrator can specify whether or not the account is permitted to have shell access. Accounts that are denied shell access can either log in to an Avaya Communication Manager administration screen or a Web page upon successful login. In both cases, the operations that these logins can perform are restricted. Generally, only people who perform hardware maintenance or software maintenance on the server need shell access permissions administered in their login accounts.

Root access

On a Linux system, the highest administrative-access level is called *root*. Direct logins to root-level accounts are not permitted on S8700 and S8300 servers. Administrative access, which requires root-level permissions, is handled through "proxy" programs that grant specific access to specific accounts. The ability to obtain full, root-level access is granted only in very special circumstances. By tightly restricting the root password, Avaya systems are less susceptible to accidental or malicious system access.

Remote access

Avaya S8700 and S8300 servers have a modem port for remote maintenance access, and for sending maintenance alarms calls. The server logins that establish this remote connection are separate from other logins that allow administrative functions. One login account can establish a connection, and once the link is established, a second login is necessary to administer the system. The dial-in line can also be restricted to:

- Disallow all incoming calls.
- Allow only one incoming call.
- Allow all incoming calls.

When the interface is set to "allow one incoming call only," the line is enabled to answer a single call. As soon as a call arrives, the line is disabled, and must be re-enabled through administration before another call will be accepted. This feature does not inhibit outgoing alarm calls, which are needed for maintenance. Normally, the line is disabled for all calls. When a maintenance activity is needed, the maintenance technician must contact the server administrator and request that the line be activated. The server administrator must then log in to the server, and enable the line for one call only. The maintenance technician then calls the server, performs the necessary maintenance, and disconnects. At this point the line is automatically disabled again. Enabling the data line for one call only is a good example of a feature that illustrates the trade-off that is required between security and convenience. Having the data line

disabled provides better security, but during diagnostic activity, when multiple calls must be made, the server administrator must be called to manually reenable the line for each call. In addition, Avaya employs Expert systems technology to contact systems automatically for monitoring and diagnostics. Disabling the data line disables this technology, which results in higher maintenance costs, and possibly longer times out of service when a failure does occur.

Secure access

Typical server access methods include telnet, Web browser (HTTP), and FTP for file transfers. Each of these mechanisms can support login authentication, but suffer a common weakness. The password that you type during login is sent in clear text, which allows someone with a network monitor/sniffer to capture the password and to gain access. These mechanisms also transmit all the session information in clear text. Some of this information might contain data such as account codes, authorization codes, or other data that might be useful to an attacker.

To overcome these problems, Avaya S8700 and S8300 servers support:

- Secure Shell Access (SSH) and Secure Copy (SCP). Provide an access mechanism for terminal access and file copy that encrypt the entire session, including the login sequence, and subsequent data transfer. ***SCP is the preferred method of transferring files.***
- Secure WEB access using the Secure Sockets Layer (SSL) with HTTPS. All Web access to an Avaya S8700 and S8300 servers is through a secure connection. Unencrypted Web access is not supported. The Avaya servers also support one-time-passwords for logins through these mechanisms, even though the exchange is already encrypted.
- FTP service that is disabled by default. Each time a file is to be transferred to the Avaya server, an administrator must log in and enable the FTP server. The file is then transferred using anonymous FTP, and the FTP server can then be disabled. Using anonymous FTP in this manner avoids the problem of sending passwords in clear text.

Monitoring and alarming

Avaya S8700/S8300 Media Servers support the following security monitoring and alarming features:

- Sessions are automatically disconnected after a period of inactivity.
- Accounts are automatically locked out for a period of time as a consequence of consecutive failed login attempts.
- Files and directories are monitored and audited by Tripwire, which maintains a cryptographically encoded signature of the files on the system, and generates alarms if any changes occur.
- All login sessions, whether successful or not, are logged.
- User activity logging.
- Security events are alarmable and reported by sending an SNMP trap to one or more destinations.

Data encryption

Attacks against a system are not limited to attempts to find holes in the access structure. Avaya S8700 and S8300 servers store backup copies of critical configuration information, including authentication and account information, on external systems. If this information is stored in clear text, and the file server on which it is stored is compromised, the servers also can be compromised. S8700 and S8300 servers can encrypt all backup data, and thus make use of the data impossible, even if access to it is possible. The user is responsible for remembering the encryption key, because Avaya cannot assist you if you forget it. Avaya also cryptographically signs all new software or firmware media to prevent malicious modification in transit. If the system detects a modification, the installation is aborted.

LAN isolation configurations

S8700 with Avaya MCC1 or SCC1 Media Gateways

An Avaya S8700 Media Server contains multiple Ethernet Network Interfaces (NICs):

- Each Avaya S8700 Media Server with Avaya MCC1 or SCC1 Media Gateway has five Ethernet interfaces (NICs), each dedicated to these specific functions:
  - The two control LANs are only used to connect between the servers and the port networks (PNs). These two LANs must be private LANs, and carry no other traffic.
  - The duplication interface is a point-to-point LAN that is only used to send information between the two servers.
  - The laptop computer interface is a point-to-point LAN that is used only for local administration and carries no other type of traffic.
  - The enterprise LAN is used for administration and time synchronization. Telephony traffic does not use this LAN. However, in this case, it is possible to subvert this security measure by interconnecting the enterprise LAN NIC with one of the other LANs shown.
- PNs contain additional Ethernet interfaces.

Figure 58, Avaya S8700 Media Server with an Avaya MCC1 or an SCC1 Media Gateway, on page 147 shows the different LANs that are possible on an S8700 server that is configured with Avaya MCC1 or SCC1 Media Gateways along with some of the common adjuncts. The enterprise LAN, adjunct LANs, and agent's LAN can all be connected together to form one network. Or these LANs can be kept physically separate for either traffic reasons or security reasons.

Figure 58: Avaya S8700 Media Server with an Avaya MCC1 or an SCC1 Media Gateway

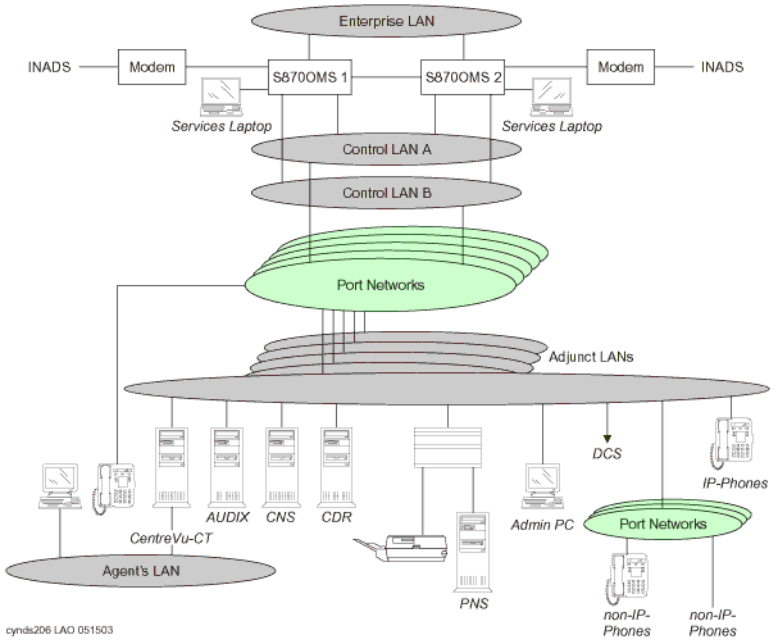

To provide the most secure environment that is possible for the system, network access should be divided into separate zones of control. These zones are sometimes referred to as *DMZs*.

- One VLAN can be administered for administrative traffic, one for call signaling, another for voice bearer traffic, and so on.
- Layer 3 boundary devices (routers, layer 3 switches, and firewalls) should be administered to enforce the corporate security policy on traffic that is destined for the Avaya S8700 Media Server, its Avaya MCC1 or SCC1 Media Gateways, or adjuncts.
- Packet filters can permit administrative access only from an administrator's PC and to deny access from the Avaya S8700 Media Server or its gateways to the corporate LAN while allowing call signaling and bearer traffic from all IP Telephones appropriate access.

Figure 59: Isolated LANs (Avaya S8700 Media Server with an MCC1 or an SCC1 Media Gateway)

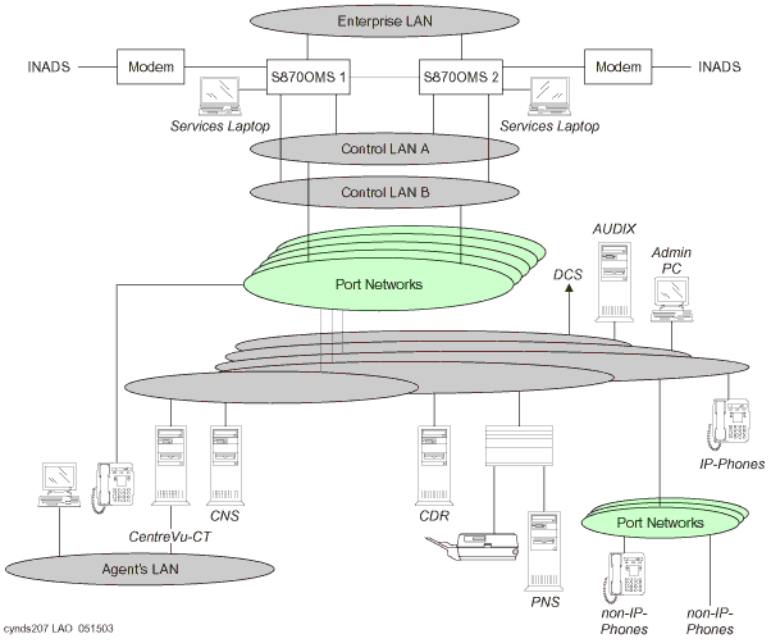

Figure 59, Isolated LANs (Avaya S8700 Media Server with an MCC1 or an SCC1 Media Gateway), on page 148 shows how Communication Manager can be configured to allow only certain types of access to specific LAN interfaces on its PNs. For example, even if you connected an administration terminal to one of the other LANs, you cannot get administration access.

S8700 with Avaya G600 Media Gateways

The Avaya S8700 Media Server with an Avaya G600 Media Gateway also have five interfaces each (Figure 60, Isolated LANs (Avaya S8700 Media Server with a G600 Media Gateway), on page 149):

- The enterprise LAN and control LANs are connected together
- There is only one control LAN.
- There are two spare NICs that are not used.

The messages between the Avaya S8700 Media Server and the Avaya G600 Media Gateways are encrypted.

Figure 60: Isolated LANs (Avaya S8700 Media Server with a G600 Media Gateway)

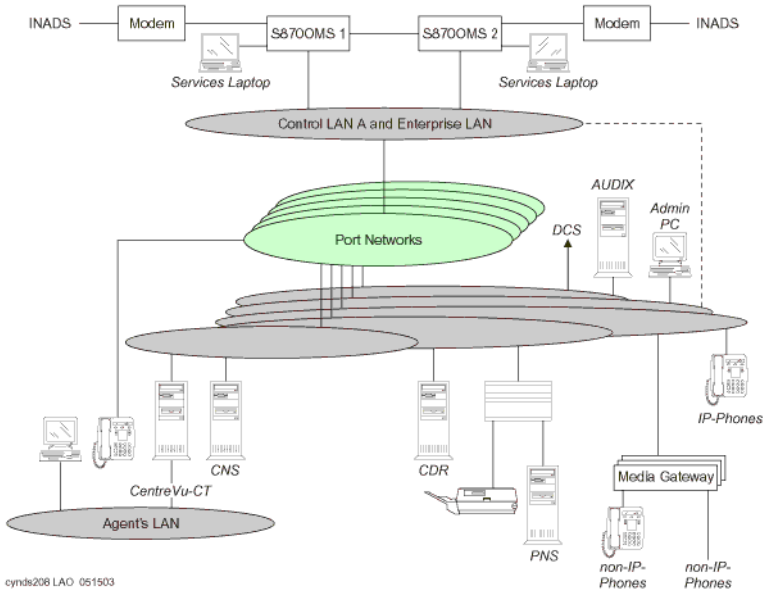

Virus and worm protection

Viruses and worms are most often targeted at Microsoft Windows operating systems or such commonly used applications as IIS, Exchange, Outlook, or Word. Because the Avaya S8700 and S8300 servers are Linux-based and do not interface with these Microsoft products, they have some degree of natural immunity. In addition, viruses and worms are most commonly delivered by e-mail, by visiting infected Web sites, or by sharing disk drives. The Avaya S8700 and S8300 servers do not

- Support incoming email and, therefore, do not forward e-mail
- Contain the Internet Explorer Web browser
- Share drives

All file transfers to the S8700 and the S8300 servers are restricted. software, and patchfiles are cryptographically signed to prevent introduction of unwanted software. In addition to this natural immunity, the files and file systems of the S8700 and the S8300 servers are monitored by Tripwire.

Testing

During the development of the Avaya S8700 and S8300 servers, or in production of upgrades to its software, Avaya subjects the system to a variety of common “attack tools” to find any overlooked or accidentally created security holes. The exact set of tools that are used varies to keep up with the technology. Common tools include nmap and nessus. Security problems found by these efforts are corrected before the product or the update is released.

Environment

Avaya S8700 and S8300 servers are as secure as reasonably possible, consistent with the operational needs of the product and business in which they are used. Security, however, does not end with the servers. These servers are connected to one or more networks that are, in turn, connected to other equipment in the enterprise.

Recommendations for network security

Avaya recommends that these servers be located behind a firewall. Where this firewall is located with respect to other LAN components must be designed on a case-by-case basis. Avaya Professional Services can assist owners in configuring their networks for both security and optimal IP Telephony operation. Other vendors also specialize in this type of consulting. Owners are advised to seek assistance if internal staff is not trained in these areas. Security holes that arise from negligence, ignorance, or oversight or the pressures of schedule or budget are all equally usable by hackers. Malicious activity is a moving target, and what is safe today might not be safe tomorrow. Avaya is committed to providing appropriate secure solutions for its products, and to continuously monitoring evolving security threats. Avaya S8700 and S8300 servers are appropriately secure against the known threats. Avaya responds quickly should new threats appear. Consult these resources for the latest security information:

- Your Avaya account team
- The Avaya support Web site:

  http://support.avaya.com/

  Click **Security Advisory** in the Technical Database list on the left side of the page.

IP Telephony circuit pack security

Avaya circuit packs such as those in the G600 have a variety of security measures that combine both voice and data security strategies in to a secure package.

The G600 uses three different Ethernet interfaces to help isolate the traffic, and protect the specific interfaces that must be secured:

- TN2312 IP Server Interface (IPSI)
- TN2302 Media Processor (MedPro)
- TN799 Control LAN (C-LAN)

TN2312 IP Server Interface (IPSI)

Topics in this section include

- Telnet
- FTP
- DHCP
- Control link

Telnet

A telnet service is currently required on the IPSI for manual administration of the IPSI (IP address, default gateway address, VLAN ID, QoS, and Ethernet settings). Telnet access to the IPSI circuit pack is through

- Standard TCP port 23, but only for connections that are physically made through its secondary services Ethernet port. When established, these Telnet accesses are directed to a command menu that supports a variety of administration tasks.
- The OS-debugging shell over port 2312. This port is always available when accessed through the local services port. Telnet access to this port on the control link is opened by a Communication Manager command, and disabled immediately after a session has closed the connection or after 5 minutes of inactivity (see also Control link).

FTP

An FTP service exists, but is disabled by default. Communication Manager must enable the FTP service, and only does so for firmware downloads. Once the FTP service is started, Communication Manager initiates the client-side of the FTP protocol, and then transfers a new firmware file to the IPSI. Once the transfer is complete, the FTP service is automatically disabled. A 5-minute time-out is enforced to guard against cases where the firmware download is started but terminated prematurely. When time-out occurs, the FTP service is disabled until a new command from Communication Manager enables it again.

DHCP

In S8700 Multi-Connect systems only, the IPSI has the ability to receive its IP address information from the S8700 server through DHCP. This DHCP service only runs on the control network, and does not connect to a customer's LAN. Avaya has also implemented mechanisms for restricting this DHCP service, so that non-IPSIs do not receive an IP address and IPSIs do not receive an address from a non-S8700 server.

Control link

In order to communicate with the S8700 server, the IPSI establishes a control link. This link is encrypted through Triple-DES (3DES) by default, although AES is also available. The control link is not open for communication to or from any other entity than the S8700 server.

TN2302 Media Processor (MedPro)

The TN2302 circuit pack is the interface to the audio gateway portion of IP Telephony. The circuit pack:

- Uses an isolated/proprietary operating system, so it is not susceptible to known viruses.
- Runs independently of administrator traffic in order to maintain an isolated security domain, protecting against attacks that exploit trusted relationships.
- Establishes audio connections and only responds to a connection when a corresponding signaling connection is established.
- Successfully survives some Denial of Service (DoS) attacks, including SynFlood, and is very resilient to flood-based attacks.

Because of the proprietary operating system, limited number of open ports, and reliance on UDP sessions, the TN2302 is very secure, and is difficult to take out of service. Regardless, the TN2302 is completely independent of the administration, maintenance, or reliability of the Avaya Media Gateways, so it cannot be used a “jumping point” to the Media Gateways.

TN799 Control LAN (C-LAN)

The C-LAN circuit pack interface not only supports signaling for IP Telephony applications, but also supports asynchronous links to INTUIY AUDIX, Call Management System (CMS), and other adjuncts. This interface

- Is independent of the Media Gateway.
- Has no IP link back to the central administration or maintenance processes of Communication Manager.
- Successfully survives DoS attacks created by the SynFlood tools.
- Maintains the IP endpoint RAS authentication sequence, a safeguard against exploiting toll services through IP endpoints.

For more information on the security of Avaya circuit packs, see:

http://support.avaya.com/elmodocs2/multivantage/95933.pdf

Toll fraud

This section contains information about Avaya’s design for preventing toll fraud, and includes these topics:

Avaya’s security design

Telecommunications systems face significant and growing problems of theft of customer services. Toll fraud, the unauthorized use of a system and its facilities by a third party, can result in substantial additional charges for telecommunications services.

Avaya makes every effort to assist customers in their battle against “hackers” through the technology that goes into every Avaya product. Avaya Communication Manager is designed with security in mind, and offers many features and capabilities to help maintain security and prevent toll fraud:

- Your company completely controls its communication facilities.
- Your company completely controls its communication’s security policy and features.
- Your company can make immediate changes at any time.

Each new release of Communication Manager addresses customer needs for even greater security capabilities, including enhancements to support the recent changes in the North American Numbering Plan.

Hacking methods

Hackers often facilitate toll fraud activity by gaining access to:

- A system's administration or maintenance port by randomly dialing thousands of telephone numbers, and then attempt to log in using default passwords. Statistical sampling indicates there is a high likelihood that customers still have one or more default passwords in place on their telecommunications system. This allows hackers to completely modify the system to allow toll fraud activity.
- A system's remote access port, and then use the remote access feature.
- A voice messaging system, and then transfer their calls to outgoing facilities.

To aid in combating these crimes, Avaya continuously works with its customers and supporting law enforcement officials to apprehend and prosecute those criminals.

Your toll fraud responsibilities

No telecommunications system can be entirely free from risk of unauthorized use. But diligent attention to system management and security can reduce that risk considerably. Often a trade-off is required between reduced risk and flexibility. The user and the administrator of the system are in the best position to determine how to tailor the system to meet their mutual needs, while protecting the system from unauthorized use. Under applicable law, customers are responsible for any toll fraud charges that occur. Because you have ultimate control over the configuration and use of the Avaya products and services that you purchase, your company bears the responsibility for fraudulent uses of those products and services. Not only can the financial loss from these calls be substantial, but operational impacts such as reduced productivity can also have an adverse effect.

Toll fraud indemnification

As part of Avaya's ongoing efforts to combat communications fraud and its threat to our business customers, Avaya has introduced an enhancement to its Service Agreement. Beginning January 1, 1996, Avaya indemnifies its customers for charges associated with fraud. This indemnification is available to all customers who are covered by warranty and/or maintain an Avaya Service Agreement for Avaya Communication Manager, INTUITY AUDIX voice messaging, and Avaya Interactive Voice Response systems.

The indemnification enhancement is offered at no additional cost to your service agreement during warranty, or as part of a multiyear Avaya Service Agreement. The only requirement is to follow and maintain the sound security practices that every business should implement. A complete list of these security practices can be obtained from your Avaya Account Team.

Additional toll fraud resources

In an effort to assist customers, Avaya has developed a variety of service offerings and provides materials to assist in helping to identify and combat toll fraud. These offerings and materials include:

- Security Audit Service
- Security Tune-up Service
- Toll Fraud Intervention Hotline
- Avaya Security Handbook

Security Audit Service

The Avaya Security Audit Service is a fee-based, consultation service that provides a security evaluation of a customer's telecommunications system. The Security Audit is conducted by a Avaya team of experts and includes:

1. Preliminary telephone interview
2. On-site or remote security audit of the equipment
3. Analysis of system vulnerability
4. Written recommendations for increasing security

Security Tune-up Service

The Security Tune-up Service is a fee-based, consultative service designed to provide an expedient, online review of the toll-fraud potential in your system. This service is provided for ACM systems and voice messaging systems. Customer support engineers who specialize in security:

1. Remotely access your system.
2. Analyze the potential risks in the system.
3. Optionally implement agreed-upon changes to secure the system.

Toll Fraud Intervention Hotline

If you suspect you are being victimized by toll fraud or theft of services and need technical support or assistance, call the Avaya Toll Fraud Intervention group toll free at

1-800-643-2353 (24 hours a day/7 days a week)

- Consultation charges may apply.
- There is no charge for intervention services performed on equipment that is covered by warranty or service agreement.

Avaya Security Handbook

The *Avaya Security Handbook* summarizes the principal steps that a system administrator can take to reduce the risk of toll fraud. This handbook complements specific documentation for Avaya products and provides a system administrator with a complete, detailed reference for planning and implementing security measures.

Voice quality network requirements

There are a number of network parameters that affect voice quality. This chapter lists some of the more important ones. The concept of quality has different meanings to different people. IP Telephony quality can be engineered to several different levels to accommodate differing business needs. A small company might choose to implement IP Telephony with very good sound instead of buying newer networking equipment to support excellent voice sound. A large call center company might want excellent voice sound as part of its corporate strategy. Avaya therefore presents options in network requirements to allow the customer to choose which "quality" level best suits their specific business needs.

Delay

Packet delay is the length of time it takes a packet to traverse the network. Each element of the network adds to packet delay including switches, routers, distance traveled through the network, firewalls, and jitter buffers (such as those built into H.323 audio applications like the Avaya IP SoftPhone™ or Microsoft NetMeeting). Router delay depends not only on hardware, but also on configurations such as access lists, queuing methods, and transmission modes. Delay (latency) can have a noticeable affect but can be controlled somewhat in a private environment (LAN/WAN) because the company or enterprise manages the network infrastructure or SLA. When using the public network, there are inherent delays that one cannot control.

Network delay

The following are suggested guidelines for one-way network delay between endpoints, meaning LAN/WAN measurements not including IP phones:

- 80ms (milliseconds) delay or less can, but may not, yield best quality.
- 80ms to 180ms delay can give business communication quality. This is much better than cell-phone quality and in fact is very well suited for the majority of businesses.
- Delays exceeding 180ms may still be quite acceptable depending on customer expectations, analog trunks used, codec type, etc.

Again, there is a trade-off between voice quality and the technical and monetary constraints with which businesses confront daily.

The ITU-T has recommended 150ms one-way delay (including endpoints) as the limit for "excellent" voice quality. This value is largely misinterpreted as the only range to calculate a network delay budget for IP Telephones. One-way delays in excess of 250ms can cause the well-known problem of "talk-over," when each person starts to talk because the delay prevents them from realizing that the other person has already started talking. Certainly long WAN transports must be considered as a major contributor to the network delay budget, as one major WAN service provider averaged 75ms delay from Los Angeles to New York. Los Angeles to Paris was found to be about 145ms. Some WAN service providers can lower delay in their network if it is negotiated and recorded as part of the companies SLA (Service Level Agreement). Even so, staying within 150ms (end to end) may not be possible.

Finally, end-to-end delay over 400ms between port networks and the S8700 Media Server can cause port network instability. A network assessment is highly recommended to measure latency (also jitter and packet loss) and make sure all values are within bounds before implementing IP Telephony.

Codec delay

Some delay may be added by various codec algorithms. G.729, for example, adds approximately 10 ms of delay compared to G.711 when configured for the same sample size and jitter buffer size. In addition, the compression algorithm in G.723.1 uses multiples of three 10 ms samples per packet, which results in increased latency over codecs configured to use two 10 ms samples per packet.

Jitter

Jitter is a measure of variance in the time it takes for communications to traverse from the sender (application) to the receiver, as seen from the application layer (from RFC_2729 Taxonomy of Communication Requirements for Large-scale Multicast Applications). Jitter is thought of as the statistical average variance in delivery time between packets or datagrams. Jitter can create audible voice-quality problems if the variation is greater than 20ms (assuming an existing 20ms packet size).

To compensate for network jitter, many vendors implement a jitter buffer in their H.323 endpoints. The purpose of the jitter buffer is to hold incoming packets for a specified period of time before forwarding them to the decompression process. A jitter buffer is designed to smooth packet flow. In doing so, it can also add packet delay Depending on the size of the jitter buffer, excessive jitter will be experienced either as packet loss (if the jitter exceeds the jitter buffer) or delay (if the jitter is less than or equal to the buffer size).

Jitter buffers should be dynamic to give the best quality, or if static, should generally be sized to twice the largest statistical variance between packets. It is not enough to just select the right size of jitter buffer, one must also pair an appropriate debuffering algorithm type with the jitter buffer.

The network topology can also affect jitter. Multiple paths between endpoints with load balancing enabled can contribute significant amounts of jitter.

These Avaya products all have dynamic jitter buffers to minimize delay by reducing the jitter buffer size as the network allows:

- Avaya G600 and Avaya G700 Media Gateways
- Avaya S8100 Media Server with G600 Media Gateway
- Avaya IP SoftPhone software
- Avaya 4600 Series IP Telephone

NOTE:
This feature can exacerbate problems in an uncontrolled network. Many good tools are commercially available to measure jitter, delay, and packet loss to help monitor and bring control to the network.

Packet loss

Packet loss occurs when packets are sent, but not received at the final destination due to some network problem. Qualifying problems caused by occasional packet loss are difficult to detect because each codec has its own packet loss concealment method. Therefore, it is possible that voice quality would be better using a compression codec (G.729A) compared to a full bandwidth G.711 codec. Several factors make packet loss requirements somewhat variable, such as the following:

- Packet loss requirements are tighter for tones (other than DTMF) than for voice. The ear is less able to detect packet loss during speech (variable-pitch), than during a tone (consistent pitch).
- Packet loss requirements are tighter for short, continuous packet loss than for random packet loss over time. Losing ten contiguous packets is worse than losing ten packets evenly spaced over an hour time span.
- Packet loss may be more noticeable for larger voice payloads than for smaller ones, because more voice is lost in a larger payload.
- Packet loss may be more tolerable for one codec over another.
- Even small amounts of packet loss can greatly affect TTY (TDD) device's ability to work properly.
- Packet loss for TCP signaling traffic increases network traffic substantially when loss is greater than 3% loss due to retransmissions.
- Packet loss perception is non-linear as shown in Figure 61, Packet loss perception, on page 157. As packet loss increases, the perception of impairment increases at a slower rate.

Figure 61: Packet loss perception

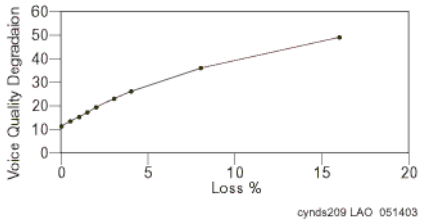

Network packet loss

The maximum loss of packets (or frames) between endpoints should be:

- 1% or less can yield best quality depending on many factors.
- 3% or less should give Business communications quality. Again, this quality is much better than cell-phone quality.
- More than 3% may be acceptable for voice but may interfere with signaling. More information on signaling bandwidth requirements can be found at this URL:

  http://www1.avaya.com/enterprise/whitepapers/S8700g600configurationguide.pdf

  Look for *Network Requirements and Configuration Guidelines for Communication Manager on S87000/G600 Issue 1.1*

Like delay, Avaya allows customers a tiered approach of packet loss to balance new network costs and limitations with business directives. Tools such as the Agilent (HP) Internet Advisor, Finisar's Surveyor Explorer, Radcom's Prism, NAI's Sniffer, and others measure packet loss. Remember that too much delay or packet mis-order can be perceived as dropped packets, and it may appear that the network is losing packets when in fact they have been discarded intentionally.

Packet loss concealment (PLC)

IP networks are characterized by unintentional packet loss. Some packet loss can be dealt with attempting to conceal the loss by generating packets to take the place of the missing packets. ITU standards G.711 Annex I and the G.729 standard define methods by which packet loss concealment can be provided. Excessive packet loss cannot be disguised and so ultimately PLC gives way to CNG if too many packets are lost in succession.

Ramping down to silence is the standard way in which PLC is performed. If just one or two consecutive packets are lost, then this strategy quite successfully covers up the break in audio. 6 is accepted as a maximum number of packets over which PLC can be sensibly applied. Any more packet loss than this should result in the activation of comfort noise generation (CNG).

The principle of independence requires that intentional packet loss is exactly recognized as such without relying upon specific behavior from other VoIP systems in the call. Intentional packet loss must be dealt with through CNG, while unintentional packet loss uses PLC to ramp down to silence (or CNG) as this kind of packet loss creates discontinuities in the conversation. Codecs that do not support PLC within the standard should use a Voice Activity Detector (VAD) to detect when a silence period is being entered. This should be communicated to the PLC layer.

Echo

The two main types of echo are acoustic and impedance, although the sources of echo can be many. Echo will result when an IP Telephony call leaves the LAN through a poorly administered analog trunk into the PSTN. Another major cause is from an impedance mismatch between four-wire and two wire systems. Echo also results when an impedance mismatch exists in the conversion between a headset and its adapter. Impedance mismatch causes inefficient energy transfer. The energy imbalance must go somewhere and so it is reflected back in the form of an echo. Usually the speaker hears an echo but the receiver does not. Echo cancellers, which have varying amounts of memory, compare the received voice with the current voice patterns. If the patterns match, the canceller cancels the echo. Echo cancellers aren't perfect, however. There will be a residual level left even in optimal operating conditions. The problem is exacerbated in VoIP systems. If the one-way trip delay between endpoints is larger than the echo canceller memory, the echo canceller won't ever find a pattern to cancel. Also, people's perception of echo increases with delay. Most echo received by the ear in under 30 ms is ignored, but echo received after 30 ms is perceived as an annoyance. Avaya's G600, G700, Avaya TN2302 Media Processor, Avaya IP SoftPhone software and Avaya 4600 Series IP Telephone all incorporate echo cancellers designed for IP Telephony to improve voice quality.

Codec discussion

Depending upon the bandwidth availability and acceptable voice quality, it might be worthwhile to select a codec that produces compressed audio:

- A G.711 codec produces audio uncompressed at 64 kbps
- A G.729 codec produces audio compressed at 8 kbps
- A G.723 codec produces audio compressed at approximately 6 kbps

Table 32, Comparison of speech coding standards (without IP / UDP / RTP overhead), on page 159 provides comparisons of several voice quality considerations associated with some of the codecs supported by Avaya products.

NOTE:
Toll-quality voice must achieve a MOS (Mean Opinion Score) of 4 or above. The MOS scoring is a long-standing, subjective method of measuring voice quality.

Table 32: Comparison of speech coding standards (without IP / UDP / RTP overhead)

Standard	Coding Type	Bit Rate (kbps)	MOS
G.711	PCM	64	4.3
G.729	CS-ACELP	8	4.0
G.723.1	ACELPMP-MLQ	6.3 5.3	3.8

Because it does not use compression (and thus loses some source information), G.711 offers the highest fidelity audio. Unfortunately, there is a tradeoff with higher bandwidth usage. G.729 offers a good compromise with lower bandwidth usage, but still good fidelity audio. It is generally used where bandwidth is scarce, such as on a WAN link.

Compressed codecs use twice as many DSP (digital signal processor) resources than uncompressed (G.711) codecs. Each call over a compressed codec uses two DSP resources – one for sampling and one for the compression. On the TN2302AP (MedPro) circuit pack there are 64 DSP resources. Thus, the number of calls supported by one MedPro is

- 64 G.711 calls
- 32 compressed calls (for example, G.729)
- Some number in between for a call mix.

The formula for calculating the number of calls one MedPro supports is

$$(\text{Number of uncompressed calls}) + 2 \times (\text{Number of compressed calls}) \leq 64$$

Generally, G.711 is used within LANs because bandwidth is abundant and inexpensive whereas G.729 is used across WAN links because of the bandwidth savings and good performing voice quality.

Transcoding

Transcoding is a voice signal converted from analog to digital or digital to analog (possibly with or without compression and decompression). If calls are routed using multiple voice codecs, as can be the case of call coverage on an branch office system back to a centralized voice mail system, the calls may experience multiple transcodings (this could include G.729 across the WAN and G.723 into the voice mailbox). Each transcoding episode results in some degradation of voice quality. These problems may be minimized by the use of the Communication Manager feature called DCS with Rerouting (Path Replacement). This feature detects that the call coming through the main ECS has been routed from one tandem ECS, through the main, and back out to a third switch. In these cases, the system then re-routes the call directly, thus replacing the path through the main system with a more direct connection. Avaya products minimize transcoding while non-Avaya products may cause slight to excessive transcoding. Shuffling and Hairpinning also reduce transcoding.

Silence suppression/VAD

Voice Activity Detection (VAD) / silence suppression is another way to save bandwidth. During a conversation, because only one party is speaking at any given time, more than 50% of the transmission is silence. Voice Activity Detection (VAD) monitors the received signal for voice activity. When no activity is detected for the configured period of time, the Avaya software informs the Packet Voice Protocol. This prevents the encoder output from being transported across the network when there is silence, resulting in additional bandwidth savings. The Avaya software also measures the idle noise characteristics of the telephony interface. It reports this information to the Packet Voice Protocol to relay this information to the remote end for comfort noise generation when no voice is present. The trade-off with silence suppression lies with the silence detection algorithm. If it is too aggressive, the beginnings and ends of words can be “clipped;” if not aggressive enough, bandwidth is wasted. Silence suppression is built-in to G.729B. It can be enabled for other codecs from within Communication Manager. Because of issues with clipping, silence suppression is generally not used (with the exception of G.729B). These Avaya products employ silence suppression to preserve vital bandwidth:

- Avaya Communication Manager
- Avaya 4600 series IP Telephone
- Avaya IP SoftPhone

For procedures to administer QoS parameters, refer to *Administration for Network Connectivity* (555-233-504).

Network management

Network management is the practice of using specialized software tools to monitor and maintain network components. Proper network management is a key component to the high availability of data networks. There are two basic network management models:

- Distributed. Specialized, nonintegrated tools (and sometimes organizations) to manage discrete components
- Centralized. Integrating network management tools and organizations for a more coherent management strategy.

This chapter outlines Avaya's network management products, and some common third-party tools. It also discusses the distributed and centralized management models, and describes how Avaya management products fit into those models. The links in Table 33, Network management topics, on page 161 take you to the corresponding topic or product description.

Table 33: Network management topics

Link to section	Link to product category	Link to product
Avaya network management products (Integrated Management)	System management products	Avaya Communication Manager Configuration Manager
		Avaya Communication Manager Fault and Performance Manager
		Avaya Communication Manager Proxy Agent
		Avaya Site Administrator
	Network management products	MultiService Network Manager
		MultiService SMON Manager
		VoIP Monitoring Manager
	Directory management products	Directory Enabled Management
		Avaya Terminal Configuration
Third-party network management products		Multi Router Traffic Grapher
		HP OpenView Network Node Manager
Management models		Distributed (component)
		Centralized (hybrid)

Avaya network management products (Integrated Management)

Avaya offers several network management products. These products are categorized as:

- System management products
- Network management products
- Directory management products

System management products

Avaya's system management products include:

- Avaya Communication Manager Configuration Manager
- Avaya Communication Manager Fault and Performance Manager
- Avaya Communication Manager Proxy Agent
- Avaya Site Administrator

Avaya Communication Manager Configuration Manager

Avaya Communication Manager Configuration Manager is a new administration product that runs on a Linux server. The Communication Manager Configuration Manager provides Web-based configuration and administration of multiple Avaya Media Services.

The intuitive wizards allow you to quickly perform complex administrative tasks including:

- Add, move, and change individual stations and execute bulk moves
- Schedule routine tasks, multiyear maintenance, and system downloads
- Import and export data
- Enhanced number portability (ENP)
- Make global changes for most system objects

The Communication Manager Configuration Manager provides a single point of entry to distributed networks and campus environments. You can launch the Avaya Site Administration and the Avaya Terminal Emulator from the Communication Manager Configuration Manager access the terminal emulation feature for direct cut-through to supported systems.

Avaya Communication Manager Fault and Performance Manager

Avaya Communication Manager Fault and Performance Manager is the next-generation product for fault and performance management. Avaya Communication Manager Fault and Performance Manager:

- Runs on a Linux server and provides Web-based access from a universal client
- Integrates with your HP OpenView running on either a Windows 2002 server or Sailors 8 server. (HP OpenView is not included with this product.)
- Integrates with a supported Network Management System (NMS) to show the hierarchical view of devices and their status.
- Collects and stores data from the network devices, and lets you generate reports and view that data in a text or table format on their screens and print reports.
- Uses the Simple Network Management Protocol (SNMP) to communicate with the managed Avaya media servers and other supported devices.

Avaya Communication Manager Proxy Agent

Avaya Communication Manager Proxy Agent provides the interface between the Communication Manager Fault and Performance Manager and the Avaya Media Servers that run Avaya Communication Manager. Avaya Communication Manager Proxy Agent:

- Runs on Linux, and acts as a protocol converter between the proprietary OSSI protocol and Simple Network Management Protocol (SNMP)
- Sends and receives alarm traps, and can filter alarms by system, type, day, and hours.
- Has a command line interface directly accesses Avaya media servers and non-IP systems.
- Is required for customers who want to receive Communication Manager alarms through SNMP.

