Pre-installation Planning Guide for Cisco Unified ICM Enterprise and Hosted

Release 8.0(1)

February 2010

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0833

Table of Contents

List of Figures

Preface

Purpose

This guide describes pre-installation requirements and issues to address in preparing for a Cisco Unified Intelligent Contact Management (Unified ICME) installation. It does not discuss, for example, pre-installation planning for Unified ICME multichannel software or for Cisco Unified Contact Center and its components such as Cisco Unified Communications Manager (Unified Communications Manager or Cisco Unified IP IVR (Unified IP IVR)). For information on Cisco Unified multichannel software, see the documentation for Cisco Unified E-Mail Interaction Manager (Unified EIM), Cisco Unified Web Interaction Manager (Unified WIM) and Cisco Media Blender (CMB). For Unified ICME, see the relevant documentation.

Audience

This guide is intended for contact center managers, system support personnel, and plant engineers who are planning and preparing contact center sites for a Unified ICM system installation. Readers should be familiar with contact center site planning and preparation issues. They should also have a basic understanding of the Unified ICM system and the components that are installed as part of the system.

Organization

This document is organized as follows:

Chapter 1, "Pre-installation Planning Overview" (page 7)	Provides an overview of the Unified ICME pre-installation planning process. This chapter includes a pre-installation document roadmap, which suggests an order to follow in using the Unified ICM pre-installation planning guides.

Chapter 2, “Cisco Unified ICM Enterprise Overview” (page 11)	Describes the role of the Unified ICM software within the Cisco Unified Contact Center Enterprise (Unified CCE). This chapter also reviews the main Unified ICM software features.
Chapter 3, “IXC Overview” (page 25)	Describes how to plan for access to the carrier’s intelligent network service. This chapter includes an overview of Unified ICME/IXC interaction and a discussion of Unified ICM -Network Interface Controller (NIC) fault tolerance.
Chapter 4, “Switch Overview” (page 31)	Provides an overview of Unified ICM PG-to-peripheral interaction.
Chapter 5, “Peripheral Gateway Configurations” (page 33)	Describes the options for configuring Peripheral Gateways in the Unified ICME.
Chapter 6, “Cisco Computer Telephony Integration Option (CTI) Planning” (page 41)	Describes the pre-installation planning for CTI, including reviewing CTI Server communications and platform options; becoming familiar with the desktop options; estimating CTI message traffic; planning fault tolerance for the CTI Server; and reviewing ACD support for client control and third-party call control.
Chapter 7, “Unified IVR Planning” (page 53)	Describes the pre-installation planning tasks for the Unified IVR option, including reviewing the options for integrating IVRs into the Uinfied ICME system, determining if any IVR programming or application development is necessary, and reviewing the PG platform requirements for Unified IVR.
Chapter 8, “ICM Application Gateway and ICM Gateway SQL Planning” (page 63)	Describes the pre-installation planning tasks for the Unified ICM Application Gateway and Unified ICM Gateway SQL options, including preparing host systems and databases; reviewing fault tolerance issues; and planning for data transfer (in the case of Gateway SQL).
Chapter 9, “Planning for ICM Platforms” (page 67)	Describes how to determine the numbers and types of Unified ICM nodes you will need.
Chapter 10, “Determining the Datacom Requirements” (page 73)	Describes how to prepare network facilities for an Unified ICM system installation, such as determining the requirements for visible and private networking, allocating IP addresses, and ordering any required network hardware.
Chapter 11, “Site Preparation” (page 109)	Presents a brief list of basic considerations for site preparation.
Chapter 12, “IP Address Worksheets” (page 111)	Provides worksheets you can use to record IP addresses for the visible and private networks.

Related Documentation

Documentation for Cisco Unified ICM/Contact Center Enterprise & Hosted, as well as related documentation, is accessible from Cisco.com at: **http://www.cisco.com/cisco/web/psa/default.html**.

- Related documentation includes the documentation sets for Cisco CTI Object Server (CTI OS), Cisco Agent Desktop (CAD), Cisco Agent Desktop - Browser Edition (CAD-BE), Cisco Unified Contact Center Management Portal, Cisco Unified Customer Voice Portal (CVP), Cisco Unified IP IVR, Cisco Unified Intelligence Center, and Cisco Support Tools.
- For documentation for these Cisco Unified Contact Center Products, go to **http://www.cisco.com/cisco/web/psa/default.html**, click **Voice and Unified Communications**, then click **Customer Contact**, then click **Cisco Unified Contact Center Products** or **Cisco Unified Voice Self-Service Products**, then click the product/option you are interested in.
- For troubleshooting tips for these Cisco Unified Contact Center Products, go to **http://docwiki.cisco.com/wiki/Category:Troubleshooting**, then click the product/option you are interested in.
- Documentation for Cisco Unified Communications Manager is accessible from: **http://www.cisco.com/cisco/web/psa/default.html**.
- Technical Support documentation and tools are accessible from: **http://www.cisco.com/en/US/support/index.html**.
- The Product Alert tool is accessible from (login required): **http://www.cisco.com/cgi-bin/Support/FieldNoticeTool/field-notice**.
- For information on the Cisco software support methodology, refer to *Software Release and Support Methodology: ICM/IPCC* available at (login required): **http://www.cisco.com/en/US/partner/products/sw/custcosw/ps1844/prod_bulletins_list.html**.
- For a detailed list of language localizations, refer to the *Cisco Unified ICM/Contact Center Product and System Localization Matrix* available at the bottom of the following page: **http://www.cisco.com/en/US/products/sw/custcosw/ps1001/prod_technical_reference_list.html**.

Product Naming Conventions

In this release, the product names defined in the table below have changed. The New Name (long version) is reserved for the first instance of that product name and in all headings. The New Name (short version) is used for subsequent instances of the product name.

Note: This document uses the naming conventions provided in each GUI, which means that in some cases the old product name is in use.

Old Product Name	New Name (long version)	New Name (short version)
Cisco IPCC Enterprise Edition	Cisco Unified Contact Center Enterprise	Unified CCE
Cisco System IPCC Enterprise Edition	Cisco Unified System Contact Center Enterprise	Unified SCCE **Note:** Cisco Unified System Contact Center Enterprise (Unified SCCE) is supported in 8.0(1); however, there is

Old Product Name	New Name (long version)	New Name (short version)
		not a separate 8.0(1) version. If you request features that are in 8.0(1), you must migrate to the Unified ICM/CCE/CCH software. Full migration information is documented in the *Upgrade Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted*.
Cisco IPCC Hosted Edition	Cisco Unified Contact Center Hosted	Unified CCH
Cisco Intelligent Contact Management (ICM) Enterprise Edition	Cisco Unified Intelligent Contact Management Enterprise	Unified ICME
Cisco Intelligent Contact Management (ICM) Hosted Edition	Cisco Unified Intelligent Contact Management Hosted	Unified ICMH
Cisco CallManager/Cisco Unified CallManager	Cisco Unified Communications Manager	Unified CM

Conventions

This manual uses the following conventions:

Convention	Description
boldface font	Boldface font is used to indicate commands, such as user entries, keys, buttons, and folder and submenu names. For example: • Choose **Edit** > **Find**. • Click **Finish**.
italic font	Italic font is used to indicate the following: • To introduce a new term; for example: A *skill group* is a collection of agents who share similar skills. • For emphasis; for example: *Do not* use the numerical naming convention. • A syntax value that the user must replace; for example: IF (*condition, true-value, false-value*) • A book title; for example: Refer to the *Cisco CRS Installation Guide*.
`window font`	Window font, such as Courier, is used for the following:

Convention	Description
	• Text as it appears in code or that the window displays; for example: **`<html><title>Cisco Systems,Inc. </title></html>`** • Navigational text when selecting menu options; for example: **`ICM Configuration Manager > Tools> Explorer Tools > Agent Explorer`**
`< >`	Angle brackets are used to indicate the following: • For arguments where the context does not allow italic, such as ASCII output. • A character string that the user enters but that does not appear on the window such as a password.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Documentation Feedback

You can provide comments about this document by sending email to the following address:

mailto:ccbu_docfeedback@cisco.com

We appreciate your comments.

Chapter 1

Pre-installation Planning Overview

Note: This manual deals with Unified ICME. For information on Unified CCE, see the Installation and Configuration Guide for Cisco Unified Contact Center Enterprise & Hosted and the Administration Guide for Cisco Unified Contact Center Enterprise & Hosted.

The Unified ICM software is a distributed application that routes telephone calls, web inquiries, and e-mail across geographically distributed contact centers. A typical Unified ICM system includes a number of computers located at different sites. A small Unified ICM system might have computers at two or three sites. A larger system might have computers at 20 sites or more.

Because the Unified ICM software works with different types of contact center equipment and sometimes one or more carrier networks, some pre-installation planning is necessary to ensure that the Unified ICM installation process proceeds smoothly and on schedule.

This chapter provides an overview of the Unified ICM pre-installation planning process. It also contains a pre-installation planning document roadmap, which suggests an order in which you can start the tasks.

The Planning Process

The Unified ICM pre-installation planning process involves coordinating and scheduling several tasks so they are completed in time for the arrival of the Unified ICM server platforms. You typically need to make preparations at each site that is to contain Unified ICM components. Some pre-installation tasks may take longer than others. Therefore, try to start the time-consuming tasks early and continue working in parallel on the other pre-installation tasks.

Coordinating and Scheduling Tasks

Cisco suggests that one person in your organization has overall responsibility for coordinating and scheduling the pre-installation planning tasks. This person can also delegate responsibility to ensure that tasks are assigned to people with the applicable expertise. For example, you might

have your MIS expert begin working with Cisco to order the server platforms. At the same time, your data communications expert can start the process of provisioning network facilities at each contact center site.

Pre-installation Document Road Map

The current document provides guidance on topics such as provisioning IXC access, preparing ACDs, and determining the Unified ICM datacom requirements. In each case, one or more pre-installation tasks are covered.

You typically start the pre-installation planning tasks in the following order:

1. Getting Started: Document current contact handling procedures. Provide configuration data for contact center sites. Understand the Unified ICM software. Review Unified ICM product options. Determine Unified ICM Configuration.

   See Chapter 2, “ICM Enterprise Edition Overview” (page 11); Administration Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted.

2. IXC Access: Review ICM/IXC interaction. Choose network link fault tolerances strategy. Review IXC access specifics. See Chapter 3, “IXC Overview” (page 25); the relevant Cisco NIC Supplement document.

3. Switch Preparation: Determine ACD requirements. Determine CTI and MIS link requirements. Order required upgrades and enhancements. See Chapter 4, “Switch Overview” (page 31); Chapter 5, “Peripheral Gateway Configurations” (page 33); the relevant Cisco ACD Supplement document(s).

4. Product Options and System Integration: Determine product option requirements. Order any required upgrades or enhancements. See Chapter 6, “CTI Planning” (page 41); Chapter 7, “IVR Planning” (page 53); Chapter 8, “ICM Application Gateway and ICM Gateway SQL Planning” (page 63); .

5. Estimating System Size: Enter data using the Unified ICM database sizing tool. Note the specifications provided by the tool. Determine the number of PCs required.

   See the discussion of the ICM Database Administration tool (ICMDBA) in the Administration Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted; Chapter 9, “Planning for ICM Platforms” (page 67).

6. Network and Site Requirements: Determine requirements for Unified ICM networking. Allocate IP addresses. Order any additional network hardware. Meet basic site requirements. Order additional cabling or other equipment required.

   See Chapter 10, “Determining the Datacom Requirements” (page 73); Chapter 11, “Site Preparation” (page 109); Chapter 12, “IP Address Worksheets” (page 111).

   For example, since the lead-time for provisioning IXC access is several weeks, this task is started early in the process. You can then proceed with tasks such as making sure your contact center equipment (ACDs, PBXs, IVRs) have the necessary software releases and

options. While that task is in progress, you can select Unified ICM product options and component platforms and begin preparing the installation sites.

7. Pre-installation EOL Component Check: ICM Installer will check if the below mentioned EOL components are installed in the \ machine before upgrading the machine:

   EOL Components

   – PG type Md110, Siemens, Rolm9005, Galaxy, G2, ACP1000, Meridian

   – MEI Server

   – Application Bridge Server

   – AIN Network Gateway

   Above mentioned components will be removed by the ICM Installer after the user confirms that they can be removed.

Caution: Application Bridge Server needs to be manually removed. It is not supported in Unified ICM 8.0. Automatic removal is not provided in this release.

NIC and ACD Supplements

The NIC Supplements are reference documents that contain specific information on how the Unified ICM Network Interface Controller (NIC) interfaces with the supported IXC carrier networks. The NIC is the software process that allows the Unified ICM system to communicate with the carrier's intelligent switching network. You may want to refer to the NIC supplements for detailed technical information when you are planning for IXC access.

There are NICs, and NIC Supplements, for each carrier supported by the Unified ICM software (AT&T, MCI, Sprint, and so on). The NIC Supplements are intended to be used as technical reference companions to the Cisco Unified ICM software documentation set.

The ACD Supplements are reference documents that contain the specific information you need to maintain Unified ICM Peripheral Gateways (PGs) in an Unified ICM environment. The PG is the Unified ICM component that provides an interface to proprietary ACD systems.

There are ACD supplements for each ACD supported by the Unified ICM software (Aspect CallCenter, Avaya ACM, Nortel Symposium, and so on). The ACD Supplements are intended to be used as the ACD-specific companions to the Unified ICM software documentation set. For example, while other Unified ICM documents such as the Configuration Guide for Cisco Unified ICM/Contact Center Enterprise and Hosted and the Scripting and Media Routing Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted, cover general topics such as configuring an overall Unified ICM system and writing scripts to route contact center requests, the ACD Supplements provide specific information on configuring certain types of PGs and making any necessary adjustments to the ACD configuration. Refer to the ACD Supplements for detailed technical information when you are determining the requirements for your ACDs.

Chapter 2

Cisco Unified ICME Overview

In the initial phase of pre-installation planning, you need to become familiar with the Unified ICM system and understand how it fits into your Unified CCE. You can then determine which products and components you want to deploy in an Unified ICM virtual contact center.

In this chapter, complete the following pre-installation tasks:

- Determine the role of the Unified ICM software in your enterprise. Understand how the Unified ICM software fits into the Unified CCE and carrier networks.
- Choose Unified ICM products. Will your system be a complete pre-routing and post-routing system? Will you have other options such as Unified ICM Gateway SQL, Cisco CTI, or Unified IP IVR?

This chapter contains the following topics:

How the ICM Software Works

The Unified ICM Enterprise Edition works with your contact center equipment and the IXC carrier network to create a virtual contact center. In the virtual contact center model, multiple distributed contact centers are linked to form one Unified CCE. The agents within the Unified CCE become members of a single team that is capable of servicing customer contacts throughout the enterprise.

ICM Call Routing

The Unified ICM software makes the best use of your contact handling resources while ensuring that each customer is directed to the most appropriate resource available. To get a better idea of how the Unified ICM software fits into the contact center and carrier environments, it might help to examine how the Unified ICM software routes telephone calls.

Figure 1: Intelligent Contact Routing (Telephone Calls)

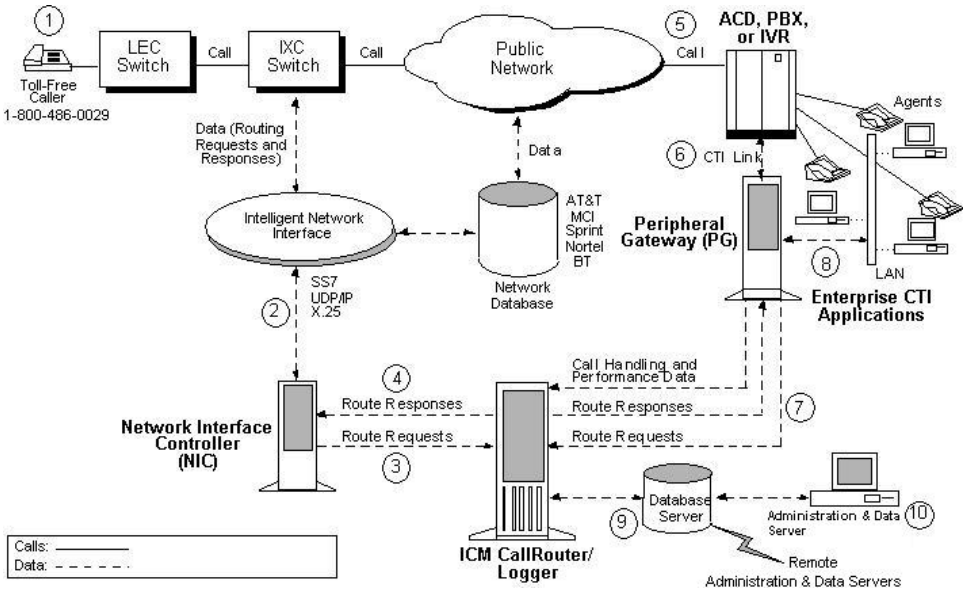

Pre-Routing

The Unified ICM software executes call routing decisions before a call terminates at a contact center. This concept is called pre-routing. As shown in the preceding figure, calls to be routed usually originate in the public telephone network as calls to a toll-free number (1).

The IXC Network

The Unified ICM software is configured in the intelligent network of the IntereXchange Carrier (IXC) to receive a route request for each designated incoming call (2). A subsystem of the Unified ICM software, called the Network Interface Controller (NIC), communicates with the carrier's network through an intelligent network interface.

Route Requests

The NIC translates the network's description of the call, including point of origin, number dialed, and any customer entered digits, into the language of the Unified ICM software. The NIC passes this call information to the CallRouter in the form of a route request (3).

Note: For clarity, the NIC is usually shown in figures as a separate computer. Actually, NICs are implemented as software on the Unified ICM software platform (usually on the CallRouter or CallRouter/Logger [Rogger] machines).

Route Responses

At this point, the Unified ICM software may query an ANI or customer profile database before returning a route response to the NIC (4). The NIC passes a destination for the call back to the IXC network. The IXC is responsible for connecting the call and maintaining the voice path.

ACDs

Each contact center has one or more Automatic Call Distributor (ACD) systems that direct incoming calls to the telephone sets of individual agents (5). The Unified ICM software maintains real-time communications with the ACDs in each contact center by using a Peripheral Gateway (PG).

Peripheral Gateway

The PG communicates with the ACD over the switch vendor's Computer Telephony Integration (CTI) link (6). To make optimal decisions, the Unified ICM software must know the latest status for every call, agent, and agent group in its network. One purpose of the PG is to extract this status information from the ACD and forward it to the CallRouter's in-memory database. (The PG can also be used as a CTI Server and as a communications interface between the Unified ICM and Interactive Voice Response (IVR) systems located at contact center sites or in the network.)

Post-Routing

In private network configurations, ACDs can also originate call routing requests. This is called post-routing. Post-routing provides the same intelligence used in pre-Routing, but applies it to calls originating from a private network of ACD, PBX, and IVR systems. The PG assists in post-routing by forwarding routing requests to the Unified ICM software and returning the target destinations to the ACD (7).

CTI Server

External server or workstation applications can subscribe with a PG that acts as a CTI Server (8). The CTI Server provides call and agent event data that can be used in screen-pops and other CTI applications. At the desktop level, the Unified ICM CTI desktop provides an environment for integrating soft-phone, screen-pop, and data entry at the agent's workstation.

Monitoring and Reporting

All event data that are gathered by the PG and router are forwarded to the Unified ICM software and stored in an industry-standard relational database (9). These data are used in real-time monitoring and historical reporting. The standard Unified ICM monitoring screens and reports

can be easily modified with Unified ICM-provided database access tools. Optionally, the data can be accessed directly with SQL or Open Database Connectivity (ODBC) tools.

Administration & Data Server

The overall operation of the Unified ICM software is monitored and controlled from an Administration & Data Server (10) . The Unified ICM software can support multiple Administration & Data Servers located throughout the contact center network.

ICM System Components and Processes

Many different Unified ICM system software components are involved in pre-installation planning. You may want to become familiar with the role of the components in the Unified ICM system.

Note: Not every component is used in every Unified ICM System.

CallRouter

This is the part of the Unified ICM system that contains the call routing logic. The Unified ICM software receives call routing requests and determines the best destination for each call. It also collects information about the entire system. The Unified ICM software serves as a real-time server by forwarding performance and monitoring information to Administration & Data servers.

Logger

The Logger is the interface between the Unified ICM software and the database manager (SQL Server). As the Unified ICM software collects performance and monitoring information about the system, it passes the information to the Logger for short-term storage in a central relational database. The Logger forwards historical information to the Historical Data Server (HDS). The HDS on the Logger maintains statistics and data for use in monitoring and reporting.

Network Interface Controller (NIC)

The NIC connects the Unified ICM software to the IXC signaling network. The NIC receives a route request from the signaling network for each incoming call and passes the request to the Unified ICM software. The Unified ICM software responds with routing information (a routing label), which the NIC passes back to the IXC signaling network.

Note: For clarity, the NIC is usually shown in figures as a separate computer. Actually, NICs are implemented as software on the Unified ICM software platform (usually on the CallRouter or CallRouter/Logger [Rogger] machines).

Peripheral Gateways

Each contact center device (ACD, PBX, or IVR) communicates with a Peripheral Gateway (PG). The PG reads status information from the device and passes it back to the Unified ICM software. The PG runs one or more Peripheral Interface Manager (PIM) processes, which are the software components that communicate with proprietary ACD systems. A single PIM is required for each peripheral to which the PG will interface. Therefore, a single PG (and its associated PIMs) can serve multiple peripherals of the same kind. For example, one PG with four Aspect ACD PIMs can serve four Aspect ACDs in the contact center.

Note: A single PG can support both ACD PIMs and IVR PIMs; however, the ACD PIMs and the IVR PIMs must all be the same type of PIM (ACD PIMs must be the same type; IVR PIMs must be the same type).

A single server can support up to two PGs. For details, refer to the Configuration Guide for Cisco Unified ICM/Contact Center Enterprise and Hosted.