Avaya Site Administrator

Avaya Site Administration is a GUI-based administration product that runs on Windows 2000. Avaya Site Administration supports:

- Non-IP systems
- Avaya media servers
- Avaya Communication Manager
- AUDIX messaging systems

You must launch the Avaya Site Administration from the Communication Manager Configuration Manager to use do the following:

- Move, add, and change information on stations and perform basic traffic analysis easily
- Access the terminal emulation feature for direct cut-through to supported systems
- Execute the “Find and Replace” and “Import and Export” features to manage subscriber data

Network management products

Avaya's network management products include:

- MultiService Network Manager
- MultiService SMON Manager
- VoIP Monitoring Manager

MultiService Network Manager

Avaya MultiService Network Manager can operate as a stand-alone Network Management System (NMS) or integrate with your HP OpenView running on either Sailors 8 or Windows 2000 server. (HP OpenView is not included with this product.) Avaya MultiService Network Manager:

- Supports data and voice devices over the network.
- Supports the native SNMP agent on Avaya Media server S8700 and S8300, and supports a system view of the converged network.
- Detects IP devices and displays them on a network map.
- Provides fault monitoring and logs SNMP traps for IP devices.
- Is the central launch point for these products:
  - Avaya Communication Manager Configuration Manager
  - Avaya Communication Manager Fault and Performance Manager
  - Avaya Site Administration
  - Avaya Terminal Emulator
  - Avaya Voice Announcement Over LAN Manager
  - Avaya ATM WAN Survivable Processor Manager

MultiService SMON Manager

Avaya MultiService SMON Manager is a monitoring product that operates with the Avaya MultiService Network Manager. A license key is required to activate the SMON Manager, which:

- Provides a complete view of all switched traffic in the network by using embedded agents and leveraging special hardware capabilities.
- Monitors the network and displays a top-down view of all traffic that is traversing the entire network of switches.

VoIP Monitoring Manager

Avaya VoIP Monitoring Manager is a client/server Voice over IP (VoIP) monitoring application that tracks the quality of voice transmissions over the network. Avaya VoIP Monitoring Manager product runs on Windows 2000, and offers these features:

- Receives quality of service (QoS) statistics from Avaya IP endpoints and displays this data in graphs and reports.
- Isolates voice quality problems and send traps to any network management system (NMS) when poor voice quality is detected. This product supports converged data and voice products.

You can launch the Avaya VoIP Monitoring Manager from the Avaya MultiService Network Manager or operate it as a standalone application.

Directory management products

Avaya's directory management products include:

- Directory Enabled Management
- Avaya Terminal Configuration

Directory Enabled Management

Avaya Directory Enabled Management is a Web-based server component that operates with the Microsoft Internet Explorer browser and runs on Windows 2000. Avaya Directory Enabled Management:

- Uses the Lightweight Directory Access Protocol (LDAP) to provide real-time, directory-based read/write access to Avaya Media Servers, Avaya Communication Manager, and INTUITY AUDIX messaging service.
- Synchronizes and manages the station data and subscriber data based on events to add, delete, and modify data. The rules engine manages the events that take place between the servers and software applications.

Avaya Terminal Configuration

Avaya Terminal Configuration is a Web-based client component that operates on both Microsoft Internet Explorer and Netscape Navigator browsers, and runs on Windows 2000 server. The Avaya Terminal Configuration runs on top of the Avaya Directory Enabled Management software, and allows you to:

- Set up preferences on your personal telephone station, and print the button labels for your telephone set.
- Control permissions to access a limited set of features.

Third-party network management products

This section describes some third-party monitoring tools that might provide benefit to companies implementing IP Telephony. Avaya is not involved with the development of these products. Inclusion on this list is not exhaustive, nor does it represent an endorsement from Avaya. Products are listed here as a convenience for our customers:

- Multi Router Traffic Grapher
- HP OpenView Network Node Manager

Multi Router Traffic Grapher

The Multi Router Traffic Grapher (MRTG) monitors the traffic load on network links, and generates HTML pages of graphic-displayed images that provide a live visual representation of this traffic. MRTG is based on Perl and C, and works under UNIX and Windows NT. The Multi Router Traffic Grapher:

- Uses SNMP to read the traffic counters of your routers, logs the traffic data, and creates graphs that represent the traffic on the monitored network connection. These graphs are embedded into Web pages. MRTG even allows you to accumulate two or more data sources into a single graph.
- Creates visual representations of the traffic seen daily, during the last week, the last month, and the last year. MRTG performs well enough to monitor 200 or more network links from any reasonably-performing PC.
- Monitors any SNMP variable that you choose. You can even use an external program to gather the data that MRTG should monitor, for example:
  - System load
  - Login sessions
  - Modem availability

For more MRTG information, see:

- http://www.mrtg.org for the main MRTG Web site. Their product is available free of charge under the terms of the GPL.
- http://www.ee.ethz.ch/stats/mrtg/ for an example.

HP OpenView Network Node Manager

HP OpenView Network Node Manager (NNM) 6.4 and Network Node Manager Extended Topology 2.0 together provide your management team with the capabilities that you need to address your key business and network challenges:

- A new approach to root-cause analysis that includes a set of easy-to-use tools to help you identify and resolve conditions before they become problems.
- ID for Networks delivers advanced capabilities for network event reduction, root-cause analysis and a new management concept called State Analysis, which actively determines the health of network protocols and complex network configurations.
- Includes out-of-the-box correlators for enhanced root-cause analysis and the new Correlation Composer to easily tailor the out-of-the-box correlators that are shipped with Network Node Manager to fit your particular needs.

- The NNM serves as a SNMP manager, trap collector, and connectivity tester. It also acts as a framework for the attachment of other programs, such as Avaya MultiService Network Manager.
- Topology discovery visually shows the interconnection of routers, switches, and endpoints.

Management models

This section contains information on the two main network management models:

- Distributed (component)
- Centralized (hybrid)

Distributed (component)

Distributed network management is the default management model for network equipment. As Figure 62, Tools for distributed network management, on page 167 shows, each device is managed separately, and can have its own management interface. There is no commonality between these interfaces. Some might be CLI-based, Web-based, or GUI-based applications. In addition, third-party tools such as MRTG complement integrated management interfaces to provide additional functionality.

Figure 62: Tools for distributed network management

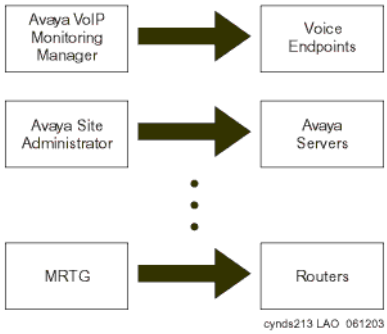

The advantage to this model is that the cost of tools is low compared to the centralized model. Many of the tools are included with the purchase of networking equipment, and many are open source. Also, many of these tools are more specialized on a specific platform or task than centralized management tools. Most Avaya Integrated management products, including Avaya VoIP Monitoring Manager, fall into this category.

There are numerous disadvantages to this model. First, this model requires more support personnel than the centralized model. Next, there are numerous interfaces that support staff must learn, which greatly increases training costs. Finally, the support person must check many places to get the full status of the network, which adds time and the likelihood of missing critical data. This model is appropriate in small to medium-sized enterprises with relatively few pieces of network equipment. It is not appropriate for most large enterprises or enterprises with complicated networks.

Centralized (hybrid)

The centralized management model strives to make all network management available in a central location. It generally begins with a framework product such as HP OpenView NNM. This framework product serves as an SNMP trap receiver for alarm data sent by networking devices. It also provides network topology discovery and availability testing.

Figure 63: Centralized management model

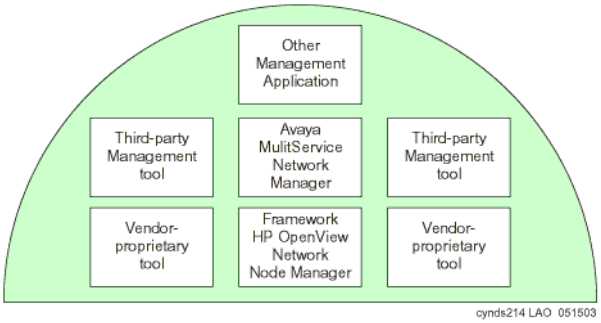

Additional management tools, such as Avaya MultiService Network Manager, attach to the framework (Figure 63, Centralized management model, on page 168). They can be launched directly from the underlying application, and can share data with it. This allows a network administrator to go to a central location for most network management and configuration tasks. Client devices are configured to send alarm and event data to the centralized manager, generally through SNMP. The management station also has the ability to periodically poll the client for specific information. This can be used to graph performance, for example. Polling can also be used for inventory management.

There are many advantages to this model:

- Because a centralized location is used, fewer administrators are required to manage a network.
- Administrators are more likely to catch critical information because it is all in one place.
- Administrators need to learn fewer interfaces, which reduces training costs.
- More advanced centralized management products offer event correlation, which increases the likelihood of proactively catching a problem before it adversely affects users.

The disadvantage to the centralized model is cost. Typically, centralized management tools cost more than distributed tools. In addition, the implementation and integration can be complex. Finally, the enterprise must adjust the manager as the network changes. If the management server is not actively maintained, it quickly falls into disuse.

In practice, it is rare for an enterprise to completely embrace the centralized model. Some applications may not “bolt on” to a particular framework, for example. Also, sometimes an enterprise writes a “homegrown” application to cover an outage with the management server. In addition, the distributed model is useful for times when the central management tool is unavailable.

This resulting hybrid management model that combines elements of centralized management with distributed management tools is most appropriate for large enterprises, or enterprises with complex networks. It is also appropriate for smaller enterprises that can justify the cost of the tools and have in-house expertise to keep the system running.

Reliability and recovery

The purpose of this chapter is to provide the reader an overview of the subject of communication-system "availability," specific to Avaya Communication Manager and Avaya Media Servers and Gateways. The discussion that follows demonstrates Avaya's long-standing commitment to high availability in hardware and software design and the architectural strength of Avaya Application Solutions.

A brief description of availability and its significance to a communications system is provided. Hardware-design considerations, software-design and recovery considerations, IP Telephone and remote media gateway recovery, and overall maintenance strategy are also described. The reliability tables specify the reliability performance of Avaya Application Solutions building blocks.

This chapter contains information on these topics:

Reliability

Customers need the full reliability of their traditional voice networks, including feature richness and robustness, and they want the option of using converged voice and data infrastructures. With the convergence of voice and data applications that run on common systems, a communications failure could bring an entire business to a halt. Enterprises are looking to vendors to help them design their converged infrastructure to meet their expected availability level.

"High availability" communications require the system to work reliably with preexisting transport infrastructures, and to integrate with a wide variety of external connectivity options. As a result, the underlying architecture should be designed to support reliable performance at every level. Avaya Communication Manager running on the Avaya S8700 and the 8300 Media Servers employs a variety of techniques to achieve this high reliability and availability.

Communication Manager is designed to automatically and continually assess performance, and detect and correct errors as they occur. The software incorporates component and subassembly self-tests, error detection and correction, system recovery, and alarm escalation paths. Its maintenance subsystem manages hardware operation, software processes, and data relationships.

Employing the TCP/IP packet-based transport architecture allows additional reliability advantages. One example is the load-sharing and fail-over ability of the principal IP resources found in the media gateways. The TCP/IP architecture also allows telephones to have a recovery mechanism of their own, so they can connect to alternate controllers if the link to their primary gatekeeper is broken.

For large systems, Avaya S8700 Media Servers provide server redundancy, with call preserving fail-over, on the strength of a Linux operating system. The Avaya S8300 Media Servers can further enhance redundancy by serving as Local Survivable Processors (LSPs) within networks. LSPs can take over segments that have been disconnected from their primary call server, and provide those segments with Avaya Communication Manager operation until the outage is resolved.

Reliability and availability

The reliability of maintained systems is often expressed in terms of availability, which is defined as the percentage of time that the system is available to most of the users. The basic formula for calculating availability is:

$$A = \frac{MTBO}{(MTBO + MTTR)}$$

where

- Mean Time Between Outage (MTBO) measures length of time between outages.
- Mean Time To Recovery (MTTR) measures the time to recover from an outage.

Table 34, Expected range of typical availability, on page 170 shows the range of availability that is typically expected of communications systems:

Table 34: Expected range of typical availability

Availability	Downtime per year	Option level	Who might need this
99.95	4 ½ hours	“Standard”	Generally accepted as the minimum standard of acceptable downtime for business
99.99	53 minutes	“High”	Businesses or organizations that highly depend on their communication system
99.999+	5 minutes or less	“Critical”	Hospitals, emergency services, high-performance call centers, critical government agencies, financial institutions

High availability – general design considerations

High availability requires dedicated design diligence at multiple layers and with several overlapping objectives (Table 35, Measurements used for assessing availability expectations, on page 171).

Table 35: Measurements used for assessing availability expectations

Design element	Measurement
Failures of each component and subsystem must be infrequent	Mean Time between Failures (MTBF)
System outages must be infrequent	Mean Time between Outages (MTBO)
When there is a failure or outage: • The impact must be minimized and isolated. • Recovery must be speedy.	Mean Time to Recovery (MTTR)
The system collects its own performance statistics	Various

Not only are failures at the device or subassembly level minimal, but also when there is a failure, the design itself helps in many ways to alleviate the impact of the failure. For example, the design tests itself frequently to detect problems before they become customer affecting. The design isolates subassemblies that are not functioning properly, and runs verification tests on them. If necessary the circuits are taken out of service, and an alarm is automatically sent, prompting the dispatch of a technician. Where necessary, the design incorporates redundancy at the device or subassembly level to add reliability where it is most needed. As an example of the level at which the maintenance architecture is thoroughly built in, consider that at least 30% of the software code for Avaya Communication Manager is dedicated to the maintenance subsystem. The firmware that runs the circuit packs, which also interacts with the maintenance software, is similarly designed.

Hardware considerations

Table 36, Comparison of circuit packs and subassemblies failure rates, on page 173 shows that Avaya circuit packs and subassemblies are extraordinarily reliable relative to the industry in general. This is not by accident. The heritage of "five 9s" (99.999%) availability results from design, manufacture, and lifetime support based on an uncompromising focus on system availability, and refined by tens of billions of hours of user experience.

This is due to highly effective knowledge and execution of "Design for Manufacture, Installation, Reliability, and Serviceability:"

- Quality control that is executed thoroughly from electrical device vendor partnerships, through every stage of the assembly process. The highest quality is pushed to the earliest step of the process possible (based on Deming's "zero defects and zero errors;" this actually reduces overall costs substantially).[1]
- Commonality that is leveraged at all levels
  - ***Piece-parts***. Many of the "workhorses" of the product are in their fifth to seventh generation of silicon integration. This keeps us on the leading edge of technology curves.
  - ***Subassemblies***. Help customers in many ways, not the least of which is investment protection. The subassemblies are also in their fifth to seventh level of renewal and refinement.
  - ***Shared designs***. Even in cases where subassemblies cannot be directly reused, common designs that have been "bullet-proofed" over years are reapplied in new configurations as appropriate.
  - ***Communication Manager***. Avaya's robust, feature-rich, "battle-hardened" software for high-reliability enterprise systems, is common across Avaya IP solutions and traditional solutions.

Table 36, Comparison of circuit packs and subassemblies failure rates, on page 173 compares average failure rates of various Avaya components with industry averages for similar components, demonstrating Avaya's commitment to reliability.

1 In his *Leading the Revolution* Gary Hamel speaks of the importance of "getting different" rather than "getting better." The "zero defects and zero errors" passion fostered by the Quality giant, Dr. Deming, in the 1980s was revolutionary. The prevailing conventional wisdom was that quality just needed to be "good enough" (whatever that is), and that to increase quality beyond "good enough" would be cost prohibitive...and provide diminishing returns. The convention wisdom missed the concept that if "causes" of quality problems are addressed, overall costs actually go down.

Table 36: Comparison of circuit packs and subassemblies failure rates[1]

Industry data		Avaya system elements	
Component	**Mean time to failure**	**Component**	**Mean time to failure**
Logic boards[2]	3-20 years	Media processor circuit pack[3]	50 years
Disks[2]	1-50 years	Protocol preprocessor circuit pack (C-LAN)	50 years
ISP server class power supply[4]	20-25 years	Digital line/trunk circuit packs	72-77 years
Power (North America)[2]	5.2 months	S8700 Server complex: “duplex”/”high & critical”	9 / 90+ years
LAN[2]	3 weeks	Power supplies	45 to 60 years
		IP Server Interface (IPSI) circuit pack	50 years
		Expansion Interface (EI) circuit pack	20 years
		(Industry) Power (North America)	5.2 months
		(Industry) LAN	3 weeks

1 All numbers assume 24 hours per day, 7 days per week usage.
2 Taken from Microsoft High Available Operations Guide.
3 Based on numerous internal Avaya studies of millions of user -hours.
4 Based on an internal survey of reputable vendors.

The data in Table 36, Comparison of circuit packs and subassemblies failure rates, on page 173 shows that in several cases Avaya’s subassemblies are so reliable that it would take twice the number of industry-typical subassemblies to reach the same availability level.

S8700 Server Complex

The high availability philosophy is scrupulously implemented the Avaya S8700 server complex:

- Linux is the operating system (OS) for many reasons:
  - — Provides access to full source for quicker bug-fix turnaround
  - — Allows easy system customization for high availability enhancements
  - — Exhibits fewer known security flaws than other OSs, and allows customization for added security features

- Two servers with a memory shadowing link allow:
  - — One processor to take over when the other fails.
  - — Simple duplication of all other server components (for example, modem, disk, or memory) to eliminate a single point of failure.
  - — Upgrades with insignificant time out of service.
- High availability enhancements:
  - — Software sanity is continuously evaluated. Any insanity (due to unexpected conditions) is detected, and the offending software is forced to go through escalating levels of recovery until finally the entire system can switch over to a standby processor.
  - — Disks are partitioned to keep most of the variable information away from the invariant, and to allow for automatic recovery if newly loaded software fails.
  - — All event logs are proactively scanned for potential service-affecting items. If found, alarms are generated. If necessary, a service dispatch is launched.
  - — Applications running on the OS are thoroughly pr-tested to assure proper performance. This OS is closed to any applications other than those provided by the manufacturer to avoid interference of operation. Alarms can be generated if any untested software is loaded on the system.
  - — Customers and technicians (in most cases) can access the operating system shell directly, providing protection from inadvertent adverse alterations to the system.
  - — Enhanced tools allow secure, remote, nonstop system debugging and unattended data collection.

Avaya S8700 Media Server

While all businesses require solid performance from their communications systems, there are increasing levels of “availability” performance needed. To meet that need, the Avaya S8700 Media Server and its associated gateways accommodate three availability levels:

	Reliability level			
Configuration	**Standard**	**High**	**Critical**	**Link to more information**
Multi-Connect	[2]	[2]	[2]	S8700 Multi-Connect hardware availability
IP-Connect	[2]			S8700 IP-Connect hardware availability

The Avaya S8700 Media Server also provides:

- Automatic restoration of the most recently saved versions of translations following a power outage. Translations are automatically shadowed onto the standby server across a high-speed fiber optic link for memory duplication.

  NOTE:
  Avaya G700 Media Gateways translations can also be copied to Local Spare Processors (LSPs) for automatic recovery in the case of network partitioning or complete central site failure.

- Scheduled backups of critical system information locally and/or at remote sites. In an emergency, multiple copies of Communication Manager translations and server configuration information are available. Saved information can be quickly restored.
- Ability to recover from software failures through server interchanges. If the active server needs to perform a non-call-preserving restart, the standby server can take over under a slightly different operating system environment with nothing more than a call-preserving warm restart. This ability is expected to enhance even traditional abilities, as it provides a fail-safe mechanism to recover from obscure, intermittent "bugs." (This is allowed by processes duplicated on each server deliberately not running in lock-step synchronization).

S8700 Multi-Connect hardware availability

Table 37, S8700 Multi-Connect availability in 3 reliability configurations, on page 175 lists the Multi-Connect hardware that is availability in the three reliability configurations.

Table 37: S8700 Multi-Connect availability in 3 reliability configurations[1]

	Standard (duplex) reliability			High reliability			Critical reliability		
Sub-system	Failures/ year	MTBO (years)	Availability[2] (%)	Failures/ year	MTBO (years)	Availability[2] (%)	Failures/ year	MTBO (years)	Availability[2] (%)
S8700 Server complex[3]	0.1095	9.1	99.995 / 0.99998	<0.01095	>91.3	>99.9995	<0.01095	>91.3	>99.9995
SCC1 or MCC1 Media Gateways[4]	0.153	6.5	99.993 / 99.997	0.0657	15.2	>99.997[5] / 99.999	<0.01095	>91.3	>99.9995
CSS Intf	0.1095	9.1	99.995 / 99.998	0.0876	11.2	>99.996 / 99.998	<0.01095	>91.3	>99.9995

1 September 30, 2002 data.

2 The lower number is the equivalent availability for MTTR of 2 hours, which is attainable with technicians and spares on site. See also Reliability and availability.

3 For high and critical reliability, the 8700 server complex has duplicated servers, dedicated Ethernet switches, and UPSs. The duplex reliability configuration consists of duplicated servers and a UPS for each, with a single, dedicated Ethernet switch.

4 Standard reliability assumes a single IPSI per PN or shared IPSI across multiple PNs. High reliability assumes duplicated IPSI per carrier in each PN. Critical reliability assumes duplicated IPSI per carrier and duplicated ATM interfaces per carrier in each PN.

5 Duplicated IPSIs help with two situations. First, a pair eliminates single point of failure for this module itself. More significantly, a pair allows dual paths through a duplicated data network.

NOTE:
This availability does not include availability of the customer's data networks, PSTN contributions, or contributions due to power outages. A conservative MTTR of 4 hours is assumed, which includes travel and repair time.

S8700 IP-Connect hardware availability

The S8700/G700 IP-Connect hardware availability analyses reflect similar availability results to those in the S8700 Multi-Connect standard configuration.

Configuration drawings

Figure 64: S8700 Media Server in a standard reliability configuration

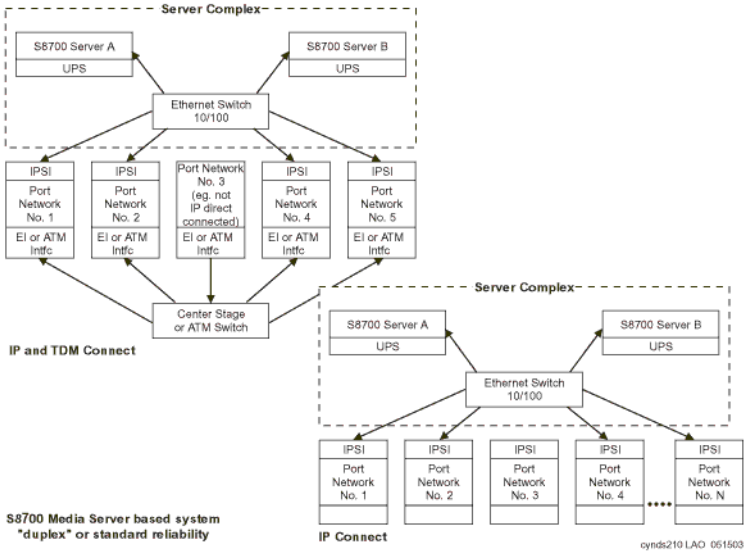

Figure 65: S8700 Media Server in a high reliability configuration

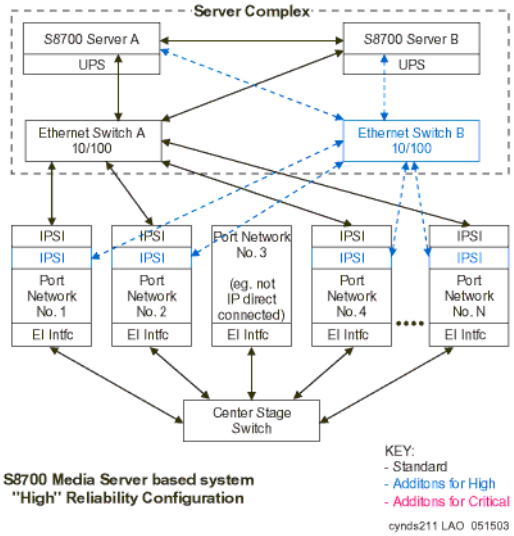

Figure 66: S8700 Media Server in a critical reliability configuration

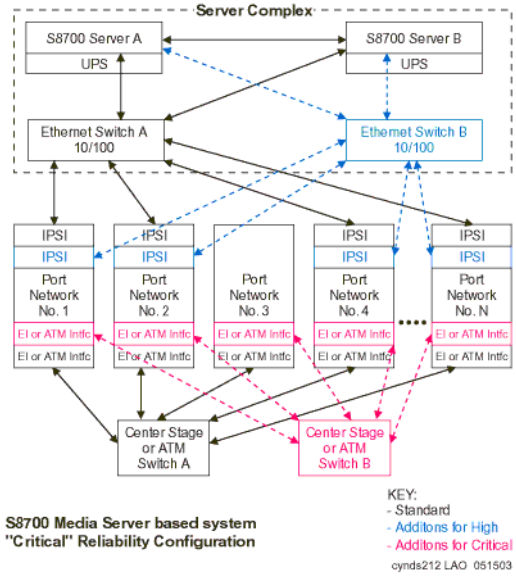

Avaya S8300 Media Server

The S8300 Media Server, like the S8700 Media Server, accommodates several levels of availability performance. In contrast to the S8700 that uses various levels of duplication for processing, and Port Network (PN) connectivity for bearer and signaling to add heightened layers of availability, the S8300 Media Server is designed to use Local Spare Processing (LSP). LSP architecture provides added availability, and survivability to a network of small to medium-sized offices. This table shows the S8300/G700 availability coverage:

Table 38: S8300/G700 reliability analysis[1]

	Standard reliability			High reliability			
Sub-system	**Failures/ year**	**MTBO (years)**	**Availability[2] (%)**	**Failures/ year**	**MTBO (years)**	**Availability[2] (%)**	**Configuration**
S8300 Media Server	1.314	0.65	99.91 / 99.95	0.219	4.6	99.99 / 99.995	**Standard**: Single Internal Call Controller (ICC) equipped Media Gateway per site.
							High: Can failover to an LSP; N+1 media gateways at each site; each site has duplicate interfaces to the data network: each IP endpoint is homed to at least 2 systems that are run by Communication Manager (S8300, or otherwise).[3]
G700 Media Gateway							**Standard**: Media Gateway interface to the call controller is supported by a non-redundant data network.
							High: Can failover to LSP upon a link failure; N+1 Media Gateways at each site with duplicate interfaces to the data network. IP phones can home to an alternate gatekeeper (S8700, S8300, or LSP).

1 September 30, 2002 data. These are conservative engineering estimates. As field data are collected, the numbers will be updated.

2 The lower number is the equivalent availability for MTTR of 2 hours, which is attainable with technicians and spares on site. See also Reliability and availability.

3 Avaya IP Telephones have multi-homing abilities, and can be configured to re-home to any Communication Manager-run system. For example, in a configuration with @8700 at a main site and S8300-ICC or G700 at remote site, the telephones at the remote site could re-home to main S8700 through separate Ethernet switches. This configuration could be said to provide 99.99% availability as well.

Avaya DEFINITY Server R

NOTE:
This availability does not include availability of the customer's data networks, PSTN contributions, or contributions due to power outages. A conservative MTTR of 4 hours is assumed, which includes travel and repair time.

Table 39: Avaya DEFINITY Server R reliability analysis[1]

	Standard reliability (single processor complex)			High reliability (duplicated processor complex)			Critical reliability (duplicated processor complex)		
Sub-system	Failures/ year	MTBO (years)	Availability[2] (%)	Failures/ year	MTBO (years)	Availability[2] (%)	Failures/ year	MTBO (years)	Availability[2] (%)
G3R Processor Complex	0.153	6.5	99.993 / 99.997	<0.01095	>91.3	>99.9995	<0.01095	>91.3	>99.9995
SCC1 or MCC1 Media Gateways	0.153	6.5	99.993 / 99.997	0.0657	15.2	0.99997 / 99.999	<0.01095	>91.3	>99.9995
CSS Intf	0.1095	9.1	99.995 / 99.998	0.0876	11.2	99.996 / 99.998	<0.01095	>91.3	>99.9995

1 Running Avaya Communication Manager (September 30, 2002 data).

2 The lower number is the equivalent availability for MTTR of 2 hours, which is attainable with technicians and spares on site. See also Reliability and availability.

Avaya DEFINITY Server SI and CSI

Table 40: Avaya DEFINITY Server SI and CSI reliability analysis[1]

	Standard reliability (single processor carrier)			High reliability (CSI only)		
Sub-system	Failures/ year	MTBO (years)	Availability[2] (%)	Failures/ year	MTBO (years)	Availability[2] (%)
SI Processor Complex	0.219	4.6	99.99 / 99.995	<0.01095	>91.3	>99.9995
SCC1 or MCC1 Media Gateways	0.153	6.5	99.993 / 99.997	0.0657	15.2	99.997

1 Running Avaya Communication Manager (September 30, 2002 data).

2 The lower number is the equivalent availability for MTTR of 2 hours, which is attainable with technicians and spares on site. See also Reliability and availability.

Maintenance architecture

Software and maintenance architecture recovery

The Communication Manager maintenance architecture is designed to detect and correct errors as they occur. This greatly reduces events that can cause a system outage. This also enables quick identification (fault isolation) to a replaceable subassembly. This automatic assessment is done constantly, in the background of normal operation, so errors can be addressed early and proactively. Component self-testing, subassembly self-testing, error detection and correction, and system reconfiguration and alarming escalation paths are all elements of this architecture. The system software is designed to recover from intermittent failures, and to continue providing service with a minimum of disruption. Firmware that runs each circuit pack does similar tasks at the module level, working tightly with the system software.

Figure 67, Maintenance management architecture, on page 180 shows the various levels of maintenance strategies that are built into the communication system.

The maintenance subsystem manages three categories of maintenance objects:

- ***Hardware maintenance objects*** are tested, and where appropriate alarmed and removed from service by the software. The error is reported to an operations center so that the object can be replaced.
- If a ***software process*** encounters trouble, it is recovered or restarted.
- ***Data relationships*** are audited and corrected.

Figure 67: Maintenance management architecture

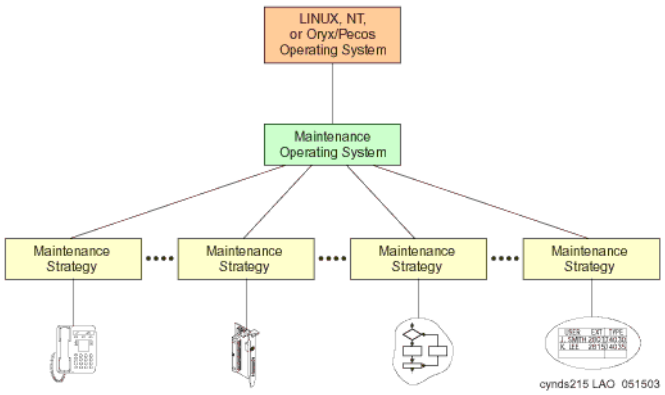

All systems are provided with remote diagnostics capability, which enables rapid troubleshooting and maintenance, in the cases where the system cannot repair itself. Studies have shown that most problems experienced by Avaya systems are self-corrected without impact to the customer. Even with the highly reliable hardware components discussed previously, this sophisticated maintenance management implementation is required to attain the 99.99 – 99.999+% availability of Avaya systems.

Software failure recovery levels

One key to rapid self-recovery of software failures is the judicious use of the appropriate level of recovery. With too little action, the attempted recovery becomes time wasting and ineffective, when stronger action should be taken. Conversely, with too much action, the recovery is unnecessarily prolonged. Figure 68, Avaya Communication Manager's recovery levels, on page 181 shows the 5 recovery levels of Communication Manager.

Figure 68: Avaya Communication Manager's recovery levels

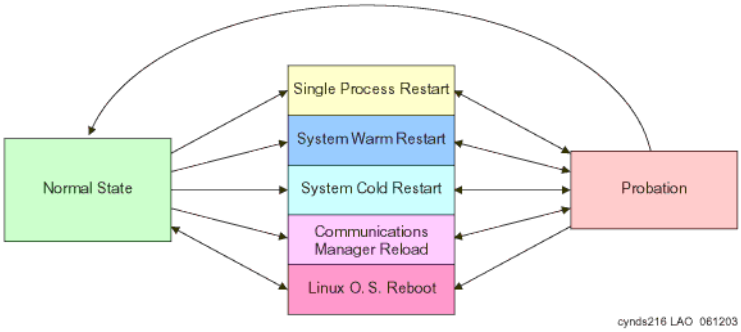

These automatic recovery levels are listed from mildest to strongest, which is also from quickest to slowest and more frequent to less frequent:

- Single process restarts
- System warm restarts
- System cold restarts
- Communication Manager reloads
- Linux operating system reboots on S8300/S8700 Media Servers

Single process restarts

Process sanity audits are performed routinely (approximately every 10 seconds) on many dozens of key software processes. In the event of a process hang, that single process will be restarted, and no call outage will result. If three single process restarts are needed within a 60-second probationary period, the third single process restart is deemed ineffective, and will instead escalate to a system warm restart.

System warm restarts

This mechanism preserves all stable and held calls, as well as feature activity data, throughout the brief recovery period. Processes are essentially started with the same data and stacks that they were using prior to the warm restart. If three warm restarts are needed within a 15-minute probationary period, the system warm restart is deemed ineffective, and will instead escalate to a system cold restart.

System cold restarts

In this recovery mechanism, processes are still started with the same data and stacks that they were using prior to the cold restart, but calls are dropped and Port Networks (PNs) are reset, followed by a port board activation phase of recovery. If three cold restarts are needed within a 15-minute probationary period, the third system cold restart is deemed ineffective, and will instead be escalated to a Communication Manager reload.

Communication Manager reloads

In this recovery mechanism, all calls are dropped, and all processes that are related to call processing are stopped and restarted. PN configuration data known as “translations” is reread from disk, and, as in system cold restarts, PNs are reset, and port boards are activated. If three Communication Manager reloads are needed within a 10-minute probationary period, the reload is deemed ineffective, and will instead be escalated to an operating system reboot.

Linux operating system reboots on S8300/S8700 Media Servers

In this recovery mechanism, all calls are dropped, all processes are killed, and the operating system is completely rebooted. Processes are then read off disk and loaded into memory, where recovery then proceeds exactly as it does in Communication Manager reloads. If the reboot fails after a recent software upgrade, another reboot is attempted, but from a disk partition containing the previous version of software.

For the S8700 Media Server, this implementation is expected to result in a weighted average of 2 minutes per year or less of call-affecting time out of service that can be attributed to software.

IP endpoint and remote media gateway recovery

Avaya’s distributed IP-based systems can also enjoy increased availability by virtue of the “alternate gatekeeper.” When IP Telephones register with Communication Manager, they are given a list of “alternate gatekeepers” to which they can re-register in the event of a failure. Thus, if a C-LAN fails or becomes unavailable, users that are registered to that C-LAN can re-home to another C-LAN that is unaffected by the failure.

IP endpoint recovery

The Avaya server is designed to have a scalable architecture with different server components. These components provide processing and relay signaling information between Communication Manager and the Avaya IP endpoints. The architecture is inherently distributed, thus allowing the system to be scalable to handle large number of endpoints, and flexible to work in different network configurations.

This distributed nature of the architecture introduces additional complexity in dealing with endpoint recovery, since failure of any element in the end-to-end connectivity path between an IP endpoint and the switch software can result in service failure at the endpoint.

The recovery algorithm that is outlined here deals with detection and recovery from the failure of signaling channels for IP endpoints. Such failures are due to connectivity outages between the server and the endpoint, which could be due to failure in the IP network or any other component between the endpoint and the server.

In the S8700 Server configurations the connectivity path between the endpoint and the server is:

Endpoint ⇔ IP network ⇔ C-LAN ⇔ PN backplane ⇔ IPSI ⇔ IP network ⇔ S8700

In this configuration, IP endpoints connect through C-LAN on the PN. The DEFINITY platforms G3r, G3si, and G3csi, which support Avaya Application Solutions features, also use C-LAN for signaling connecting to IP endpoints.

A C-LAN provides two basic reliability functions:

- A C-LAN hides server interchanges from the IP endpoints. The signaling channels of the endpoints remain intact during server interchanges, and do not have to be reestablished with the new active server.
- A C-LAN terminates TCP keepalive messages from the endpoints, and thus frees the server from handling frequent keepalive messages.

Recovery algorithm

The recovery algorithm is designed to minimize service disruption to an IP endpoint in the case of a signaling channel failure. When connectivity to a gatekeeper is lost, the IP endpoint progresses through three phases:

- Recognition of the loss of the gatekeeper
- Search for (discovery of) a new gatekeeper
- Re-registration

When the IP endpoint first registers with the C-LAN, the endpoint receives a list of alternate gatekeeper addresses from the DHCP server. The telephone uses the list of addresses to recover from a signaling link failure to the C-LAN/gatekeeper.

When the telephone detects failure of signaling channel (H.225/Q.931), its recovery algorithm depends on the call state of the endpoint:

- If the user of the telephone is on a call, and the telephone loses its call signaling channel, the telephone waits until the call is hung up before proceeding with the new registration, and does not try to reestablish the link with the server. As a result, the call is preserved, but call features are not available for the remainder of the call.
- If the user of the telephone is not on a call, the telephone closes its signaling channel, and searches for a gatekeeper.

To reestablish the link, the endpoint tries to register with the first C-LAN or S8300 Server on the list. If it is unsuccessful, it tries to register to the next gatekeeper address from the alternate gatekeeper list. When it does successfully register, the server updates its station mapping to reflect this new connection.

In the S8300/G700 configuration, the IP endpoint connects directly to the S8300 Server (there is no C-LAN.) The connectivity path between the endpoint and the server is:

Endpoint ⇔ IP network ⇔ S8300

To discover connectivity failure, keepalive messages are exchanged between the IP end point and the server. When the endpoint discovers that it no longer has communication with its primary gatekeeper, it looks at the next address on its list. If the next address is for an LSP, the LSP accepts the registration and begins call processing.

While the LSP is not call preserving, the fail-over from primary gatekeeper to LSP is an automatic process, and does not require human intervention. The fail-back from LSP to primary gatekeeper, however, is not currently automatic, and requires a system reset on the LSP. During the fail-back to the primary gatekeeper, all calls are dropped, with the exception of IP-to-IP calls.