Administration & Data Server

The Administration & Data Server is the human interface to the Unified ICM software. It serves as a control console from which you can monitor agent and contact center activity and change how the Unified ICM software routes calls. For example, you can use the Administration & Data Server to configure the Unified ICM contact center data, create call routing scripts, and monitor and report on the Unified ICM system or some part of the system. Administration & Data Servers can be located anywhere, as long as they have LAN or WAN connections to the Unified ICM software.

Administration & Data Servers have several roles: Administration, Real-time data server, Historical Data Server, and Detail Data Server. A Unified ICM deployment must have Administration & Data Servers to fill these roles. The servers may be deployed in the following combinations to achieve the needed scalability with the minimum number of servers:

- Administration Server and Real-time Data Server (AW)
- Configuration only Administration Server
- Administration Server, Real-time and Historical Data Server and Detail Data Server (AW-HDS-DDS)
- Administration Server and Real-time and Historical Data Server (AW-HDS)
- Historical Data Server and Detail Data Server (HDS-DDS)

An Administration Client (formerly known as a “client AW”) serves the administration role, but is deployed as a client to an Administration Server for scalability. The Administration Client may view and modify the configuration, and receive real-time reporting data from the Administration & Data Server, but does not store the data itself, and does not have a database. Each Administration & Data Server must be installed on a separate server for production systems

to ensure no interruptions to the real-time call processing of the Call Router and Logger processes. For lab or prototype systems, the Administration & Data Server (with the WebView Server option) can be installed on the same server as the Call Router and Logger.

Historical Data Server

Administration & Data Servers need to access historical data (half hour data, call detail, and so on) for historical reporting in the Script Editor or in third-party tools. At least one real-time Administration & Data Servers, in a system, must be installed with a Historical Data Server (HDS) to support reporting and long-term historical data storage. The HDS IP address requirements are identical to those of a standard Administration & Data Server.

ICM Reporting

The Unified ICM Reporting solution provides an interface to access data describing the historical and real-time states of the system.

The reporting solution consists of the following components:

- Unified IC or WebView —reporting user interfaces
- Configuration and Reporting Data — contained on Administration & Data Server(s)

Reporting concepts and data descriptions are described in Reporting Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted; this description is independent of the reporting user interface being used.

Cisco Unified Intelligence Center

Cisco Unified Intelligence Center (Unified IC) is an advanced reporting product used for CCE and other products. This platform is a web-based application offering many Web 2.0 features, high scalability, performance, and advanced features such as the ability to integrate data from other Cisco Unified Communications products or third-party data sources. Unified IC incorporates a security model which defines different access and capabilities for specific users. Unified IC Standard is included with Unified ICM. Unified IC Premium is an optional product with additional features. Unified IC must be installed on a separate server; it cannot be co-resident with other Unified ICM components.

For a complete description of both Unified IC products see Reporting Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted.

WebView

WebView is a web-based application for Unified ICM reporting. WebView performs the basic operations of gathering user input, querying the databases and presenting the requested data for both real time and historic data.

Note: WebView does not support historical data which is collected in 15 minute intervals (a feature new in Release 8.0(1)). WebView only supports historical reporting for data collected in half hour intervals. WebView also does not supply reports for the call type/skill group data (new in Release 8.0(1)). Unified IC must be used for 15 minute data intervals or call type/skill group reporting.

WebView has the ability to export data, launch scheduled reports, save and share report settings, and mark favorite reports. It also has features to display service affecting events reported by the system. WebView can be installed on an Administration & Data Server. To increase scalability, Webview can be installed on a standalone server.

The WebView architecture is described in the WebView Installation and Administration Guide, available at **http://www.cisco.com/en/US/products/sw/custcosw/ps4145/prod_installation_guides_list.html** . For a description of all of the reports provided with WebView, refer to the WebView Template Reference Guide for Cisco Unified Contact Center Enterprise & Hosted, available at **http://www.cisco.com/en/US/products/sw/custcosw/ps4145/products_user_guide_list.html**.

ICM Options and Related Products

The Unified ICM software can be set up with a variety of options, such as adding software to perform database lookups or performing secondary call routing once a call has terminated at an ACD. In some cases, the Unified ICM software is an integral part of other Cisco contact center products, such as the IP Contact Center (IPCC). You may want to review the Unified ICM software options and related products to learn about the different ways the Unified ICM software can be deployed in a Unified CCE.

Pre-Routing

Pre-Routing allows the Unified ICM software to execute routing decisions before a call terminates at a contact center. With pre-routing, the Network Interface Controller (NIC) receives the route request from the IXC and passes the call information to the Unified ICM software. The Unified ICM software processes the route request through a call routing script, which defines how the call should be routed. The Unified ICM software returns a route response to the NIC, which in turn forwards it to the IXC. The route response contains the call's final destination.

In pre-routing, the Peripheral Gateway's role is to keep the Unified ICM software informed of the real-time status of switches, calls, and agents in the Unified CCE. The Unified ICM software uses this real-time data to make an informed call routing decision.

Pre-Routing systems require the following components:

- Network Interface Controller (NIC)
- CallRouter
- Logger
- Administration & Data Server

- WebView Server
- Peripheral Gateway (PG)

The pre-routing capabilities are enabled through the Network Interface Controller (NIC) and the CallRouter processes. NICs are implemented as software on the Unified ICM software platform (for example, on the CallRouter or Logger machines).

The Unified ICM routes calls within the public network based on several dynamic variables. You can use any combination of the following variables to route calls:

Agent availability	Day of week
Agent skills	Number dialed
Caller-entered digits	Origin of call
Cost of the call	Cost of the transaction
Customer database lookup	Scheduled agents
Customer-defined business rules	Time of day

Calls are routed in the most efficient manner possible given the current contact center load conditions.

Post-Routing

In a traditional time-division multiplexing (TDM) environment, post-routing systems have software that allows the CallRouter to make secondary routing decisions after a call has been received at a contact center. In post-routing, the ACD or IVR submits a route request to the Unified ICM software. The Unified ICM software executes scripts to process the routing request and return a destination address to the ACD. The Unified ICM software then directs the ACD to transfer the call to an agent, skill group, or service, either in the same contact center or at a different contact center. In making a post-routing decision, the Unified ICM software can use the same information and script it uses in pre-routing. In other words, the same call routing intelligence that is used in the pre-routing of calls is applied to calls that are interflowed between contact center sites, transferred between agents, or transferred into or out of IVRs.

Pre- and Post-Routing Systems

A pre- and post-routing Unified ICM system is a complete intelligent call routing, monitoring, and reporting system. The Unified ICM software can execute routing decisions before a call terminates at a contact center. It can also make secondary routing decisions after a call has been received at a contact center. A Pre- and post-routing system can be expanded with optional features such as Unified ICM Application Gateway, Unified ICM Gateway SQL, Unified ICM IVR interface, and CTI Server to create an intelligent call routing and management solution in which all the elements of the Unified CCE play a role in intelligent routing.

Computer Telephony Integration (CTI)

Cisco CTI software provides an interface between the Unified ICM software and agent desktop and server applications. The CTI software works with a PG's ACD and IVR interface software and all associated ACDs to track events and transactions and forward call- and transaction-related data to an agent's desktop computer.

The CTI software has full third-party call control features that allow agents and integrated desktop applications to perform tasks such as transferring calls, conferencing calls, and setting call data all within an enterprise framework. Voice and data collected by an agent at the desktop can be transferred in the form of a screen-pop among agents and across different ACD platforms. This allows customer and transaction data to accompany a call from an IVR or web server to the agent and from site-to-site as required. The Unified ICM system can also use CTI data to determine call destinations based on factors such as customer value, business objectives, market penetration, and personalized service.

CTI Server

CTI Server, the basic server component of Cisco CTI, enables the Unified ICM software to deliver agent, call, and customer data in real-time to a server and/ or workstation application as events occur throughout the life of a call. The CTI Server is a software process that runs on a Peripheral Gateway (PG).

It is a gateway into the Unified ICM software's data and services.

- Pre-route indications identify a caller and provide associated attributes to applications while the call is still in the public or private network and before the caller is connected to an agent, web server or IVR.
- Call events are provided throughout all stages of the call flow, from the moment a call arrives at an answering location (ACD, PBX, IVR, web server) until the caller hangs up.
- Agent work state changes are reported as they occur.

Cisco CTI Object Server (CTIOS)

CTI Object Server (CTIOS) is a high-performance, scalable, fault-tolerant server-based solution for deploying CTI applications. It serves as a single point of integration for third-party applications, including Customer Relationship Management (CRM) systems, data mining, and workflow solutions.

CTIOS is a client of CTI Server, and has a single all-events connection to Cisco CTI Server. In turn, CTIOS accepts client connections using session, agent, and call interfaces. These interfaces are implemented in .NET, COM, Java, and C++, allowing for a wide range of application development uses. The interfaces are used for call control, to access data values, and to receive event notifications.

CTIOS configuration and behavior information is managed at the CTIOS server, simplifying customization, updates, and maintenance. Servers can be accessed and managed remotely. Thin-client and browser-based applications that do not require Cisco software on the desktop can be developed and deployed with CTIOS.

CTIOS incorporates the following major components:

- CTIOS Toolkit
- Client Interface Library
- CTIOS Combo Desktop for Agents and Supervisors

Note: Refer to the Cisco CTIOS Software documentation for more information.

IVR Interface

This option allows for running a Voice Response (IVR) system. The IVR interface software runs on a PG platform. It allows the Cisco Unified ICM software to route calls to targets on IVRs and collect data from IVRs for use in call routing, real-time monitoring, and historical reporting.

The IVR interface can also provide queuing at a network-based or premises-based IVR. With this feature, calls can be directed to an IVR queue when no other appropriate answering resource is available. The IVR interface is not specific to a particular IVR system or manufacturer. It is based on an open IVR model. Many IVR systems support Cisco's Open IVR Interface Specification, including Unified CVP.

The Cisco Customer Voice Portal integrates with both traditional time-division multiplexing (TDM) and IP-based contact centers to provide a call-management and call-treatment solution with a self-service IVR option that can use information available to customers on the corporate Web server. With support for automated speech recognition (ASR) and text-to-speech (TTS) capabilities, callers can obtain personalized information and can conduct business without interacting with a live agent.

Note: Unified CVP was previously called Internet Service Node (ISN).

For a list of IVRs that support this interface, contact your Cisco representative.

Note: You can integrate IVR systems into the Cisco Unified ICM software in several different ways. Chapter 7, "IVR Planning" (page 53) provides more information on IVR integration along with examples of how you might integrate IVRs with the Cisco Unified ICM system.

ICM Application Gateway

The Cisco Unified ICM Application Gateway option allows the Cisco Unified ICM software to interact with a host system that is running another contact center application. Within the Cisco Unified ICM software, the Gateway feature is implemented as an Application Gateway node in a call routing script. You add an Application Gateway node to a script to instruct the system

to execute an external application. This allows the script to evaluate responses from the external application and base subsequent routing decisions on the results produced by the application.

The Gateway option allows the Cisco Unified ICM system to interface with any external application, not just database applications.

You can use the Gateway option within the Cisco Unified ICM system to:

- Allow other applications to select a call's destination.
- Control or trigger external applications through Cisco Unified ICM call routing scripts.
- Pass data to and collect data from other contact center applications.

For example, a simple Gateway application might return a variable to the CallRouter that identifies the caller as having a premium account. The routing script can use this information to control where and how the call is routed. Optionally, the Cisco Unified ICM can pass the retrieved information to the site that is receiving the call. Data such as account numbers, dates, billing phone numbers, and addresses can be passed along with the call to an answering resource.

Note: Chapter 8, "ICM Application Gateway and ICM Gateway SQL Planning (page 63)" provides more information on planning for the Gateway feature.

ICM Gateway SQL

Cisco Unified ICM Gateway SQL allows the Cisco Unified ICM software to query an external SQL Server database and use the data in call routing. If you have databases that contain customer account or profile information, you might want to perform database lookups to assist in call routing. The database lookups can be based on Calling Line ID (CLID), Dialed Number (DN), or Caller Entered Digits (CED) such as account or social security numbers.

A typical Gateway SQL application might prioritize callers. For example, a call routing script might use the caller's CLID to access a database and retrieve data about the caller such as the caller's average monthly bill. Based on this information, the routing script would route the caller to the most appropriate answering resource.

The following figure shows a basic Gateway SQL configuration. Note that this configuration requires an additional database server on which you can load the external SQL Server database and data.

Figure 2: Gateway SQL Configuration

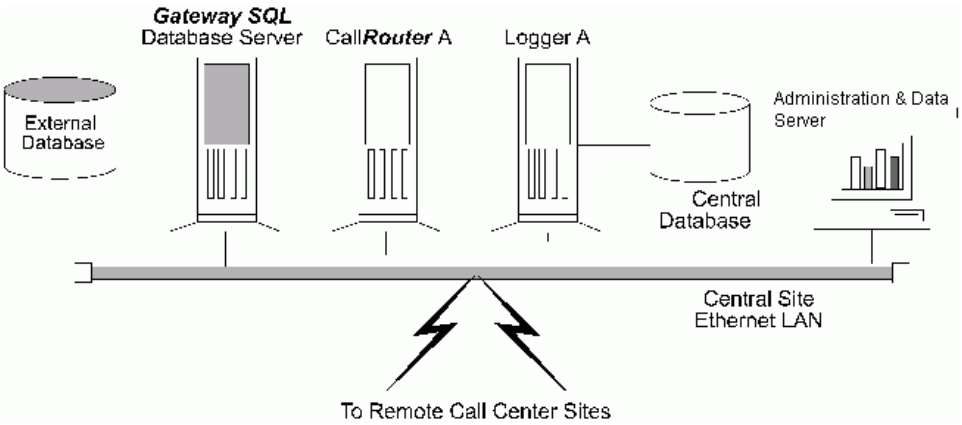

Note: You need to perform some pre-installation planning if you are going to use the Cisco Unified ICM Gateway SQL option. Chapter 8, "ICM Application Gateway and ICM Gateway SQL Planning" (page 63) provides more information on planning for the Cisco Unified ICM Gateway SQL feature.

Internet Script Editor

Internet Script Editor is an application you can use to work with routing and administration scripts. It provides the same functionality as the Cisco Unified ICM Script Editor software, without the need for an Administration & Data Server.

Internet Script Editor works through the IIS Web server on Cisco Unified ICM software, using HTTP to communicate with the Cisco Unified ICM software.

The Internet Script Editor and the Cisco Unified ICM Script Editor GUIs are essentially the same. The menus, toolbars, palette, and work space are utilized in the same manner in both applications. The differences between the two occur primarily in the method by which each application communicates with the Cisco Unified ICM software.

ICM Multichannel Software

The Cisco multichannel software provides a flexible, integrated architecture to support a variety of agent and customer interactions for a contact center. The contact center manager can configure agents to handle voice, Web collaboration, text chat, and e-mail requests and have the agents switch between these media types on a task-by-task basis. The manager can also configure agents to support only one media type. Customers can choose the medium that is most comfortable and convenient for them.

Requests are routed by the Cisco Unified ICM system using the same kind of business rules applied to contacts arriving from a carrier network. Every request is delivered to the most appropriate agent anywhere in the enterprise.

Note: For information on Cisco Unified multichannel software, see the documentation for Unified Email Interaction Manager, Unified Web Interaction Manager, and Cisco Media Blander.

Cisco Unified Contact Center

Cisco Unified Contact Center combines Cisco's IP telephony products and Cisco Unified ICM software to create an IP-based contact management solution. Cisco Unified Contact Center provides a migration path to an IP-based contact center by supporting integration with legacy call center platforms and networks. With Cisco Unified Contact Center, agents can use Cisco IP phones to receive both time-TDM and VoIP phone calls. Capabilities of the Cisco Unified Contact Center include intelligent call routing, ACD functionality, network-to-desktop CTI, Unified IP IVR integration, call queuing, and consolidated reporting.

Cisco Unified Contact Center is based mainly on two Cisco products: Unified CM and ICM software. Unified CM provides traditional PBX telephony features in an IP telephony environment. Unified ICM software provides enterprise-wide management and distribution of voice and data from ACDs, Unified IP IVR systems, small office/home office (SOHO) agents, and desktop applications. Cisco Unified IP phones and Unified IP IVRs (as well as traditional TDM IVRs) are also part of the Cisco Unified Contact Center.

Note: For information on Unified CCE, see the Installation Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted and the Administration Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted.

Chapter 3

IXC Overview

If pre-routing is to be performed, the Unified ICM Enterprise Edition software requires access to the IntereXchange Carrier intelligent call routing network. Each interexchange carrier offers intelligent network services that allow customer-premises equipment to participate in network-level call routing. The Unified ICM software connects to one or more networks by using a Cisco Network Interface Controller (NIC).

Specifically, this chapter helps you to complete the following tasks:

- Choose your carrier(s). Cisco supports network interfaces with several carriers. You can use one or more carriers with the Unified ICM software.
- Choose the types of network link fault tolerance to apply. It is important to apply fault tolerance in the network interface and the links to the carrier's intelligent network.
- Order intelligent network service. Once you review the requirements for your specific Cisco NIC, order intelligent network service and work with the carrier and Cisco to bring the service on line.

This chapter contains the following topics:

ICM Software and IXC Interaction

The Network Interface Controller (NIC) is the interface between the Unified ICM software and the IXC intelligent network. The NIC communicates with the IXC network by using network control links. These links are typically offered as part of the carrier's intelligent network service.

Cisco provides a NIC to interface to the specific carrier network. For example, if you have Sprint toll-free service, your Unified ICM system is equipped with a Cisco-supplied Sprint NIC.

The Sprint NIC allows the Unified ICM to interface with the Sprint intelligent network service. If you use both AT&T and Sprint as carriers, your Unified ICM system is equipped with AT&T and Sprint NICs. The following figure shows the interaction between the IXC network and the Unified ICM NIC.

Figure 3: Network Interface Controller

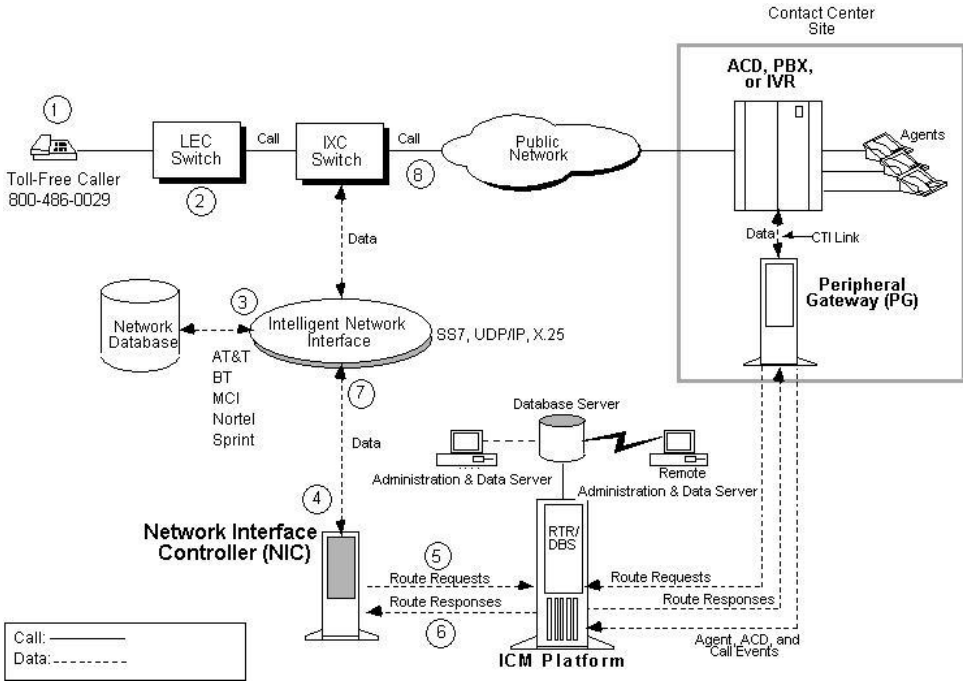

For clarity, the NIC in the figure above is shown as a separate computer. Actually, NICs are implemented as software on the Unified ICM software platform (for example, on the CallRouter or Logger machines).

An Unified ICM Network Gateway can be implemented for SS7 networks. The Unified ICM Network Gateway is implemented as a separate node on the Unified ICM signaling access network. When this node is implemented, the NIC software can be installed on the CallRouter machine. For Sigtran SS7 networks, a Sigtran Gateway can be deployed on either the CallRouter machine or a separate machine; the NIC software is installed on the CallRouter machine.

The specific flow of messages to and from the NIC within the Unified ICM software and the IXC network is shown by the circled numbers in the above diagram, and is explained in the following sections below.

Toll-Free Caller

As shown in the preceding figure, the flow of messages between the network and the Unified ICM begins when a caller dials a toll-free number (1).

LEC-to-IXC

The Local Exchange Carrier (LEC) determines which interexchange carrier (IXC) is providing transport for that particular number and forwards the call to the IXC switch (2).

Network Query

The IXC switch holds the call momentarily while it queries a network database to determine where to route the call (3).

ICM NIC

The network database forwards the query to the NIC and requests an intelligent routing decision (4).

NIC-to-CallRouter

The NIC software process receives the request, translates it into a standard format, and forwards it to the Unified ICM CallRouter process (5)

Best Destination Returned

The Unified ICM software selects the appropriate call routing script, assesses the skills and current real-time status of agents throughout the contact center network, and returns the best destination address back to the NIC (6).

IXC Network

The NIC sends the destination address to the IXC network (7).

Connecting the Call

The network instructs the originating IXC switch to connect the call to the destination specified by the Unified ICM software (8). The total time taken by the carrier to connect the call varies. However, the additional time added by the Unified ICM software to process the route request is typically less than half a second.

Carrier Connections

The table titled **Interexchange Carrier Connections** summarizes the basic supported carrier connections and their corresponding Unified ICM software routing client (NIC) and network transport protocol. Note that the SS7IN NIC is used for a number of carrier SS7 INAP interfaces.