To discover connectivity failure, two kinds of keepalive procedures are currently implemented between the IP endpoints and the Avaya server:

- RAS keepalive
- TCP keepalive

RAS keepalive

This procedure is used for the H.225.0/RAS signaling channel, and depends on the exchange of RAS keepalive messages between an IP endpoint and the server. The Avaya server uses this procedure to determine if a currently registered endpoint is still available. The endpoint uses the procedure to determine whether the server is reachable or not.

A registered H.323 endpoint periodically sends keepalive messages for the purpose of performing lightweight re-registration with the switch.

TCP keepalive

This procedure is implemented for the H.225.0/Q.931 call signaling channel. In this procedure, the endpoint periodically sends TCP keepalive messages to determine if the TCP peer is reachable or not. The TCP peer can be a C-LAN board (in the case of an S8700 Server), or the Media Server (in the case of the S8300 Server). The TCP keepalive message is only sent when the TCP connection is idle.

IP Telephones and softphones rely on RAS and TCP keepalive procedures to monitor the signaling connectivity status with ACM Server. For the S8300 Media Server, there is no TCP keepalive from S8300 to the endpoints.

G700 Media Gateway recovery

If the link between the remote media gateway and the gatekeeper is broken, or the gatekeeper is down, the LSP activates and assume call processing for the media gateway. The main gatekeeper can be the

- S8700, also known as an external communication controller (ECC), see S8700/G700 configuration
- S8300, also known as the internal communication controller (ICC), see S8300/G700 configuration

The strategy by which the media gateways change control from the primary to the LSP gatekeeper is driven by the gateway using the alternate gatekeeper list.

S8700/G700 configuration

In this configuration, the connectivity path between the G700 Media Gateway and the server is:

G700 ⇔ IP network ⇔ C-LAN ⇔ PN backplane ⇔ IPSI ⇔ IP network ⇔ S8700

To discover gatekeeper/media gateway connectivity failure, there are keepalive messages between the gateway and the media server (S8700 or S8300). The G700 alternate gatekeeper list is divided into primary and secondary addresses. All primary addresses receive priority weighting over the secondary addresses. Normal practice is to designate all C-LANs for the primary controller as the primary gatekeeper addresses, and all LSPs as secondary addresses. This practice gives the Media Gateway the best possible chance of registering with the primary controller before registering with the LSP, and entering into survivable mode. In the S8700/G700 configuration, up to 50 LSPs back up the gateways that are controlled by the S8700 Server.

When it discovers the link failure between the gateway and the primary gatekeeper, the Media Gateway sends three keepalive messages, and waits for responses to these messages from the primary gatekeeper. When there is no response, the socket that the gateway is communicating through closes due to a controller failure, and the TCP connection is closed. The gateway realizes the failure immediately, and enters fail-over mode.

Once the G700 Media Gateway enters the fail-over mode, it begins to search for the alternate gatekeeper. The Media Gateway continues down the list, and once it receives a response from a controller, the gateway performs a reboot of the Media Gateway Processor (MGP). Following the MGP reboot, the Media Gateway checks again with the first address in its list to find out if during the reboot the primary controller recovered, and begins processing registrations once again.

When the call processing control transfers to the LSP, all calls are dropped, with the exception of IP-to-IP calls. These calls are preserved, with the exception of features for that call, such as conferencing, call waiting, and call transfer.

S8300/G700 configuration

In this configuration, the connectivity path between the G700 Media Gateway and the S8300 Media Server is:

Endpoint ⇔ IP Network ⇔ S8300 Server

The link failure discovery and recovery process is the same as above, except there are no C-LAN addresses in the alternate gatekeeper list. In the S8300/G700 configuration, up to 10 LSPs can back up the media gateways that are controlled by the S8300 Server.

Section 3. Getting the IP network ready for telephony

IP Telephony network engineering overview

This section contains information about these topics:

- Delay, loss, and jitter measurements
- QoS and CoS
- Network topologies
- Bandwidth estimates for signaling and bearer traffic
- Network design recommendations
  - Overview
  - Voice quality
  - Best practices
  - Common issues

Network design recommendations

In the early days of local area networking, network designers used hubs to attach servers and workstations, and routers to segment the network into manageable pieces. Because of the high cost of router interfaces and the inherent limitations of shared-media hubs, network design was generally well done. In recent years, with the rise of switches to segment networks, designers were able to hide certain faults in their networks and still get good performance. As a result, network design was often less than optimal. IP Telephony places new demands on the network. Suboptimal design cannot cope with these demands. Even with switches installed, a company must follow industry best practices to have a properly functioning voice network. Because most users do not tolerate poor voice quality, administrators should implement a well-designed network before they begin IP Telephony pilot programs or deployments.

Overview

Industry best practices dictate that a network be designed with consideration of the following factors:

- Reliability and redundancy
- Scalability
- Manageability
- Bandwidth

Voice mandates consideration of the following additional factors when designing a network:

- Delay
- Jitter
- Loss
- Duplex

In general, these concerns dictate a hierarchical network that consists of at most three layers (Table 41, Layers in a hierarchical network, on page 189):

- Core
- Distribution
- Access

Some smaller networks can collapse the functions of several layers into one device.

Table 41: Layers in a hierarchical network

Layer	Description
Core	The core layer is the heart of the network. The purpose of the core layer is to forward packets as quickly as possible. The core layer must be designed with high availability in mind. Generally, these high-availability features include redundant devices, redundant power supplies, redundant processors, and redundant links. Today, core interconnections increasingly use Gigabit Ethernet.
Distribution	The distribution layer links the access layer with the core. The distribution layer is where QoS feature and access lists are applied. Generally, Gigabit Ethernet connects to the core, and either Gigabit Ethernet or 100base-TX/FX links connect the access layer. Redundancy is important at this layer, but not as important as in the core.
Access	The access layer connects servers and workstations. Switches at this layer are smaller, usually 24 to 48 ports. Desktop computers and workstations are usually connected at 10 Mbps (or 10 Mbps), and servers are connected at 100 Mbps (or 1 Gbps). Limited redundancy is used. Some QoS and security features can be implemented in the access layer.

For IP Telephony to work well, WAN links must be properly sized with sufficient bandwidth for voice and data traffic. Each voice call uses between 6.3 Kbps and 80 Kbps, depending on the desired codec, quality, and header compression used. G.729, which uses 24 Kbps of bandwidth, is one of the most promising standards today. Traditional telephone metrics, such as average call volume, peak volume, and average call length, can be used to size interoffice bandwidth demands. See Traffic engineering for more information.

Quality of Service (QoS) also becomes increasingly important with WAN circuits. In this case, QoS means the classification and the prioritization of voice traffic. Voice traffic must be given absolute priority through the WAN. If links are not properly sized or queuing strategies are not properly implemented, the quality and the timeliness of voice and data traffic will be less than optimal.

Three WAN technologies are commonly used with IP Telephony:

- ATM
- Frame Relay
- Point-to-point (PPP) circuits

These technologies all have good throughput, low latency, and low jitter. ATM has the added benefit of enhanced QoS. Frame Relay and PPP links are more economical, but lack some of the traffic-shaping features of ATM.

Of the three technologies, Frame Relay is the most difficult WAN circuit to use with IP Telephony. Congestion in Frame Relay networks can cause frame loss, which can significantly degrade the quality of IP Telephony conversations. With Frame Relay, proper sizing of the committed information rate (CIR) is critical. In a Frame Relay network, any traffic that exceeds the CIR is marked as discard eligible, and is discarded at the option of the carrier if it experiences congestion in its switches. Because voice packets must not be dropped, CIR must be sized to maximum traffic usage. Also, Service Level Agreements (SLAs) must be established with the carrier to define maximum levels of delay and frame loss, and remediation if the agreed-to levels are not met.

Network management is another important area to consider when implementing IP Telephony. Because of the stringent requirements imposed by IP Telephony, it is critical to have an end-to-end view of the network, and ways to implement QoS policies globally. Products such as HP OpenView Network Node Manager, Avaya™ Integrated Management, Concord NetHealth, and MRTG help administrators maintain acceptable service. Outsource companies are also available to assist other companies that do not have the resources to implement and maintain network management.

Voice quality

Voice quality is always a subjective topic. Defining “good” voice quality varies with business needs, cultural differences, customer expectations, and hardware and software. The requirements set forth are based on the ITU-T and EIA/TIA guidelines and extensive testing at Avaya Labs. Avaya requirements meet or exceed most customer expectations. However, the final determination of acceptable voice quality lies with the customer’s definition of quality, and the design, implementation, and monitoring of the end-to-end data network.

Quality is not one discrete value where the low side is good and the high side is bad. A trade-off exists between real-world limits and acceptable voice quality. Lower delay, jitter, and packet loss values can produce the best voice quality, but also can come with a cost to upgrade the network infrastructure to get to the low values. Another real-world limit is the inherent WAN delay over an IP trunk that links the west coast of the United States to India. This link could add a fixed delay of 150 milliseconds (ms) into the overall delay budget.

Perfectly acceptable voice quality is attainable, but will not be “toll” quality. Therefore, Avaya presents a tiered choice of elements that make up the requirements.

The critical objective factors in assessing IP Telephony quality are delay, jitter, and packet loss. To ensure good and consistent levels of voice quality, Table 42, Factors that affect voice quality, on page 191 lists Avaya’s suggested network requirements. These requirements are true for both LAN only and for LAN and WAN connections.

Table 42: Factors that affect voice quality

Network factor	Measurement[1]
Delay (one-way between endpoints)	• A delay of 80 ms or less can, but may not, yield the best quality. • A delay of 80 ms to 180 ms can yield business-communication quality. Business-communication quality is much better than cell-phone quality, and is well-suited for the majority of businesses.[2] • Delays that exceed 180 ms might still be quite acceptable depending on customer expectations, analog trunks used, codec type, and so on.
Jitter (variability of the delay between endpoints)	• 20 ms, or less than half the sample size, for the best quality. **NOTE:** This value has some latitude, depending on the type of service that the jitter buffer has in relationship to other router buffers, the packet size used, and so on.
Packet loss (maximum packet/frame loss between endpoints)	• <1% can yield the best quality, depending on many factors. • <3% should give business-communications quality, which is much better than cell-phone quality.[2] • >3% might be acceptable for voice, but might interfere with signaling.

1 All measurement values are between endpoints because this document assumes that IP Telephony is not yet implemented. All values therefore reflect the performance of the network without endpoint consideration.

2 Also, “business-communication quality” is defined as less than toll quality, but much better than cell-phone quality.

For more information see Voice quality network requirements.

Best practices

To consistently ensure the highest quality voice, Avaya highly recommends consideration of the following industry best practices when implementing IP Telephony. Note that these suggestions are options, and might not fit individual business needs in all cases.

- **QoS/CoS**. QoS for voice packets is obtained only after a Class of Service (CoS) mechanism tags voice packets as having priority over data packets. Networks with periods of congestion can still provide excellent voice quality when using a QoS/CoS policy. The recommendation for switched networks is to use IEEE 802.1p/Q. The recommendation for routed networks is to use DiffServ Code Points (DSCP). The recommendation for mixed networks is to use both. Port priority can also be used to enhance DiffServ and IEEE 802.1p/Q. Even networks with plentiful bandwidth should implement CoS/QoS to protect voice communications from periods of unusual congestion, such as a computer virus might cause. See Implementing Communication Manager on a data network for more information.

- **Switched network**. A fully switched LAN network is a network that allows full duplex and full endpoint bandwidth for every endpoint that exists on that LAN. Although IP Telephony systems can work in a shared or hub-based LAN, Avaya recommends the consistently high results that a switched network lends to IP Telephony.
- **Network assessment**. A Basic Network Readiness Assessment Offer from Avaya is vital to a successful implementation of IP Telephony products and solutions. Contact an Avaya representative or authorized dealer to review or certify your network. Network assessment offer explains the options that are available with this offer.
- **VLANs**. Placing voice packets on a separate VLAN or subnetwork from data packets is a generally accepted practice to reduce broadcast traffic. When data is on a shared LAN, this practice also reduces contention for the same bandwidth as voice. Note that Avaya IP Telephones provide excellent broadcast storm protection. Other benefits become available when using VLANs, but there can be a substantial cost with initial administration and maintenance. Section 3.2.1.2 “Using VLANs” explains this concept further.

Common issues

Some common negative practices that can severely impact network performance, especially when using IP Telephony, include:

- **A flat, nonhierarchical network**, for example, cascading small workgroup switches together. This technique quickly results in bottlenecks, because all traffic must flow across the uplinks at a maximum of 1Gbps, versus traversing switch fabric at speeds up to 256 Gbps. The greater the number of small switches or layers, the greater the number of uplinks, and the lower the bandwidth for an individual connection. Under a network of this type, voice performance can quickly degrade to an unacceptable level.
- **Multiple subnets on a VLAN**. A network of this type can have issues with broadcasts, multicasts, and routing protocol updates. This practice can have a significant negative impact on voice performance, and complicate troubleshooting.
- **A hub-based network**. Hubs in a network create some interesting challenges for administrators. It is advisable not to link more than four 10baseT hubs or two 100baseT hubs together. Also, the *collision domain*, the number of ports that are connected by hubs without a switch or router in between, should be kept as low as possible. Finally, the effective (half-duplex) bandwidth that is available on a shared collision domain is approximately 35% of the total bandwidth that is available.
- **Too many access lists**. Access lists slow down a router. While access lists are appropriate for voice networks, care must be taken not to apply them to unnecessary interfaces. Traffic should be modeled beforehand, and access lists applied only to the appropriate interface in the appropriate direction, not to all interfaces in all directions.

Avaya recommends caution when using the following:

- Network Address Translation (NAT). Most implementations that use IP Telephony endpoints behind NAT fail because many H.323 messages (the protocol carrying the voice information) contain multiple instances of the same IP address in a given message, but NAT is unlikely to find and translate all of them. See NAT on page 217 for more information on using NAT with IP Telephony. Avaya products work seamlessly with static NAT implementation, even if that NAT is not H.323-aware.

- **Analog dial-up**. Be careful in using analog dial-up (56 K) to connect two locations. Upstream bandwidth is limited to a maximum of 33.6 K, and in most cases is less. This results in insufficient bandwidth to provide toll-quality voice. Some codecs and network parameters provide connections that are acceptable, but consider each connection individually.
- Virtual Private Network (VPN). Large delays are inherent in some VPN software products due to encryption, decryption, and additional encapsulation. Some hardware-based products, including Avaya VPN products, encrypt at near wire speed, and can be used. In addition, if the VPN is run over the Internet, sufficient quality for voice cannot be guaranteed unless delay, jitter, and packet loss are contained within the parameters that are listed above. See VPN for more information.

Network design

This chapter discusses the network design process for IP Telephony. This chapter focuses on LAN issues, IP addressing, deployment of IP terminals, WAN issues, VPNs, and NAT.

LAN

This section covers LAN issues, including speed and duplex, inline power, hubs versus switches, and so on.

General guidelines

Because of the time-sensitive nature of IP Telephony applications, IP Telephony should be implemented on an entirely switched network. Ethernet collisions, which are a major contributor to delay and jitter, are virtually eliminated on switched networks. Additionally, the C-LAN, MedPro, and IP Telephones should be placed on a separate subnetwork or VLAN (that is, separated from other non-IP Telephony hosts). This separation provides for a cleaner design where IP Telephony hosts are not subjected to broadcasts from other hosts, and where troubleshooting is simplified. This separation also provides a routed boundary between the IP Telephony segments and the rest of the enterprise network, where restrictions can be placed to prevent unwanted traffic from crossing the boundary. When personal computers are attached to IP Telephones, the uplink to the Ethernet switch should be a 100-Mbps link, so that there is more bandwidth to be shared between the telephone and the computer.

Sometimes enterprises are unable to follow these guidelines, and Avaya's solutions can be made to work in some less-than-ideal circumstances. If IP Telephones will share a subnetwork with other hosts, the IP Telephones should be placed on a subnetwork of manageable size (24-bit subnet mask or larger, with 254 hosts or less), with as low a rate of broadcasts as possible. If the broadcast level is high, remember that 100-Mbps links are less likely to be overwhelmed by broadcast traffic than 10-Mbps links. Perhaps a worst-case example is the scenario where Avaya IP Telephones are deployed on a large subnetwork that is running IPX or other broadcast-intensive protocol, with broadcasts approaching 500 per second. Although the performance of the IP Telephones and the voice quality can be satisfactory in this environment, this type of deployment is strongly discouraged.

Ethernet switches

The following recommendations apply to Ethernet switches to optimize operation with Avaya endpoints. Theses recommendations are meant to provide the simplest configuration by removing unnecessary features.

- Enable spanning tree fast start feature or disable spanning tree at the port level. The Spanning Tree Protocol is a Layer 2 loop-avoidance protocol. When a device is first connected (or reconnected) to a port that is running spanning tree, the port takes approximately 50 seconds to cycle through the Listening, Learning, and Forwarding states. This 50-second delay is neither necessary nor desired on ports that are connected to IP endpoints. Instead, enable a fast start

feature on these ports to put them into the Forwarding state almost immediately. If this feature is not available, disabling spanning tree on the port is an option that should be considered. Do not disable spanning tree on an entire switch or VLAN.

- Disable Cisco features. Cisco features that are not required by Avaya endpoints include channeling, cdp, and inline power. These features are nonstandard mechanisms that are relevant only to Cisco devices, and can sometimes interfere with Avaya devices. The CatOS command **set port host <mod/port>** automatically disables channeling and trunking, and enables portfast. Execute this command first, and then manually disable cdp and Cisco inline power. Then manually enable 802.1Q trunking as necessary.
- Properly configure 802.1Q trunking on Cisco switches. When trunking is required on a Cisco CatOS switch that is connected to an Avaya IP Telephone, enable it for 802.1Q encapsulation in the nonegotiate mode (**set trunk <mod/port> nonegotiate dot1q**). This causes the port to become a plain 802.1Q trunk port with no Cisco autonegotiation features. When trunking is not required, explicitly disable it, because the default is to autonegotiate trunking.

Speed and duplex

One major issue with Ethernet connectivity is proper configuration of speed and duplex. A significant amount of misunderstanding exists in the industry as a whole with regard to the autonegotiation standard. Table 43, Speed/duplex matrix, on page 196 is a quick reference for how speed and duplex settings are determined and typically configured. It is imperative that the speed and duplex settings be configured properly.

Table 43: Speed/duplex matrix

Device 1 configuration	Device 2 configuration	Result
Autonegotiate	Autonegotiate	100/full expected and often achieved, but not always stable. Suitable for user PC connections, but not suitable for server connections or uplinks. May be suitable for a single IP Telephony call, such as with a softphone. Not suitable for multiple IP Telephony calls, such as through a MedPro circuit pack.
Autonegotiate	100/half	100/half stable. Device 1 senses the speed, and matches accordingly. Device 1 senses no duplex negotiation, so it goes to half duplex.
Autonegotiate	10/half	10/half stable. Device 1 senses the speed and matches accordingly. Device 1 senses no duplex negotiation, so it goes to half duplex.
Autonegotiate	100/full	Device 1 goes to 100/half, resulting in a duplex mismatch, which is undesirable. Device 1 senses the speed, and matches accordingly. Device 1 senses no duplex negotiation, so it goes to half duplex.
100/full	100/full	100/full stable. Typical configuration for server connections and uplinks.
10/half 100/half	10/half 100/half	Stable at respective speed and duplex. Some enterprises do this on user ports as a matter of policy for various reasons.

A duplex mismatch condition results in a state where one side perceives a high number of collisions, while the other side does not. This results in packet loss. Although it degrades performance in all cases, this level of packet loss might go unnoticed in a data network because protocols such as TCP retransmit lost packets. In voice networks, however, this level of packet loss is unacceptable. Voice quality rapidly degrades in one direction. When voice quality problems are experienced, duplex mismatches are the first thing to look for.

VLANS

VLANS are an often-misunderstood concept. This section begins by defining VLANS, and then addresses configurations that require the Avaya IP Telephone to connect to an Ethernet switch port that is configured for multiple VLANs. The IP Telephone is on one VLAN, and a personal computer that is connected to the telephone is on a separate VLAN. Four sets of configurations are given: Avaya Cajun P330 v3.2.8 and later, Avaya Cajun P330 pre-3.2, Cisco CatOS, and some Cisco IOS.

VLAN defined

With simple Ethernet switches, the entire switch is one Layer 2 broadcast domain that usually contains one IP subnetwork (Layer 3 broadcast domain). Think of a single VLAN (on a VLAN-capable Ethernet switch) as being equivalent to a simple Ethernet switch. A VLAN is a logical Layer 2 broadcast domain that typically contains one IP subnetwork. Therefore, multiple VLANs contain logically separated subnetworks. This arrangement is analogous to multiple switches being physically separated subnetworks. A Layer 3 routing process is required to route between VLANs, just as one is required to route between subnetworks. This routing process can take place on a connected router or a router module within a Layer 2/Layer 3 Ethernet switch. If no routing process is associated with a VLAN, devices on that VLAN can only communicate with other devices on the same VLAN.

For more information, use the links below to see Avaya's white paper, "LANs and VLANs: A Simplified Tutorial."

- Avaya Associates use this link (http://gozer.dr.avaya.com/)
- Business Partners use this link (www.avaya.com)

The port or native VLAN

Port VLAN and native VLAN are synonymous terms. The IEEE 802.1Q standard and most Avaya switches use the term *port VLAN*, but Cisco switches use the term *native VLAN*. Issue the **show trunk** command on P330s and CatOS Catalysts to see which term is used in the display output.

Every port has a port VLAN or a native VLAN. Unless otherwise configured, it is VLAN 1 by default. It can be configured on a per-port basis with the commands in Table 44, Commands to configure a port VLAN or a native VLAN, on page 197.

Table 44: Commands to configure a port VLAN or a native VLAN

Avaya P33xT v3.2.8 and later	Cisco CatOS
set port vlan <id> <mod/port>	**set vlan <id> <mod/port>**

All untagged Ethernet frames (with no 802.1Q tag, for example, from a personal computer) are forwarded on the port VLAN or the native VLAN. This is true even if the Ethernet switch port is configured as an 802.1Q trunk, or otherwise configured for multiple VLANs. For more information, see VLAN binding feature (P330 v3.2.8)).

Trunk configuration

A trunk port on an Ethernet switch is one that is capable of forwarding Ethernet frames on multiple VLANs through the mechanism of VLAN tagging. IEEE 802.1Q specifies the standard method for VLAN tagging. Cisco also uses a proprietary method called ISL. Avaya products do not interoperate with ISL.

A trunk link is a connection between two devices across trunk ports. This connection can be between a router and a switch, between two switches, or between a switch and an IP Telephone. Some form of trunking or forwarding multiple VLANs must be enabled to permit the IP Telephone and the attached personal computer to appear on separate VLANs. The commands in Table 45, Administration commands for VLAN trunking, on page 198 enable trunking.

Table 45: Administration commands for VLAN trunking

Avaya P33xT v3.2.8 and later	Cisco CatOS
set trunk <mod/port> dot1q	**set trunk <mod/port> nonegotiate dot1q**
By default, only the port VLAN or the native VLAN is enabled on the trunk port. Another set of commands is required to specify other allowed VLANs.	By default, all VLANs (1 to1005) are enabled on the trunk port. VLANs can be selectively removed with the command **clear trunk <mod/port> <vid>**.

Note that Cisco *can* remove VLANs from a trunk port. This is a highly desirable feature because only two VLANs at most should appear on a trunk port that is connected to an IP Telephone. That is, broadcasts from nonessential VLANs should not be permitted to bog down the link to the IP Telephone. The pre-3.2 version of P330 code did not have the capability to clear off unwanted VLANs. Enabling 802.1Q trunking enabled all the VLANs. Newer versions of P330 code can limit the VLANs on a trunk, but in doing so they alter the previous trunking behavior.

VLAN binding feature (P330 v3.2.8)

With both Cisco and the pre-3.2 P330 code, the default behavior of trunking is to permit all VLANs. With the new P330 code, the default behavior is to permit only the port VLAN or the native VLAN, and to block all other VLANs. Additional VLANs are added to a port using the VLAN binding feature. In addition, the port does not need to be a trunk at all to forward multiple VLANs. For one application, connecting to an Avaya IP Telephone, the port *must* not be a trunk (do not issue the set trunk command).

To enable VLAN binding:

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

   Static option

   a. Put the port in bind-to-static mode by typing **set port vlan-binding-mode <mod/port> static**.

   b. Statically add another VLAN in addition to the port VLAN or the native VLAN by typing **set port static-vlan <mod/port> <vid>**.

Configured option

a. Add a VLAN to the configured VLAN list by typing **set vlan <id>**.

b. Type **show vlan** to see entire list.

c. Apply the configured VLANs to the port, and permit only those VLANs (bind-to-all permits all VLANs and not just the configured) by typing **set port vlan-binding-mode <mod/port> bind-to-configured**

3. For simplicity, Avaya recommends using the static option for IP Telephony. 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 Telephone with an attached personal computer, trunking must not be enabled so that none of the egress frames are tagged. This is necessary because most personal computers cannot understand tagged frames.

Setting the priority without trunking or VLAN binding (single-VLAN scenario)

With Avaya, it is possible to set the Layer 2 priority on the IP Telephone, even if the telephone 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 VLAN or the native VLAN configured. This is useful if the telephone and the attached personal computer are on the same VLAN (same IP subnetwork), but the telephone traffic requires higher priority. Enable 802.1Q tagging on the IP phone, set the priorities as desired, and set the VID to zero. Per the IEEE standard, a VID of zero assigns the Ethernet frame to the port VLAN or the native VLAN.

Cisco switches behave differently in this scenario, depending on the hardware platforms and OS versions. Table 46, Cisco hardware characteristics, on page 199 shows Avaya's laboratory test results with a sample of hardware platforms and OS versions.

Table 46: Cisco hardware characteristics

Hardware platform / operating system	Laboratory test results
Catalyst 6509 with 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 with CatOS 6.3(3)	Did 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 with 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 the OS version, and consult the Cisco documentation, or call TAC.

NOTE:
Setting a Layer 2 priority is only useful if QoS is enabled on the Ethernet switch. Otherwise, the priority-tagged frames are treated no differently than clear frames. See Appendix C. Multi-VLAN example for an example.

IP addressing

This section discusses IP addressing concepts, including classful addresses, RFC 1918, CIDR/VLSM, and DHCP.

Overview

An IP (v4) address is a 32-bit network address (Layer 3 on the OSI model). An IP address is usually written in dotted-quad notation. Dotted-quad notation consists of four integer fields that range from 0 to 255, and are separated by periods.

An IP address consists of a network portion and a host portion. The boundary that separates the network portion and the host portion of the address is defined by the subnet mask. The subnet mask is another 32-bit address (again usually written in dotted-quad notation) with a consecutive string of 1s followed by a consecutive string of 0s when written in binary. All bit positions of the IP address that are covered by a 1 in the subnet mask are network address bits. Bit positions that are covered by 0s represent the host address.

Some standard subnet masks have been assigned. For addresses that begin with 0 to 127, the default subnet mask is 255.0.0.0, in which 8 bits are used for the network, and 24 bits are used for the host. This is known as a Class A address. For addresses that begin with 128 to 191, the default subnet mask is 255.255.0.0, in which 16 bits are used for the network, and 16 bits are used for the host. This is known as a Class B address. Finally, for addresses that begin with 192 to 223, 24 bits are used for the network, and 8 bits are used for the host. This is known as a Class C address.

In recent years, additional techniques, including Variable Length Subnet Masks (VLSM) and Classless InterDomain Routing (CIDR), have extended subnetting techniques to make more efficient use of address space. CIDR introduced the concept of supernets, which is a technique for aggregating a range of older classful address blocks (for example, Class C) under a single network mask. VLSM provides a technique for allocating subnets of varying size out of a classful address block. Prior to this point, once a subnet mask was applied to a network, the same mask had to be applied to all subnetworks.

Some addresses, defined by RFC 1918, are available for private use. Each address class range includes one group of addresses. The available addresses are:

- Class A: 10.0.0.0 through 10.255.255.255 (mask 255.0.0.0)
- Class B: 172.16.0.0 through 172.31.255.255 (mask 255.240.0.0)
- Class C: 192.168.0.0 through 192.168.255.255 (mask 255.255.0.0)

These addresses can be allocated by companies and individuals in any way. Be aware, however, of the following caveats:

- These addresses are not routable across the Internet. If an organization that uses RFC 1918 addresses wants to connect to the Internet, that organization must use Network Address Translation (NAT).
- If a company is connecting its network to another company, it must take care that their RFC 1918 addresses do not overlap. Overlapping address ranges prohibit unimpeded communication across affected networks.

DHCP

Dynamic Host Configuration Protocol (DHCP) is a tool that automates the assignment of IP addresses. DHCP is a successor of BOOTP, the Bootstrap Protocol. Avaya IP Telephones can use DHCP to learn their IP addresses, default gateways, call controller, tftp server, QoS settings, and other parameters.

DHCP is a broadcast protocol. which means that request messages from DHCP clients such as Avaya IP Telephones are seen by all devices on the local network, but are not forwarded to additional subnetworks. If the DHCP server is present on a different network, DHCP forwarding must be enabled on the router. DHCP forwarding converts the broadcast message into a unicast message, and forwards the message to the configured DHCP server. DHCP forwarding is offered on most routers and layer 3 switches, including those offered by Avaya and Cisco.

For more information, see Appendix D. DHCP / TFTP or the *Avaya IP Telephone LAN Administrator's Guide*.

Recommendations for IP Telephony

Avaya recommends using a separate subnetwork for voice. Isolating voice traffic from data traffic allows protection from viruses, excessive broadcast traffic, and security threats that are caused by malicious users or external intruders. For most IP Telephony implementations, using RFC 1918 (private) address space is acceptable. Generally, Voice over IP (VoIP) is not deployed across the public Internet. Therefore, providing addresses in the private range saves public IP addresses, and provides a layer of security protection by denying connections directly in from the Internet. Should a public Internet connection prove necessary, Avaya recommends setting up a C-LAN and a Media Processor card in a demilitarized zone (DMZ) off the firewall, and using Communication Manager as a proxy server between the internal and external networks.

Avaya also recommends using DHCP to configure IP Telephones. Using DHCP reduces administration to a single point, and reduces the incidence of typographical errors that could cause configuration problems. Avaya includes special configuration options for IP Telephones in Option 176. Microsoft and ISC (Linux and Unix) DHCP servers support this option. Additional methods exist for configuring IP Telephones if a particular DHCP server does not support Option 176. Contact your Avaya representative for more information.

IP terminals deployment

IP Telephone

This sections covers some general information regarding the IP Telephone. See the following resources for more detailed information:

- *4600 Series IP Telephone Installation Guide*
- *4600 Series IP Telephone LAN Administrator's Guid*e

Both documents can be found at:

Link to 4600 Series IP Telephone documents (support.avaya.com)

The current GA firmware releases can be obtained at the Avaya Support Center (support.avaya.com).

The information that is covered in this section may or may not be covered in the resources that are listed above. It might also be necessary to read the "*4600 Series…*" guides above to fully understand the information covered in this section.

Basics

Speed and duplex

Newer Avaya IP Telephones such as the 4620 and the 4602SW contain a 10/100 Ethernet switch. This switch is set to autonegotiate speed and duplex by default. As was stated above, the closet Ethernet switch to which the IP Telephone is attached should be set to autonegotiate, as well. Locking down the closet switch to full duplex will lead to packet loss, and thus result in problems with voice quality.

Older Avaya IP Telephones such as the 4606, 4612, 4624, and 4630 contain a 10/100 hub. The integrated hub in the IP Telephone operates at 10 mbps, or 100 mbps half duplex. There are generally no speed or duplex issues with the IP Telephone. When connected to an Ethernet switch port that is configured to autonegotiate, the Ethernet switch port stabilizes at 100/half. The exception to this is if a personal computer is attached to the telephone that is capable of only 10 mbps. In this case, case all three devices stabilize at 10/half. If no personal computer is to be attached to the telephone, or if the attached computer will always be capable of 100 mbps operation, it is good practice to lock down the Ethernet switch to 100/half. If a personal computer might be attached to the telephone, and there is a chance that the computer might have a 10-mbps NIC, leave the Ethernet switch port in autonegotiate mode.

Note about Single-Speed Bus: Dual-speed hubs and switches must inherently buffer and discard traffic because of the inconsistent flows (one port receives at 100 mbps, but the other can only send at 10 mbps). The Avaya IP Telephone is designed with a single-speed bus in the hub, and does not perform these functions. Instead, these functions are transferred to the enterprise Ethernet switch, where they really belong. Although the IP Telephone can accommodate a second user device (the telephone itself being the first), its primary function is not that of an enterprise network device.

30A base switch

The 30A base switch is a three-port switch that is integrated into the base stand of 4612 and 4624 sets. The pigtail cable attaches to the uplink port of the IP Telephone. The other two ports are an uplink port to connect to the enterprise Ethernet switch, and a user port to connect to a personal computer, just like the IP Telephone. Both ports to require standard Ethernet cables, just like the IP Telephone. Each port supports 10/100 capability at full duplex or half duplex. The ports are in autonegotiate mode, and cannot be configured. Therefore, the attached devices must also be in autonegotiate mode, or they must be fixed to 100/half or 10/half. Limited experience has shown that the 30A base switch functions adequately with the attached devices in autonegotiate mode. Because the 30A base switch is not an enterprise-class switch, it is best to have the speed and duplex on both ports be the same. Otherwise, the 30A base switch is required to buffer and discard frames, which it can do but not as well as an enterprise Ethernet switch.

The 30A base switch has built-in QoS, and gives strict priority to the traffic of the IP Telephone traffic on the uplink port. That is, when the IP Telephone and the personal computer are both transmitting, the telephone traffic is given strict priority out the uplink port to the enterprise Ethernet switch. This is not an issue for the computer, because under normal conditions the IP Telephone transmits less than 100 Kbps of audio traffic. Prioritization of traffic downstream from the enterprise Ethernet switch to the 30A base switch must be handled by the enterprise Ethernet switch.

Sequence of operation

The following are key boot-up events, listed in order, that can help to verify proper operation of the IP Telephone. This list includes only key events, and might not be comprehensive. Note also that the telephone may go blank between events. In such cases, wait a few seconds or more for an indication from the telephone as to what event is taking place.

1. Initial startup. At power-up or manual reset, the telephone goes through a short initial startup procedure. The display shows Restarting… (if the telephone was intentionally restarted with Hold RESET#), and then Loading… and Starting…
2. DHCP. The telephone queries the DHCP server for an IP address and other needed information. The following packets are transmitted: DHCP Discover from telephone to broadcast; DHCP Offer from server to broadcast, or relay agent to telephone; DHCP Request from telephone to broadcast; and DHCP ACK from server to broadcast, or relay agent to telephone. Note that this step is bypassed if the telephone is manually configured with all the necessary information.

   Also note that a protocol analyzer that is attached to the PC port of an Avaya IP Telephone with a:

   - Switch (4602SW, 4620), sees only broadcast packets.
   - Hub (4606, 4612, 4624, 4630), sees all packets.
3. TFTP ping. The telephone pings the TFTP server for verification purposes.
4. Request file “46XXUPGRADE.SCR” (all caps) and others from TFTP server. This text script file tells the telephone what boot code (“bbla0_##.bin”) and application code (“def##r#_##.bin”) are needed. If the telephone does not have the current codes, it requests them from the TFTP server. A brand new telephone makes all three requests, because telephones come from the factory with no code, or outdated code. When captured using a protocol analyzer, all three requests show up as intuitive TFTP messages that reveal the file name that is being requested or transferred. Note that there is a loading period after each code is received for the first time. Note also that the file names are case sensitive on some servers (Unix), and not on others (Microsoft).
5. Ext and Password prompts. The telephone prompts for the extension and the password if there are no previously stored values.

6. Registration. The telephone registers with a media controller after the codes are successfully loaded. This registration happens very quickly, and does not show up on the display. However, the following packets can be captured using a protocol analyzer: RAS-Gatekeeper Request (GRQ) from the telephone to the media controller; RAS-Gatekeeper Confirm (GCF) from the media controller to the telephone; RAS-Registration Request (RRQ) from the telephone to the media controller (not necessarily the same one that the GRQ was sent to); and RAS-Registration Confirm (RCF) from the media controller to the telephone.
7. Telephone is operational. The administered display shows up on the telephone, and the extension LED illuminates.
8. Keepalive messages. These messages are sent by each telephone to the media controller at time intervals that are determined by Avaya Communication Manager, and based on the number of registered sets. On a protocol analyzer, the keepalive message shows up as a RAS-Registration Request (RRQ) message with the keepalive bit set in the RAS header. Each request message is answered by the media controller with a RAS-Registration Confirm (RCF) message.
9. Unregistration messages. If the Avaya Media Server intentionally unregisters a set, or if the set intentionally unregisters itself, the message sent by either the media controller or the set is a RAS-Unregistration Request (URQ). The acknowledgment message is RAS-Unregistration Confirm (UCF). All unregistration requests should be confirmed. Future releases will include various URQ types, whereas currently there is only one type.

Connecting a personal computer to an IP Telephone

On the back of the IP Telephone, the port with the icon that looks like a terminal is the user port. (The port with the icon that looks like a network jack is the uplink port, which connects to the Ethernet switch.) Use discretion when connecting a personal computer to the telephone, and remember that its primary function is not that of an enterprise network device. For example, do not connect an enterprise server to the telephone. Such high-traffic servers require their own separate connections to the enterprise Ethernet switch. Also, do not connect a personal computer to the telephone at 10 mbps if that computer routinely runs high-volume transactions. The telephone itself operates well at 10 mbps, and the computer itself may also operate adequately at 10 mbps. But the two combined can cause the computer to overwhelm the 10-mbps link at the expense of audio quality. Connecting a user computer to the telephone at 100 mbps works very well.

An IP Telephone and an attached PC on different VLANs

The third scenario for attaching a PC to the telephone (the first two were covered in the previous subsection) is to have the telephone and the PC on separate VLANs. This requires a trunk port, or some other multi-VLAN port, on the Ethernet switch. One of the VLANs is the port/native VLAN, and the clear Ethernet frames (ones with no 802.1Q tag) from the PC reside on this VLAN. The IP Telephone must tag its traffic with the ID of the VLAN to which it belongs. The Hold QOS# options are exactly the same as described in the previous section, except that now the VID must not be zero. The Layer 2 and Layer 3 priority options may or may not be implemented. The *IP Telephony Implementation Guide* contains more detail about how to implement this third scenario.