Table 1: Interexchange Carrier Connections

Routing Client	Connection to ICM
AT&T	AT&T Network (SS7 INAP Gateway)
CRSP	Call Routing Service Protocol (UDP)

Routing Client	Connection to ICM
CWC	Cable & Wireless Gateway (SS7 Gateway)
GKTMP	Gatekeeper GKTMP interface (TCP/IP)
ICRP	Intelligent CallRouter Protocol (UDP)
INCRP	NAM/ICM Gateway Call Routing Protocol interface (UDP)
MCI	MCI Network (TCP/IP)
Nortel	Nortel Network (SS7 INAP Gateway)
NTL	NTL Network (TCP/IP)
Sprint	Sprint Network (X.25)
SS7IN	Generic / Extensible SS7 INAP (SS7 INAP Gateway)
Stentor	Stentor Adv Toll-free Gateway (HyperStream, TCP/IP)
UniSource INAP	Unisource (SS7 INAP Gateway)

Applying Fault Tolerance in NICs

You may already have a strategy for fault tolerance for some parts of the Unified ICM system. For example, you might have decided to use a duplexed, distributed Unified ICM central controller and duplexed PGs at each call center. It is just as important to apply fault tolerance to the NICs and intelligent network access links. Without a connection to the carrier's intelligent network, the Unified ICM system cannot perform pre-routing. If these links are lost, calls are typically routed according to the default routing plans set up in the carrier network.

Note: For more information on Unified ICM system fault tolerance, see the Administration Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted.

Goals for NIC Fault Tolerance

The goal in applying NIC fault tolerance is to add levels of protection that successively eliminate single points of failure.

Cisco requires an order of importance to follow when choosing the types of fault tolerance to apply in the carrier network-to-ICM system connection:

- First, use **redundant links** from the Cisco NIC to the carrier's intelligent network.
- Next, if you have redundant links, provision those links on **diverse facilities**. This adds another level of fault tolerance to your network connection.
- For NICs that run on the Unified ICM CallRouter platform, the NIC processes are duplexed when the CallRouter is duplexed.

  The types of NIC fault tolerance you apply have a bearing on the number of links you need to provision for IXC intelligent network access.

Link Redundancy

Cisco requires that you configure redundant links to the IXC network. In other words, rather than having a single link from the NIC to the IXC intelligent network, provision two links. Having just one link to the IXC network represents a single point of failure (that is, an area or node in the system that, should it fail, could cause the system to stop routing calls).

By using redundant links, you increase the reliability of the IXC network connection and add an important level of fault tolerance to the system. The following figure shows a simplexed Unified ICM central controller and NIC with redundant links to the IXC network.

Figure 4: Redundant Links

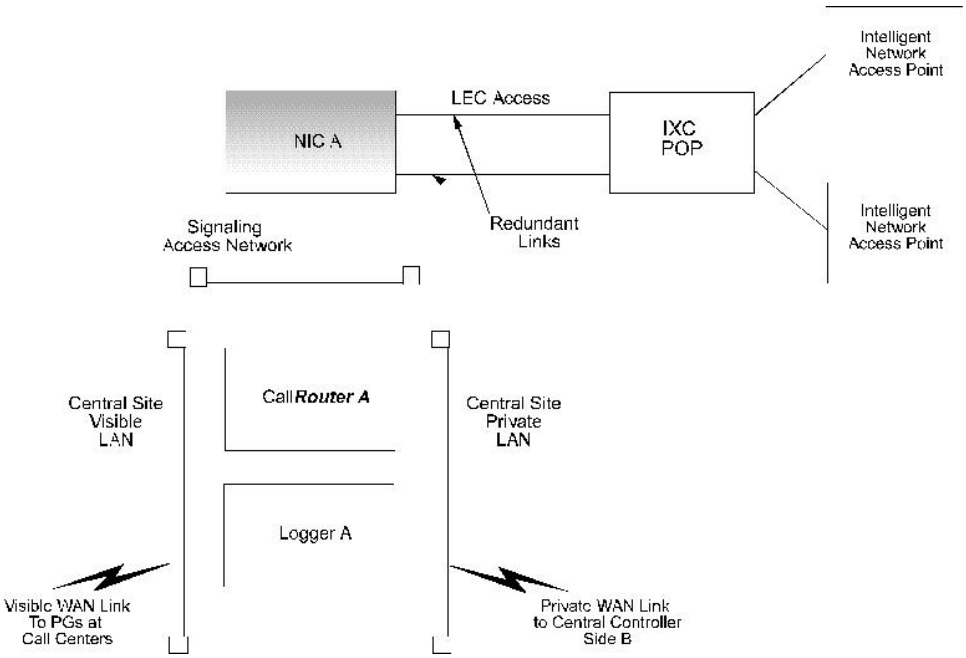

In the preceding figure, single points of failure still exist because the NIC, CallRouter, and Logger are simplexed. The simplexed central controller and NIC configuration is shown here only as an example. This type of simplexed configuration is used only for non-critical systems that can tolerate potentially long interruptions in service (for example, in lab or demo systems).

The major IXCs support redundant links to their intelligent networks. Contact your carrier for more information on access link options.

Route Diversity

For even more protection against network outages, Cisco requires that the network links be provisioned on diverse network facilities. By having diverse links, you further reduce the risk that another single point of failure (in this case, the failure of a circuit) might cause you to lose the connection to the IXC network. For example, you might provision one access link on one T1 circuit and provision the other access link on a different T1 circuit. By having diverse links, you protect against network failures in which an entire circuit is lost.

The following figure shows a simplexed Unified ICM system with redundant links and route diversity:

Figure 5: Redundant Links and Route Diversity

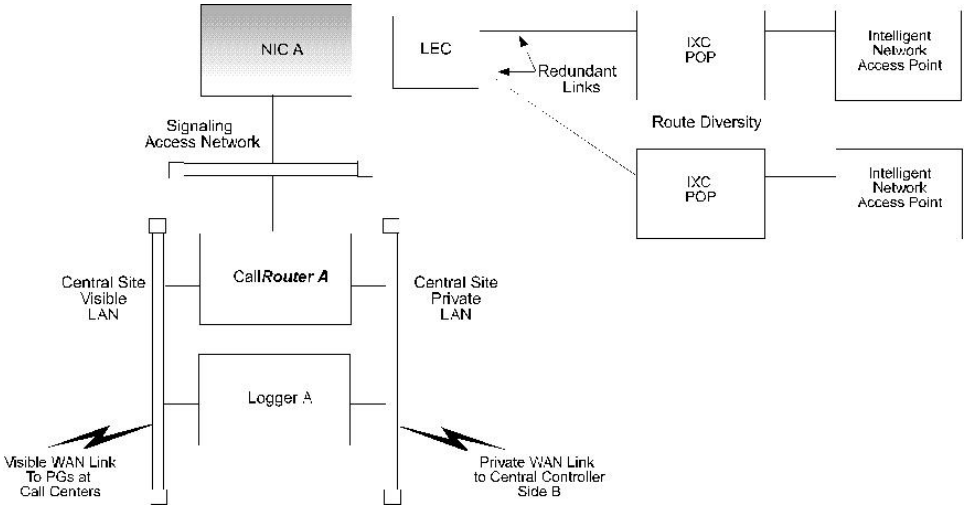

This example provides more fault tolerance by protecting against circuit failure or the loss of an IXC Point Of Presence (POP). Although the NIC is at one location, the redundant links connect to two different POPs. If one IXC POP is taken out of service (for example, in the event of a natural disaster), one link can still access the IXC network through the other POP.

The major carriers provide options for route diversity. Check with your carrier to discuss having the links handled by different POPs. You need to make sure that both the IXC and the Local Exchange Carrier (LEC) are using diverse circuits. Your LEC may impose some limitations on link diversity from the NIC to the IXC POP (that is, over the “last mile”). These limitations often depend on whether the call center is located in a metropolitan or rural area.

Chapter 4

Switch Overview

Each contact center device (ACD, PBX, or IVR) communicates with an Unified ICM Peripheral Gateway (PG). The PG reads status information from the device and passes it back to the Unified ICM software. The PG runs one or more Peripheral Interface Manager (PIM) processes, which are the software components that communicate with proprietary ACD systems. One PIM is required for each peripheral to which the PG will interface. So if you have two identical ACDs, your PG will require two PIMs.

A single PG can serve multiple peripherals of the same kind. For example, one computer with an Aspect PG and several Aspect PIMs can serve several Aspect ACDs in the contact center. Another PG and PIM on the same computer might serve an IVR.

Note: A single PG can support both ACD PIMs and IVR PIMs, though the ACD PIMs must all be of the same kind.

This chapter provides an overview of how the PG interfaces with ACDs in a contact center environment.

This chapter contains the following topics:

PG-to-Peripheral Connections

Each contact center peripheral (ACD, PBX, or IVR) requires a connection to a Cisco Peripheral Gateway (PG). The Peripheral Gateway provides a software interface between ACD, PBX, and IVR systems and the Unified ICM routing software.

The PG connects to a peripheral via the peripheral's computer telephony integration (CTI) link. In some cases, the PG also connects to the peripheral's MIS subsystem. The MIS subsystem

may be on a separate hardware platform or it may be integrated with the ACD, PBX, or IVR. The relationship of the Peripheral Gateway to an ACD system is shown in the following figure.

Figure 6: Peripheral Gateway ACD/PBX Interface

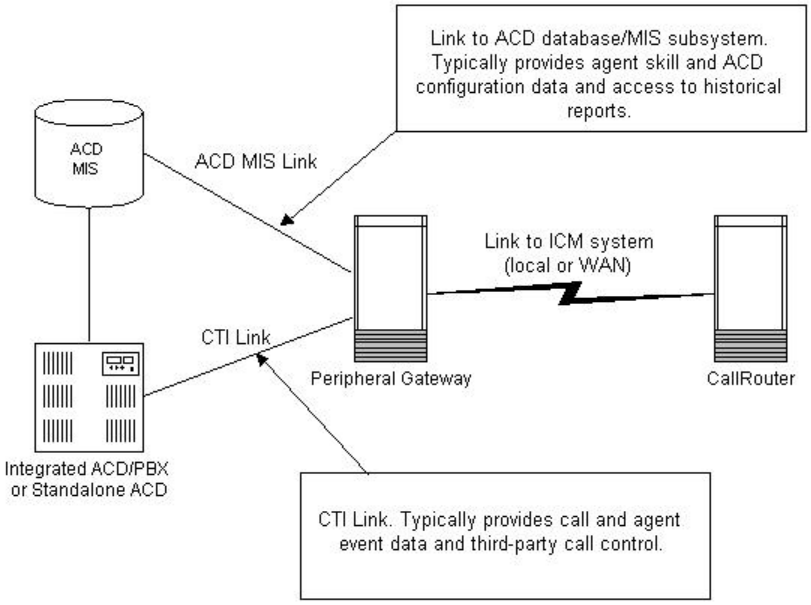

Through the CTI link, the PG monitors changes in agent status, calculates call handling performance statistics, and forwards events to the CallRouter. The MIS connection provides additional information such as the mapping of individual agents to skill types and the current status of agents (either by themselves or relative to a given agent group or skill group). Typical agent states include Logged In, Ready, Talking In, Talking Out, Work Not Ready, and so on. The MIS link may also provide the Unified ICM system with ACD configuration data and historical reports.

Each PG has one or more connections to the peripheral. The type of connection used depends on the type of peripheral. For example, some ACDs use a TCP/IP Ethernet connection, while others require X.25 links. Refer to the Cisco Unified ICM Software Supported Switches (ACDs) documentation for more information.

Supported ACD Switches

To ensure that your ACD software version is compatible with Unified ICM software, refer to the Cisco Unified ICM ACD PG Supportability Matrices document **http://www.cisco.com/univercd/cc/td/doc/product/icm/icmentpr/acddoc/icmacdmx.pdf** This document contains the latest information on Unified ICM switch support.

Note: For more details on how ACDs interface to the Unified ICM software, see the appropriate Cisco Unified ICM software ACD Supplement. The ACD Supplements provide more technical details on the ICM-to-ACD interface than is provided in this document.

Chapter 5

Peripheral Gateway Configurations

As part of planning for ACDs and PGs, you need to decide whether your Peripheral Gateways will be simplexed or duplexed. Simplexed means that one PG is used. Duplexed means that two essentially identical PGs are used, one as a backup system. (Both PGs run simultaneously, with some processes active on both PGs. See the discussion in Peripheral Gateway Fault Tolerance (page 34).)

The following figure shows examples of simplexed and duplexed contact center configurations. Typically, duplexed PGs are installed for fault tolerance.

Figure 7: PG Contact Center Configurations

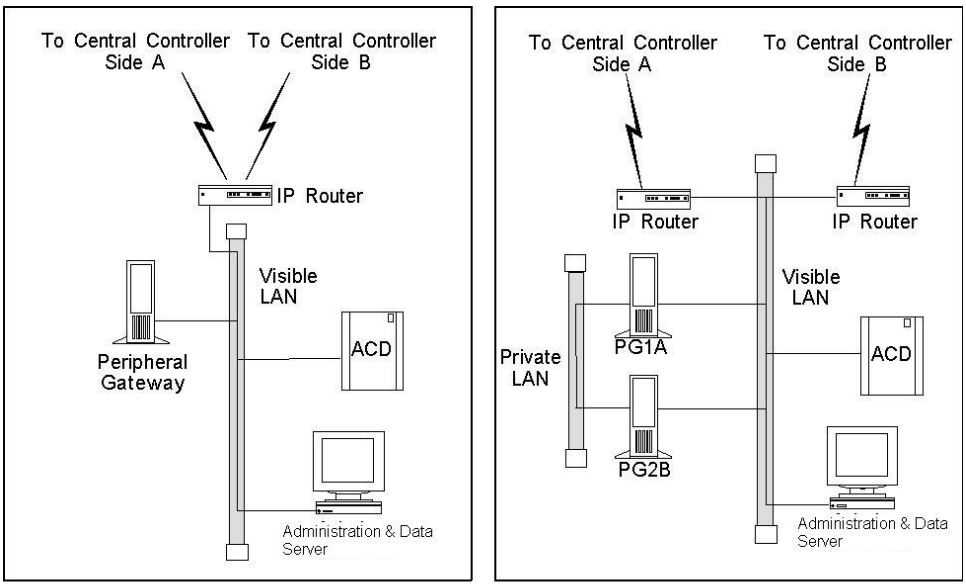

Simplexed PG **Duplexed PG**

Note: Some ACDs can connect directly to the Unified ICM visible LAN. Others connect to the PG via serial or other types of communication links.

The Peripheral Gateway reads information from one or more peripherals at a contact center and sends status information back to the Unified ICM CallRouter. A peripheral might be an ACD, IVR, PBX, or another device that distributes calls within the contact center. If the Unified ICM system is performing post-routing, the PG also sends route requests to the CallRouter and receives routing instructions in response.

This chapter contains the following topics:

Peripheral Gateway Fault Tolerance

Duplexed PGs are usually implemented to provide fault tolerance in the Unified ICM software's communication with peripherals. The duplexed PGs use a private network. The PG private network is used to synchronize certain processes within a duplexed PG pair. It is also used in "heartbeat detection," which is a process by which each PG sends a heartbeat packet every 100ms to keep track of the "health" of the other PG.

PGs use a combination of the hot standby and synchronization approaches to fault tolerance. In the hot standby approach, one set of processes is called the primary, and the other is called the backup. In this model, the primary process performs the work at hand while the backup process is idle. In the event of a primary process failure, the backup process is activated and takes over. In a duplexed PG system, the Peripheral Interface Manager (PIM) processes use the hot standby approach to fault tolerance.

In the synchronization approach, the critical process is duplicated on separate computers. There is no concept of primary and backup. Both process sets run in a synchronized fashion, processing duplicate input and producing duplicate output. Each synchronized process is an equal peer. Cisco refers to these equal peers as a synchronized process pair. In a duplexed PG system, the Open Peripheral Controller (OPC) process operates as a synchronized process pair.

The following figure shows how hot standby and synchronization are employed in a duplexed Peripheral Gateway.

Figure 8: PG Fault ToleranceACD2PG

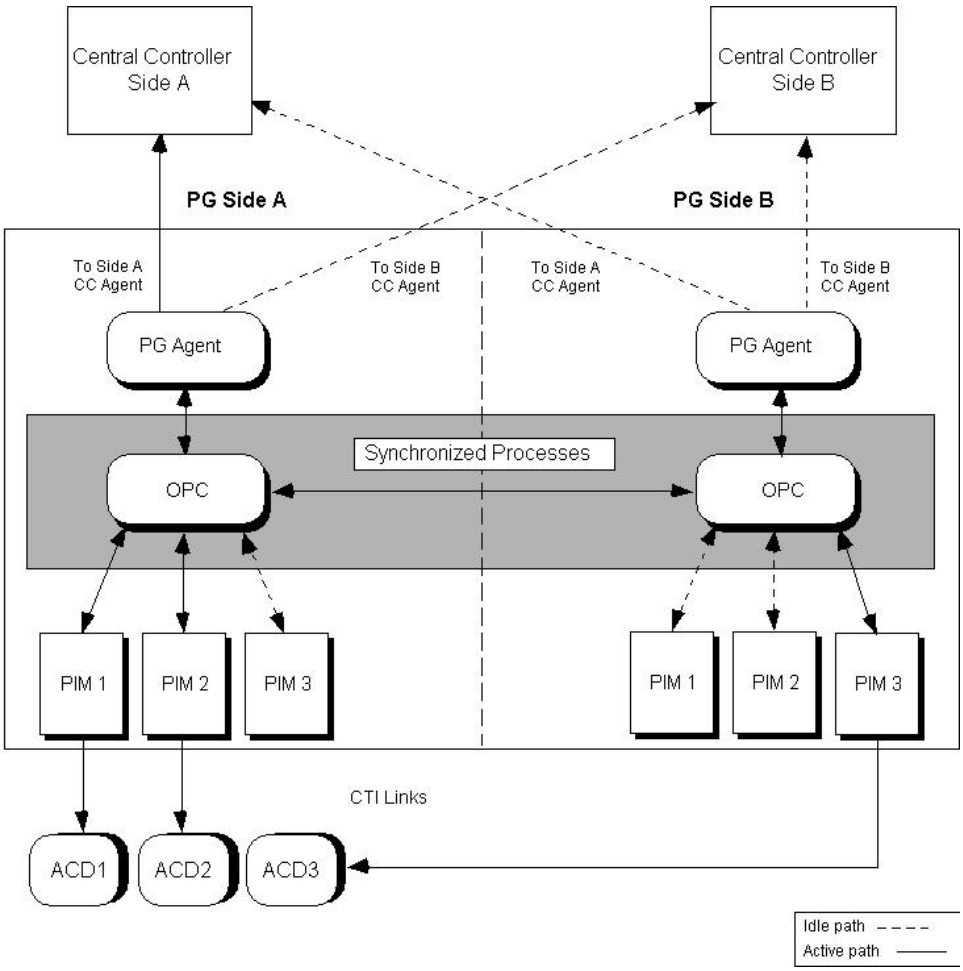

\

The OPC processes communicate with each other through a private network connection and the Cisco Message Delivery Service (MDS). The MDS provides a synchronizer service which combines the input streams from the PIMs and PG Agents on both sides of the PG to ensure that both OPC processes see exactly the same input.

The OPC process is responsible for activating PIMs and PG Agents on each side of the duplexed PG. The OPC process also supplies uniform message sets from various PG types to the Unified ICM central controller.

The PIMs manage the interface between different types of ACDs and the OPC. PIMs are duplicated on each side of the system and operate in hot standby mode. A PIM can be active on either side of the duplexed PG, but not on both sides at the same time. For example, in the preceding figure PIMs 1 and 2 are active on Side A; PIM 3 is active on Side B. The duplexed OPCs communicate with each other through the MDS to ensure that a PIM is active only on one side at a time.

The duplexed PG architecture protects against a failure on one side of the PG. For example, if an adapter card controlling access to an ACD fails, a hot standby PIM can use the alternate PIM activation path. As shown in the preceding figure, PIM3 has been activated from Side B of the

PG. This might be in response to an adapter failure between the Side A PIM3 and ACD3. In this type of failure scenario, the PG is able to maintain communication with the attached ACD.

Only one PG Agent actively communicates with a side of the central controller. When messages arrive at the central controller, they are delivered to both sides by the central controller Synchronizer process. The PG maintains idle communication paths to both sides of the central controller in case a switch-over to the other side of the central controller or PG is necessary.

PG Platform Options

A maximum of two PGs can run on a single hardware platform. A single PG can serve only one type of ACD, but can also contain one or more VRU PIMs and/or Media Routing PIMs provided that the server hardware has the capacity to support the aggregate processing load. For a single hardware platform to serve two different types of ACDs, you need two PGs—one for each peripheral type. The following figure shows some possible PG options.

Figure 9: PG Platform Examples

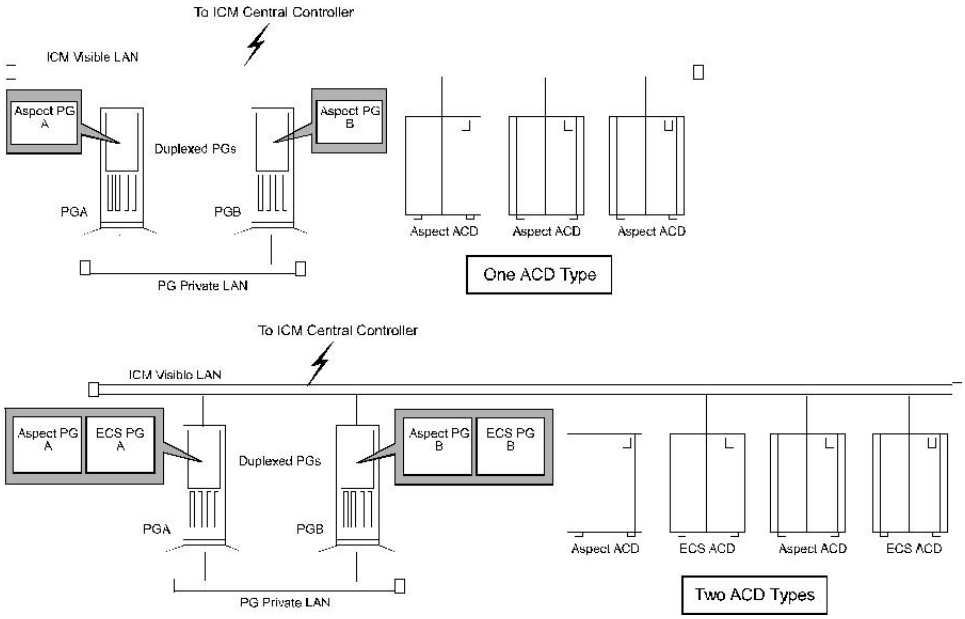

As shown in the preceding figure, you might have an Aspect PG on PGA and an Aspect PG on PGB. This duplexed PG pair could serve multiple Aspect ACDs. One Aspect Peripheral Interface Manager (PIM) would be added through Peripheral Gateway Setup for each Aspect ACD to be connected to this PG. In this example, three Aspect PIMs would be installed on each PG. The PIM is the Unified ICM software interface between the PG and different types of contact center peripherals. One PIM is required for each peripheral connected to a PG.