An IP Telephone and an attached PC on the same VLAN

Three variations exist for attaching a personal computer to the telephone. The first two involve having both the telephone and the computer on the same VLAN, which is the port VLAN or the native VLAN. Refer to the I*P Telephony Implementation Guide* for a primer on VLANs. In the first scenario, traffic from both the telephone and the PC have no CoS tagging. In this case, no special configurations are necessary. Just attach the telephone to an access port (one with only the port VLAN or the native VLAN configured), and attach the computer to the telephone.

The second scenario is similar to the first, except that traffic from the telephone is tagged with Layer 2/Layer 3 priority while remaining on the port VLAN or the native VLAN. The telephone must be configured to tag its Ethernet frames and/or IP packets with the desired priority. This is normally set by DHCP or TFTP by specifying the following parameters:

- 802.1Q. On/off for 802.1Q tagging. Turn this on if Layer 2 priority tagging is desired. Otherwise, turn this off.
- L2 audio. Layer 2 CoS tag for Ethernet frames that contain audio packets. Set this to a value between 0 and 7. This value is sent to the telephone by Avaya Communication Manager, as configured on the IP Network Region form.
- L2 signaling. Layer 2 CoS tag for Ethernet frames that contain signaling packets. Set this to a value between 0 and 7. This value is sent to the telephone by Avaya Communication Manager, as configured on the IP Network Region form.
- LAN ID. Must be set to zero (0) for this scenario. A VID of zero indicates that the Ethernet frame belongs on the port VLAN or native VLAN. The VID has no effect when 802.1Q tagging is disabled. Cajun switches require no special configuration for this scenario. Cisco switches, however, behave differently for different hardware platforms and OS versions. Table 46, Cisco hardware characteristics, on page 199 shows Avaya laboratory test results on a sample of hardware platforms and OS versions.
- L3 audio. Layer 3 DSCP for audio IP packets. Set this to a value between 0 and 63. This value is sent to the telephone by Avaya Communication Manager, as configured on the IP Network Region form.
- L3 signaling. Layer 3 DSCP for signaling IP packets. Set this to a value between 0 and 63. This value is sent to the telephone by Avaya Communication Manager, as configured on the IP Network Region form.

Remember that for the CoS tags to have any effect, the corresponding QoS configurations must be implemented on the necessary network devices. Remember also that improperly enabling Layer 2 and Layer 3 tagging can break processes that were working without tagging. Quality of Service guidelines contains more information on CoS and QoS.

DHCP and TFTP

Dynamic Host Configuration Protocol (DHCP) provides a way to assign configuration parameters to clients on a TCP/IP network automatically. This minimizes the maintenance of a network of 4600 Series IP Telephones by removing the need to assign and maintain IP addresses and other parameters for each IP Telephone on the network individually.

Trivial File Transfer Protocol (TFTP) provides a way to transfer files that does not require user intervention. TFTP is used by Avaya IP Telephones to download their configuration files, and the latest firmware.

Software checklist

Ensure that you own licenses to install and use the DHCP server software and the TFTP server software.

WARNING:
The circuitry in the 4600 Series IP Telephones reserves IP addresses of the form 192.168.2.x for internal communications. The telephones do not properly use the addresses you specify if the addresses are from that range.

Required network information

DHCP is the control point where an enterprise controls its IP Telephones. Before administering DHCP and TFTP, complete the information that is outlined below to ensure that you have the necessary information regarding your network. There can be more than one gateway, TFTP server, subnet mask, and C-LAN in your configuration. You need a copy of this table for each DHCP server.

Release 1.5 and later of the 4600 Series telephones supports the ability to specify a list of IP addresses for a gateway/router, TFTP server, and one or more C-LAN circuit packs. Each list can contain up to 127 total ASCII characters, with IP addresses that are separated by commas with no intervening spaces. When you specify IP addresses for the TFTP server or call server, you can use either dotted decimal format (“xxx.xxx.xxx.xxx”) or DNS names to identify the addresses. If you use DNS, note that the system value DOMAIN is appended to the IP addresses that you specify. If DOMAIN is null, the DNS names must be fully qualified, in accordance with IETF RFCs 1034 and 1035. For more information about DNS, see Appendix D. DHCP / TFTP and the *IP Telephone LAN Administrators Guide*.

You can install both the DHCP server and the TFTP server on the same machine.

Before installing each DHCP server, obtain the following required network information:

- Gateway/router IP addresses
- TFTP server IP addresses
- Subnet mask
- Media controller (C-LAN circuit pack) IP addresses
- Media controller (C-LAN circuit pack) port
- TFTP server path
- Telephone IP address range
  - From:
  - To:
- DNS server addresses

The TFTP server file path is the “root” directory that is used for all transfers by the server. This is the default directory which all files will be uploaded to, or downloaded from. In configurations where the upgrade script and the application files are in the default directory, TFTP server path should not be used.

Avaya can use a special option, Option 176, to pass these values. Avaya has done significant testing of and had good success with Option 176 on the Microsoft Windows 2000 DHCP server and the ISC DHCP server (common on Linux and Unix platforms). Results from other DHCP servers may vary. A typical Option 176 string looks like the following string:

"MCIPADD=#.#.#.#,MCPORT=1719,TFTPSRVR=#.#.#.#,L2Q=1,L2QVLAN=0"

where

- MCIPADD is the IP address of the C-LAN
- MCPORT is the UDP port that is used for telephone registration
- TFTPSRVR is the TFTP server that the telephone uses to look for firmware and configuration upgrades
- L2Q is 802.1Q. 1 is on, 0 is off
- L2QVLAN is the VLAN that the telephone uses. Vlan ID 0 is a special vlan ID that tells the next Layer 2 switch to replace the 0 tag with the native vlan ID on that ingress port.

See Appendix D. DHCP / TFTP for more information.

Powering IP Telephones

Introduction

The Avaya 4600 Series IP Telephones were designed to use flexible powering methods. Some of these powering solutions require the use of special cables that are designed specifically for the Avaya 4600 Series telephones.

Background

To meet the critical needs of the business, two generations, Gen-1 and Gen-2, of the 4606, 4612, and 4624 IP Telephones were developed. The second-generation IP Telephones, Gen-2, added Power over Ethernet (PoE) to the capabilities of the original IP Telephone. Either local or centralized power can be provided to the IP Telephones (4606, 4612, and 4624 models only) by one of the following four methods:

- Power over Spare Pairs pins 4/5 (GRD) and 7/8 (-48 volts) of an RJ45 connector
- Power over Data/Signal Pairs pins 1/2 (-48 volts) and 3/6 (GRD) of an RJ45 connector
- Power over, Traditionally, pins 7 (-48 volts) and 8 (GRD) of an RJ45 connector
- Power through the barrel connector on the bottom of the telephone

Types of IP Telephone power

Centralized power

IEEE, the standards body that governs PoE, has not ratified a final position on PoE. However, a working draft (Rev.3.0) of IEEE 802.3af has been in place to establish guidelines for this area since November 2001, and was updated to Rev.3.2 in September, 2002. With PoE,
(IEEE Draft 802.3af standard), both power and data are carried over one CAT 5 Ethernet cable. Deploying the IP Telephones using PoE eliminates the need for a local power supply, AC adapter, and cables. Thus, power can be provided from the wiring closet or the switch room, where it can be easily connected to a UPS system.

The key technical characteristics of the IEEE Draft 802.3af standard for PoE are:

- Power Sourcing Equipment (PSE) output voltage is 44 VDC to 57 VDC.
- Power Sourcing Equipment (PSE) output current is 350 mA, maximum.
- Power Sourcing Equipment (PSE) power is 15.4 watts, maximum.
- Powered Device (PD) power draw allowed is 12.95 watts, maximum.
- Powered Device (PD) is ready to accept power from either set of pairs:
  - Spare Pairs (pins 4/5 and 7/8)
  - Signal/Data Pairs (pins 1/2 and 3/6)
- The method of signature detection is the “Resistor” concept.
- Mid-Span supplies power on the Spare Pairs (pins 4/5 and 7/8).
- End-Span supplies power on either the Signal/Data Pairs (pins 1/2 and 3/6), or the Spare Pairs (pins 4/5 and 7/8).
- The power detection and the power feed operate on the same set of pairs.

For more information on the IEEE Draft 802.3af standard (technically known as “DTE Power through MDI Task Force”), see:

Link to IEEE Draft: “DTE Power through MDI Task Force

Local power

Local power is the power that is supplied at the immediate location of the telephone. Local power requires a 120/240 VAC outlet that is located within 6 feet of the telephone. Power is provided to the IP Telephone through a power supply with either a CAT 5 LAN cable or a barrel connector or a special split cord. Each power supply has a different power range in which it can operate.

Configuring the IP Telephones for power

The Avaya 4600 Series IP Telephones are comprised of the following models:

- 4602 IP Telephone (no barrel connector, one RJ45 jack, no switch, no hub)
- 4606 IP Telephone (barrel connector, two RJ45 jacks, built-in hub)
- 4612 IP Telephone (barrel connector, two RJ45 jacks, built-in hub)
- 4620 IP Telephone (no barrel connector, two RJ45 jacks, built-in switch)
- 4624 IP Telephone (barrel connector, two RJ45 jacks, built-in hub)
- 4630 IP Telephone/Screenphone (barrel connector, 2 RJ45 jacks, built-in hub)

Model 4630

The 4630 IP Telephone is popularly known as the *IP Screenphone*. Currently, the IP Screenphone consumes more power than the IEEE limits, and therefore requires local power. The IP Screenphone must be powered locally through the barrel connector on the bottom of the telephone using the power supply that is provided with the unit. When using the 30A Switch with the IP Screenphone, a special split cord is required.

Models 4602 and 4620

The 4602 and the 4620 IP Telephones accept power only through the RJ45 Jack on the telephone, using either of the following types of power supplies:

- **Centralized power supply**
  - — Avaya P333T-PWR switch
  - — 1152A1 Mid-Span Power Distribution Unit
- **Local power supply**
  - — 1151B1 “Desktop” Power Supply
  - — 1151B2 “Desktop” Power Supply with Battery Backup

Models 4606, 4612, and 4624

The following section discusses the ways in which local or centralized power can be applied to the IP 4606, 4612, and 4624 telephones.

> **NOTE:**
> Legacy Power is the least preferred method for powering these telephones.

Centralized power For centralized power (PoE), use either an Avaya P333T-PWR switch for new or *Greenfield* installations, or an 1152A1 Mid-Span Power Distribution Unit for legacy systems. The Avaya P333T-PWR switch with Power over LAN capability can use either the data (pins 1/2 and 3/6) or spare (pins 4/5 and 7/8) pairs for power feeding. A Category 5 Ethernet cable from an Ethernet LAN Switch carrying DATA is connected to the “data” port of the 1152A1 Mid-Span Power Distribution Unit (PDU). Power is then injected from the “data & power” port over the spare pair (pins 4/5 and 7/8) on a Category 5 Ethernet cable that is connected to the IP Telephone.

Local power. The 1151B1 and 1151B2 switching power supplies are the preferred global solution for local power, and replace the 1151A1 and 1151A2 units, respectively. In addition to being the preferred solution, the 1151B1 and 1151B2 local power supply units eliminate the need for any special split cord.

Legacy power. Many existing IP Telephone installations within the United States and Canada use the IP Phone Aux power supply with a barrel connector. This local power supply is also known as a “leader” power supply or a wall power supply. Power is supplied through the barrel connector to a jack on the bottom of the telephone. The 1151A1 or 1151A2 power supply (commonly referred to as a *brick transformer*) is an alternative method of powering that requires one or more special split cords. Many existing IP Telephone installations in international regions use the 1151A1 or 1151A2 power supply with a required special split cord. If a 30A switch, three-port switched hub is used with the 1151A1 or the 1151A2 local power supply, two special cords are required (these cables are included with the 30A switch). The 30A switch is applicable to the 4624 and 4612 models, but not the 4606 because of its small footprint (the 30A switch does not fit in the base of 4606 model). The 1145B Bulk Power Supply is an existing prestandard solution (that is, the 1145B power supply was developed before the IEEE Draft 802.3af specifications) that provides centralized power over pins 7 and 8 using a special cord with over-current protection, and alarm LEDs for each output. The 1145B uses a locking station cord to guard against damage, compared to the signature detection method of the IEEE Draft 802.3af specifications.

WAN

This section covers such WAN issues as

- Serialization delay
- Protocols
- Point-to-point versus multipoint
- Queuing

General guidelines

Because of the high costs and lower bandwidths available, there are some fundamental differences in running IP Telephony over a WAN versus a LAN. Because of the resource scarcity, it is important to consider network optimizations and proper network design, because problems are more likely to manifest themselves in a WAN environment.

QoS

In particular, QoS becomes more important in a WAN environment than in a LAN. In many cases, transitioning from the LAN to the WAN reduces bandwidth by approximately 99%. Because of this severe bandwidth crunch, strong queuing, buffering, and packet loss management techniques have been developed. These are covered in more detail in the Quality of Service guidelines chapter.

Recommendations for QOS

In general, for the WAN, Avaya recommends tagging IP Telephony bearer and signaling packets with DiffServ Code Point (DSCP) 46 (Expedited Forwarding). This tagging can be administered in Avaya IP Telephones, Communication Manager, and circuit packs. At the routers, Avaya recommends using strict priority queuing for voice packets, and weighted-fair queuing for data packets. Voice packets should always get priority over non-network-control data packets. This type of queuing is called Class-Based Queuing (CBQ) on Avaya data networking products, or Low-Latency Queuing (LLQ) on Cisco routers.

Codec selection and compression

Because of the limited bandwidth that is available on the WAN, using a compressed codec allows much more efficient use of resources without a significant decrease in voice quality. Avaya recommends that IP Telephony implementations across a WAN use the G.729 codec with 20-ms packets. This configuration uses 24 Kbps (excluding Layer 2 overhead), 30% of the bandwidth of the G.711 uncompressed codec (80 Kbps). For more information on bandwidth, see IP Bandwidth and Call Admission Control on page 133.

To conserve even more bandwidth, RTP header compression (cRTP) can be used on point-to-point links. cRTP reduces the IP/UDP/RTP overhead from 40 bytes to 4 bytes. With 20-ms packets, this translates to a savings of 14.4 Kbps, making the total bandwidth required for G.729 approximately 9.6 Kbps. The trade-off for cRTP is higher CPU utilization on the router. The processing power of the router determines the amount of compressed RTP traffic that the router can handle. Avaya testing indicates that a typical small branch-office router can handle 768 Kbps of compressed traffic. Larger routers can handle greater amounts. cRTP is available on Avaya and Cisco routers.

Serialization delay

Serialization delay refers to the delay that is associated with sending bits across a physical medium. Serialization delay is important to IP Telephony because this delay can add significant jitter to voice packets, and thus impair voice quality. See Layer 3 QoS on page 220 for techniques to minimize serialization delay.

Network design

Routing protocols and convergence

When designing a IP Telephony network across a WAN, some care should be taken when selecting a routing protocol or a dial-backup solution. Different routing protocols have different convergence times, which is the time that it takes to detect a failure and route around it. While a network is in the process of converging, all voice traffic is lost. Routing protocol convergence is covered in more detail in section 3.5.

The selection of a routing protocol depends on several factors:

- If a network has a single path to other networks, static routes are sufficient.
- If multiple paths exist, is convergence time an issue? If so, EIGRP and OSPF are appropriate.
- Are open standards-based protocols required? If so, OSPF and RIP are appropriate, but not EIGRP or IGRP, which are Cisco proprietary.

In general, Avaya recommends the use of OSPF when routing protocols are required. OSPF allows for relatively fast convergence, and does not rely on proprietary protocols.

In many organizations, because of the expense of dedicated WAN circuits, dial-on-demand circuits are provisioned as backup if the primary link fails. The two principal technologies are ISDN (BRI) and analog modem. ISDN dial-up takes approximately 2 seconds to connect, and offers 64 Kbps to 128 Kbps of bandwidth. Analog modems take 60 seconds to connect, and offer up to 56 Kbps of bandwidth. If G.729 is used as the codec, either technology can support IP Telephony traffic. If G.711 is used as the codec, only ISDN is appropriate. Also, because of the difference in connect times, ISDN is the preferred dial-on-demand technology for implementing IP Telephony.

Multipath routing

Many routing protocols, such as OSPF, install multiple routes for a particular destination into a routing table. Many routers attempt to load-balance across the two paths. There are two methods for load balancing across multiple paths. The first method is per-packet load balancing, where each packet is serviced round-robin fashion across the two links. The second method is per-flow load balancing, where all packets in an identified “flow” (source and destination addresses and ports) take the same path. IP Telephony does not operate well over per-packet load-balanced paths. This type of setup often leads to “choppy” quality voice. Avaya recommends that in situations with multiple active paths, per-flow load balancing is preferable to per-packet load balancing. This behavior is enabled by default on Avaya products. On Cisco routers, the command for this is “ip route-cache,” applied per interface.

Frame Relay

The nature of Frame Relay poses somewhat of a challenge for IP Telephony. This section presents an overview of Frame Relay, and then discusses an issue that affects IP Telephony across Frame Relay links.

Overview

Frame Relay service is composed of three elements: the physical access circuit, the Frame, Relay port, and the virtual circuit. The physical access circuit is usually a T1 or fractional T1 and is provided by the local exchange carrier (LEC) between the customer premise and the nearest central office (CO). The Frame Relay port is the physical access into the Frame Relay network, a port on the Frame Relay switch itself.

The access circuit rate and the Frame Relay port rate must match. The virtual circuit is a logical connection between Frame Relay ports that can be provided by the LEC for intra-lata Frame Relay, or by the inter-exchange carrier (IXC) for inter-lata Frame Relay. The most common virtual circuit is a permanent virtual circuit (PVC), which is associated with a committed information rate (CIR). The PVC is identified at each end by a separate data-link connection identifier (DLCI) in Figure 69, Data-link connection identifiers over an interexchange carrier Frame Relay network, on page 212.

Figure 69: Data-link connection identifiers over an interexchange carrier Frame Relay network

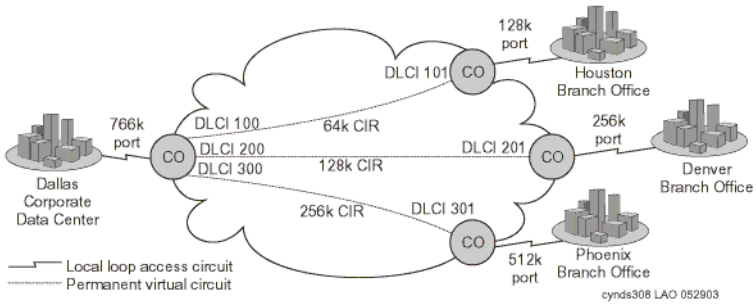

This hypothetical implementation shows the Dallas corporate office connected to three branch offices in a common star topology (or hub and spoke). Each office connects to a LEC CO over a fractional T1 circuit, which terminates onto a Frame Relay port at the CO, and onto a Frame Relay capable router at the customer premise. The port rates and the access circuit rates match. PVCs are provisioned within the Frame Relay network between Dallas and each branch office. The CIR of each PVC is sized so that it is half the respective port rate, which is a common implementation. Each branch office is guaranteed its respective CIR, but it is also allowed to burst up to the port rate without any guarantees.

The port rate at Dallas is not quite double the aggregate CIR, but it does not need to be, because the expectation is that not all three branch offices will burst up to the maximum at the same time. In an implementation like this, the service is probably negotiated through a single vendor. But it is likely that Dallas and Houston are serviced by the same LEC, and that the Frame Relay is intra-lata, even if it was negotiated through an IXC, such as AT&T, WorldCom, or Sprint. The service between Dallas and the other two branch offices, however, is most likely inter-lata.

An issue and alternatives

The obstacle in running IP Telephony over Frame Relay involves the treatment of traffic within the CIR and outside of CIR, commonly termed the “burst range.”

Figure 70: Committed information rate (burst range)

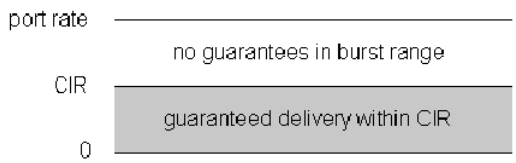

As Figure 70, Committed information rate (burst range), on page 213 shows, traffic up to the CIR is guaranteed, whereas traffic beyond the CIR usually is not. This is how Frame Relay is intended to work. CIR is a committed and reliable rate, whereas burst is a bonus when network conditions permit it without infringing upon the CIR of any user. For this reason, burst frames are marked as discard eligible (DE), and are queued or discarded when network congestion exists. Although experience has shown that customers can achieve significant burst throughput, it is unreliable and unpredictable, and not suitable for real-time applications like IP Telephony.

Therefore, the objective is to prevent voice traffic from entering the burst range and being marked DE. One way to accomplish this is to prohibit bursting by shaping the traffic to the CIR and setting the excess burst size (B_e – determines the burst range) to zero. However, this also prevents data traffic from using the burst range.

Additional Frame Relay information

One interesting piece of knowledge is that most IXCs convert the long-haul delivery of Frame Relay into ATM. That is, the Frame Relay PVC is converted to an ATM PVC at the first Frame Relay switch after leaving the customer premise. It is not converted back to Frame Relay until the last Frame Relay switch before entering the customer premise. This is significant because ATM has built- in Class of Service (CoS). A customer can contract with a carrier to convert the Frame Relay PVC into a constant bit rate (CBR) ATM PVC. ATM CBR cells are delivered with lower latency and higher reliability.

Finally, under the best circumstances, Frame Relay is still inherently more susceptible to delay than ATM or TDM. Therefore, after applying the best possible queuing mechanism, one should still expect more delay over Frame Relay than is present over ATM or TDM.

VPN

Many definitions exist for Virtual Private Networks (VPNs). VPNs refer to encrypted tunnels that carry packetized data between remote sites. VPNs can use private lines, or use the Internet through one or more Internet Service Providers (ISPs). VPNs are implemented in both dedicated hardware and software, but can also be integrated as an application to existing hardware and software packages. A common example of an integrated package is a firewall product that can provide a barrier against unauthorized intrusion, as well as perform the security features that are needed for a VPN session.

The encryption process can take from less than 1 millisecond (ms) to 1 second or more, at each end. Obviously, VPNs can represent a significant source of delay, and therefore have a negative affect on voice performance. Avaya VPN products encrypt traffic with less than 1ms of delay, and thus are appropriate for IP Telephony. Also, because most VPN traffic runs over the Internet and there is little control over QoS parameters for traffic crossing the Internet, voice quality may suffer due to excessive packet loss, delay, and jitter. Users might be able to negotiate a service-level agreement with the VPN provider to guarantee an acceptable level of service. Before implementing IP Telephony with a VPN, users should test their VPN network over time to ensure that it consistently meets the requirements that are specified in the *Avaya IP Voice Quality Network Requirements Document Summary.*

For more information, see:

- IP Voice Quality Network Requirements Website
- IP Voice Quality Document (.PDF)

Convergence advantages

For increasing numbers of enterprises, the VPN carries not only data, but voice communications. Though voice communication over IP networks (IP Telephony) creates new quality of service (QoS) and other challenges for network managers, there are compelling reasons for moving forward with convergence over maintaining a traditional voice and data infrastructure:

- A converged infrastructure makes it easier to deploy eBusiness applications, such as customer care applications, that integrate voice, data, and video.
- Enterprises can reduce network costs by combining disparate network infrastructures, and eliminating duplicate facilities.
- A converged infrastructure can increase the efficiencies of the IT organization.
- Long distance charges can be reduced by sending voice over IP networks.

Voice over IP VPN is emerging as a viable way to achieve these advantages. The emergence of public and virtual private IP services promises to make it easier for customers, suppliers, and businesses to use data networks to carry voice services. As with any powerful new technology, however, VPNs require skilled management to achieve top performance. The highest network performance becomes imperative when the VPN network must deliver high-quality voice communication. Not all IP networks can meet these quality requirements today. For instance, the public Internet is a transport option for voice communication only when reduced voice performance is acceptable, and global reach has the highest priority. When high voice quality is a requirement, ISPs and Network Service Providers (NSPs) can provide other VPN connections that meet required Service Level Agreements (SLAs).

Managing IP Telephony VPN issues

The following issues are among the issues that managers must consider in providing voice communications over IP VPN networks:

- Communications security
- Special requirements for voice communication over VPNs
- Class of Service (CoS) and Quality of Service (QoS) management
- Network Address Translation (NAT) and IP Telephony protocol compatibility
- Models for network management and outsourcing
- Vendor capabilities

This section provides an overview of the considerations managers face in each of these areas.

Communication security

The public nature of the Internet, its reach, and its shared infrastructure provide cost savings when compared to leased lines and private network solutions. However, those factors also contribute to make Internet access a security risk. To reduce these risks, network administrators must use the appropriate security measures.

It is important to note that a managed service can be implemented either as a premises-based solution or a network-based VPN service. A premises-based solution includes customer premises equipment (CPE) that allows end-to-end security and Service Level Agreements (SLAs) that include the local loop. These end-to-end guarantees of quality are key differentiators. A network-based VPN, on the other hand, is provisioned mainly by equipment at the service provider's point-of-presence (PoP), so it does not provide equivalent guarantees over the last mile. For a secure VPN that delivers robust, end-to-end SLAs, an enterprise must demand a premises-based solution that is built on an integrated family of secure VPN platforms.

The "private" in virtual private networking is also a matter of separating and insulating the traffic of each customer traffic so that other parties cannot compromise the confidentiality or the integrity of data. IPSec tunneling and data encryption achieves this insulation by essentially carving private end-to-end pipes or "tunnels" out of the public bandwidth of the Internet, and then encrypting the information within those tunnels to protect against someone else accessing the information. In addition to IPSec, there are two standards for establishing tunnels at Layer 2. These are the Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP), neither of which includes the encryption capabilities of IPSec. The value of IPSec beyond these solutions is that IPSec operates at IP Layer 3. It allows for native, end-to-end secure tunneling and, as an IP-layer service, it also promises to be more scalable than the connection-oriented Layer 2 mechanisms.

Also, note that IPSec can be used with either L2TP or PPTP, since IPSec encrypts the payload that contains the L2TP/PPTP data. Indeed, IPSec provides a highly robust architecture for secure wide-area VPN and remote dial-in services. It is fully complementary to any underlying Layer 2 network architecture, and with its addition of security services that can protect the VPN of a company, IPSec marks the clear transition from early tunneling to full-fledged Internet VPN services.

An issue, however, is the fact that different implementations of IPSec confer varying degrees of security services. Products must be compliant with the latest IPSec drafts, must support high- performance encryption, and must scale to VPNs of industrial size.

Finally, a VPN platform should support a robust system for authentication of the identity of end users, based on industry standard approaches and protocols.

Firewall technologies

To reduce security risks, appropriate network access policies should be defined as part of business strategy. Firewalls can be used to enforce such policies. A firewall is a network interconnection element that polices traffic the flows between internal (protected) networks and external (public) networks such as the Internet. Firewalls can also be used to "segment" internal networks.

The application of firewall technologies only represents a portion of an overall security strategy. Firewall solutions do not guarantee 100% security by themselves. These technologies must be complemented with other security measures, such as user authentication and encryption, to achieve a complete solution.

The three technologies that are most commonly used in firewall products are packet filtering, proxy servers, and hybrid. These technologies operate at different levels of detail, and thus they provide varying degrees of network access protection. That means that these technologies are not mutually exclusive. A firewall product may implement several of these technologies simultaneously.

Network management and outsourcing models

While enterprises acknowledge the critical role that the Internet and IP VPNs can play in their strategic eBusiness initiatives, they face a range of choices for implementing their VPNs. The options range from enterprise-based or "do-it-yourself" VPNs that are fully built, owned, and operated by the enterprise, to VPNs that are fully outsourced to a carrier or other partner. In the near term, it is generally believed that enterprise-operated and managed VPN services will hover around a 50/50 split, including hybrid approaches.

Increasingly, enterprises are assessing their VPN implementation options across a spectrum of enterprise-based, carrier-based/outsourced, or hybrid models. Each approach offers a unique business advantage.

- **Enterprise based**. This option operates over a public network facility (most commonly the Internet) using equipment that is owned and operated by the enterprise. Its greatest benefit to the enterprise is the degree of flexibility and control it offers over VPN deployment, administration, and adaptability or change.
- **Fully outsourced**. This managed service could be implemented by a collection of partners, including an ISP and a security integration partner. Its advantages include quick deployment, easy global scalability, and freedom from overhead network management.
- **Shared management**. With this hybrid approach, a partner can take responsibility for major elements of infrastructure deployment and management, but the enterprise retains control over key aspects of policy definition and security management.

Conclusion

Moving to a multipurpose packet-based VPN that transports both voice and data with high quality poses a number of significant management challenges. Managers must determine whether to operate the network using an enterprise-based model, an outsourced or carrier-based model, or a hybrid model. They must settle security issues that involve several layers of the network. And they must ensure that they and their vendors can achieve the required QoS levels across these complex networks. Yet the advantages of converged, multipurpose VPNs remain a strong attraction. The opportunity to eliminate separate, duplicate networks and costly dedicated facilities, avoid costly public network long distance charges, and reduce administrative overhead provides a powerful incentive. Most important, by helping integrate voice and data communication, multimedia messaging, supplier and customer relationship management, corporate data stores, and other technologies and resources, converged networks promise to become a key enabler for eBusiness initiatives.

NAT

IP Telephony does not work well with networks that use NAT (Network Address Translation) because most NAT implementations do not support H.323 protocols. The destination IP address is encapsulated in more than one header, including the Q.931, H.225, and IP headers. NAT changes only the address in the IP header, which results in a mismatch that prohibits the control of calls. Avaya suggests using a firewall to guard against intruders, but the firewall should not provide NAT functions for IP Telephony packets unless it is Avaya Q.931 friendly. Figure 71, IP Telephony without NAT, on page 217 shows an approved sample implementation of a firewall that uses selective NAT. With Avaya Communication Manager 1.3, products work seamlessly with many static NAT applications, even if they are not H.323 aware.

Figure 71: IP Telephony without NAT

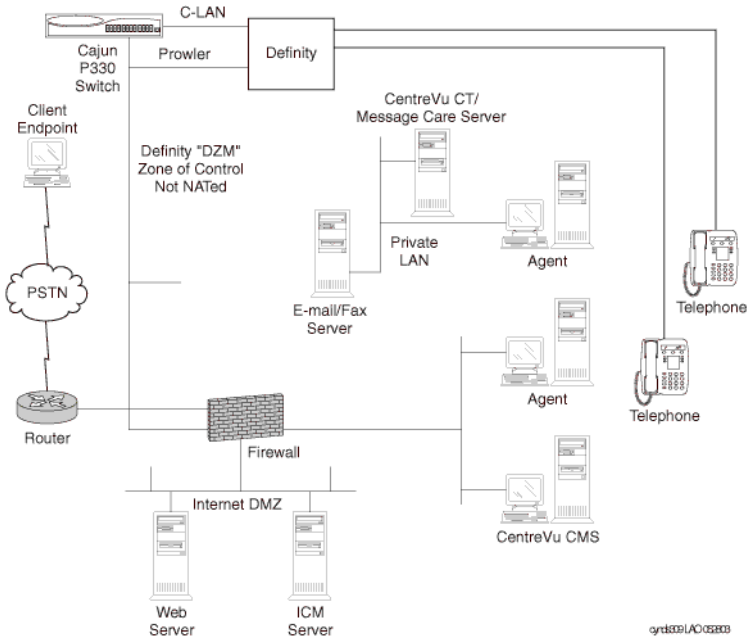

Quality of Service guidelines

This chapter contains guidelines for deploying Quality of Service (QoS) for an IP Telephony network. This chapter begins with an overview of Class of Service (CoS) versus QoS.

Class of Service refers to mechanisms that tags traffic in such a way that the traffic can be differentiated and segregated into various classes. *Quality of Service* refers to what the network does to the tagged traffic to give higher priority to specific classes. If an endpoint tags its traffic with Layer 2 802.1p priority 6 and Layer 3 Differentiated Services Code Point (DSCP) 46, for example, the Ethernet switch must be configured to give priority to value 6, and the router must be configured to give priority to DSCP 46. The fact that certain traffic is tagged with the intent to give it higher priority does not necessarily mean it will receive higher priority. CoS tagging does no good without the supporting QoS mechanisms in the network devices.

CoS

IEEE 802.1p/Q at the Ethernet layer (Layer 2) and DSCP at the IP layer (Layer 3) are two standards-based CoS mechanisms that are used by Avaya products. These mechanisms are supported by the IP Telephone, the S8300 Media Server, and the C-LAN and MedPro circuit packs. Although TCP/UDP source and destination ports are not CoS mechanisms, they can be used to identify specific traffic, and can be used much like CoS tags. Other non-CoS methods to identify specific traffic are to key in on source and destination IP addresses and specific protocols, such as RTP. The MedPro circuit pack and IP Telephones use RTP to encapsulate audio.

Note that the 802.1Q tag changes the size and the format of the Ethernet frames. Because of this, many switches must be explicitly configured to accept 802.1Q tagged frames. Otherwise, these switches might reject the tagged frames. The two fields to be concerned with are the Priority and Vlan ID (VID) fields. The Priority field is the “p” in 802.1p/Q, and ranges in value from 0 to 7. (*802.1p/Q* is a common term that is used to indicate that the Priority field in the 802.1Q tag has significance. Prior to real-time applications, 802.1Q was used primarily for VLAN trunking, and the Priority field was not important.) The VID field is used as it always has been, to indicate the VLAN to which the Ethernet frame belongs.

The IP header with its 8-bit Type of Service (ToS) field, which was, and in some cases still is, originally used. This original scheme was not widely used, and the IETF developed a new Layer 3 CoS tagging method for IP called Differentiated Services (DiffServ, RFC 2474/2475). DiffServ uses the first 6 bits of the ToS field, and ranges in value from 0 to 63. Table 47, Comparison of DSCP with original TOS, on page 219 shows the original ToS scheme and DSCP in relation to the 8 bits of the ToS field.

Table 47: Comparison of DSCP with original TOS

8-bit ToS field							
IP precedence bits			**ToS bits**				**0**
0	1	2	3	4	5	6	7
DSCP bits						0	0

Ideally, any DSCP value should map directly to a precedence and traffic parameter combination of the original scheme. This is not always the case, however, and it can cause problems on some older devices.

On any device, new or old, having a nonzero value in the ToS field cannot hurt if the device is not configured to examine the ToS field. The problems arise on some legacy devices when the ToS field is examined, either by default or by enabling QoS. These legacy devices (network and endpoint) might contain code that only implemented the precedence portion of the original ToS scheme, with the remaining bits defaulted to zeros. This means that only DSCP values that are divisible by 8 (XXX000) can map to the original ToS scheme. For example, if an endpoint is tagging with DSCP 40, a legacy network device can be configured to look for precedence 5, because both values show up as 10100000 in the ToS field. However, a DSCP of 46 (101110) cannot be mapped to any precedence value alone. Another snag is if the existing code implemented precedence with only one traffic parameter permitted to be set high. In this case, a DSCP of 46 still does not work, because it requires 2 traffic parameter bits to be set high. When these mismatches occur, the older device mighty reject the DSCP tagged IP packet, or exhibit some other abnormal behavior. Most newer devices support both DSCP and the original ToS scheme.

Layer 2 QoS

On Avaya and Cisco switches, IP Telephony traffic can be assigned to higher priority queues. The number and the sizes of queues and how the queues function are device dependent, and beyond the scope of this document.

However, in general, a fixed number of queues exist, and the queues are usually not configurable. If the queues are configurable, it is typically not recommended. Older or lower end switches commonly have only two queues or none at all. Newer or higher-end switches commonly have four or eight queues, with eight being the maximum because there are only eight Layer 2 priority levels. When configured to do so, the Ethernet switch can identify the high-priority traffic by the 802.1p/Q tag, and assign that traffic to a high-priority queue. On some switches, a specific port can be designated as a high-priority port, which causes all traffic that originates from that port to be assigned to a high-priority queue.

Layer 3 QoS

It is usually more complicated to implement QoS on a router than on an Ethernet switch. Unlike Ethernet switches, routers do not just have a fixed number of queues. Instead, routers have various queuing mechanisms. For example, Cisco routers have standard first-in first-out queuing (FIFO), weighted fair queuing (WFQ), custom queuing (CQ), priority queuing (PQ), and low-latency queuing (LLQ). LLQ is a combination of priority queuing and class-based weighted fair queuing (CBWFQ), and it is Cisco's recommended queuing mechanism for real-time applications such as IP Telephony. Each queuing mechanism behaves differently, is configured differently, and has its own set of queues.

First, the desired traffic must be identified using DSCP, IP address, TCP/UDP port, or protocol. Then the traffic must be assigned to a queue in one of the queuing mechanisms. Then the queuing mechanism must be applied to an interface.

The interface itself might also require additional modifications, independent of the queuing mechanism, to make QoS work properly. For example, Cisco requires traffic shaping on Frame Relay and ATM links to help ensure that voice traffic is allotted the committed or guaranteed bandwidth. Cisco also recommends link fragmentation and interleaving (LFI) on WAN links below 768 kbps, to reduce serialization delay. Serialization delay is the delay that is incurred in encapsulating a packet and transmitting it out the serial interface. It increases with packet size, but decreases with WAN link size. The concern is that large low-priority packets induce additional delay and jitter, even with QoS enabled. This is overcome by fragmenting the large low-priority packets and interleaving them with the small high-priority packets, thus reducing the wait time for the high-priority packets. Table 48, Serialization delay matrix, on page 221 lists serialization delay for a variety of packet sizes and line speeds. The formula for determining serialization delay is:

$$\text{Serialization delay} = \frac{\text{Packet size (bits)}}{\text{Line speed}}$$

Table 48: Serialization delay matrix

	Packet size					
WAN line speed	**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 µs	1.28 ms	2.56 ms	5.12 ms	10.24 ms	15 ms

Because of all these configuration variables, properly implementing QoS on a router is no trivial task. However, QoS is needed most on the router where, because most WAN circuits terminate on routers.