In a mixed contact center environment, you might want to run two different types of PGs on a single hardware platform. For example, you might want to put an Aspect PG and a DEFINITY ECS PG on the same computer. In this way, one hardware platform could serve two types of ACDs provided that the hardware platform has the necessary memory and CPU capacity to support the aggregate processing load.

Considerations for PGs and PIMs

Here are some points to remember when planning for PGs and PIMs:

Points to remember:

- Maximum PGs on a platform. A maximum of two PGs can run on a single hardware platform. These may be of the same or different types. For example, on a single machine you could have an Aspect PG and an Avaya PG, or you could have two Avaya PGs.

- PIMs and peripherals. You need one PIM for each peripheral that will be connected to the PG. The PIMs are installed along with the PG software by using the Peripheral Gateway Setup tool.

- A single PG serves peripherals of the same type. A single PG (and its associated PIMs) can serve only ACDs of the same type. For example, an Aspect PG with four PIMs can serve only four Aspect ACDs. It cannot serve three Aspect ACDs and an Avaya DEFINITY ACD. You can put VRU and Media Routing PIMs on the same PG as an ACD, but all VRU PIMs must service VRUs of the same type.

- Using two PGs on a platform. Before you commit to installing two PGs on a single computer, consider the expected call load for the ACDs that will be connected to the PGs.

  Note: Along with call load, you should also consider the number of CTIOS agents, number of VRU ports, as factors in determining server capacity.

  You need to be sure that the computer has enough memory and processing power to handle the expected call load. In addition, you should ensure that the bandwidth in the network between the PG and the Unified ICM central controller is enough to handle the route request and event traffic that will be generated by the PGs. (These same considerations apply when using multiple PIMs on a PG, but to a lesser extent.)

- **Properly sizing the PG hardware platform.** To properly size the PG hardware platform(s) and to determine which PG configuration is appropriate for your application see th*e Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1).* Cisco offers standard and high-end PG hardware platforms to suit more demanding contact center applications.

- **CTI Server and an ACD PG on the same platform.** You must install CTI Server and an IVR or ACD PG on the same hardware platform. The PG may run multiple PIMs. (The same considerations described earlier in “Using two PGs on a platform,” also apply to the CTI Server-PG configuration.)

- **IPCC Gateway.** In Unified ICM Enterprise, the IPCC Gateway PG allows the Unified ICM to pre-route calls to IPCC call centers and can also post-route IPCC calls. The IPCC Gateway feature allows IPCC Enterprise or IPCC Express to act as enhanced ACDs connected to the Unified ICM . Refer to the Cisco Contact Center Gateway Deployment Guide for Cisco Unified ICME/CCE/SCCE/CCX for more information.

Standard PG Configuration

In most PG configurations, the PG is located with the ACD at a contact center site. The PG communicates with the central controller via the Unified ICM visible network WAN links. These WAN links can be a dedicated circuit, or—if Quality of Service (QoS) is implemented—the corporate WAN can be used. When Administration & Data Servers are located with PGs and ACDs at the contact center site, the WAN links to the central controller can be shared by both PGs and Administration & Data Servers. If the PG is collocated with the Unified ICM central controller, the PGs connect directly to the Unified ICM visible LAN. The following figure shows an example of a standard PG configuration.

Figure 10: Standard PG Configuration (Duplexed PGs)

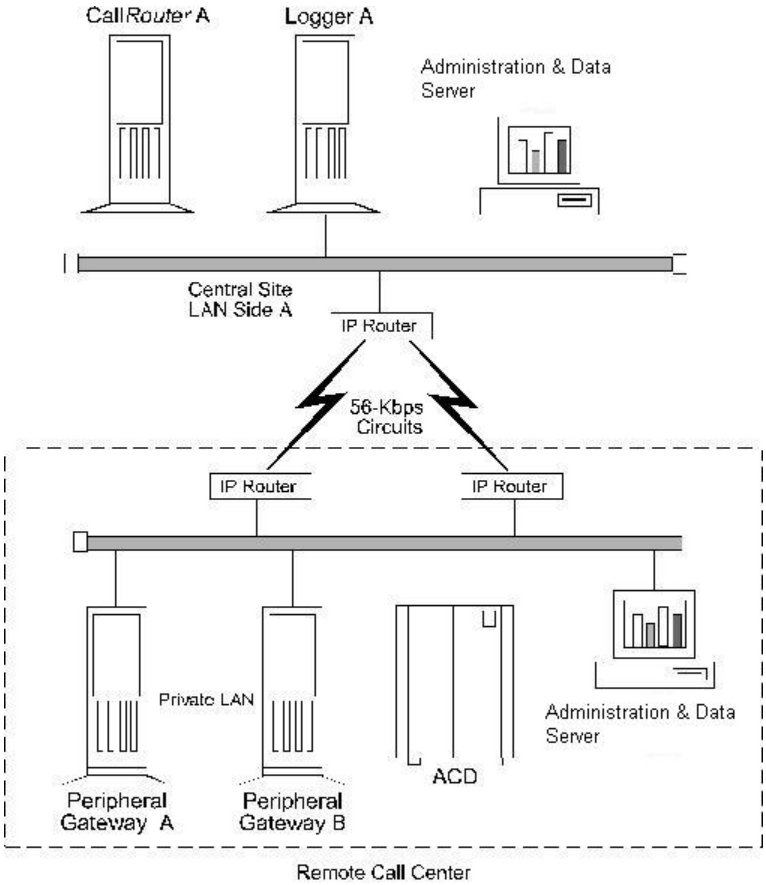

Remote ACD and IVR Connection to PGs

Some ACDs allow a remote connection to the Unified ICM Peripheral Gateway. In a remote ACD configuration, the PGs are located at the central site along with the CallRouter, Logger, and NIC. The ACD is located at a remote contact center site.

For information on remote PG support, see the ACD Supplement for the particular ACD. Generally speaking, Alcatel, Aspect, Avaya, NEC, Siemens, Symposium ACDs are supported

over the WAN. However, in all cases, you must check with the ACD manufacturer for any WAN limitations.

The IVR PG can communicate remotely with IVRs via a TCP/IP network. However, you must ensure that the network link between the PG and IVR system provides enough bandwidth to support the call load for the VRU.

Multiple PGs Connecting to a Single ACD

It may be necessary to connect multiple PGs to the same ACD. This type of configuration is required when multiple Unified ICM customers need to share the same service bureau ACD. In order for this configuration to be possible, the ACD must allow multiple CTI applications to share its CTI link(s). Support for multiple PG connections varies depending on the ACD platform. Please see the appropriate Cisco Unified ICM software ACD Supplement and contact your ACD vendor to determine the availability of this functionality.

Chapter 6

CTI Planning

Cisco CTI software provides an interface between the Unified ICM software and agent desktop and server applications. The CTI software works with a Peripheral Gateway's ACD and IVR interface software and associated ACDs to track call events and transactions and forward call- and transaction-related data to an agent's desktop computer.

Pre-installation planning for CTI involves several tasks:

- Review CTI Server communications and platform options.
- Become familiar with the desktop options available with CTI Server.
- Estimate CTI message traffic.
- Plan fault tolerance for the CTI Server.
- Review ACD support for client control and third-party call control.

This chapter contains the following topics:

CTI Server

CTI Server, the basic server component of Cisco CTI, enables the Unified ICM software to deliver agent, call, and customer data in real-time to a server and/or workstation application as events occur throughout the life of a call. The CTI Server is a software process that runs on a

Peripheral Gateway (PG). It is the CTI gateway into the Unified ICM software's data and services.

The following figure shows a sample CTI Server system. CTI Servers may be running at one or several call centers in the enterprise.

Figure 11: CTI Server Overview

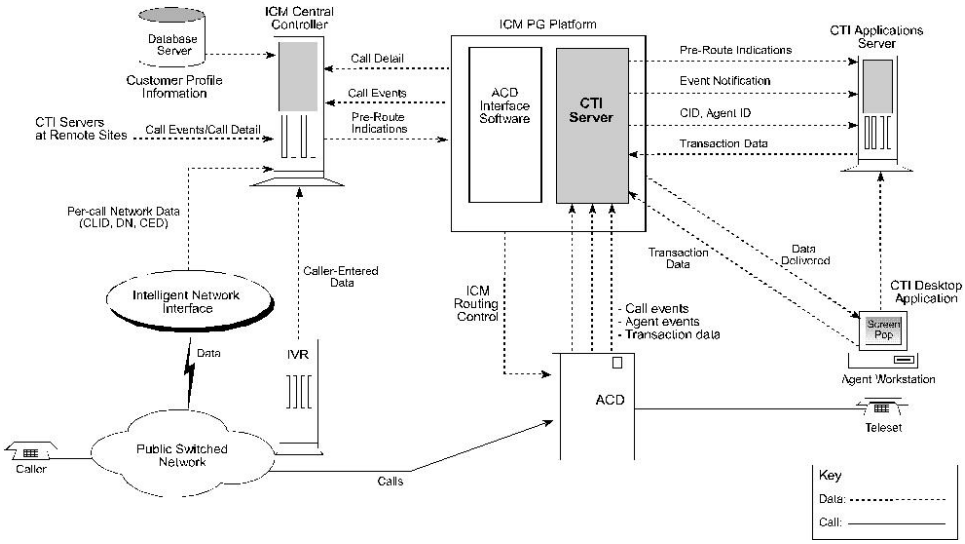

One function of the CTI Server is to forward pre-route indications to CTI application servers. Pre-route indications identify the caller and provide CTI applications with other call attributes while the call is still in the public or private network (that is, before the call is connected to an agent or IVR resource).

CTI Server also reports call events and agent work state changes as they occur through each stage of the call flow—from the moment a call arrives at an answering resource (ACD, PBX, IVR), until the caller hangs up. In a desktop application environment, call event information is delivered to the targeted agent desktop at the same time the call is delivered.