QoS guidelines

There is no all-inclusive rule regarding the implementation of QoS because all networks and their traffic characteristics are unique. It is good practice to baseline the IP Telephony response on a network without QoS, and then apply QoS as necessary. Avaya Network Consulting Services can help with baselining services. Conversely, it is bad practice to enable multiple QoS features simultaneously, not knowing what effects, if any, each feature is introducing.

Generally, for newer network equipment, best practices involve enabling Layer 3 (DiffServ) QoS on WAN links traversed by voice. Tag voice and data with DiffServ Code Point 46 (Expedited Forwarding), and set up a strict priority queue for voice. If voice quality is still not acceptable, or if QoS is desired for contingencies such as unexpected traffic storms, QoS can then be implemented on the LAN segments as necessary.

There is one caution to keep in mind about QoS with regard to the processor load on network devices. Simple routing and switching technologies have been around for many years and have advanced significantly. Packet forwarding at Layer 2 and Layer 3 is commonly done in hardware (Cisco calls this fast switching, with “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 process, it inherently requires more processing resources from the network device. Many of the newer devices can handle this additional processing in hardware, and maintain speed without a significant processor burden. However, to implement QoS, some devices must take a hardware process and move it to software (Cisco calls this process “switching”). 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 policies are implemented on WAN links, the following general points for Cisco routers are offered to increase the level of confidence that QoS remains in hardware (consult Cisco to be sure):

- Newer hardware platforms are required: 2600, 3600, 7200, and 7500
- Newer interface modules (WIC, VIP, and so on) are required. Consult Cisco to determine which hardware revision is required for any given module.
- Sufficient memory is required: device dependent.
- Newer IOS is required: 12.0 or later.

Avaya Layer 3 switches and the X330 WAN module support both 802.1 p/Q and DiffServ QoS.

Several things should be examined whenever QoS is enabled on a network device. First, the network administrator should examine the processor load on the device, and compare it to levels before QoS was enabled. It is likely that the levels will have gone up, but the increase should not be significant. If it is, then it is likely that the QoS process is being done by software. Also, the processor load must remain at a manageable level (50% average, 80% peak). If the processor load is manageable, then the IP Telephony response (for example, voice quality) should be checked to verify that it has improved under stressed conditions (for example, high congestion). If the IP Telephony response has improved, the other applications should be checked to verify that their performances have not degraded to unacceptable levels.

IEEE 802.1 p/Q

Figure 72: 802.1Q tag

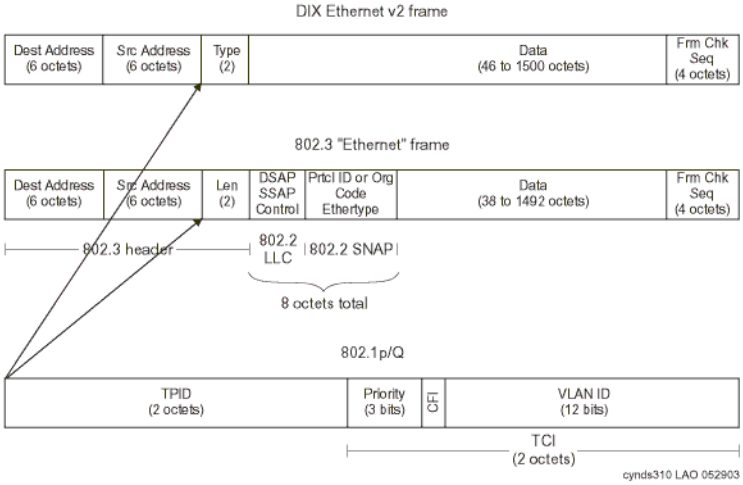

The IEEE 802.1Q standard is a Layer 2 tagging method that adds 4 bytes to the Layer 2 Ethernet header. IEEE 802.1Q defines the open standard for VLAN tagging. Two bytes house 12 bits that are used to tag each frame with a VLAN identification number. The IEEE 802.1p standard uses 3 of the remaining bits in the 802.1Q header to assign one of 8 different classes of service. Communication Manager users can add the 802.1Q bytes and set the priority bits as desired. *Avaya suggests that a priority of 6 be used for both voice and signaling.* The Avaya line of data switches can switch frames with or without these VLAN headers, with no configuration time spent. IEEE 802.1p and IEEE 802.1Q are OSI layer 2 solutions, and work on frames.

Because 802.1Q is a Layer 2 (Ethernet) standard, it only applies to the Ethernet header. At every Layer 3 boundary (router hop), the Layer 2 header, including CoS parameters, is stripped and replaced with a new header for the next link. Thus, 802.1Q does not enable end-to-end QoS.

Recommendations for end-to-end QoS

When end-to-end QoS is desired, Avaya recommends using DiffServ, a Layer 3 CoS method. Modern can map DiffServ Code Points (DSCP) to 802.1p priority values, so 802.1p tags can be recreated on each Ethernet link. This functionality is supported in Avaya Layer 3 switches, and the X330 WAN module.

IEEE 802.1p states a standard according to which these bits are used for CoS. The precedence is listed in Table 49, IEEE 802.1 precedence and service mapping, on page 224.

Table 49: IEEE 802.1 precedence and service mapping

User priority	Service mapping
000	Default, assumed to be best effort
001	Reserved, less than best effort
010	Reserved
011	Reserved
100	Delay sensitive, no bound
101	Delay sensitive, 100 ms bound
110	Delay sensitive, 10 ms bound
111	Network control

DiffServ

The Differentiated Services (DiffServ) prioritization scheme redefines the existing TOS byte in the IP header (Figure 73, Differentiated Services (DiffServ) TOS byte, on page 224) by combining the first 6 bits into 64 possible combinations. This use of the TOS byte is still evolving, but can be used now by Communication Manager, IP Telephones, and other network elements such as routers and switches in the LAN and WAN.

Figure 73: Differentiated Services (DiffServ) TOS byte

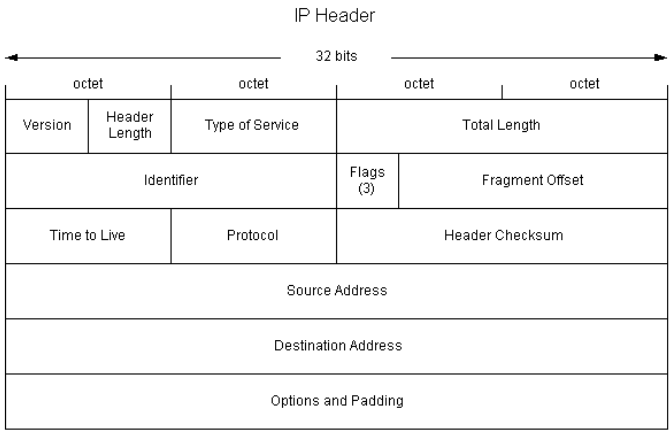

A DiffServ Code Point (DSCP) of 46 (101110), referred to as expedited forwarding (EF), is suggested for the proper treatment of voice and signaling packets. However, with Communication Manager, one can set any DSCP value as desired to work with a company's QoS scheme. Some common DiffServ Code Points are defined in RFCs 2474 and 2475. Although DSCPs are specified in IETF RFCs, the treatment of packets that are tagged with DiffServe is implementation- dependent.

Note that older routers might require a DSCP setting of 40 (101000), which is backward compatible to the original TOS byte definition of critical. But again, Avaya products and software allows users to set any of the 64 possible DSCP values to work with your voice quality policy. The TOS byte is an OSI model Layer 3 solution, and works on IP packets on the LAN and possibly the WAN, depending upon the service provider.

Table 50: Original TOS specification

Bit description	Value	Use
Bits 0-2IP precedence	000	Routine
	001	Priority
	010	Immediate
	011	Flash
	100	Flash Override
	101	CRITIC/ECP
	110	Internetwork control
	111	Network control
Bit 3 delay	0	Normal
	1	Low
Bit 4 Throughput	0	Normal
	1	High
Bit 5 reliability	0	Normal
	1	High
Bit 6 monetary cost	0	Normal
	1	Low
Bit 7 reserved		Always set to 0

RSVP

Resource Reservation Protocol (RSVP) is a protocol that hosts can use to request specific QoS parameters through the network for a particular application data stream. A host can request guaranteed service through a network. If all routers have RSVP support enabled, and if there exists sufficient unreserved bandwidth, a reservation is established throughout the network. If insufficient bandwidth exists, the reservation fails and notifies the hosts. At that point, hosts can choose to send traffic without a reservation, or drop the connection.

RSVP is supported in Communication Manager beginning with Release 1.3. RSVP can be enabled per network region on the network region form. If RSVP is enabled, endpoints including IP Telephones and media processors attempt to establish a reservation for each call. If the reservation fails, Avaya endpoints still try to place a call, but lower the DiffServ priority of the call to the better-than-best-effort (BBE) DSCP that is defined on the network region form. By default, this value is 43.

If RSVP is enabled on a network region, it is very important that it also be enabled on associated routers. If not, all RSVP reservations fail, and all voice traffic in that region is marked with the BBE DSCP, which will generally receive degraded service versus the EF (DSCP 46) DiffServ Code Point.

Queuing methods

This section discusses common queuing methods and their appropriateness for voice.

WFQ

Weighted fair queuing (WFQ) is similar to first in, first out (FIFO) queuing, except that it grants a higher weight to small flows, and flows that are marked with higher DiffServ or IP TOS priorities. This queuing strategy does allow smaller (for example, telnet) and higher-priority (for example, IP Telephony) protocols to squeeze in before high-flow (for example, ftp) packets, but does not starve off any traffic. By itself, it is not appropriate for IP Telephony traffic because high-flow traffic can still delay IP Telephony traffic, and cause unacceptable latency and jitter.

PQ

Strict priority queuing (PQ) divides traffic into different queues. These queues are usually high, medium, normal, and low, based on traffic type. This form of queuing services the queues in order of priority, from high to low. If there is a packet in the high-priority queue, it will always be serviced before the queue manager services the lower-priority queues. With priority queuing, however, it is possible to starve out lower-priority flows if sufficient traffic enters the
high-\priority queue. This mechanism works very well for IP Telephony traffic (where IP Telephony bearer and signaling are inserted in the high-priority queue), but might work less well for routine data traffic that is starved out if sufficient high-priority traffic arrives.

Round-robin

Round-robin (sometimes called *custom*) queuing sorts data into queues, and services each queue in order. An administrator manually configures which type of traffic enters each queue, the queue depth, and the amount of bandwidth to allocate to each queue.

Round-robin queuing is not particularly suited to IP Telephony. It does not ensure strict enough priority to voice packets, so they may still wait behind other traffic flows in other queues. Latency and jitter can be at unacceptable levels.

CB-WFQ / LLQ / CBQ

Class-Based Weighted Fair Queuing (CB-WFQ) with Low-Latency Queuing (LLQ), which is sometimes called Class-Based Queuing (CBQ), combines the above-mentioned queuing mechanisms. Generally, there is one strict-priority queue, several round-robin queues, and weighted fair queuing for the remainder. This queuing mechanism works very well for converged networks. IP Telephony bearer and signaling packets receive the priority they need, while there remains an equitable mechanism for distributing remaining bandwidth. In addition, limits can be set on the high-priority queue to prevent it from using more than a specified amount of bandwidth. Bandwidth that is reserved for the high-priority queue will be given to other queues if insufficient traffic enters the high-priority queue.

RED / WRED

Although they are not queuing methods *per se*, Random Early Detection (RED) and Weighted Random Early Detection (WRED) are important queue management techniques. RED and WRED work by randomly discarding packets from a queue. RED takes advantage of the congestion control mechanism of TCP. By randomly dropping packets prior to periods of high congestion, RED causes the packet source to decrease its transmission rate. Assuming that the packet source is using TCP, it will decrease its transmission rate until all the packets reach their destination, which indicates that the congestion is cleared. Some implementations of RED, called Weighted Random Early Detection (WRED), combines the capabilities of the RED algorithm with IP Precedence. This combination provides for preferential traffic handling for higher-priority packets. It can selectively discard lower-priority traffic when the interface begins to get congested, and provide differentiated performance characteristics for different classes of service.

RED and WRED are useful tools for managing “data” traffic, but should not be used for “voice.” Because IP Telephony traffic runs over UDP, because IP Telephony protocols do not retransmit lost packets, and because IP Telephony transmits at a constant rate, the IP Telephony queue should never be configured for WRED. WRED only adds unnecessary packet loss, and consequently reduces voice quality.

Traffic shaping and policing

Traffic shaping is a mechanism to reduce the rate at which data is transmitted over an interface. When people discuss traffic shaping, they are usually referring to the related technology of traffic policing. Policing works by either reprioritizing excess traffic to a lower queue, or discarding it. As with RED, discarding TCP traffic has the effect of throttling the stream by forcing window size to shrink, and decreasing its transmission rate. Because RTP is a fixed-bandwidth application, discarding RTP packets reduces voice quality without altering the transmission rate. Reprioritizing voice traffic removes the strict priority protection that reduces latency and jitter, and offers the highest voice quality. Thus, in most cases, it is beneficial to use the QoS mechanisms listed above, rather than traffic shaping and policing, to offer the highest quality for voice.

Frame Relay traffic shaping

Traffic shaping is important in technologies that implement virtual circuits (VCs), such as Frame Relay or ATM, where the Committed Information Rate (CIR) might be less than the physical speed of the interface, the port speed. In such scenarios, it is possible for traffic to burst above the CIR. Depending on the Service Level Agreement (SLA), a carrier might mark excess traffic as Discard Eligible (DE), and either delay or discard it if congestion is detected within the network of the carrier. This behavior is unacceptable for voice traffic, which must minimize delay and jitter to achieve optimal voice quality. To solve this issue, Frame Relay traffic shaping gives an administrator tools to limit the transmission rate on a Frame Relay virtual circuit to the CIR.

A popular misconception is that voice traffic can be confined to the CIR, while data traffic can be allowed to burst. Unfortunately, that is not how Frame Relay works. There is not a QoS mechanism for Frame Relay that is negotiated between service providers and customers. Service providers view all traffic equally, and mark any packet that exceeds the CIR as DE, even if the packet is high-priority voice. Thus, the only way to guarantee optimal performance for voice traffic is to restrict the traffic rate to the CIR.

On a Cisco router, there are several steps to take to ensure proper handling for voice:

1. Disable Frame Relay adaptive shaping. This technique reduces the CIR in response to backwards explicit congestion notification (BECN) messages from the service provider. Because traffic is being transmitted at the CIR in the first place, it does not need to be throttled.
2. Set cir and mincir to the negotiated CIR. If FRF.12 fragmentation is implemented, reduce the cir and mincir values slightly to account for the fragment headers.
3. Set be, the excess burst rate, to 0
4. Set bc, the committed burst rate, to cir/100. This accounts for at most a 10-ms serialization delay.
5. Apply this map class to an interface, subinterface, or VC.

Thus, the complete configuration for Frame Relay traffic shaping looks like:

```
map-class frame-relay NoBurst
no frame-relay adaptive shaping
frame-relay cir 384000! (for a 384K CIR)
frame-relay mincir 384000
frame-relay be 0
frame-relay bc 3840


interface serial 0
frame-relay class NoBurst
```

Fragmentation

One large cause of delay and jitter across WAN links is serialization delay, or the time that it takes to put a packet on a wire. For example, a 1500-byte FTP packet takes approximately 214 ms to be fed onto a 56-Kbps circuit. For optimal voice performance, the maximum serialization delay should be close to 10 ms. Thus, it can be problematic for a voice packet to wait for a large data packet over a slow circuit. The solution to this problem is to fragment the large data packet into smaller pieces for propagation. If a smaller voice packet comes in, it can be squeezed between the data packet fragments and be transmitted within a short period of time. The sections that follow discuss some of the more common fragmentation techniques.

MTU

The maximum transmission unit (MTU) is the longest packet (in bytes) that can be transmitted by an interface without fragmentation. Reducing the MTU on an interface forces a router to fragment the large packet at the IP level. This allows smaller voice packets to squeeze through in a timelier manner.

The drawback to this method is that it increases overhead and processor occupancy. For every fragment, a new IP header must be generated, which adds 20 bytes of data. If the MTU is 1,500 bytes, the overhead is approximately 1.3%. If the MTU is shortened to 200 bytes, however, the overhead increases to 10%. In addition, shortening the MTU to force fragmentation increases processor utilization on both the router and the end host that needs to reassemble the packet.

For these reasons, shortening the MTU is only recommended as a last resort. The techniques described later in this section are more efficient, and should be used before changing the values of the MTU. When changing the MTU, size it such that the serialization delay is less than or equal to 10 ms. Thus, for a 384-kbps circuit, the MTU should be sized as follows: 384000 bps *0.01 second (10 ms)/8 bits/byte = 480 bytes. As the circuit size diminishes, however, care should be taken to never reduce the MTU below 200 bytes. Below that size, telephony signaling and bearer (voice) packets can also be fragmented, which reduces the link efficiency and degrades voice performance.

LFI

Link Fragmentation and Interleaving (LFI) is an enhancement to Multilink PPP (MLP) that fragments packets at the Layer 2 (PPP) level. Fragmenting at the IP layer, as with MTU reduction, forces the addition of a new 20-byte IP header and an 8-byte PPP header. However, fragmenting at the data link (PPP) layer only forces generation of an 8-byte PPP header, which greatly increases the efficiency of the link.

Avaya recommends use of LFI functionality instead of MTU manipulation when transmitting IP Telephony packets over PPP links. As with MTU, Avaya recommends sizing packets so that the serialization delay is approximately 10 ms or less.

FRF.12

FRF.12 is a Frame Relay standard for fragmentation. It works for Frame Relay in the same way that LFI works for PPP, with similar increases in efficiency over MTU manipulation. When implementing a Frame Relay network, Avaya recommends using FRF.12 for fragmentation, and sizing the fragments so the serialization delay is no more than 10 ms.

cRTP

RTP header compression is a mechanism that reduces the protocol overhead that is associated with IP Telephony audio packets. It is a function of the network, and not a function of the IP Telephony application. Along with the benefits of using RTP header compression, there are also cautions. This section discusses both.

Application perspective

The following table shows the anatomy of a 20-ms G.729 audio packet, which is recommended for use across limited bandwidth WAN links. Notice that two-thirds of the packet is consumed by overhead (IP, UDP, and RTP), and only one-third is used by the actual audio.

IP header	UDP header	RTP header	20 ms of G.729 audio
20 B	8 B	12 B	20 B

It is important to understand that all 20-ms G.729 audio packets, regardless of the vendor, are constructed like this. Not only is the structure of the packet the same, but the method of encoding and decoding the audio itself is also the same. This sameness is what allows an Avaya IP Telephone to communicate directly with a Cisco IP Telephone, or any other IP Telephone, when using matching codecs. The packets from the application perspective are identical.

Network perspective

RTP header compression is a mechanism that routers use to reduce the 40 bytes of protocol overhead to approximately 2 to 4 bytes. Cisco routers use this mechanism, as does the Avaya X330WAN router, which is a module for the P330 chassis. RTP header compression can drastically reduce the IP Telephony bandwidth consumption on a WAN link when using 20-ms G.729 audio. When the combined 40-byte header is reduced to 4 bytes, the total IP packet size is reduced by 60% (from 60 bytes to 24 bytes). This equates to reducing the total IP Telephony WAN bandwidth consumption by roughly half, and it applies to all 20-ms G.729 audio packets, regardless of the vendor.

Recommendations for RTP header compression

Enterprises that deploy routers that are capable of this feature might be able to benefit from it. However, Cisco recommends caution in using RTP header compression on its routers because it can significantly tax the processor if the compression is done in software. Depending on the processor load before compression, enabling RTP header compression can significantly slow down the router, or cause the router to stop completely. For best results, use a hardware/IOS/interface module combination that permits the compression to be done in hardware.

RTP header compression has to function with exactness or it will disrupt audio. If for any reason the compression at one end of the WAN link and decompression at the other end do not function properly, the result can be intermittent loss of audio or one-way audio. This has been very difficult to quantify, but there is some anecdotal evidence that cRTP sometimes leads to voice-quality issues. One production site in particular experienced intermittent one-way audio, the cause of which was garbled RTP audio samples inserted by the cRTP device. When, for experimentation purposes, RTP header compression was disabled, the audio problems went away.

The test

This section details the results of a simple RTP header compression test that was conducted in a laboratory environment. Although this test was conducted using Cisco routers, the expected behavior is the same for any router that performs this function as specified in RFC 2508. This test was performed in the laboratory configuration that is shown in Figure 74, Equipment configuration for RTP header compression test, on page 230.

Figure 74: Equipment configuration for RTP header compression test

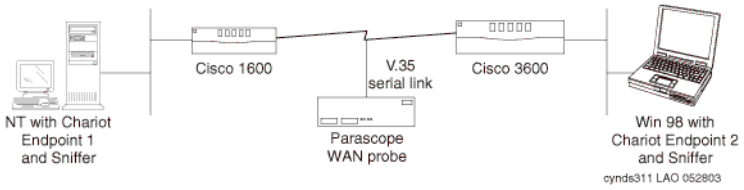

In Figure 74, Equipment configuration for RTP header compression test, on page 230:

- NetIQ Chariot v4.0 was used to simulate IP Telephony calls between the two endpoints. Chariot v4.0 accurately simulates the characteristics of various codecs, and uses a 40-byte IP/UDP/RTP header.
- Sniffer Pro v3.50.02 was used to capture the sent and received packets.
- The Cisco 3600 had IOS v12.1(2)T, and the Cisco 1600 had IOS v12.0(12).
- The Frederick Engineering Parascope WAN probe was tapped into the V.35 serial link to take bandwidth measurements.
- This test was performed using PPP encapsulation on the WAN link.

A single call was placed between the Chariot endpoints using various codecs, all sending 20-ms voice packets. Table 51, Test call (20ms-packets) results, on page 231 shows the results with and without RTP header compression. *Note that these are rough measurements*.

Table 51: Test call (20ms-packets) results

Codec	Payload bytes per packet	Packets per second	Avg WAN BW consumption (kbps)		% reduction
			without compression	with compression	
G.711 (64 kbps)	160	50	84	68.5	~18%
G.729A (8 kbps)	20	50	27.5	13	~53%
G.723.1 (5.3 kbps)	20	33	18	9	~50%
G.723.1 (6.3 kbps)	24	33	19	10	~47%

For each codec, there was an attempt to verify that the audio packets were received intact. This was done by spot checking the audio packets before and after compression, using two Sniffer protocol analyzers. For every codec except G.711, the RTP header and payload were identical before and after compression. With G.711, however, the received packets had the PADDING flag set in the RTP header, although the flag was not set when the packets were transmitted. The PADDING flag indicates the presence of padding octets at the end of the RTP payload, which cannot be true for G.711.

Configuration

To configure RTP header compression on a Cisco router,

1. Specify the number of RTP connections that can be compressed (cache allocation). In interface configuration mode, the command is **ip rtp compression-connections <*number*>**, where
   - The default for **<*number*>** is 32, and each call requires two connections.
   - The configurable range is 3 to 256 for PPP and HDLC using IOS v11.3 and later.
   - The configurable range is 3 to 1000 for PPP and HDLC using IOS v12.0(7)T and later.
   - For Frame Relay, the value is fixed at 256.
2. The command to turn on compression is **ip rtp header-compression** in interface configuration mode. It must be implemented at both ends of the WAN link. When the command was entered into the router, ip tcp header-compression was also installed automatically. When either command was removed, the other was automatically removed.

See the Cisco documentation for more specific configurations on other types of WAN links (that is, Frame Relay and ATM). Configuration for the X330WAN router is very similar to Cisco, and is well documented in the X330WAN User Guides. For this documentation, see the P330 section at:

http://support.avaya.com

Examples

This section contains sample commands for QoS implementation on Avaya products and Cisco products.

Example 1

Cisco router configuration for point-to-point WAN links

There is a three-step process to turn on QoS on a Cisco router:

1. Set up a class map that defines “interesting traffic” to be prioritized.
2. Select a queuing strategy. In this case, use a policy map to set priority. Set up a route map that sets the priority level (critical).
3. Apply the policy map to an interface.

In Figure 75, High-quality service across a congested WAN link, on page 233, set priority-aware Class-Based Weighted Fair Queuing (CB-WFQ) with Low Latency Queuing (LLQ). Although there are more aggressive QoS strategies, they can have a severe impact on data performance. Those other strategies, including Priority Queuing, Custom Queuing, and RSVP, can be implemented at a later date, if conditions warrant. This is a good starting point.

Figure 75, High-quality service across a congested WAN link, on page 233 is used as a reference point. The objective is to assure high quality of service to IP Telephony applications across the congested WAN link.

Figure 75: High-quality service across a congested WAN link

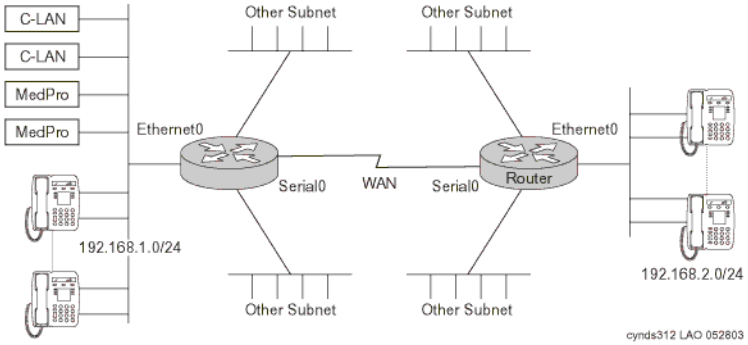

CB-WFQ/LLQ is a priority-aware queuing strategy that has a strict priority queue for voice packets, and does round-robin queuing for other types of traffic. Nonprioritized traffic is still forwarded, however, so this should not interfere with a customer's data network. Use weighted random early detect to manage the fair queue.

The actual router configuration used for this testing follows. First, set the endpoints to tag interesting traffic as DSCP 46. Cisco routers support DiffServ in IOS 12.0 and later. Next, set up a class map to match traffic that is marked with DSCP 46. Once traffic is defined by the class map, set policies for it using a policy map. For the policy map to take effect, it has to be applied to an interface. Queue packets on the outgoing interface. In the sample configuration, 768 K of bandwidth is reserved for RTP. This value should be set at or above the maximum bandwidth to be used for IP Telephony. In our case, 768 K supports 9 calls using G.711, or 31 calls using G.729. This example should work well in most cases using Cisco routers with point-to-point WAN links. Networks that use Frame Relay might need additional steps.

Assumptions for Example 1

Suppose all endpoints are capable of tagging with DSCP 46, which is the default for audio. This would be true in a Communication Manager system with *TN799DP C-LAN circuit packs running firmware v5 or later.* Previous firmware versions and the TN799C circuit pack cannot tag at Layer 2 or Layer 3. A matching set of configurations is applied to both routers.

Administration commands for Example 1

Table 52: Administration commands for Example 1

Command		Meaning
1	**class-map match-any VoIP**	Create a class map called “VoIP.”
2	**match ip dscp 46**	Any packet with DSCP 46 is in the class “VoIP.”
3	**policy-map voipQoS**	Create a policy map called “voipQoS.”
4	**class VoIP priority 768**	Give strict priority to packets in the class “VoIP” on up to 768 k of this WAN link.
5	**class class-default fair-queue**	Put everything else in the default class, and transmit it out the default queue in a fair queue fashion.
6	**random-detect dscp-based**	If the default queue starts to get full, randomly discard packets in this queue based on DSCP. The lower values are discarded first.
7	**interface Serial0 description T1 ip address 172.16.0.1 service-policy output voipQoS**	Apply the “voipQoS policy” outbound on this interface.

Example 2

C-LANS cannot tag their traffic

Assumptions for Example 2

- The C-LANs 192.168.1.10 and .11 cannot tag their traffic (TN-799C or earlier).
- The configuration commands in Table 53, Administration commands for Example 2, on page 235 are applied only to the left router.

Administration commands for Example 2

Table 53: Administration commands for Example 2

Command		Meaning
1	**access-list 101 permit ip host 192.168.1.10 192.168.2.0 0.0.0.255**	The command “access-list 101...” permits any IP traffic from the 2 C-LANs to the 192.168.2.0/24 network. There is an implicit “deny any” at the end of this access list.
2	**access-list 101 permit ip host 192.168.1.11 192.168.2.0 0.0.0.255**	
3	**class-map match-any untaggedVoIP**	Create a class map called “untaggedVoIP.”
4	match access-group 101	Packets that match access list 101 are in the class that is “untaggedVoIP.”
5	**policy-map setDSCP**	Create a policy map called “setDSCP.”
6	**class untaggedVoIP set ip dscp 46**	For all packets in the class “untaggedVoIP,” set the DSCP to 46.
7	**interface Ethernet 0/0 service-policy input setDSCP**	Apply the “setDSCP” policy inbound on this interface.

Example 3

More restrictions on the traffic

Assumptions for Example 3

- DSCP 46 is used throughout to simplify the access list.
- A somewhat matching set of configurations is applied to both routers.

Administration commands for Example 3

Table 54: Administration commands for Example 3

	Command	Meaning
1	**access-list 101 permit ip 192.168.1.0 0.0.0.255 192.168.2.0 0.0.0.255 dscp 46**	Left router
2	**access-list 101 permit ip 192.168.2.0 0.0.0.255 192.168.1.0 0.0.0.255 dscp 46**	Right router The “access list 101...” permits any IP traffic that is tagged with DSCP 46 between the two VoIP subnets. There is an implicit deny any at the end of this access list.
3	**class-map match-any VoIP**	Create a class map called “VoIP.”
4	**match access-group 101**	Only packets matching access list 101 are in the class VoIP; this is more restrictive than matching any packet with DSCP 46 or 34.
		The remainder of the configuration is identical to Example 1.
5	**policy-map voipQoS**	Create a policy map called “voipQoS.”
6	**class VoIP priority 768**	Give strict priority to packets in the class “VoIP” on up to 768k of this WAN link.
7	**class class-default fair-queue**	Put everything else in the default class and transmit it out the default queue in a fair queue fashion.
8	**random-detect dscp-based**	If the default queue starts to get full, randomly discard packets in this queue based on DSCP (lower values get discarded first).
9	**interface Serial0 description T1 ip address 172.16.0.1 service-policy output voipQoS**	Apply the “voipQoS” policy outbound on this interface. If any of the endpoints are incapable of tagging, the “dscp 46” can be removed from access list 101. Then any traffic between the two IP Telephony subnetworks, regardless of the tag, is in the class “VoIP.”

Converged infrastructure LAN switches

P580/P882 family

QoS Process for ingress traffic process

The QoS process for ingress traffic involves the following steps:

1. Identifying the priority, also called *class*, of the frame or the packet. The switch can identify the priority of the frame or the packet by using one or more of the following criteria:
   - The priority of the physical port that on which the switch received the frame or the packet
   - The Cisco ISL tag priority
   - The 802.1p tag priority (default)
   - The source or destination MAC address
   - The DiffServ code point
   - The IP protocol (assigned by means of an ACL rule)
   - The source or destination IP address (assigned by means of an ACL rule)
   - The source or destination TCP port or UDP port (assigned by means of an ACL rule)
2. Storing the frame or the packet in one of eight ingress queues.

   The switch stores the frame or the packet in the queue that matches the priority that was identified in Step 1.
3. 3. Forwarding the frame or the packet from the ingress queue to its destination.

   If you enable policing for the queue, the switch forwards ingress traffic that falls within the maximum bit rate that you set, and drops ingress traffic that exceeds the maximum bit rate.

Example

Assign a priority of 6 to a VoIP flow that is destined to an S8100/G600 telephone switch. Corporate policy requires policing the port that receives the VoIP data to 5 Mbps.

1. Set up an ACL rule that associates a priority of 6 with the destination IP address of the VoIP flow. VoIP traffic cannot tolerate latency or frame loss, so it needs a high priority to ensure its timely delivery.

   NOTE:
   Priority 6 is as an example only. Actual implementations can vary.
2. Enable policing on the port that receives the VoIP flow, and set the guaranteed bit rate to 5 Mbps.
3. The switch stores packets that match the ACL rule in queue 6.

   The switch stores packets in the queue that matches their priority.
4. The switch forwards the VoIP traffic in queue 5 as long as its bit rate does not exceed 5 Mbps. If the bit rate of the queue exceeds 5 Mbps, the switch drops the excess traffic.

Traffic classification

The switch assigns traffic to one of eight queues according to the priority, or class, of the traffic. Priorities range from 0 to 7, with 7 being the highest priority. You can set the switch to classify traffic by the priority assigned to the following characteristics:

- Layer 2
  - Physical port on which the frame or the packet is received
  - Cisco ISL tag or 802.1p tag
  - Source MAC address
  - Destination MAC address
- Layer 3
  - DSCP in the packet
  - New DSCP that replaces the original DSCP, and that you specify.
  - IP protocol. which is assigned by means of an ACL rule
  - Destination IP address, which is assigned by means of an ACL rule
  - Source IP address, which is assigned by means of an ACL rule
- Layer 4
  - Destination TCP or UDP port, which is assigned by means of an ACL rule
  - Source TCP or UDP port, which is assigned by means of an ACL rule

Default priority

By default, the switch uses the priority from the 802.1p tag field, if present, to classify a frame.

If none of the QoS default settings are changed and the frame does not have an 802.1 tag or a Cisco ISL tag, the switch assigns the priority of the physical port to the packet. Each physical port has a default priority of 3.

However, the priority of the 802.1 tag and the Cisco ISL tag take precedence over the priority of the physical port, so the switch uses the priority of the physical port only if:

- No tags are present in the frame.

  or

- You have set the physical port to ignore priorities in tags.

For more information on configuring P580/882 switches for QoS, see:

Link to P580 / 882 User Guide

P330 family

By default, P330 LAN switches accept 802.1Q frames on all ports, and use the 802.1p priority tags. There are two queues within the P330s. The high-priority queue represents 802.1p values 4 to7. The low-priority queue represents 802.1p values 0 to 3. In addition to accepting the values tagged by an endpoint, tagging can be done per port with the command:

set port priority <*high*|*low*>

X330 WAN Module

The new X330WAN versions contain a predefined queue management strategy for IP Telephony that is called CBQ. Use the following procedure to activate CBQ:

Table 55: X330 WAN Module administration commands

	Command	Meaning
1	**set qos policy-source local**	Define DSCP-CoS mapping.
2	**! no external policy source**	
3	**ip access-list-name 100 voice**	
4	**ip access-list-dscp operation 100 34 fw7**	The X330 WAN has four queues with eight behaviors. fw6 and fw7 are different behaviors within the top strict priority queue.
5	**ip access-list-dscp operation 100 46 fw6**	
6	**ip access-list-dscp trust 100 trust-cos-dscp**	Trust packet tagging.
7	**interface FabricFastEthernet 1**	Activate the above mapping on ingress traffic to the Fabric Fast Ethernet interface
8	**ip access-group 100 in**	Apply the ACL to traffic that is arriving on the FabricFastEthernet port.
9	**exit**	?
10	**interface Serial 1**	Activate “VoIP queue management mode” on the serial interface.
11	**voip-queue**	Turn on QoS for VoIP.
12	**exit**	?
13	**interface Serial 1**	Enable cRTP (optional).
14	**ip rtp header-compression**	

Implementing Communication Manager on a data network

S8700 Multi-Connect

Figure 76: S8700 Multi-Connect system

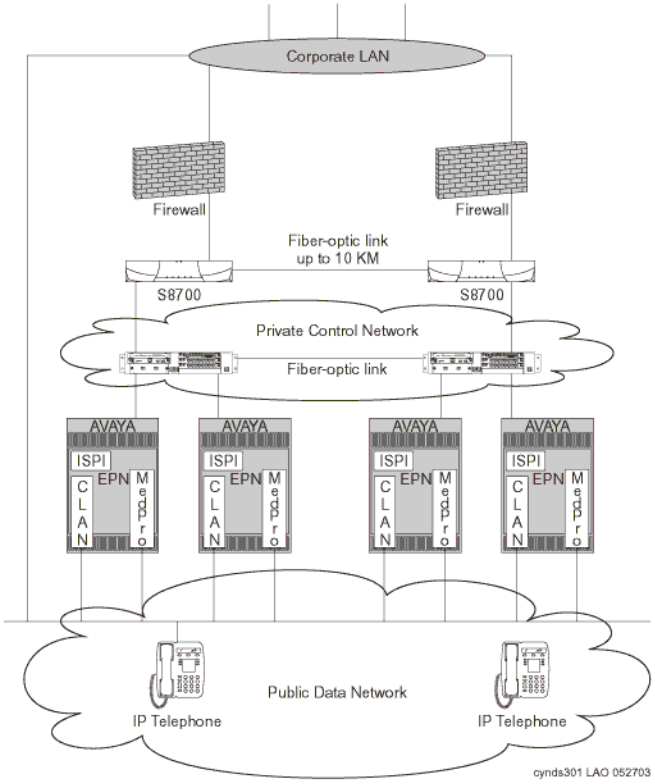

As shown in Figure 76, S8700 Multi-Connect system, on page 241, the S8700 Multi-Connect system is relatively straightforward to implement on the data network. It consists of an S8700 server pair and some number of IPSIs, depending on the number of port networks. The connection between the IPSIs and the S8700 pair is done on a private LAN utilizing one (or more) switches. This is one of the simpler network configurations, as control traffic is completely isolated from the production data network.

IPSI configuration

In a Multi-Connect system, IPSI configuration is automatic. Upon power-up, IPSIs send out a DHCP request on the control network. The S8700s, then, provide DHCP addresses and server configuration information to the IPSIs.

Server separation

The S8700 servers can be separated by a maximum distance of 10 km. This limitation is based on the fiber channel interface used for memory shadowing. When contemplating a 10 km server separation, it is important to plan for the Ethernet separation, as well. Both servers' arbitration link and control network links MUST be on the same IP subnet. Although the use of media converters to convert the 10/100BaseTX Ethernet to 100BaseFX is an approved configuration, it is not recommended. Media converters are not reliable and are difficult to troubleshoot. As shown in Figure Figure 76, S8700 Multi-Connect system, on page 241, the recommended alternative is to use two P130 switches with 100BaseFX (or gigabit) interfaces. Each server would connect to one switch, and the two switches would be linked by a fiber optic connection. This solution is much more reliable and easier to troubleshoot.

Security

Because the S8700 servers allow for (and recommend) a modem for Avaya services access, placing a firewall between of the S8700 servers and enterprise network is advisable. In a Multi-Connect system, this is straightforward. Place a firewall (or router access-lists) in front of the S8700 pair's administrative interfaces. Permit administrative traffic (https/ssh/telnet/ftp) from administrator PCs into the S8700s. Also permit the return traffic. Deny all other traffic. The S8700 Multi-Connect system is architected such that it is impossible for traffic originating on the S8700 servers to traverse the control network and enter the customer's network through the port network.

Other IP interfaces

The C-LAN and Media Processor connect directly to the customer's data network (that is, not the control network). They must be reachable by IP Telephones on the network, so they should be placed in the voice vlan, should one exist, or should at least be reachable by all subnets containing voice endpoints. The architecture of the system is such that traffic entering either the C-LAN or MedPro cannot cross into the control network.

The IPSI connects to the control network and provide an interface between the S8700 servers and the port network. It does not need to be reachable from the enterprise network.

S8700 IP connect

Figure 77: S8700 IP-Connect system

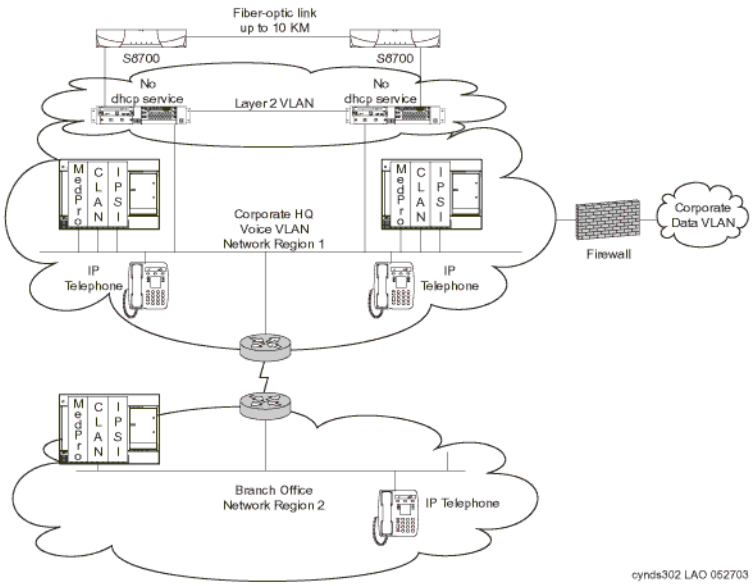

As shown in Figure 77, S8700 IP-Connect system, on page 243, the S8700 IP-Connect system also consists of a pair of S8700 servers with IPSIs per port network. The main difference between a Multi-Connect system and an IP connect system is that in IP-Connect systems, control traffic is not isolated to a private control network, but traverses the customer's network, instead. Unlike Multi-Connect, IP-Connect supports port networks across a WAN.

IPSI config

DHCP is not available in an IP-Connect system. IPSI cards must be administered manually through the services ethernet interface on the front of the card.

Network design

Proper network design is very important in an IP-Connect system. Because port network control traffic is flowing across a customer's network, and not isolated on a control network, suboptimal network design may lead to stability problems on the port networks.

Figure 78: Server farm placement in network configuration

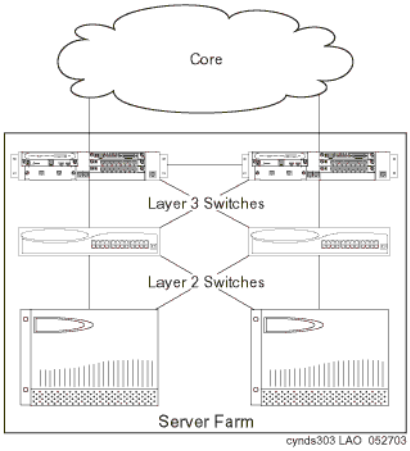

In general, both S8700 Media Servers and G600 Media Gateways should be treated as servers for the purpose of network placement. They should be connected into a server farm utilizing redundant switches. If the network path between the S8700 media servers and the G600 gateways is disrupted for more than a few seconds, the port network could reset, disrupting calls to the PSTN. Because stability is a concern, G600s should not be connected on the same subnet as user PCs. One such configuration of G600 media gateways is shown in Figure 78, Server farm placement in network configuration, on page 244.

Provisioning Network Regions

Network regions are Communication Manager constructs for selecting codecs and for grouping resources by location or network topology. When determining MedPro resources, for example, Communication Manager will try to select them in the same network region as the IP Telephone attempting to use them. In addition, codec sets are negotiated by endpoints based on whether they are in the same or different network regions. It is common to specify G.711 operation within a network region and G.729 between regions. For voice quality reasons, then, it is advisable to never have a network region extend across a WAN connection. Thus, when deploying an IP-Connect system, it is common to have one network region for the central headquarters location and separate network regions for each remote branch office. MedPros should be configured for whichever network region in which they are physically located. This concept applies to VoIP resources on the G700 media gateways, as well. The IP Telephony resources in a G700 deployed in a branch office should be configured for the network region of the remote branch office, and not the network region of the MedPros located in headquarters.

QoS

Because G600s can be separated from S8700s by a WAN link, and because network issues affect G600 stability, it is important to properly enable QoS, specifically DiffServ. QoS is most important across WAN links, but may be important in LAN environments, as well. If a network has already been configured for DiffServ, Avaya servers and gateways can utilize the policy already existing on the network.

Recommendations for QoS DiffServ

If the network has not already been configured for DiffServ, Avaya recommends keeping the policy simple: set the DiffServ Code Point (DSCP) to 46 (Expedited Forwarding, or EF) for control, signaling, and voice bearer traffic. Use CBQ or strict priority queuing (with IP Telephony traffic inserted into the highest priority queue) to offer the appropriate level of service.

Security

Because the S8700 servers have modems for Avaya services access, it is advisable to use an external firewall or router access-lists to filter traffic to and from the S8700 servers. Appendix B. Access list lists port ranges used by Avaya products, and can be used to harden a VoIP system.

S8700 / S8300 LSP

The S8700 server pair can also be used to control a G700 media gateway. For survivability, an S8300 server running in Local Survivable Processor (LSP) mode can take over control of up to 50 G700s and all associated endpoints in case of a network failure.

Security

Because both the S8700 servers and the S8300 server can support modems, it is important to consider using access-lists and firewalls to isolate servers and media gateways from the rest of the data network. Appendix B has information necessary for establishing access-lists or setting up a firewall policy. Avaya has worked on hardening Avaya Application Solutions products against penetration and denial of service attacks. For more information see Security.

G700 connections to the C-LAN

A G700 uses H.248 signaling with its controller. If the G700 is homed off of an S8700 server pair, it must be able to reach a C-LAN for its signaling connection. The G700 gateway does not communicate directly with an S8700. This restriction does not apply when controlled by an S8300 Media Server. The G700 would communicate directly with the S8300 server, and not a C-LAN, in that case.

LSP-to-S8700 connection

In order to ensure up-to-date translations, an S8300 Media Server configured as an LSP must be able to communicate directly over the IP network with the S8700 server pair. Communication Manager uses rsync to synchronize translations between the S8700 servers and their LSPs. This communications channel does not use a C-LAN. The S8300 must communicate with a C-LAN circuit pack for keepalive traffic.

S8300 / G700 (ICC)

The S8300 server, operating as primary controller, will control a G700 media gateway. It supports all of the same features as an S8700 server pair and is inserted directly into the G700. All H.248 signaling occurs across the Ethernet backplane in the G700.

Native NIC

The S8300 server supports C-LAN functionality natively within the server. IP Telephones, IP trunks, and G700 gateways can connect to the S8300 server interface directly for H.323 and H.248 signaling.

Stacking

The G700 can be stacked with other switches in the Avaya P330 family. It uses an 8-Gbps stacking cable connecting the switches. As with other members of the P330 family, it can be managed by the switch designated as the stack master, freeing customers from managing each switch separately. It can take advantage of all of the features of P330 switches, including Layer 3 switching when stacked with a P330R switch. In addition, the G700 works with all of the X330 expansion modules, including gigabit Ethernet, IEEE 802.3af inline power, and 100BaseFX Ethernet. It also supports the X330 WAN blade, allowing cost-effective routing of T-1/E-1 data traffic.

Sample Multi-Connect deployment

Figure 79: S8700 /G700 Multi-Connect system (gateways deployed at remote offices)

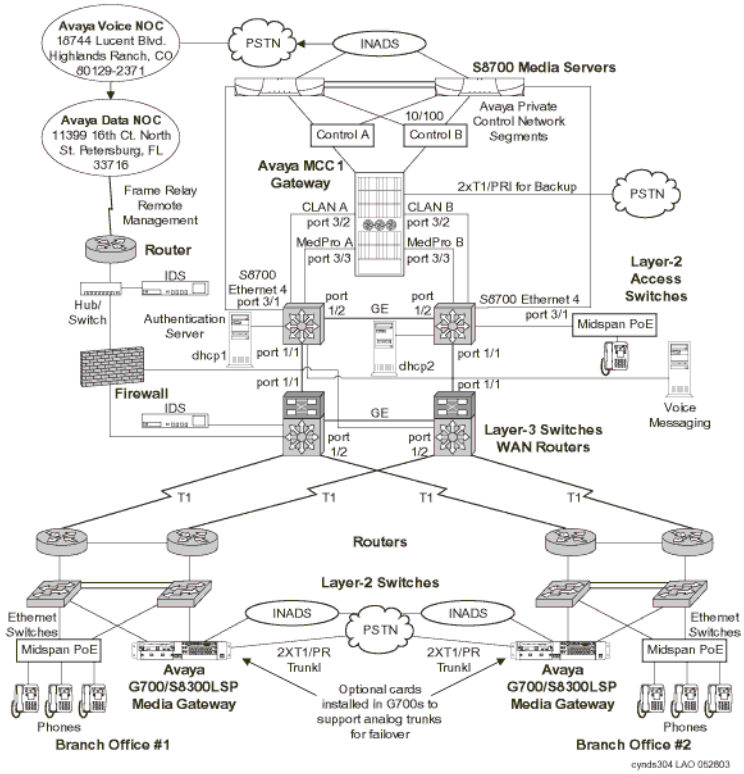

Figure 79, S8700 /G700 Multi-Connect system (gateways deployed at remote offices), on page 247 illustrates a typical S8700 Multi-Connect system with G700 gateways deployed at two remote offices. It demonstrates a number of the features previously discussed:

- The network as designed is highly resilient. There are redundant routers and switches at every level. In addition, the routers and Layer 3 switches are running VRRP to offer redundant default gateways to endpoints.
- The G700s in the branch offices use LSPs for local survivability, should connectivity to headquarters be disrupted.
- The phones are powered by midspan power units connected to UPSs, reducing the likelihood of power outages to the phones.
- This network is very secure. Access-lists (set up in accordance with Appendix B) are applied on the routers and Layer 3 switches. In addition, firewalls and intrusion detection sensors have been deployed. The control networks are physically disconnected from the building network, and only the C-LANs, MedPros, and S8700 administrative interfaces are connected to the building network (through access-lists). This does provide sufficient connectivity, however, for the G700s to reach the C-LANs and the LSPs to reach the S8700s.
- DiffServ (L3) QoS has been applied on the T-1 circuits. In this case, 75% of the bandwidth has been reserved for voice and voice signaling (DSCP 46) and 25% has been left for everything else. Should voice traffic not fill the voice queue, the excess bandwidth is passed to the data queue. This level of queuing may not be appropriate for all networks, but is often a good place to start.

- 802.1Q has been enabled on the Ethernet segments. Voice and signaling traffic are both tagged at 6, the value reserved for voice (and the second-highest level of Layer 2 QoS, leaving the highest for network control). Because 802.1Q tags are stripped at every router hop, they are regenerated at the far side of WAN links by mapping 802.1p tags based on Layer 3 DiffServ tags.
- Three network regions have been set up: one for headquarters and one for each branch office. The codec sets have been established such that intra-region traffic uses G.711, while inter-region traffic uses G.729.

Network recovery

Conventional wisdom holds that network reliability is typically 3-9s (99.9%) on a LAN, and 2-9s (99%) on a WAN. The leading causes of network failure are a WAN link failure, administrator error, cable failure, issues that involve connecting new devices or services, and malicious activity, including DoS attacks, worms, and viruses. Somewhere lower down on the list are equipment failures. To achieve the highest levels of availability, it is important that a strong change control policy and network management strategy be implemented.

There are numerous techniques for improving the reliability of data networks, including spanning tree, self-healing routing protocols, network management, and change control. This chapter discusses a few of those techniques.

Change control

Change control describes a process by which an organization can control nonemergency network changes, and reduce the likelihood of administrator errors that cause network disruption. It involves carefully planning for network changes (including back-out plans), reviewing proposed changes, assessing risk, scheduling changes, notifying affected user communities, and performing changes when they will be least disruptive. By implementing a strict change control process, organizations can reduce the likelihood of administrator errors, which are a major cause of network disruption, and increase the reliability of their networks. is Appendix A. Change control contains more information on change control.

Layer 2 mechanisms to increase reliability

Spanning tree

IEEE 802.1D spanning tree is an Ethernet loop avoidance protocol. It allows network managers to connect redundant network links within their networks. Prior to the advent of spanning tree, loops within a switched Ethernet network would forward traffic around the loop forever, which saturated the network and prevented new traffic from getting through. Spanning tree selects one switch as a root and creates a loop-free topology connecting to the root. If loops are discovered, one switch blocks that port until its alternate path to the root is disrupted. Then the blocked port is brought back into service. There are several drawbacks to spanning tree:

- By default, all switches have the same priority, which means that root bridge selection is random. This can be suboptimal in a network.
- Spanning tree is slow to converge. It typically takes at least 50 seconds from link failure for a backup link to become active. As Layer 2 complexity increases, so does convergence time.
- Although there are mechanisms for speeding up spanning tree, most are proprietary.
- Traditional spanning tree is not VLAN aware. Thus, it will block links even if VLAN provisioning would have prevented a loop.

To solve these issues, the IEEE has recently introduced 802.1s and 802.1w enhancements. 802.1w introduces rapid spanning tree protocol (RSTP). RSTP uses active handshaking to speed up convergence times. 802.1s introduces multiple spanning trees (MST), which is a way of grouping different VLANs into different spanning tree instances. These features might not be present in data network switches yet, but look for them soon.

Link Aggregation Groups

Link Aggregation Groups (LAGs) are a mechanism for combining multiple real interswitch links (typically four, Avaya products are configurable from two to eight) into one point-to-point virtual interswitch link. The advantage of this mechanism over spanning tree is that an organization can have the redundant links in if a failure occurs in one of the LAG links, the two switches will quickly discover it, and remove the failed link from the LAG., which reduces the convergence time to nearly instantaneous. Not all implementations interoperate, so care must be taken when the LAG connects switches from multiple vendors. Also, LAG links are a point-to-point technology. They cannot be used to connect a backup switch in case the primary fails. When available, this is a very good mechanism for improving the resiliency of LANs.

Layer 3 availability mechanisms

Routing protocols

Routing protocols allow routers to dynamically learn the topology of the network. Should the topology of the network change, routing protocols update their internal topology table, which allows them to route around failure.

There are two types of routing protocol, distance vector and link state. Distance vector protocols, including RIP and IGRP, exchange their entire routing table periodically. To each route, they add their metric (for RIP, this is “hop count”) and insert it in the routing table. If updates fail to arrive before the router’s timer expires, it purges the route and looks for another path. These protocols are usually slow to converge. See Table 56, Sample convergence times (single link failure), on page 252.

Link-state protocols, such as OSPF, take a more holistic view of the network. They compute the entire topology of the network and insert the best path to a destination in the routing table. Link state protocols exchange their routing tables only once, when routers first establish a relationship. After that, they only send updates. They also send hello messages periodically to ensure that the other routers are still present. Link state protocols converge much more quickly than distance vector protocols, and thus are generally better suited to networks that require high availability.

VRRP and HSRP

Virtual Router Redundancy Protocol (VRRP) and the related Cisco proprietary Hot Standby Router Protocol (HSRP) provide a mechanism to deal with router failure without disrupting endpoints on the network. In essence, these protocols work by assigning a virtual IP address and MAC address for the routers. This address is given to endpoints as their default gateway. The two routers send periodic hello messages marked with a priority value between each other. The high-priority router assumes the virtual address, and traffic flows through it. If the primary router fails or its capabilities become degraded (such as if a WAN link fails), the secondary router takes over. This is a useful mechanism to protect endpoints from router failures, and works with IP Telephony endpoints.

Multipath routing

Modern routers and Layer 3 switches allow multiple routes for a particular destination to be installed in the routing table. Depending on the implementation, this can be as high as six routes. Some implementations require that all routes that are inserted in the routing table have the same metric, while others allow unequal metric routing. In cases where the metric for all installed routes are the same, the router will load balance traffic evenly across each path. When the metric for multiple routes vary, the traffic is load balanced in proportion to the metric (in other words, if one path is “twice as good” as another, two-thirds of the traffic travels down the good path, and one-third of the traffic selects the other one). Asymmetric routing is suboptimal for voice, so route-caching (described earlier) should be considered in this environment.

In addition to using all (up to 6) active paths and optimally using available bandwidth, multipath routing greatly improves convergence time. As soon as a router detects a path failure, it remove it from the routing table, and sends all traffic over the remaining links. If this is a physical link failure, the detection time is nearly instantaneous. Therefore, Avaya recommends the use of multipath routing, where available, across multiple links to a particular location.

Additional mechanisms

Dial backup

One cost-effective technique for installing backup WAN links is to use dial backup. This can be done using either ISDN-BRI or analog lines. ISDN lines typically take 2 seconds to connect, while 56-k analog modems take approximately 1 minute. While this strategy is effective for data traffic, it is less effective for voice. First, the bandwidth may have been greatly reduced. If this is the case, the number of voice channels that can be supported might have been reduced proportionally. Also, if QoS is not properly applied to the backup interface, high packet loss and jitter can adversely affect voice quality. Finally, the time that is required to establish the new link can be up to 1 minute, which disrupts active calls. However, providing that these considerations are taken into account, proper QoS is applied, and a compressed codec is chosen, dial backup can be an effective solution for two to four users.

Convergence times

Convergence is the time that it takes from the instant a failure occurs in the network until a new path through the network is discovered, and all routers or switches are aware of the new path. Convergence times vary, based on the complexity and size of a network. Table 56, Sample convergence times (single link failure), on page 252 lists some sample convergence times that are based on a single link failing in a relatively simple network. They reflect update and/or hello timers expiring. Dialup "convergence" times reflect the time that it takes to dial, connect, and authenticate a connection. These times do not take into account LAG, fast spanning tree, or multipath routing, which speed up convergence. This table shows the importance of carefully planning for fail-over in a network. For example, both OSPF and EIGRP (Layer 3) protocols converge faster than spanning tree (Layer 2). When designing a highly available data network, it is more advantageous to use Layer 3 protocols, especially link-state (OSPF) or hybrid (EIGRP) protocols, than Layer 2 (spanning tree).

Table 56: Sample convergence times (single link failure)

Protocol	Approximate convergence time (seconds)
EIGRP (Cisco)	2
OSPF	6 to 46
RIP	210
IGRP (Cisco)	400
Spanning tree (Layer 2)	50+
ISDN dialup (connect + authentication)	2
56-k dialup (connect + authentication)	60

Network assessment offer

Avaya Network Consulting Services (NCS) supports a portfolio of consulting and engineering offers to help plan and design

- IP Telephony
- Data Networking Services
- Network Security Services.

How to contact the NCS

- On the Web -- http://ncs.avaya.com/
- E-Mail: ncs@avaya.com
- Phone: +1 866-832-0925

Problems with data networks

Many customer IP infrastructures appear to be stable and perform at an acceptable levels but have performance and stability issues that create problems for Avaya IP Telephony. While the customer network appears to be ready for full-duplex IP Telephony, Avaya cannot assure performance and quality without a Network Assessment.

Avaya network assessment solutions

The Network Assessment services for Avaya IP Telephony consist of 2 phases:

- Customer Infrastructure Readiness Survey (CIRS) is a high-level LAN/WAN infrastructure evaluation that determines the suitability of an existing network for IP Telephony.

  The CIRS Report includes detailed technical information about any problems that are discovered in the customer infrastructure. It also includes performance predictions based on network and system administration standards.

  If the survey discovers significant network issues, these must be remedied before deploying any Avaya IP Solution. Customers can resolve the problems independently and follow up with another CIRS (at an additional charge) or to move ahead with the Network Analysis Network Optimization (NANO) service.

  NOTE:
  The CIRS is available in the U.S. and Canada through direct and indirect channels.

- Network Analysis Network Optimization (NANO) is typically the second phase in the Network Assessment for IP Telephony solutions. The NANO takes information gathered from the CIRS, performs problem diagnosis and provides functional requirements for the network to implement Avaya IP Telephony.

  A NANO is required when the CIRS indicates that the customer's network as it is configured cannot support the proposed IP Telephony application at the desired performance levels. Sometimes customers already know that their existing network is not configured to support Avaya IP Telephony, and they can order a NANO without first completing a CIRS. The assessment requires that the customer complete a CIRS-like analysis as the first phase). Customers may also request a NANO to optimize their network.

Customer Infrastructure Readiness Survey (CIRS)

The CIRS offer is a scheduled remote network evaluation that is valuable for all customers that are expanding of their communication capabilities.

The CIRS evaluates the customer's current network environment by

- Maximizing the available resources.
- Identifying additional resources that are required to support the proposed IP Solution.

The outcome of the CIRS is a road map that identifies the gaps in the existing network today, but it does not provide step-by-step configuration instructions on how to deploy the solution. The CIRS is performed remotely and must be scheduled with the NCS team 2 weeks before implementing Avaya IP Telephony.

What if my network functions well today?

Even if your network appears to perform acceptably, IP Telephony taxes network resources and performance because IP Telephony requires dedicated bandwidth and is more sensitive to network problems than data applications. Table 57, Customer Infrastructure Readiness Survey (CIRS) components, on page 255 shows the CIRS components and the depth of Avaya's network analyses.

Table 57: Customer Infrastructure Readiness Survey (CIRS) components

Component	Who does this?	What does this do? What are the results?	What happens with the results?
Network Topology Report	Customer	Describes your network configuration.	Topology report integrated with all other CIRS components.
Site Configuration Survey	Customer	Data for individual customer site; high-level health check.[1]	Professional Service (PS) engineer reviews data and recommends where to deploy Protocol Analysis software.
Vital Agent analysis	Customer downloads application to identified desktops	Vital Agent collects data about the customer's network traffic.	PS Engineer analyzes the traffic data and determines where to deploy the Performance Analysis software.
Protocol Analysis	Avaya	Measures the data exchange between the local host and remote resources and tests the current load for appropriate bandwidth.	Deeper data analysis (NCS engineer)

1 If the network fails to meet the minimum criteria required for network throughput, configuration, or additional resources are needed, Avaya recommends the more in-depth analysis of the Network Analysis Network Optimization (NANO).

Site Configuration Survey

The ECLIPS Site Configuration Survey (SCS) is an detailed customer-view of the their network. This survey is required as part of the Customer Infrastructure Readiness Survey process. The ECLIPS SCS questionnaire must be filled out as completely as possible, and can require the Account Team's regional Sales Engineering resources to assist. In addition to the SCS the customer must provide a topology map of their existing network (LAN/WAN and hardware/ software configuration listings). When the SCS is complete, the customer provides both the SCS and Topology Maps with detailed descriptions of topology components (routers, Ethernet switches, PSTN linked systems, firewalls, servers, E-mail systems, etc.). This information goes to the CVDN Professional Service engineering team that reviews this information and prepares for the Vital Agent Analysis, the second component of the CIRS.

Vital Agent analysis

Vital Agent is a high-level analysis tool that passively monitors and reports throughput and performance statistics and errors and reports any problems that the host computer encounters. The customer must install and run the Vital Agent software on all desktops targeted for Avaya IP Telephony.

If the customer has a somewhat standardized network infrastructure, Avaya can waive the need to install this application on every desktop and instead to run this utility only on key desktops.

The Vital Agent software gathers data for up to 5 consecutive business days after which the customer sends the data file to the CVDN Professional Service engineers for analysis. The CVDN then determines if the proposed Avaya IP Telephony application can perform acceptably over the customer's network.

If a problem is uncovered as a result of the survey, the CVDN Professional Service engineering team notifies the Account Team and includes detailed technical information regarding the problem. The customer has two choices:

- Resolve the problems independently and then re-run the survey afterward;
- Hire Avaya to perform an on-site Network Analysis Network Optimization (NANO) evaluation.

Network Analysis Network Optimization (NANO)

The Network Analysis Network Optimization (NANO) offer includes

- Scheduled on-site evaluations
- Traffic simulation
- Network testing
- Analysis of the results
- Recommendations to resolve any network throughput issues

In order to reap the benefits of IP Telephony, customers must either possess or acquire a keen understanding of their network and its performance capabilities. This ensures that the transfer of information between systems and processes is not compromised and that the network infrastructure remains stable.

The NANO results are documented in a Network Assessment Report that identifies the root cause of the network issues and provides the customer with recommendations on how to resolve those issues to support the implementation of the IP Telephony solution. CVDN Professional Services utilizes proven methodologies performed by a staff of highly-experienced, certified network engineers. These engineers are capable of addressing the customer's critical business needs in complex, multimedia, and multivendor environments.

Use these links for more information about the NANO components:

- The NANO process
- Customer responsibilities
- Discovery
- Element monitoring
- Synthetic IP Telephony measurements

- Remote analysis
- Report generation
- Customer deliverables

The NANO process

To begin the NANO process, the customer must have completed the

- Customer Infrastructure Readiness Survey (CIRS). If a customer has already concluded that their network is not ready for the implementation of Avaya IP Telephony, they can skip the CIRS.
- Site Configuration Survey (SCS).
- Network topology map.

During a NANO, data collection utilities and network simulation tools are loaded onto a customer's network at pre-determined endpoints. Traffic with similar characteristics injected onto the network and monitored for performance under load conditions. After the performance analysis, a comprehensive report documenting network performance, problem areas, and suggested resolutions is given to the customer. The CVDN Professional Services organization can also provide a separate proposal to assist the customer in configuration and integration/administration engineering services to prepare the network for the proposed Avaya IP Telephony application.

Table 58, Network Analysis Network Optimization components, on page 257 shows the NANO components and the information exchange between Avaya and the customer.

Table 58: Network Analysis Network Optimization components ***1 of 2***

Component	Who does this?	What does this do? What are the results?	What happens with the results?
Network Topology Report	Customer (may already be part of CIRS)	Describes your network configuration.	Topology report integrated with all other NANO components.
Site Configuration Survey	Customer (can already be part of CIRS)	Data for individual customer site; high-level health check.	Professional Service (PS) engineer reviews data and recommends where to deploy Protocol Analysis software.
Traffic Injection Monitoring Data Collection	Avaya	Determines endpoints (with SNMP agents installed) for data collection. Monitors each network segment for busy hour traffic	Data analyzed to determine the highest level phone quality (starting at 64Kbps) and working through lower quality levels.

1 of 2

Table 58: Network Analysis Network Optimization components ***2 of 2***

Component	Who does this?	What does this do? What are the results?	What happens with the results?
Additional tests	Avaya	Summary of Impact of Delay	
		Packet Loss and Jitter on Quality of Service on Voice Quality	
		Summary of Quality using Avaya's Specification for Delay, Loss, and Jitter	
		Impact of Quality of Service on Voice Quality	
		Summary of Quality of Real World Pilot	Summarizes the entire network analysis.
		Layer 3 Traffic Analysis	

2 of 2

Customer responsibilities

In order to successfully complete a NANO the customer must:

- Provide technical resource personnel who are well-versed in the network infrastructure.
- Provide complete access to the network.
- Provide passwords for networking equipment.
- Provide access to personnel for interviews.
- Update or provide network topology maps.
- Identify a place on the network for test equipment.
- Define times to complete network testing.

Discovery

- Perform interviews with IT staff to determine application and network performance expectations
- Locate and identify all SNMP enabled devices
- Identify hosts on each subnet
- Identify all routers, switches, and hubs
- Manual identification of all non-SNMP enabled devices
- Identify operating system of each Host found
- Map hosts to communication paths between hosts
- Generate Layer 3 topology map to compare with CIRS

- Install endpoints for testing
- Review WAN-specific circuits, bandwidths, DLCI/PVC configurations, and channeled T1 configurations
- Review the customer's Layer 2 architecture

Element monitoring

- Monitor router status through SNMP (port utilization, MIB II errors)
- Capture all network device SNMP data real-time into database
- CPU utilization capturing per host being used for testing
- Monitor LAN switch utilization, MIB II errors

Synthetic IP Telephony measurements

- Inject busy hour IP Telephony call traffic simulation into live network segments
- Random CODECs and injection points between pre-defined end points/hosts
- Injections initially within single facilities, replicated across WAN end points as appropriate
- Capture of all test data into database real-time

Remote analysis

- Analysis of element/endpoint data by router, time period, and other performance variables
- Analysis of element/endpoint data by switch, time period, other performance variables
- Analysis of IP Telephony call data by IP endpoint pair, time period, and other performance variables, then integrated with SNMP data
- Generation of graphs representing usage for all endpoint data
- "What if" analysis of IP Telephony codecs to determine best match for performance and call quality

Report generation

- Summary of IT and Voice team's interviews: perceived expectations and requirements as related to proposed applications and network performance levels
- Physical topology map on all devices discovered and monitored on the network
- Analysis of WAN circuits: current status and recommendations for support of proposed Avaya IP Telephony
- Traffic analysis reports, including archive on CD-ROM of all captured data for all segments monitored and injected with simulated busy hour IP Telephony calls
- Recommendations of Avaya Engineering team to resolve infrastructure problems discovered and/or make-ready for proposed Avaya IP Telephony
- Summary reports of segment utilization, errors, and dropped packets
- Summary E-Model calculations for different CODEC reports per segment/per layer
- Summary reports for the Level 3 QoS audits (if performed)
- Summary reports for different network layers' performance

Customer deliverables

- Avaya networking experts perform discovery of the customer's network and document findings in a NANO Report delivered to the customer.
- Accurate network topology
- Measurements of actual usability performance levels, throughput performance of the LAN, and server utilization
- Results of traffic simulation on the network at projected volumes
- Define problem areas, causes, and functional requirement recommendations to be implemented in the network design

Troubleshooting

This chapter contains an overview of the troubleshooting process, and common tools and techniques for isolating and solving problems.

Methodology

In general, two steps are needed to resolve an IP Telephony problem:

- Identify the location of the problem (IP Telephone, network, PBX, and so on), by using alarms and the state information of devices and any information provisioned or gathered by a network management system.
- Repair the problem (for example, rebooting), correcting parameter provisioning, reinstalling software or firmware, or replacing hardware.

It is important to pinpoint the location of the problem as precisely as possible so that any repair actions require minimal effort to reduce the costs and minimize the impact on noncorrupted service.

Many strategies can identify the location of a IP Telephony problem. For example, one could pinpoint the location of a problem in the following ways:

- Bottom up, protocol layer after protocol layer, starting at the physical layer.
- First analyze the perceived voice impairments (echo, delay, and voice clipping) if any, and then analyze signaling and network impairment problems.
- Start with a solution that is most likely to resolve the problem, followed by less likely solutions if necessary.
- Start by looking at large patterns:
  - Has the customer had a network voice readiness assessment done? If not, the network may not be compatible with the voice network readiness guidelines for Avaya products.
  - Has the network been changed after a network assessment been performed? Again, any network modifications should follow the network readiness guidelines.
  - Do other IP Telephones on the same subnetwork/VLAN, floor, switch port, router MedPro, C-LAN, network region, campus, software or firmware version, or Communication Manager version have the same problem? Similar problems with multiple IP Telephones might indicate shared resource problems such as power problems, Ethernet switch or IP router problems, or remote connectivity WAN problems. It may also indicate software or firmware version problems.
  - Does the problem repeat at a specific time of day? At specific times, the network load may be higher, which causes Communication Manager to run out of IP Telephony resources.

- Look for other simple solutions such as:
  - — If only one telephone has a problem:
    - If exchanging the IP Telephone solves the problem, then the IP Telephone is likely the source of the problem, unless the problem is intermittent.
    - If the problem is solved when the IP Telephone is connected to a different Ethernet switch port or IP router port, then the IP Telephone is not the problem.
  - — Are compatible codecs used?

In general, the fastest strategy to resolve IP Telephony problems is based on a combination of the last three strategies. The problem-specific Symptom Resolution Notes are created based on the combination of these strategies, found in Appendix E. Troubleshooting.

Sometimes, additional information needs to be obtained about the state of the network, such as router and switch port statistics, or about router access control lists. This information can be obtained by directly logging into the IP network, or by using a protocol analyzer tool to monitor traffic.

Avaya tools

Avaya provides several tools for troubleshooting IP Telephony systems. These tools include VoIP Monitoring Manager (part of Avaya Integrated Management) that monitors and records statistics from all IP Telephony calls, list trace commands that enable customers to observe call progress, and status station/trunk commands that give a snapshot of the observed state of the element in question. For very complex problems Avaya Services can diagnose problems using Message System Tracer to trace system calls within Communication Manager.

VoIP Monitoring Manager

VoIP Monitoring Manager can be used to troubleshoot voice-quality problems. VoIP Monitoring Manager captures RTCP messages between all the IP endpoints that are involved with a call. From the RTCP messages, it calculates packet loss, jitter, and round-trip delay. These results are graphed for active calls, or can give an overall status for historical calls. This is an excellent tool for troubleshooting voice-quality problems because it displays statistics (either real-time or historical) that it gathers from the endpoints themselves, and thus are not based on the location of the VoIP Monitoring Manager server or the configuration of the data network. By analyzing patterns in the data, VoIP Monitoring Manager can also help identify problems in the data network.

Troubleshooting commands

There are several commands that are helpful in troubleshooting IP Telephony problems. Table 59, Troubleshooting commands and their use, on page 263 lists those commands and their usage.

Table 59: Troubleshooting commands and their use

Command	Use
list trace station	This command traces the behavior of a particular station. It shows off-hook status, call setup and teardown messages, call routing, and call performance (for IP sets only). Every 10 seconds, it displays packet loss and jitter statistics for the previous 10 seconds, to assist in voice- quality troubleshooting. This is useful for voice-quality troubleshooting, or calls that fail to set up properly.
list trace tac	This command operates similar to **list trace station**, but it operates on trunks. In addition to call setup, teardown, and routing, it also lists voice-quality statistics in 10- second increments. This is useful for troubleshooting call routing problems or voice-quality problems across IP trunks.
list trace ras	This command allows an administrator to watch the state of the RAS messages that Communication Manager is processing. This can either be limited to a single station or expanded to the whole system. It shows registration, keepalive, and unregistration requests. This is useful when IP Telephones are rebooting spontaneously or fail to register.
status station	This command shows a snapshot of the state of an individual station. It lists registration status, the C-LAN and media processor or IP endpoint that is connected to an IP station, and lists 10 seconds of voice-quality (packet loss and jitter) information. It also shows whether the call is shuffled, hairpinned, or connected to the TDM bus.
status trunk	This command shows a snapshot of the state of an individual trunk. It lists the far-end C-LAN and media processor or IP endpoint that is connected to an IP trunk, and lists 10 seconds of voice-quality (packet loss and jitter) information. It also shows whether the call is shuffled, hairpinned, or connected to the TDM bus.

Common issues

In general, a number of common voice-quality problems can be preempted by performing a network assessment before implementing IP Telephony (see Network assessment offer).

Also, IP Telephone, C-LAN, and MedPro firmware releases are offered free of charge at:

support.avaya.com

It is advisable to upgrade to the latest GA firmware versions for a particular platform as a first step in troubleshooting voice-quality issues.

No dial tone

For the sake of this discussion, the “no dial tone” problem refers to the case where the light on the IP Telephone is on and the display is working, but no dial tone is heard after the IP Telephone goes off-hook. No dial tone occurs when connectivity between the MedPro and the IP Telephone is interrupted, or when insufficient DSP resources are available on the MedPro. No dial tone can also indicate incompatibilities in Network Region configuration. In some cases, no dial tone indicates a duplex mismatch between the MedPro and the Ethernet switch. Tools that help diagnose the no dial tone problem include VoIP Monitoring Manager, list trace station, status station, ping, and trace-route. See Appendix E. Troubleshooting for more information.

Talk path

One-way talk path is experienced where only one party on a call can hear the other. No-way talk path is the problem where neither party can hear the other, but the call is still connected. Talk path issues often relate to network connectivity issues. Both telephones might have a path to the MedPro, but might not have a route to each other, or might be blocked by a firewall. Also, talk-path problems could indicate a shortage of DSP resources on the MedPro. Disabling shuffling is a good way to help diagnose talk-path problems. Tools that help diagnose talk-path problems include VoIP Monitoring Manager, list trace station, status station, list measurements ip dsp-resource, ping, and trace-route. See Appendix E. Troubleshooting for more information.

Poor audio quality

Many problems can fall into the category of poor quality audio. In particular, these problems can be experienced as clipping of the beginning or ends of words, pops, or crackles.

Poor quality audio is generally caused by network problems. In particular, these problems indicate packet loss on the data network. Common solutions for such problems include applying or tuning QoS parameters and checking for duplex mismatch issues. Tools that help diagnose problems with audio quality include VoIP Monitoring Manager, list trace station, and status station. See Appendix E. Troubleshooting for more information.

Dropped calls

A dropped call is a call that is terminated by a mechanism that is outside of user control. For example, a call might be dropped without anyone hanging up. Dropped calls sometimes indicate a connectivity problem on the signaling channel. Such occurrences can be intermittent, and thus difficult to diagnose. If dropped calls do occur frequently, they can be diagnosed using **list trace station** and **list mst,** or by checking the denial event log. See Appendix E. Troubleshooting for more information.

Echo

Echo describes the phenomenon where a voice signal is reflected back to the speaker at an audible level so that it interferes with the ability to have a normal conversation with another party.

In recent years, echo has mostly been imperceptible in circuit-switched networks due to their low delay and the deployment of echo cancellers. IP calls can experience a much larger delay, and therefore echo can be much more noticeable.

Echo can be created in two ways:

- Acoustically, in a telephone handset, a telephone that is operating in speakerphone mode, a speakerphone, a headset, or a multimedia laptop computer or desktop computer with a headset or an integrated or separate microphone and speaker. In particular, speakerphones or telephones that are operating in speakerphone mode provide a high level of acoustical echo return signal. The level of acoustic echo is determined by the acoustics of the environment (such as wall and ceiling reflection), the degree to which loudspeaker and microphone are directed towards each other, and the directional acoustic characteristics of the microphone.
- Electrically, by impedance mismatches in 2-to-4 wire hybrids on analog line or trunk cards, or electrical cross-talk interference in wires or headset adapters.

In general, the perception of echo is call dependent. The perceived echo problems for calls that are made over a WAN are normally much larger compared with calls that are made over a LAN because of the larger delay in WAN-connected systems.

As echo is not caused by an IP network (although it is exacerbated by delay), so its resolution will not be covered in detail in this document. In general, there are three strategies for dealing with echo:

- Tune the network to reduce delay.
- Deploy echo cancellers.
- Tune the loss plan that is associated with the problem area.

When echo is experienced, the problem is generally resolved at the far-end of the link. For more information, see *Avaya IP Voice Quality Network Requirements*.

Appendix A. Change control

This appendix contains an overview of the change control process, why it is important, and the trade-offs that are associated with it.

Introduction

This section provides a template for change management that promotes high-availability networks. Specifically, the template provides the critical steps for creating a change management process, a high-level process flow for planned change management, an emergency change process flow, and a general method to evaluate the success of your process.

Critical steps for creating a change management process

Change management has two basic components

- Planning
- Managing

Planning

Change planning identifies the risk level that is associated with a change, and builds change planning requirements to ensure a successful change. The main steps for change planning are to:

- Assign all potential changes a risk level prior to scheduling the change.
- Document at least three risk levels for:
  - Software and hardware upgrades
  - Topology changes
  - Routing changes
  - Configuration changes
  - New deployments

  Assign higher risk levels to nonstandard adds, moves, or changes. The high-risk change process that you document must include laboratory validation, vendor review, peer review, and detailed configuration and design documentation.

- Create solution templates for deployments that affect multiple sites. Include information about:
  - Physical layout
  - Logical design
  - Configuration
  - Software versions
  - Acceptable hardware chassis and modules
  - Deployment guidelines
- Document your network standards for:
  - Configuration
  - Software version
  - Supported hardware
  - Domain Name System (DNS)
  - Device naming
  - Design
  - Supported services

Managing

Change management is the process that approves and schedules the change to ensure the correct level of notification and minimal user impact. The main activities involved in change management are to:

- Assign an individual to act as a change controller. This individual is responsible to:
  - Receive and review change requests
  - Manage change process improvements
  - Moderate change management review meetings
  - Act as liaison for user groups
- Hold periodic change review meetings. Include personnel from the following functional areas:
  - System administration
  - Application development
  - Network operations
  - Facilities groups
  - General users
- Document change input requirements, including:
  - Change owner
  - Business impact
  - Risk level
  - Reason for change
  - Success factors
  - Backout plan
  - Testing requirements

- Document change output requirements, including updates to:
  - DNS
  - Network map
  - Template
  - IP addressing
  - Circuit management
  - Network management
- Define a change approval process that verifies validation steps for higher-risk change.
- Hold postmortem meetings for unsuccessful changes to determine the root cause of the failure.
- Develop an emergency change procedure to ensure or restore an optimal solution.

High-Level process flow

The different steps to follow during a network change are represented in Figure 80, Process flow during a network change, on page 269. Each process (box) in the flowchart is discussed below.

Figure 80: Process flow during a network change

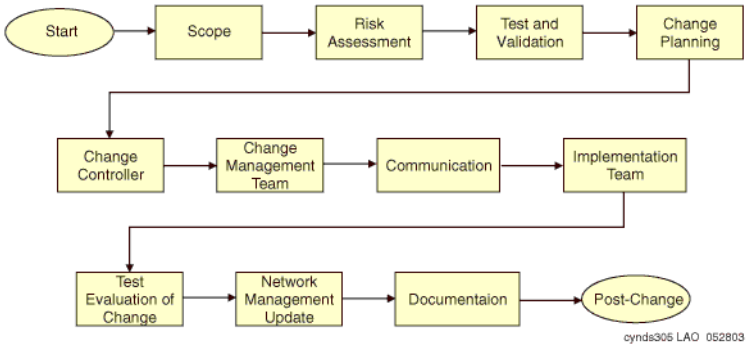

Scope

Any proposed change should include a complete technical definition, and the intent or purpose of the change. The change should also include information that describes what business units, user groups, servers, and applications might be affected, both during the change period and after deployment. Generally, most changes fall into one of the following categories:

- Network expansion
- Addition of LAN segments at existing sites
- Addition of new sites
- Connection to existing networks
- Connection to the Internet

- Corporate mergers and acquisitions
- Design and feature enhancements
- Software release upgrade
- Host software
- Distributed client software
- Configuration changes
- Support for additional protocol(s)
- Implementation of enhanced features

Risk assessment

Every network change has an associated risk level that can be assessed by modeling the change in a laboratory environment or with a network modeling tool. It might be helpful to assign one of the following risk categories to each change request:

- **High risk.** These network changes have the highest impact on user groups or particular environments, and might affect an entire site. Backing out of the change can be time consuming and difficult. Research high-risk changes using the available tools, and implement the change in conjunction with Avaya Services personnel. Ensure that management is aware of the change and its implications, and notify all users.
- **Moderate-risk.** These network changes can have a critical impact on user environments or affect an entire site, but backing out of the change is a reasonably attainable scenario. You should research moderate-risk changes the Avaya Support Centre Web site, and possibly review the change with Avaya Services personnel. Avaya recommends notifying all users of a moderate-risk change.
- **Low-risk.** These network changes have minor impact on user environments, and backing out of the change is easy. Low-risk changes rarely require more than minimal documentation. User notification is often unnecessary.

Additional risk levels night help identify the correct level of testing and validation prior to a change. Table 60, Test and validation risk levels, on page 270 defines five different risk levels that might help identify testing and validation requirements.

Table 60: Test and validation risk levels ***1 of 2***

Risk level	Definition
1	High potential impact to a large number of users (500+) or business-critical service. Introducing a new product, software, topology, or feature means network downtime.
2	High potential impact to large number of users (500+) or business-critical service. Large increases in traffic or users, backbone changes, or routing changes might require network downtime.

1 of 2

Table 60: Test and validation risk levels ***2 of 2***

Risk level	Definition
3	Medium potential impact to smaller number of users or business service. Any nonstandard change, such as a new product, software, topology, features, or the addition of new users, increased traffic, or nonstandard topology may require some network downtime.
4	Low potential impact, including adding new standard template network modules (building or server switches, IP Telephones, trunks, or routers); bringing up new remote offices or additional proven access services; and all risk level 3 changes that have been tested in the production environment. Change might require some network downtime.
5	No user or service impact, including adding individual users to the network, and standard configuration changes such as password, banner, Simple Network Management Protocol (SNMP), or other standard configuration parameters. No expected network downtime.

2 of 2

Test and validation

After the risk level of the potential change has been assessed, the appropriate amount of testing and validation can be applied. Table 61, Testing and validation recommendations, on page 271 demonstrates how testing and validation may be applied to the five-level risk model.

Table 61: Testing and validation recommendations

Risk level	Recommendations
1	Requires laboratory validation of the new solution, including documented testing, validation, and what-if analysis showing the impact to existing infrastructure; completion of an operations support document, backout plan, and implementation plan; and adherence to the change process. Recommend solution pilots and a preliminary design review prior to testing.
2	Requires laboratory what-if analysis to determine the impact to the existing environment with regard to capacity and performance; test and review of all routing changes; backout plan, implementation plan, and adherence to change process; and design review for major routing changes or backbone changes.
3	Requires engineering analysis of the new solution, which may require laboratory validation; implementation plan and adherence to change process.
4	Requires implementation plan and adherence to change process.
5	Optional adherence to change process.

For changes with risk levels 1 through 3, two types of laboratory validation are important:

- Feature and functionality testing
- What-if analysis

Feature and functionality testing

Feature and functionality testing requires that you validate all configurations, modules, and software with laboratory-generated traffic to ensure that the solution can handle the expected traffic requirements. Create a test plan that validates configuration parameters, software functionality, and hardware performance. Be sure to test behavior under real-world conditions, including spanning-tree changes, default gateway changes, routing changes, interface flaps, and link changes. Also validate the security and network management functions of the new solution.

What-if analysis

What-if analyses seek to understand the affect of the change on the existing environment. For example, if you add a new feature to a media gateway, the what-if analysis should determine the resource requirements of that feature on the media gateway. This type of testing is normally required when adding additional features, users, or services to a network.

Change planning

Change planning is the process of planning a change, including identifying requirements, ordering the required hardware and software parts, checking power budgets, identifying human resources, creating change documentation, and reviewing technical aspects of the change and change process. You should create change planning documentation such as maps, detailed implementation procedures, testing procedures, and backout procedures. The level of planning is usually directly proportional to the risk level of the change. A successful project should have the following goals for change planning:

- Ensure all resources are identified and in place for the change.
- Ensure a clear goal has been set and met for the change.
- Ensure the change conforms to all organizational standards for design, configuration, version, naming conventions, and management.
- Create a backout procedure.
- Define escalation paths.
- Define affected users and times when the network will be out of service for notification purposes.

Change planning includes the generation of a change request, which should be sent to the change controller.

Recommendations for change request information

Avaya recommends including the following information on the change request form:

- Name of person requesting change
- Date submitted
- Target date for implementing the change
- Change control number (supplied by the change controller)
- Help desk tracking number (if applicable)
- Risk level of the change

- Description of the change
- Target system name and location
- User group contact (if available)
- Lab tested (yes or no)
- Description of how the change was tested
- Test plan
- Backout plan
- If successful, will the change migrate to other locations (yes or no)
- Prerequisites of other changes to make this change successful

The technical description of the change is an important aspect of the change request, and may include the following: current topology and configuration, physical rack layouts, hardware and hardware modules, software versions, software configuration, cabling requirements, logical maps with device connectivity or VLAN connectivity, port assignments and addressing, device naming and labeling, DNS update requirements, circuit identifiers and assignments, network management update requirements, out-of-band management requirements, solution security, and change procedures.

In addition, a change request should reference any standards within your organization that apply to the change. This helps to ensure that the change conforms to current architecture or engineering design guidelines or constraints. Standards can include the following: device and interface naming conventions, DNS update requirements, IP addressing requirements, global standard configuration files, labeling conventions, interface description conventions, design guidelines, standard software versions, supported hardware and modules, network management update requirements, out-of-band management requirements, and security requirements.

Change controller

A key element to the change process is the change controller. The change controller is usually an individual within your IT organization who acts as a coordinator for all change process details. Normal job functions of the change controller include:

- Accepting and reviewing all change requests for completeness and accuracy
- Running periodic (weekly or biweekly) change review meetings with change review board personnel
- Presenting complete change requests to the change review board for business impact, priority, and change readiness review
- Preventing potential conflict by maintaining a change schedule or calendar
- Publishing change control meeting notes and helping communicate changes to appropriate technology and user groups
- Helping ensure that only authorized changes are implemented, that changes are implemented in an acceptable time frame in accordance with business requirements, that changes are successful, and that no new incidents are created as a result of a change

In addition, the change controller should must metrics for the purpose of improving the change management process. Metrics can cover any of the following:

- Volume of change processed per period, category, and risk level
- Average turnaround time of a change per period, category, and risk level
- Number of relative changes amended or rejected per period and category
- Number of relative change backouts by category
- Number of relative changes that generate new problem incidents
- Number of relative changes that do not produce the desired business results
- Number of emergency changes implemented
- Degree of client satisfaction

Change management team

You should create a change management team that includes representation from networking operations, server operations, application support, and user groups within your organization. The team should review all change requests and approve or deny each request based on completeness, readiness, business impact, business need, and any other conflicts.

The team should first review each change to ensure that all associated documentation is complete, based on the risk level. The team can then investigate the business impact issues and business requirements. The final step is to schedule the change. Once a change has been approved, the change management team is also responsible for communicating the change to all affected parties. In some cases, user training might also be needed.

> **NOTE:**
> The change management team does not investigate technical accuracy of the change. Technical experts who better understand the scope and the technical details should complete this phase of the change process.

Communication

Once a change is approved, the next step is to communicate details of the change by setting expectations, aligning support resources, communicating operational requirements, and informing users. The risk level and potential impact to affected groups, as well as scheduled network outages as a result of the change, should dictate the communication requirements.

Avaya recommends creating a matrix to help define who will be affected by a change, and what the potential time out of service might be for each application, user group, or server. Remember that different groups might require varying levels of detail about the change. For instance, support groups might receive communication with more detailed aspects of the change, new support requirements, and individual contacts, while user groups might receive only a notice of the potential time that the network will be out of service, and a short message that describes the business benefit.

Implementation team

You should create an implementation team that consists of individuals with the technical expertise to expedite a change. The implementation team should also be involved in the planning phase to contribute to the development of the project checkpoints, testing, backout criteria, and backout time constraints. This team should guarantee adherence to organizational standards, update DNS and network management tools, and maintain and enhance the tool set that is used to test and validate the change.

Specifically, the implementation team should fully understand the following testing questions, and should include them in the change documentation prior to approval by the change control board:

- How thoroughly should we test the change?
- How will we roll out the test?
- How long will testing last, and at what point can we make the decision that the change is implemented successfully?

The implementation team should also be fully aware of all backout criteria, time constraints, and procedures. The team should answer the following questions as part of the change documentation for high-risk change prior to approval by the change control board:

- How is the change to be removed?
- At what point is the decision made to back out of the change?
- What information should be gathered before backout occurs to determine why the change needed to be backed out or why it affected the network adversely?

During the implementation of any change, it is crucial to follow the change management team recommendations on how to make the change. If anything is performed on the network that deviates from the recommendations, the implementation team should document and present these steps to the change controller when the change is completed.

Test evaluation of change

Testing and verification can be critical to a successful change. You should identify testing steps after defined change checkpoints and final change completion. In addition, allocate sufficient time for testing, both during and following the implementation and backout, if necessary. In some cases, you can do testing prior to the change when new service is involved, such as new circuits or links that are not currently in production. The following additional testing and verification procedures may be pertinent to a network change:

- Extended pings for connectivity and performance (may require many to many)
- Traceroutes
- End-user station network and application testing
- Test calls or traffic generation for performance-related changes
- Bit error rate tester (BERT) for new circuits
- Display errors
- Log file verification
- List trace verification
- Display or status command verification
- Network management station availability and verification

After achieving some level of comfort with the change, evaluate what was accomplished:

- Does the change make sense?
- Did the change address the network problem?
- What should be done differently the next time that a change is warranted?

Network management update

Operational readiness requires that you update all network management tools, device configuration, and DNS to reflect the change. Your organization might also have tools for fault management, configuration management, availability measurement, inventory management, billing, and security that require updates. The following are some typical network management update requirements following change:

- Removal of DNS entries and network management system (NMS) management for devices that were removed from the network.
- Standard SNMP configuration entered on devices, including community string, location, support contact, syslog server, trap server, and SNMP server host.
- Trap source, syslog source, and SNMP source configured for loopback.
- Fault management tool update.
- Inventory management tool update.
- Circuit pack and media gateway addresses with DNS name (following naming standard)

Documentation

Possibly the most important requirement in any network environment is to have current and accurate information about the network available at all times. During the process of changing the network, it is critical to ensure that documentation is kept up to date. Network documentation should include the following:

- Detailed physical layer drawing that displays all network devices that have a medium risk (or higher) on the network. The drawing should include rack layouts, cable connections, and devices.
- Detailed network layer drawing of all network devices that have a medium risk (or higher) on the network. The drawing should include addresses, and IP subnetwork and VLAN information.
- Out-of-band management access maps and documentation.
- Solution templates.
- Detailed numbering plans and assignments.
- Detailed dial plan and call routing information.
- VLAN numbering plans and assignments.
- Network Region assignments.
- Naming standards for all network devices.
- Software code and hardware types that are currently implemented and supported.
- Protocol filtering criteria and methodologies.
- Routing protocols standards and supported modifications from default settings.
- Global configuration standards.
- Inventory database for all physical connectivity and contact information.

In addition, Avaya recommends that you develop a matrix that contains information about user groups, the applications they require, and the servers (addresses and locations) that host these applications. This information is necessary to ensure that users continue to have the level of access and performance they require during and after the change. In addition, previously used test plans assist in simplifying future changes, and they may assist in troubleshooting problems that occur because of a change.

High-Level process flow for emergency change management

Unfortunately, not all situations that occur in a network environment are conducive to the extensive research and planning described in the previous section. Sometimes you must make more immediate changes to restore network connectivity following a network outage.

The procedures that you put in place to handle emergency changes should be flexible enough to facilitate rapid resolution of the problem, including documentation of who is authorized to make emergency changes to the network, and how to contact these individuals. You should either have a sufficient number of people who can resolve network emergencies, or those people should be easily accessible at all times to prevent a roadblock in the problem resolution process.

It is critical to maintain both communication and the integrity of documentation through an emergency change. This is the time when documentation is needed most, so documenting the steps that are taken to resolve the problem is very important.

Finally, when considering changes, you should think about not only whether the change will resolve the existing problem, but also whether the change will cause other network problems. Steps that are critical for an emergency change process are shown in the process flow below.

Issue determination

It is usually obvious when an emergency change is required. However, exactly what change is required may not be obvious. For Avaya equipment, you should include the appropriate Avaya Services personnel in the troubleshooting process. In many cases, problems with Avaya equipment will be fixed by Avaya Services expert systems, or a technician will be dispatched before users are aware of the problem.

When taking corrective action, it is imperative that you implement only one change at a time. Otherwise, if the problem is resolved by multiple changes, it is impossible to pinpoint which change actually fixed the problem. Or worse, if other problems are introduced, it is impossible to determine which change was the cause of the new fault. Each change should go through the full process outlined above before you begin on the next change. If a change is shown to have no effect, you should back out of it before you begin the next change. The single exception is when the initial change is a prerequisite to the next change that is under consideration.

Limited risk assessment

In most cases, the amount of risk assessment done in an emergency situation is directly proportional to the scope of the change, and inversely proportional to the effect of the network outage. For example, the scope of changing a Communication Manager release is much greater than that of changing a protocol address. Similarly, the same change would go through increased scrutiny if a single user is unable to access the network rather than if an entire site loses connectivity.

Ultimately, risk assessment is the responsibility of the support person who implements the change. For this, the engineer should rely on personal experience, as well as that of associated support personnel. Many of the ideas given in the section on Planned Change Management can be adapted to the emergency change environment, but on a more limited scale. For instance, you can use the Avaya Support Centre Web site (support.avaya.com/), or even use a limited test bed simulation, depending on your situation.

Finally, as part of the limited risk assessment, you should determine which users might be affected by the change.

Communication

Although it is not always be possible to notify all users of all changes (especially in emergency situations), the users certainly appreciate any warning that you can provide. You should also communicate the details of any emergency changes with the change manager, and allow the change manager to maintain metrics on emergency changes and root causes. The information may also affect the scheduling or the rollout of future changes.

Documentation

Updating documentation is critical to ensure valid, up to date information. During unplanned changes, it can be easy to forget to make updates because of the frantic nature of emergencies. However, undocumented change solutions often result in increased time out of service if the solution is unsuccessful.

It helps to document changes before they are made in emergency situations from a central location, perhaps at the change manager level. If a central support organization does not exist to document changes before they occur, different individuals might make changes at the same time, not knowing about each other’s activities. The following types of documentation often require updates during a change: drawings, IP/VLAN database, engineering documents, dial plan, troubleshooting procedures, and server/application/user matrices.

Implementation

If the process of assigning risk and documentation occurs before the implementation, the actual implementation should be straightforward. Beware of the potential for changes to come from multiple support personnel without their knowing about each other’s changes. This scenario can lead to increased potential time out of service and misinterpretation of the problem.

Test and evaluation

In this phase, the person who initiated the change is responsible for ensuring that the emergency change had the desired affect and if not, restarting the emergency change process. Steps to take in the investigation of the change include the following:

- Observe and document the impact of the change on the problem.
- Observe and document any foreseen or unforeseen side effects of the change.
- Determine whether the problem is resolved, and if so, make sure all necessary documentation and network management updates occur to properly reflect the change.
- If the change is unsuccessful, back out, and continue the emergency change process until the problem is resolved or a workaround is in place.

Once the change is deemed successful, send all emergency change documentation to the change controller for review and documentation by the change control team. The change controller and change review team should perform a post mortem on the problem to determine potential improvements to prevent future emergency changes of this type. You should also send the information to engineering or architecture groups for review, and allow them the opportunity to change solution templates, standard software versions, or network designs to better meet the goals or requirements of your organization.

Performance indicators for change management

Performance indicators provide the mechanism for you to measure the success of your change management process. We recommend that you review these indicators monthly to ensure that change planning and change management are working well.

- Change management metrics by functional group
- Targeting change success
- Change history archive
- Change planning archive
- Configuration change audit
- Periodic change management performance meeting

Change management metrics by functional group

Change management metrics by functional group include the percentage and quantity of change success by functional group and risk level. Emergency changes should be identified separately in the metrics by functional group, including the success rate for attempted fixes. Functional groups include any IT teams making changes, possibly including telephony administration, network administration, database groups, application teams, and facilities. Risk level is important, because generally higher-risk changes fail or create incidents. You might define change failure as any change that is backed out or causes a problem that results in time out of service for the users.

Determining change-related incidents can be difficult. You should contact the user who is identified on the change request form following the change to get an understanding of change success. The change controller might also have a help-desk database available that includes problems closed because of change-related issues.

Targeting change success

To target change success, you should start with a baseline of change management metrics. The change controller can then identify potential issues and set overall goals. A reasonable overall goal for change success in high-availability networks should be 99% across all functional groups. If your organization is experiencing a higher rate of change failure, the rate should be targeted for improvement.

Change history archive

The change controller is also responsible for archiving the change history. Creating a spreadsheet with functional group success and failure columns and month rows is sufficient for archival. Change history archives can help identify current issues that are based on past change rates and available resources. The information can also be used to investigate change rates in general for overall planning purposes.

Change planning archive

The change controller should archive change planning documentation, such as network engineering documents, to create a reference of examples for future successful projects. If the change controller notices change problems, the controller can refer to the change planning document to investigate how well the particular issue was documented before the change. Over time, the change controller might ask to have additional information added to future change planning documents for higher-risk changes to help ensure success.

Periodic performance meeting

Each month, it is important to review the metrics that you collect, including the following:

- Change quantity and risk level
- Change failure quantity and post mortems
- Emergency changes and post mortems
- Change management goals
- Undocumented changes

The functional manager should review the metrics, and report to the appropriate teams for improvement.

Appendix B. Access list

This appendix provides guidelines for configuring access lists to facilitate basic Avaya IP Telephony functionality.

The ports used by the Avaya call server are fairly fixed and well known. The ports used by the endpoints are more variable and random. As a result, it is simpler to tailor access lists based on call server ports. Table 62, Access list guidelines to support IP Telephony, on page 281 contains access list guidelines for supporting IP Telephony.

Table 62: Access list guidelines to support IP Telephony ***1 of 3***

Action	From	TCP/UDP port or protocol	To	TCP/UDP port or protocol	Notes
Permit	Any C-LAN	UDP 1719	Any endpoint	UDP any	The C-LAN uses UDP port 1719 for endpoint registration (RAS).
Permit	Any endpoint	UDP any	Any C-LAN	UDP 1719	
Permit	Any C-LAN	TCP 1720	Any endpoint	TCP any	The C-LAN uses TCP port 1720 for H.225 call signaling.
Permit	Any endpoint	TCP any	Any C-LAN	TCP 1720	
Permit	Near-end C-LAN	TCP 1720	Far-end C-LAN	TCP 1720	This is to facilitate IP trunking between two Avaya call servers, and must be done for each IP trunk.
Permit	Far-end C-LAN	TCP 1720	Near-end C-LAN	TCP 1720	
Permit	Any MedPro	UDP port range on IP Network Region form	Any endpoint	UDP any	This is one way to facilitate audio streams between MedPros and endpoints.
Permit	Any endpoint	UDP any	Any MedPro	UDP port range on IP Network Region form	
Permit	Any MedPro	UDP port range on IP Network Region form	Any endpoint	UDP any	This is another way to facilitate RTP/RTCP audio streams between MedPros and endpoints.

1 of 3

Table 62: Access list guidelines to support IP Telephony ***2 of 3***

Action	From	TCP/UDP port or protocol	To	TCP/UDP port or protocol	Notes
Permit	Any endpoint	UDP any	Any endpoint	UDP any	This is to facilitate RTP/RTCP audio streams between direct IP-IP (shuffled) endpoints.
Permit	Any R300	UDP 1900-2100 RTP/RTCP	Any MedPro or endpoint	UDP varies --	The R300 uses this UDP port range for audio, which can be used to further restrict the access list if desired.
Permit	Any MedPro or endpoint	UDP varies RTP/RTCP	Any R300	UDP 1900-2100 --	
Permit	Any R300	UDP 1900-2100 RTP/RTCP	Any R300	UDP 1900-2100 --	
Permit	Any IP Telephone (hardphone)	UDP any	DNS server(s)	UDP 53 (dns)	These are all services used by the IP Telephone. TFTP is difficult to isolate to a port range. The GET and PUT requests from the client go to the UDP port 69 on the server, but all other messages go between random ports.
Permit	DNS servers	UDP 53 (dns)	Any IP Telephone (hardphone)	UDP any	
Permit	Any IP Telephone (hardphone)	UDP 68 (bootpc)	DHCP server(s)	UDP 67 (bootps)	
Permit	DHCP servers	UDP 67 (bootps)	Any IP Telephone (hardphone)	UDP 68 (bootpc)	
Permit	Any IP Telephone (hardphone)	TFTP	TFTP server(s)	--	
Permit	TFTP servers	TFTP	Any IP Telephone (hardphone)	--	
Permit	SNMP management stations	UDP any	Any IP Telephone (hardphone)	UDP 161 (snmp)	

2 of 3

Table 62: Access list guidelines to support IP Telephony ***3 of 3***

Action	From	TCP/UDP port or protocol	To	TCP/UDP port or protocol	Notes
Permit	Any IP Telephone (hardphone)	UDP 161 (snmp)	SNMP management stations	UDP any	
Permit	Any Avaya device	ICMP Echo	Any	--	Avaya devices ping other devices for various reasons. For example, C-LANs ping endpoints for management purposes; MedPros ping C-LANs to gauge network performance across an IP trunk; IP Telephones ping TFTP servers for verification purposes.
Permit	Any	ICMP Echo Reply	Any Avaya device	--	

3 of 3

Table 63, Access list guidelines for Avaya S8700 and S8300 Media Servers, on page 283 contains access list guidelines that pertain to Communication Manager platforms, including the S8700 and S8300 Media Servers. The S8700 enterprise interface, which is the one that is connected to the enterprise network (versus the control network), is eth4 on Multi-Connect systems and eth0 on IP-Connect systems.

Table 63: Access list guidelines for Avaya S8700 and S8300 Media Servers ***1 of 2***

Action	From	TCP/UDP port or protocol	To	TCP/UDP port or protocol	Notes
Permit	S8700 enterprise interface	TCP any	S8300 LSP	TCP 514	This allows the S8700 to synchronize translations with the S8300 LAN Spare Processor (LSP). A TCP session is initiated from the S8700 to the S8300 TCP port 514. A second session is then initiated from the S8300 to the S8700 TCP port range 512-1023.
Permit	S8300 LSP	TCP 514	S8700 enterprise interface	TCP any	
Permit	S8300 LSP	TCP any	S8700 enterprise interface	TCP 512-1023	

1 of 2

Table 63: Access list guidelines for Avaya S8700 and S8300 Media Servers ***2 of 2***

Action	From	TCP/UDP port or protocol	To	TCP/UDP port or protocol	Notes
Permit	S8700 enterprise interface	TCP 512-1023	S8300 LSP	TCP any	
Permit	Avaya Site Administration workstation	TCP any	S8700-ent-intf or S8300	TCP 5023	This allows an administrator to log in through Avaya Site Administration to a call server.
Permit	S8700-ent-intf or S8300	TCP 5023	Avaya Site Administration workstation	TCP any	
Permit	Web admin station	TCP any	S8700-ent-int or S8300	TCP 80	This allows secure and unsecure web access to a call server. The call server redirects unsecure sessions to https.
Permit	S8700-ent-int or S8300	TCP 80	Web admin station(s)	TCP any	
Permit	Web admin station	TCP any	S8700-ent-int or S8300	TCP 443	
Permit	S8700-ent-int or S8300	TCP 443	Web admin station(s)	TCP any	
Permit	S8700-ent-int or S8300	UDP any	DNS server(s)	UDP 53 (dns)	Optional services used by S8700 and S8300.
Permit	DNS server(s)	UDP 53 (dns)	S8700-ent-int or S8300	UDP any	
Permit	S8700-ent-int or S8300	UDP any	NTP server(s)	UDP 123 (ntp)	
Permit	NTP server(s)	UDP 123 (ntp)	S8700-ent-int or S8300	UDP any	
Permit	G700	TCP any	S8300 or other call server	TCP 2945	H.248 signaling between G700 Media Gateway and S8300 or other call server. G700 initiates the session.
Permit	S8300 or other call server	TCP 2945	G700	TCP any	
Permit	Call server	IP any	IPSI board	IP any	There are too many system control messages and services between the call server and IPSI board to filter each one individually.
Permit	IPSI board	IP any	Call server	IP any	

2 of 2

Appendix C. Multi-VLAN example

Figure 81, Sample Multi-VLAN scenario for Cajun P330 code 3.2.8 and Cisco, on page 285 is a sample multi-VLAN scenario. Suppose there is a Cisco router that is connected to a P330 switch that contains two VLANs, one for the VoIP devices and one for the personal computers. To conserve ports and cabling, the computers are connected to the telephones and the telephones are connected to the P330 switch.

Figure 81: Sample Multi-VLAN scenario for Cajun P330 code 3.2.8 and Cisco

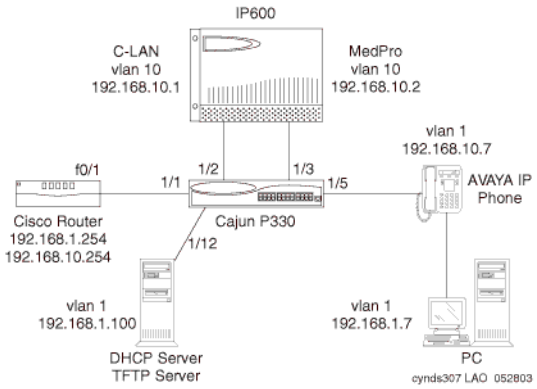

Table 64: Command set and explanations for multi-VLAN example ***1 of 4***

Command	Notes
Cisco router configuration	
interface FastEthernet0/1	
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	Forwards DHCP requests to the DHCP server.
P330 configuration (bind-to-static option)[1]	

1 of 4

Table 64: Command set and explanations for multi-VLAN example ***2 of 4***

Command	Notes
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	
set port vlan 10 1/2	Port/native VLAN changed to 10 on this port.
set port spantree disable 1/2	
set port vlan 10 1/3	
set port spantree disable 1/3	
set port vlan-binding-mode 1/5 static	Port in static binding mode by default, but command shown.
set port static-vlan 1/5 10	In addition to v1, v10 statically bound to port, but not a trunk port.
set port spantree disable 1/5	Port 1/12 for the DHCP/TFTP server already has port/native VLAN 1.
P330 configuration (bind-to-configured option)	
set vlan 1 (VLAN 1 configured)	
set vlan 10 (VLAN 10 configured)	
set port vlan-binding-mode 1/1 bind-to-configured	Port bound to configured VLANs 1 and 10.
set trunk 1/1 dot1q	Port connected to Cisco router is an 802.1Q trunk port.
set port spantree disable 1/1	
set port vlan 10 1/2 (port/native VLAN changed to 10 on this port)	
set port spantree disable 1/2	
set port vlan 10 1/3	
set port spantree disable 1/3	
set port vlan-binding-mode 1/5 bind-to-configured	Bound to configured VLANs but not a trunk port.
set port spantree disable 1/5	
If the P330 switch were a Cisco CatOS switch instead	
First, invoke the **set port host** command on all user ports, and then proceed as follows.	

2 of 4

Table 64: Command set and explanations for multi-VLAN example ***3 of 4***

Command	Notes
set vlan 1005 1/1	Cisco switches do not tag the native VLAN, but the router expects a tag on VLAN 1, so the native VLAN is changed to some unused VLAN.
set trunk 1/1 on dot1q	Port connected to Cisco router is an 802.1Q trunk port.
clear trunk 1/1 2-9,11-1004	Unnecessary VLANs removed; 1, 10, and 1005 remain.
set vlan 10 1/2	Port/native VLAN changed to 10 on this port.
set vlan 10 1/3	
set trunk 1/5 nonegotiate dot1q	Plain 802.1Q trunk port with no Cisco negotiation features.
clear trunk 1/5 2-9, 11-1005	Unnecessary VLANs removed; 1 and 10 remain.
Optional command using auxiliaryvlan on the telephone port instead of explicit trunking	
set port auxiliaryvlan 1/5 10	VLAN 10 is the auxiliaryvlan; only VLANs 1 and 10 on this port; port is an 802.1Q trunk port, though not explicitly configured.
If the P330 switch were a Cisco IOS switch instead	
interface FastEthernet0/1	
switchport trunk encapsulation dot1q	Port connected to Cisco router is an 802.Q trunk port.
switchport trunk native vlan 1005	Cisco switches do not tag the native VLAN, but the router expects a tag on VLAN 1, so the native VLAN is changed to some unused VLAN.
switchport trunk allowed vlan 1,10,1005	VLANs 1, 10, and 1005 allowed on trunk.
switchport mode trunk	
spanning-tree portfast	
interface FastEthernet0/2	
switchport access vlan 10	Port/native VLAN changed to 10 on this port.
spanning-tree portfast	
interface FastEthernet0/3	
switchport access vlan 10	
spanning-tree portfast	
interface FastEthernet0/5	
switchport trunk encapsulation dot1q802.1Q trunk port	

3 of 4

Table 64: Command set and explanations for multi-VLAN example ***4 of 4***

Command	Notes
switchport trunk native vlan 1	Since most PCs do not understand the tag, the Cisco native VLAN must be set as the PC's VLAN. VLAN 1 is already the native VLAN, but command is shown.
switchport trunk allowed vlan 1,10	VLANs 1 and 10 allowed on trunk.
switchport mode trunk	
spanning-tree portfast	
Optional commands using the voice vlan on the telephone port.[2]	
interface FastEthernet0/5	
switchport trunk encapsulation dot1q	
switchport trunk native vlan 1	
switchport voice vlan 10	VLAN 10 is the voice vlan; unsure if this removes all other VLANs from trunk or not.
	4 of 4

1 All ports have port/native VLAN 1 by default.

2 There really is no reason to do this unless a Cisco telephone will use this port. The configuration is not simpler, as with the CatOS switch and auxiliaryvlan.

IP Telephone configuration

This procedure applies regardless of the Ethernet switch that is being used. Initially placing the IP Telephone on VLAN 10 requires two DHCP scopes, one for VLAN 1 and another for VLAN 10, with identical SSON 176 parameters.

1. Run the telephone through its normal boot-up sequence.

   It will come up with an IP address on VLAN 1 - the port/native VLAN - assuming that the DHCP scope is set up properly.

2. If IP600 R9.2 and IP Telephone R1.1, after the telephone is up and operational on VLAN 1 press **Hold QOS #**.

3. Enable 802.1Q, set the priorities as desired, set the VID to 10, and save the values.

4. Press **Hold RESET #**.

5. Answer **No** to resetting the values.

6. Answer **Yes** to restarting the telephone.

   The telephone comes up with an IP address on VLAN 10, assuming DHCP is relayed and set up properly. From then on the telephone will always come up on VLAN 10 without further manual intervention until the stored values are manually reset. This is because all manually entered values remain in the telephone's NVRAM until manually reset.

With IP600 R9.5 and IP Telephone R1.51, the Hold QOS# values that were set manually in the previous IP Telephone release can be sent to the telephone by the DHCP server, using SSON 176. On the VLAN 1 DHCP scope, add L2Q=1 and L2QVLAN=10 to the existing SSON 176 comma-separated string. For example:

MCIPADD=#.#.#.#,MCPORT=1719,TFTPSRVR=#.#.#.#,L2Q=1,L2QVLAN=10

This causes the telephone to release the VLAN 1 address after the first DHCP sequence, and then enter a second DHCP sequence with tagging enabled to obtain a VLAN 10 address. Because the L2Q parameters are not manually set in this scenario, and thus are not stored in NVRAM, the telephone requires the VLAN 1 DHCP scope every time it reboots. The L2Q parameters should not be added to the VLAN 10 DHCP scope. This is so that in the event a telephone is connected to a port that has VLAN 10 as the port/native VLAN, it will not receive instructions from the DHCP scope to enable tagging. In such a case the telephone would not require tagging to function on VLAN 10, and tagging could result in an incompatibility with the Ethernet switch.

PC configuration

The PC can be statically addressed with a VLAN 1 address, or it can receive a VLAN 1 address through DHCP. No special configurations are required.

Appendix D. DHCP / TFTP

DHCP

This section provides information on possible DHCP servers and generic information on administering a DHCP server.

Required information

Before installing a DHCP server you will need the following required network information:

- Router IP address
- TFTP server IP address
- Subnet mask
- C-LAN IP address(es)
- Communication Manager C-LAN port. Although this may be a value between 0 and 65535, the default value is 1719 and should not be changed unless this conflicts with an existing port assignment.
- TFTP server file path
- Telephone IP address range (both From and To)
- DNS Server addresses (if applicable)

Choosing a DHCP configuration

This section concentrates on the simplest case of the single LAN segment. Extrapolate the information that is provided here for more complex LAN configurations.

WARNING:
Before you start, it is important that you understand your current network configuration. An improper installation can cause network failures or reduce the reliability and performance of your network. See the Network assessment offer for more information about Avaya's comprehensive network performance assessment.

DHCP software alternatives

Two DHCP software alternatives are common to Windows operating systems:

- Windows NT 4.0 DHCP Server
- Windows 2000 DHCP Server

Any other DHCP application might work.

It is the customer's responsibility to install and configure the DHCP server correctly. This appendix is limited to describing generic administration for Avaya IP Telephones.

DHCP generic setup

Set up of a DHCP server involves the following phases:

1. Install the DHCP server software according to vendor instructions.
2. Configure the DHCP server with the available IP addresses for the Avaya IP Telephones.
3. Administer the lease duration (“Infinite” is recommended).
4. Administer the gateway (router) IP addresses.

   If more than one address is listed, the total list may contain up to 127 total ASCII characters, with IP addresses separated by commas with no intervening spaces.
5. Configure the Subnet mask.
6. Administer Option 6 (DNS servers address list).

   If more than one address is listed, the total list may contain up to 127 total ASCII characters, with IP addresses separated by commas with no intervening spaces. At least one address in Option 6 must be a valid, nonzero, dotted decimal address. Otherwise, DNS will fail.
7. Administer Option 15 (DNS Domain Name).

   This string should contain the domain name to be used when DNS names in system parameters are resolved into IP addresses. This domain name is appended to the DNS name before the IP Telephone attempts to resolve the DNS address. Option 15 is necessary if you want to use a DNS name for the TFTP server. Otherwise, you might specify a DOMAIN as part of TFTP customization, as indicated in DNS Addressing on page 4-24.
8. Administer Option 66 (TFTP Server Name).

   NOTE:
   Note: Microsoft DHCP servers support only dotted decimal format for TFTP addresses, not symbolic names. Option 66 need not be used if the TFTP server is identified in the Site Specific Option string (Option 176). However, to simplify configuration, we recommend that you use Option 66. If you use both Option 66 and Option 176 to identify TFTP servers, the values in Option 176 will override the values in Option 66.
9. Administer the IP Telephone-specific DHCP options specifying information, such as TFTP server and DEFINITY C-LAN IP addresses. Use the site-specific option (SSON) at #176. The value for this option should be set to either of the following strings:

   MCIPADD=xxx.xxx.xxx.xxx,MCPORT=yyyy,TFTPSRVR=zzz.zzz.zzz.zzz,TFTPDIR=< path>

   or

   MCIPADD={list of DNS names},MCPORT=yyyy,TFTPSRVR={list of DNS names},TFTPDIR=<path>

   Where xxx.xxx.xxx.xxx is one or more IP addresses for the C-LAN IP circuit pack, yyyy is the C-LAN port (1719), zzz.zzz.zzz.zzz is one or more IP addresses for TFTP servers, and <path> is the location of the location of the upgrade script and application files on the TFTP server.

Each list can contain up to 127 total ASCII characters, with IP addresses separated by commas with no intervening spaces, and with quotes on either end (see the example in the NOTES below). If you use DNS, note that the system value DOMAIN is appended to the hostname that you specify. If DOMAIN is null, the DNS names must be fully qualified. In configurations where the upgrade script and application files are in the default directory, the TFTPDIR=<path> should not be used. You do not have to use Option 176. For example, if the DNS server is specified in Option 6, and the Domain Name is specified in Option 15, you can use the configured names "AvayaTFTPServer" and "Avaya Call Server" for TFTPSRVR and MCIPADD, respectively. The Call Server Name, TFTP Server Name, and SMTP Server Name must each be no more than 32 characters in length. Examples of good DNS administration include the following:

- Option 6: aaa.aaa.aaa.aaa
- Option 15: dnsexample.yourco.com
- Option 66: tftpserver.yourco.com,zzz.zzz.zzz.zzz
- Option 176: MCIPADD=xxxx.xxx.xxx.xxx

Depending on the DHCP application you choose, be aware of the fact that the application most likely will not immediately recycle expired DHCP leases. An expired lease may remain reserved for the original client for a day or more (for example, Windows NT DHCP reserves expired leases for about 1 day). The intent of this reservation period is to protect a client's lease in case the client and the DHCP server are in two different time zones, the computers' clocks are not in synch, or the client is not on the network when the lease expires.

The implication of this fact can be seen in the following example. Assume two IP addresses (hence two possible DHCP leases) and three IP Telephones, two of which are using the two available IP addresses. When the lease expires for the first two telephones, the third will not be able to get a lease (even if the other two telephones have been removed from the network), until the reservation period expires.

The IP Telephone sets the indicated system values to the values of the indicated fields of the DHCPACK message (Table 65, DHCP setting of system values, on page 293).

Table 65: DHCP setting of system values

System value	Set to
IPADD	The yiaddr field
NETMASK	Option #1 (if received)
GIPADD	The first four octets of Option #3 (if received)
TFTPSRVR	The first four octets of the siaddr field

Windows NT 4.0 DHCP server

This section contains details on how to verify and configure the DHCP server included in the Windows NT 4.0 server operating system.

Use Verifying the DHCP server installation below to verify whether the DHCP server is installed. If it is not, install the DHCP server. If it is installed, then proceed with the Initial configuration and the Creating a DHCP scope for the IP Telephones sections.

Verifying the DHCP server installation

Use the following procedure to verify whether the DHCP server is installed.

1. Select **Start->Settings->Control Panel**.
2. Double-click the **Network** icon.
3. Verify that **Microsoft DHCP Server** is listed as one of the Network Services on the Services Tab.
4. If it is listed, continue with the Initial configuration section below. If it is not listed, install the DHCP server.

Initial configuration

The Windows NT 4.0 DHCP server configuration involves setting up a scope for the IP Telephone. A DHCP scope is essentially a grouping of IP devices (in this case IP Telephones) running the DHCP client service in a subnet. The scope is used to define parameters for each subnet. Each scope has the following properties:

- A unique subnet mask used to determine the subnet related to a given IP address.
- A scope name assigned by the administrator when the scope is created.
- Lease duration values to be assigned to DHCP clients with dynamic addresses.

In addition, the DHCP server can assign configuration parameters to a client, and these can be specified for each individual DHCP scope. Setting up of the Windows NT 4.0 DHCP server, requires these steps:

1. Creating a DHCP scope for the IP Telephones
2. Editing custom options
3. Adding the DHCP option
4. Activating the leases

Creating a DHCP scope for the IP Telephones

Use the following procedure to create a DHCP scope for the IP Telephones.

1. Select **Start->Programs->Admin Tools->DHCP Manager**.
2. Expand **Local Machine** in the DHCP Servers window by double clicking on it until the + sign changes to a - sign.
3. Select **Scope->Create**.
4. Define the range of IP addresses used by the IP Telephones.
    - The Start Address should be the first IP address to be used for the IP Telephones.
    - The End Address should be the last IP address to be used for the IP Telephones.
    - Subnet Mask should be set to the value assigned by the network administrator.

5. Perform these steps to exclude any IP addresses that you do not want to be assigned to IP Telephones within the range specified by the Start and End Addresses.
   a. Enter the first IP address in the range that you would like to exclude in the Start Address field under Exclusion Range.
   b. Enter the last IP address in the range that you would like to exclude in the End Address field under Exclusion Range.
   c. Click the **Add** button.
   d. Repeat steps a. through c. for each IP address range that you would like to exclude.

   Example

   Suppose the ranges of the IP addresses that are available for your IP Telephone network are:

   - 135.254.76.7 to 135.254.76.80
   - 135.254.76.90 to 135.254.76.200
   - 135.254.76.225 to 135.254.76.230

   Your start address and end address should then be 135.254.76.7 and 135.254.76.230, respectively.

   You should exclude the ranges 135.254.76.81 to 135.254.76.89 and 135.254.76.201 to 135.254.76.224.

   NOTE:
   Avaya recommends that the IP Telephones be provisioned with sequential IP addresses.

6. Under **Lease Duration**, select the **Limited To** option and set the lease duration to the maximum.
7. Enter a sensible name for the Name field, such as “DEFINITY IP Telephones.”
8. Click **OK**.

   A dialog box prompts you: ‘**Activate the new scope now?**’
9. Click **No**.

Editing custom options

Use the following procedure to edit custom options:

1. Select **DHCP Options->Defaults**.
2. Click **New**.
3. Enter “46XXOPTION” for your custom in the Add Option Type dialog.
4. Select **Data Type of String,** and enter **176** in the Identifier field.
5. Click **OK**.

   The DHCP Options menu is displayed.
6. Select the **Option Name** for 176 and set the **value string**.
7. Click **OK**.
8. Select **003 Router** from the list for the Option Name field.
9. Click **Edit Array**.
10. Enter the Gateway IP address for the New IP Address field.
11. Click **Add,** and then click **OK**.

Adding the DHCP option

Use the following procedure to add the DHCP option:

1. Highlight the scope that you just created.
2. Select Scope under DHCP OPTIONS.
3. Select the 176 option that you created from Unused Option List.

   Avaya recommends that the IP Telephones be provisioned with sequential IP addresses. You will activate the scope when all options have been set.
4. Click **Add**.
5. Select **option 003** from the Unused Options List.
6. Click **Add**.
7. Click **OK**.
8. Chose the **Global parameter** under DHCP Comments.
9. Select the **176 option** that you created from the Unused Option List.
10. Click **Add**.
11. Click **OK**.

Activating the leases

To activate the leases, click **Activate** under the Scope Menu, and the icon for the scope should light.

Verifying your configuration

This section describes how to verify that the 46XXOPTIONs are correctly configured for the Windows NT 4.0 DHCP server.

Verify the default option (176 46XXOPTION)

Use the following procedure to verify the default option:

1. Select **Start>Programs>Admin Tools>DHCP Manager**.
2. Expand **Local Machine** in the DHCP Servers window.
3. In the DHCP Servers frame, click the scope for the IP Telephone.
4. Select **Defaults** from the DHCP_Options menu.
5. In the Option Name list, select **176 46XXOPTION**.
6. Verify that the Value String box contains the correct string.
7. If not, update the string, and click **OK** twice.

Verify the scope option, 176 46XXOPTION

Use the following procedure to verify the scope option:

1. Select **Scope** under DHCP OPTIONS.
2. In the Active Options scroll list, click on **176 46XXOPTION**.
3. Click **Value**.
4. Verify that the Value String box contains the correct string from the DHCP generic setup section.

   If not, update the string and click **OK**.

Verify the global option, 176 46XXOPTION

1. Select **Global** under DHCP OPTIONS.
2. In the Active Options list, click **176 46XXOPTION**.
3. Click **Value**.
4. Verify that the Value String box contains the correct value from the DHCP generic setup section.

   If not, update the string and click the **OK** button.

Windows 2000 DHCP server

This section describes the configuration of the DHCP server in Windows 2000.

Verifying the DHCP server installation

Use the following procedure to verify whether the DHCP server is installed:

1. Select **Start>Program>Administrative Tools>Computer Management**.
2. Under Services and Applications in the Computer Management tree, find DHCP.
3. If DHCP is not installed, install the DHCP server. Otherwise skip directly to Creating and configuring a DHCP Scope for instructions on server configuration.

Creating and configuring a DHCP Scope

Use the following procedure to create and configure a DHCP scope:

1. Select **Start >Programs >Administrative Tools>DHCP**.
2. In the console tree, click the DHCP server to which you want to add the DHCP scope for the IP Telephones. This is usually the name of your DHCP server.
3. Select **Action>New Scope** from the menu.

   Windows displays the New Scope wizard to guide you through rest of the setup.
4. Click **Next**.

   The Scope Name dialog box is displayed.
5. Enter a name for the scope in the Name field.
6. Enter a brief comment in the Description field.

7. Click **Next** when finished.

   The IP Address Range dialog box is displayed.

8. Define the range of IP addresses used by the IP Telephones.

   The Start IP Address should be the first IP address available to the IP Telephones. The End IP Address should be the last IP address available to the IP Telephones.

9. Define the subnet mask in one of two ways:

   - The number of bits of an IP address to use for the network/subnet IDs
   - The subnet mask IP address in dotted-quad notation

   Enter only one of these values.

10. Click **Next** when finished.

    The Add Exclusions dialog box is displayed.

11. Exclude any IP addresses in the range specified in the previous step that you do not want to be assigned to an IP Telephone.

    a. Enter the first IP address in the range that you want to exclude in the Start Address field under Exclusion Range.

    b. Enter the last IP address in the range that you want like to exclude in the End Address field under Exclusion Range.

    c. Click the **Add** button.

    d. Repeat steps a. through c. for each IP Address range that you want to exclude.

    Example

    Suppose the ranges of IP addresses available for your IP Telephone network are:

    - 135.254.76.7 to 135.254.76.80
    - 135.254.76.90 to 135.254.76.200
    - 135.254.76.225 to 135.254.76.230

    Your Start IP Address and End IP Address entered on the IP Address Range dialog box should then be 135.254.76.7 and 135.254.76.230, respectively.

    On the Add Exclusions dialog box, you should exclude the following ranges:

    - 135.254.76.81 to 135.254.76.89
    - 135.254.76.201 to 135.254.76.224

12. Click **Next** when all the exclusions have been entered.

    The Lease Duration dialog box is displayed.

13. Enter **30 days** in the lease duration for all telephones that will receive their IP addresses from the server. This is the duration after which the IP address for a device expires and needs to be renewed by the device.

14. Click **Next**.

    The Configure DHCP Options dialog box is displayed.

    You can add additional exclusion ranges later by right clicking on the Address Pool under the newly created scope and select the New Exclusion Range option.

15. Click **No, I will activate this scope later**.

    The Router (Default Gateway) dialog box is displayed.

16 For each router or default gateway, enter the IP address and click **Add**.

17 When you are finished, click **Next**.

The Completing the New Scope Wizard dialog box is displayed.

18 Click **Finish**.

The new scope is added under the server in the DHCP tree. It is not yet active and will not assign IP addresses.

19 Highlight the newly-created scope, and select **Action->Properties** from the menu.

20 Under Lease duration for DHCP clients, select **Unlimited** and then click **OK**.

WARNING:
IP Address leases are kept active for varying periods of time. To avoid having calls terminated suddenly, make the lease duration unlimited.

Adding DHCP options

Use the following procedure to add DHCP options to the scope you created in the previous procedure:

1 On the DHCP window, right-click the **Scope Options** folder under the scope you created in the last procedure.

A menu is displayed.

2 Click **Configure Options**.

The Scope Options dialog box is displayed.

3 In the General tab page, under the Available Options, select **066 Boot Server Host Name Options**.

The String Value dialog box is displayed.

4 Enter the TFTP Server addresses in the string value.

Use the same TFTPSRVR value format as discussed in the TFTP Generic Setup section. For example, if you had a TFTP server at IP address zzz.zzz.zzz.zzz and a second TFTP server at address tftpserver.yourco.com, in the string value field enter "zzz.zzz.zzz.zzz,tftpserver.yourco.com"

5 Also under the Available Options, select **176 Site-Specific Options**.

6 Click **Add**, and then click **Edit Array**.

The IP Address Array Editor dialog box is displayed.

7 Enter the IP Addresses for the TFTP Servers that support the IP Telephones.

8 Click **OK**.

The Predefined Options and Values dialog box is displayed.

9 Click **OK**.

The Predefined Options and Values dialog box is closed, leaving the DHCP dialog box enabled.

10 Expand the newly created scope to reveal its Scope Options.

11 Click **Scope Options**, and select **Action>Configure Options** from the menu.

12 In the General tab page, under the Available Options, select **176 Site-Specific Options**.

13 In the Data Entry box, enter the DHCP IP Telephone option string as described in the DHCP generic setup section.

14 From the list in Available Options, select **003 Route**r.

15 Enter the gateway (router) IP address.

16 Click **Add**.

17 Click **OK**.

Activating the New Scope

Use the following procedure to activate the new scope.

1 In the DHCP console tree, click the IP Telephone Scope created.

2 From the Action menu, select **Activate**.

The small red down arrow over the scope icon disappears, indicating that the scope has been activated.

NOTE:
You can enter the text string directly on the right side of the Data Entry box under the ASCII label.

TFTP

This section describes how to set up a TFTP server for downloading software updates to the Avaya IP Telephones.

WARNING:
The files defined by the TFTP server configuration must be accessible from all IP Telephones. Ensure that the filenames match the names in the upgrade script, including case, since some TFTP servers are case sensitive.

TFTP Generic Setup

The following phases are involved in setting up a TFTP server.

- Install the TFTP server software. The section below describes how to configure Avaya's TFTP application.
- Configure the file path parameter to the directory where the files are to be stored. For increased security, it is also recommended that you disable the ability to upload to the server. This option may be not available to all TFTP servers.
- In addition, you may want to enable the transfer size option (tsize) if your TFTP server supports it. This allows the IP Telephone to display the progress of the transfer by displaying the total number of data blocks.
- Download the upgrade script file and application file from the Avaya Web site (www.avaya.com/support) to the directory as specified by the file path.

Avaya TFTP (Suite Pro) configuration

Use the following procedure to configure the Avaya TFTP server:

1. Select **Start->Programs->Avaya TFTP Server >TFTPServer32** to run the TFTP Suite Pro server.

   The TFTP server starts.

   WARNING:
   You must restart Avaya TFTP manually every time you reboot your TFTP server.

2. Select **System->Setup**.

   On the Outbound tab (page 1) the Outbound path should be the TFTP file path.

3. Select **Enable Path**.
4. Under the Options tab, select **No Incoming**.
5. Under the Client Limits tab, Set the Maximum Simultaneous Clients to infinite by dragging the slide bar all the way to the right.
6. Place the 46xxupgrade.scr file in the file path directory. (The filename 46xxupgrade.scr is an example, not the filename you will use.)

Appendix E. Troubleshooting

This troubleshooting section addresses the following topics:

No Dial Tone

Terminology

- **No Dial Tone**. The light on the IP Telephone is on and the display is working, but no dial tone is heard after the IP Telephone goes off-hook.
- **Message Sequence Trace** (MST). A Communication Manager tool that captures the message flow between software processes, and also between software processes and the system I/O.

Symptom resolution procedure

1. Has a network assessment ever been done and has the network not been modified after the assessment?

   Y. There may be a network or MedPro problem. All possibilities need to be explored, go to Step 3.

   N. The network may not be compliant with the Avaya's network requirements. If the problem cannot be resolved by using the steps described below, then a (re-)assessment may need to be done, go to Step 2.

2. Look at the large pattern first: do other IP Telephones experience the same problem (assuming multiple IP Telephones are installed)?

   Y. There may be a C-LAN, a MedPro, or a network problem. All possibilities need to be explored, go to Step 3.

   N. Go to Step 3 because it could still be that the IP Telephone is the only one connected registered with the C-LAN or only one assigned to the MedPro that has the problem.

3. Because there is a problem with many IP Telephones, is the C-LAN that the IP Telephone is registered to operational?

   a. Execute the **status station <*extension #*>** command.

   b. Scroll to page 3 (Call Control Signaling).

   In the `Switch Port` field look up the slot location of the C-LAN circuit pack that is responsible for the IP Telephone, for example 07D1703.

c Verify that the IP Telephone is registered properly with Communication Manager by checking the `Registration Status` field on that page. If the IP Telephone is not registered, then ensure that it is registered.

d Execute the **test board 07D17** command.

This should indicate (all tests should pass) that the C-LAN board is operational according to the software. If any test fails then refer to the *Avaya Communication Manager Hardware Maintenance Manual*.

e Execute the **status link** command and ensure that the link is in service.

If the link is out of service, then check the Installation Manual to make sure the C-LAN has been installed correctly.

If both the **test board** and the **status link** command do not show any problems with the C-LAN board, then go to Step 4.

4 Because there is a problem with many IP Telephones, is the MedPro circuit pack operational?

a Go off-hook on the IP Telephone.

b Execute the **status station <*extension #*>** command.

c Scroll to page 3 (Call Control Signaling).

In the Audio Channel section if the `Switch Port` field contains a port location, then go to Step d; otherwise go to Step e.

d We have a MedPro port that is dynamically allocated to the IP call. Go to page 5 of the Station form and check the `Last Tx Sequence` field that shows the RTP sequence number of the last packet sent by the MedPro to the IP Telephone. This sequence number should increase at a regular rate when you run the **status station** command repeatedly. If it does not increase, then there is likely a MedPro hardware or firmware problem. *Use the Communication Manager Maintenance Manual* to resolve the issue. If packets are being transmitted normally, then go to Step 5.

If the audio channel on the Station form is blank, this might be due to an inability of the MedPro to allocate resource for the call.

e Execute the **list measurements ip dsp-resource** command to determine whether there are sufficient MedPro resources in the system.

Check for denials, blockage and out-of-service condition. If any of those measurements are greater than 0, this may indicate that any of the following problems might exist on the MedPro:

- The MedPro may have run out of DSP resources. After some users will have disconnected the problem will resolve itself. If this is a regular problem, another MedPro board needs to be installed.
- The firmware should be FW46 or later. Replace the firmware if needed.
- One of the DSPs may be bad or there could be firmware problem. This can be checked in the hardware error log by executing the **display errors** command. Escalate the problem to the next Avaya Maintenance tier.
- Communication Manager might not be able to find a MedPro in the network region where the IP Telephone resides.

f If there are no MedPro problems, then go to Step 5.

5 Can the MedPro ping the IP Telephone?

a Execute the **status station <*extension #*>** command.

b Scroll to page 3 (Call Control Signaling).

In the `Switch Port` field look up the slot location of the C-LAN circuit pack that is responsible for the IP Telephone, for example 07D1717.

c Get the IP address of the IP Telephone from the `Set-end IP Addr` field.

Note that hereafter, to simplify the description, it is assumed that this address is 135.9.42.105.

d Execute the **ping ip-address 135.9.42.105 board 07D17** command.

Y. The MedPro receives echo replies from the IP Telephone, thus there is network connectivity between the MedPro and the telephone. The IP Telephone might be faulty. Replace the IP Telephone with another one to verify this. If this still does not solve the problem, go to Step 7.

N. The IP Telephone is invisible to the MedPro. Go to Step 6.

6 Find out where the ping from the MedPro terminated.

a Execute the trace-route ip 135.9.42.105 board 07D15 command.

If network connectivity cannot be established between the MedPro and the telephone, one hop will be delineated with "3 *."

b Begin analyzing the network at the previous router (the last IP address displayed).

7 Are the transmission speed and transmission duplex (HDX, FDX) of the MedPro and the Ethernet switch compatible?

a Check this by verifying the layer 1 port statistics on the Ethernet switch connected to the MedPro.

Look for Frame check sequence errors, late collisions, and runts.

Y. Go to Step 8.

N. Change the port settings on the Ethernet switch and/or the IP Interfaces form (**change ip-interfaces**) in Communication Manager to make speed and duplex compatible.

NOTE:
If one side's duplex is set to **autonegotiate**, the other side must also be set to **autonegotiate** or **half**. Locking one side to full duplex will cause errors.

If this resolves the problem then no further steps need to be taken; otherwise go to Step 8.

8 Are the transmission speed and transmission duplex (HDX, FDX) of the IP Telephone and the Ethernet switch compatible?

a Verify the Layer 1 port statistics on the Ethernet switch connected to the IP Telephone (frame check sequence errors, late collisions, and runts).

NOTE:
The switch port *must* be set to **autonegotiate** or **half duplex** or there will be a duplex mismatch.

Y. Go to Step 9.

N. Change the port settings to make speed and mode compatible. If this resolves the problem then no further steps need to be taken, otherwise go to Step 9.

9 Does the C-LAN receive an off-hook indication from the IP Telephone?

a Check this by running an MST (message sequence tracer) trace for the IP Telephone.

Y. The IP Telephone at least generates messages. Go to Step 10.

N. The IP Telephone might be the problem. Replace the IP Telephone with another IP Telephone and check if this solves the problem. If it does, then order a new IP Telephone. If it does not solve the problem, then escalate the problem to the next Avaya Maintenance tier.

10 There must be a network problem. Compliance with the Avaya network requirements might be an issue as well, and a (re-)assessment may need to be done. Install a protocol analyzer in the network to capture live traffic and analyze the network in further detail.

Talk path

Terminology

General

- **One-way talk path**. The unidirectional voice audio path from one IP Telephone to another.
- **No-way talk path**. No voice talk path exists between IP Telephones.

Communication Manager

- **Shuffling**. The process of rerouting the voice channel connecting two IP endpoints so that the voice completely goes through an IP network without using intermediate Communication Manager MedPro resources.
- **Hairpinning**. Rerouting of a voice channel connecting two IP endpoints so that the voice goes through the MedPro circuit pack in IP format, without having to go through the gateway's Time Division Multiplexing (TDM) bus. Only the IP and RTP packet headers are changed as the packet goes through the MedPro. This requires that both endpoints use the same codec.

Symptom resolution procedure

Three possible problem locations can be identified if users report a one-way or no-way talkpath between IP Telephones:

- The network
- The MedPro circuit pack (if the call is not shuffled)
- The IP Telephone

For the resolution of this symptom, first disable shuffling (if turned on). See later for the details. This forces traffic to use the media processor, and simplifies the analysis of the network. Then, among other steps, check whether audio/dial-tone can be received by the IP Telephones involved in the call. If necessary, the media processor can check the connectivity of the IP Telephones and their local subnetwork using pings. Layer 1 errors can also be checked.

1 Has a network assessment ever been done and has the network not been modified after the assessment?

Y. There may be a network problem, a MedPro problem, or the IP Telephone may have outdated software. All possibilities need to be explored, go to Step 2.

N. The network may not be compliant with the Avaya's network requirements. If the problem cannot be resolved by using the procedures described below, a (re-)assessment may need to be done, go to Step 2.

2 Do other IP Telephones one the same VLAN/subnet/floor experience the same problem?

Y. There may be a network problem, or multiple IP Telephones may have outdated software. If the IP Telephone firmware version is prior to R1.5 there is a Communication Manager patch #3483 for 1-way talk path problems. Obtain through Avaya Tier3 maintenance support. If this solves the problem then no further steps are needed, otherwise go to Step 3.

N. Go to Step 3.

3 Is the call shuffled?

a Run the **status station <*extension#*>** command.

b Scroll to page 3 (Call Control Signaling).

If the `Audio Connection Type` field is

- **ip-direct**, then it is shuffled.
- **ip-tdm** or **ip-hairpin,** then it is not shuffled.

Y. Turn off shuffling with the **change station <*extension #*>** command. On the form, the field `Direct IP-IP Audio Connections` should be set to **n**.

If this resolves the problem, then there is a network problem that prevents the two IP Telephones from communicating directly. See the note below and go to Step 8.

If this does not resolve the problem, there could be a network problem or a MedPro problem. Although a network problem is still most likely, keep shuffling disabled and go to Step 4.

NOTE:
The remote ping and remote traceroute commands can be used to help pinpoint the location in the network where shuffled calls experience problems.

N. Go to Step 4.

4 Does the IP Telephone receive dial-tone?

a On both IP Telephones put the call on hold by picking up the handset and listening to any dial-tone received.

Y. Go to Step 5.

N. Go to the No Dial Tone section.

5 Are there any Communication Manager errors logged for MedPro or the IP Telephone?

a Run the **display errors** command.

Check the hardware error log and the denial event log for errors against the IP Telephone with the particular extension.

Y. Use the information in the error log and the *Communication manager Maintenance Manual* to correct the errors. If this solves the problem, no further steps are needed. Otherwise, go to Step 6.

N. Go to Step 6.

6 Is voice audio received by the MedPros from both IP Telephones in the call?

a Execute the **status station <*extension #*>** command.

b Scroll to page 5 (Network Status).

Look at the `Last Rx/Tx Sequence` field data. These RTP sequence numbers should increase upon repeatedly executing the **status station <*extension #*>** command.

Alternatively, Avaya's VoIP Monitoring Manager can be used to verify proper traffic flow.

Y. Go to Step 4.

N. The IP Telephone is not sending audio or the network is blocking audio packets. Exchange the IP Telephone to see if resolves problem.

If this resolves the problem, then replace the IP Telephone.

If it does not resolve the problem, then there is a network problem that the customer needs to resolve.

7 Is the MedPro operating correctly and does it have sufficient MedPro audio resources?

a Take an IP Telephone off-hook.

b Execute the **status station <*extension#*>** command.

c Scroll to page 3 titled (Call Control Signaling).

In the audio channel section if the `Switch Port` field contains a port location then go to Step d. Otherwise go to Step e.

d We have a MedPro port that is dynamically allocated to the IP Telephone call. Go to page 5 of the Station form and check the `Last Tx Sequence` field. This field shows the RTP sequence number of the last packet sent by the MedPro to the IP Telephone. This sequence number should increase at a regular rate when you run the **status station <*extension#*>** command repeatedly. If it does not increase then there is likely a MedPro hardware or firmware problem. Use the *Communication Manager Maintenance Manual* to resolve the issue. If packets are being transmitted normally, go to Step 8.

e If the audio channel on the status station form is blank, this might be due to an inability of the MedPro to allocate resource for the call. Run the **list measurements ip dsp-resource** command to determine whether there are sufficient MedPro resources in the system. Check for denials, blockage and out-of-service condition. If any of those measurements are greater than 0, this may indicate that any of the following problems may exist on the MedPro:

- The MedPro may have run out of DSP resources. After some users will have disconnected the problem will resolve itself. If this is a regular problem, another MedPro board needs to be installed.
- The firmware should be FW46 or later. Replace the firmware if needed.
- One of the DSPs may be bad or there could be firmware problem. This can be checked in the hardware error log by executing **display errors** command. Escalate the problem to the next Avaya Maintenance Tier.
- Communication Manager might not be able to find a MedPro in the network region where the IP Telephone resides.

If there are no MedPro problems, then go to Step 8.

8 Can the IP Telephone that experiences the 1-way problem or both IP Telephones that experience the no-way problem be pinged from the MedPro?

a Run the **status station <*extension #* >** command.

b Scroll to page 3 (Call Control Signaling). The `Switch Port` field gives the slot location of the MedPro circuit pack that is responsible for the IP Telephone, for example, 07D1717.

c Obtain the IP address of the IP Telephone from the `Set-end IP Addr` field. Hereafter, to simplify the description, it is assumed that this address is 135.9.42.105.

d Execute the command **ping ip-address 135.9.42.105 board 07D17**.

Y. The IP Telephones can be pinged from the MedPro, go to Step 10.

N. The IP Telephones cannot be pinged from the MedPro. Go to Step 9.

9 Find out where the ping terminated.

a Execute the **trace-route ip 135.9.42.105 board 07D17** command.

The customer needs to resolve the network problem in the router that terminated the trace-route command. Go to Step 12 after the problem has been resolved.

10 Is the call going through a firewall/ACLs?

a Check if the call would have to traverse a firewall by determining if it is destined to another remote network.

Y. Relax the packet/port filtering constraints in the firewall if they are too strict. If this works then go to Step 12. Otherwise, go to Step 11.

N. Go to Step 11.

11 Are there layer 1 errors detected in the IP Telephone, the intermediate switches/ routers or in the MedPro?

a Log into the switches and routers using telnet or SNMP access.

Check the port statistics.

NOTE:
Some customers will not allow this. In such case, the customer should be requested to provide this information.

Y. There is a network problem. The customer needs to fix the wiring.

N. Put a Protocol analyzer on both ends of the call by using switch port mirroring to see where packets are being dropped and resolve the problem. Go to Step 12 after the problem has been resolved.

12 If desired, return to the original state again by turning shuffling/hairpinning on if necessary. However, returning to a shuffled state may bring the problem back.

a Run the **change station <*extension #*>**. The `Direct IP-IP Audio Connections` and ***IP Audio Hairpinning*** should be set to **y**.

Choppy voice

Terminology

- **Choppy voice**. A voice audio signal that is impaired.
- **Clipping**. missing pieces in the received voice signal, especially at the beginning or end of words.
- **Pops**. Sudden interruptions of the voice by a popping sound.
- **Crackles**. Intermittent samples of noise and silence.

All these phenomena could be caused by packet loss or excessive jitter (perceived as packet loss).

Symptom resolution procedure

Several kinds of calls can be distinguished:

- IP Telephone - LAN - IP Telephone
- IP Telephone - LAN - PBX - DCP telephone
- IP Telephone - LAN - PBX - central office - telephone

1 Has a network assessment ever been done and has the network not been modified after the assessment?

Y. There might be a MedPro, IP Telephone or network problem, or the IP Telephone might have outdated software. All possibilities need to be explored, go to Step 2.

N. The network may not be compliant with the Avaya's network requirements. If the problem cannot be resolved by using the procedures described below, an assessment or reassessment might need to be done, go to Step 2.

2 Look at the large pattern first: do other IP Telephones one the same VLAN/subnetwork/floor experience the same problem?

Y. There may be a network problem, or multiple IP Telephones may have outdated software. All possibilities need to be explored, go to Step 3.

N. Go to Step 3.

3 Is a separate VLAN or subnetwork used for voice?

a The customer should check this on the Ethernet switches.

Y. Go to Step 5.

N. Go to Step 4.

4 Is the number of broadcasts messages lower than 1000 messages per second (this is the number that can safely be handled by the IP Telephone)?

a Check this by using the network management system or by hooking up a protocol analyzer to the network. If this cannot be checked through the network management system, go to the subsequent steps first, as it takes a relatively large effort to hook up a protocol analyzer.

Y. Go to Step 5.

N. There is a network problem. The customer should put the voice traffic (audio and signaling) on a separate VLAN with 802.1p priority 6 (the priority value reserved for voice and other real-time traffic).

5 Is the Ethernet switch connected to the MedPro set to (auto=speed)/(auto=mode) negotiation.

a Check this on the switch by logging in to it. The MedPro supports auto/auto by default.

Y. Go to Step 6.

N. Change the switch setting to auto/auto. If this is not possible, set the MedPro speed and duplex to match the switch port using the **change ethernet-options** command.

6 Is the Ethernet switch connected to the IP Telephone transmitting in HDX mode?

a Log in to the switch.

The 4606, 4612, 4624, and 4630 IP Telephones are only capable of HDX transmission. The 4602 and 4620 IP Telephones do support full-duplex mode, but require the attached Ethernet switch to be set in autonegotiate mode.

Y. Go to Step 7.

N. Change the switch setting to HDX (or auto for the 4602 or 4620). If this solves the problem, no further steps need to be taken. Otherwise, go to Step 7.

7 Are 802.1p QoS and IP DiffServ properly and consistently used in the switches, routers, the MedPro and the C-LAN?

Check if the QoS usage is consistent by examining the following:

a At an IP Telephone press the keypad button sequence **Hold Q O S #** and use the **#** key to walk through the menu to verify if the following recommended values are used for traffic priorities:

- Layer 2 Audio (802.1p) value = **6**.
- Layer 3 Audio DSCP value = **40** or **46**.
- Layer 3 Signaling DSCP value = **40** or **46**.

b In Communication Manager execute the **status station <*extension#*>** to determine the C-LAN circuit pack to which the IP Telephone is registered.

c Run the **display ip-interfaces** command to find the network region for that C-LAN circuit pack.

d Run the **display ip-network-region** command to check the QoS settings for the region.

e Switches and routers: log in the switches and routers

Y. Go to Step 8.

N. Turn 802.1p QoS and IP DiffServ tagging on with consistent values across the network by provisioning the recommended values in the switches, routers and IP Telephones. No further steps need to be taken if this solves the problem. Otherwise, go to Step 8.

8 Does the call traverse a WAN link? Does it have sufficient bandwidth and QoS/ packet fragmentation?

a Log on to the WAN routers and verify if the available bandwidth is sufficient to support voice.

NOTE:
Avaya recommends using G.729, which requires 24 Kbps (uncompressed, excluding Layer 2 overhead). IP packet fragmentation should be turned on when no DiffServ QoS facilities are available. On Avaya and Cisco routers it is possible to minimize bandwidth for audio usage by using the CRTP (compressed RTP).

Y. Escalate the problem to the next Avaya maintenance tier. Potentially it may be needed to perform a network assessment or a reassessment.

N. Go to Step 9.

9 Is the voice codec set to G.729 for calls across a WAN?

a This can be checked with an active call going on by running the **status station <*extension#*>** command.

b Scroll to page 3 (Call Control Signaling).

In the Audio Channel section it should indicate **G.729** as the encoder used.

Y. Go to Step 10.

N. Change the voice codec to G.729 (which is a lower bandwidth encoder than G.711, but still provides high quality) by executing the **change ip-codec-set** command and by putting G.729 at the top of the codec list. If this solves the problem, no further steps need to be taken. Otherwise, go to Step 10.

10 Is the end-to-end packet loss less than 1%?

Packet loss greater than 1% may be perceived as poor voice quality. IP Telephony packet loss can be measured using several different tools:

a The **list trace station** and **status station** commands show packet loss experienced by the MedPro.

b Avaya VoIP Monitoring Manager can measure packet loss experienced by IP Telephones as well as media processors.

c A protocol analyzer can capture packet streams between endpoints and identify packet loss.

Y. There is a network problem. The customer should explore the possibility to upgrade to a WAN link with the appropriate bandwidth and quality to ensure that it is compliant with the Avaya network requirements, possibly by establishing a new Service Level Agreement (SLA) with a network service provider. A network assessment or reassessment might need to be done.

N. There may still be a network problem. Escalate the problem to the next Avaya maintenance tier.

Dropped calls

Terminology

- **Dropped call.** A call that is terminated by a mechanism outside the control of the parties on the call.

Symptom resolution procedure

1 Does reconnecting the call solve the problem?

Y. There may have been an intermittent network problem. No further actions need to be taken unless this happens frequently. In the latter case, go to Step 2.

N. Install the latest software/firmware on the IP Telephone. Download the latest firmware from http://support.avaya.com and install it on your TFTP server. To transfer the software to the phone, type **Hold-R-E-S-E-T-#** on the phone. This reboots the IP Telephone and downloads a new version from the tftp server. If this resolves the problem, then no further steps need to be taken; otherwise go to Step 2.

2 Has a network assessment ever been done and has the network not been modified after the assessment?

Y. There may be a network problem, a MedPro problem or a C-LAN problem. All possibilities need to be explored, go to Step 1.

N. The network may not be compliant with the Avaya's network requirements. If the problem cannot be resolved by using the steps described below, a (re-)assessment may need to be done, go to Step 3.

3 Look at the large pattern first: do other IP Telephones experience the same problem?

Y. There may be a network problem, a MedPro problem, or a C-LAN problem. All possibilities need to be explored, go to Step 4.

N. Go to Step 4.

4 Perform traditional troubleshooting to determine whether Communication Manager or the IP Telephone drops the call. For example, this can be done by:

- Executing the command **list trace station <*extension #*>** command.
- Running an MST trace on the IP Telephone.
- Checking the denial event log.

If this does not solve the problem, then there is a network problem. Compliance with the Avaya network requirements may be an issue as well, and an assessment or a reassessment may need to be done.