CTI Server Communications

The CTI Server uses TCP/IP Ethernet for communication with clients. Multi-protocol IP routers may be used to provide connectivity to clients on other types of LANs. The same LAN that is used for the Peripheral Gateway's visible network interface can also be used for CTI client-to-server communications.

CTI Server Platform Options

The CTI Server runs on a machine that is also running a Cisco ACD (or VRU) PG process. The shared platform option is shown in the following figure.

Figure 12: Shared CTI Server Platform

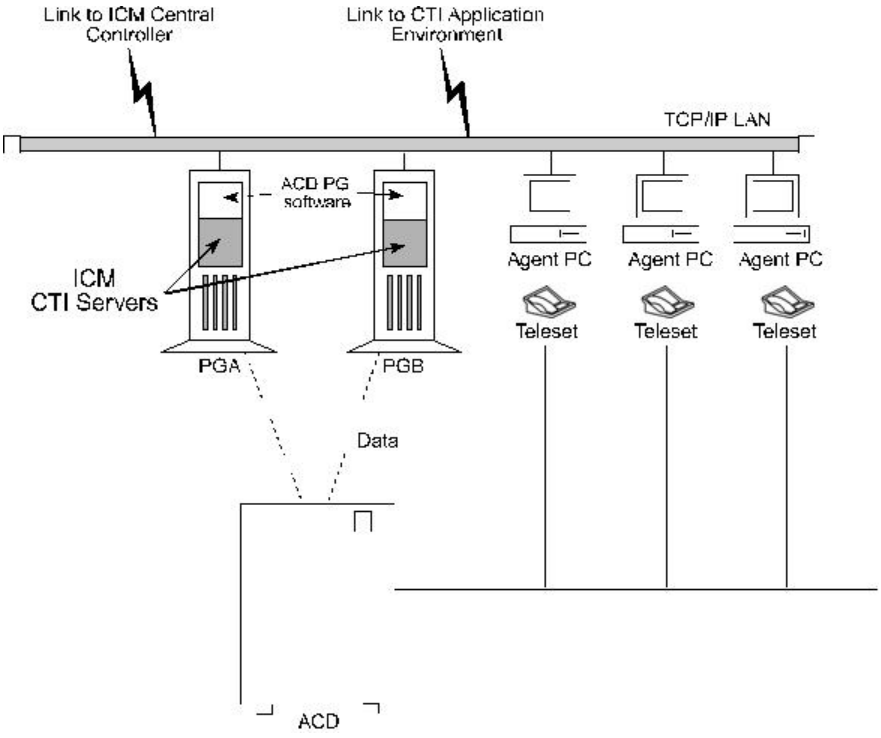

CTI Server Fault Tolerance

You can implement the CTI Server in a duplexed, fully fault-tolerant configuration. In a duplexed configuration, the CTI Server is installed on a pair of server platforms. In the event of a failed CTI client connection, the client process can automatically reestablish a connection to either side of the duplexed CTI Servers. The call's CTI client history list and any updates to call variables remain in effect when the connection is reestablished. The following figure shows a duplexed CTI Server configuration.

Figure 13: Duplexed CTI Server

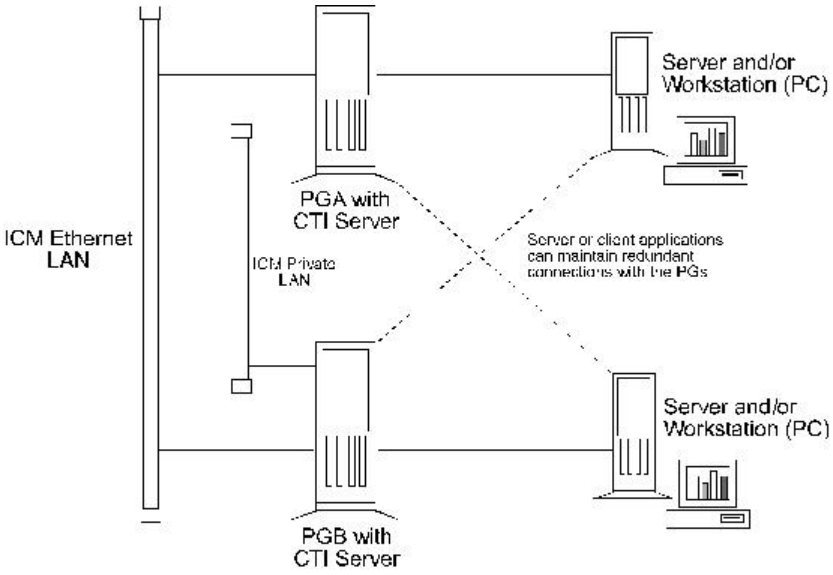

Cisco CTI Object Server (CTIOS)

CTI Object Server (CTIOS) is a high-performance, scalable, fault-tolerant server-based solution for deploying CTI applications. CTIOS serves as a single point of integration for third-party applications, including Customer Relationship Management (CRM) systems, data mining, and workflow solutions.

Configuration data is managed at the server, which helps simplify customization, updates, and maintenance of CTI applications. Servers can be accessed and managed remotely. Thin-client and browser-based applications that do not require Cisco software on the desktop can be developed and deployed with CTIOS.

CTIOS incorporates the following major components:

- CTIOS Toolkit
- Client Interface Library
- CTIOS Combo Desktop for Agents and Supervisors

CTIOS is a client of CTI Server. It has a single all-events connection to Cisco CTI Server. In turn, CTIOS accepts client connections using session, agent, and call interfaces. These interfaces are implemented in .NET, COM, Java, C++, and C, allowing for a wide range of application development uses. The interfaces are used for call control, to access data values, and to receive event notifications.

For complete and current information about the number of agents supported for CTIOS and other hardware configurations, see the *Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1).*

For new installations, the CTIOS Server should be co-resident with the PG.

For more information on CTIOS, refer to the Cisco Unified ICM Software CTIOS document set.

CTI Server Client Application Models

You can use either of two client models to integrate call center applications with the Unified ICM : agent workstation and CTI Bridge.

Agent Workstation (Desktop) Application

In the agent workstation model, the client is an application running on a personal computer on an agent's desktop. This client is interested in the call data and call events related to a single agent teleset. The agent workstation application might also be interested in agent state changes.

Typically, when the agent workstation application is informed of an incoming call, it will likely use the call data collected by the Unified ICM system to retrieve caller-specific data from a database. This data is presented on the agent workstation screen at approximately the same time that the incoming call is connected to the agent.

While handling the call, the agent may wish to update some of the call data. For example, an agent who is processing an insurance claim may make some adjustments to the call data; an update ensures that the changes are not lost before the call is transferred to a second agent.

Upon completion of the call, the client may be used by the agent to add call-specific, wrap-up information to the Termination_Call_Detail record logged in the Unified ICM central database. This wrap-up data may be a key value that can help locate more detailed transaction information in some other database. If the agent should conference with or transfer the call to another agent on the same ACD with a CTI client workstation, then that agent's CTI client also receives the incoming call data, including any updates made by the first agent. If the transfer or conference involves an agent on another ACD, the call data is provided to the remote CTI client if a translation route is used.

CTI Bridge (All Events) Application

CTI Bridge applications are interested in all call and agent state events that occur on the ACD, unlike agent workstation applications that are interested only in the events associated with a particular teleset. The CTI Bridge application is a user-written program that converts or adapts some or all of the CTI Server messages into another format; a single CTI Bridge application provides such services for multiple agent desktops. The CTI Bridge application can be designed to interface with CTI Servers or similar applications on systems that are already in use in the call center.

Some examples of CTI Bridge applications include:

- Message converter applications. For example, an application may convert the CTI Server message set to the message set of a foreign telephony server.
- Server-to-server communication applications. For example, an application may enable the CTI Server to speak directly to a help desk application's middle tier server.

  In a CTI Bridge configuration, a CTI Bridge application provides the connection between an existing desktop CTI application and the Unified ICM .

Figure 14: CTI Bridge Model

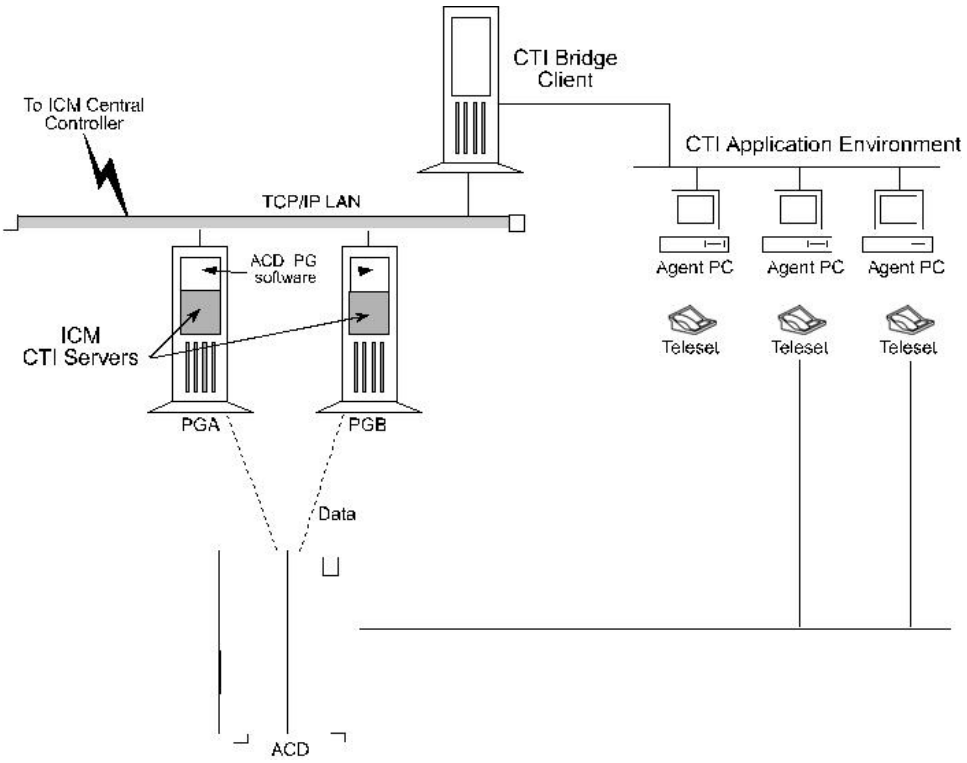

Note: All of the functionality found in the agent workstation (desktop) model is also available in the CTI Bridge application model. However, the CTI Bridge application must be written to support this functionality.

CTI Server Network and Database Planning

Some pre-installation planning is necessary to prepare your CTI desktop and network environment for the introduction of a CTI Server.

Review the Desktop Network Environment

The machine running CTI Server connects to the CTI desktop environment via an Ethernet LAN. Therefore, the CTI desktop environment must reside on an Ethernet LAN. Other networks, such as Token-Ring, may require additional network hardware if they are to be connected to a CTI Server.

Review Network Security Issues

You need to be sure that the CTI desktop environment IP routing scheme is compatible with the Unified ICM system and CTI Server. For example, you might currently have a firewall set up on the CTI environment LAN. If there is a firewall, you may need to change your system access setup or network configuration.

Address Desktop Software Roll-out and Distribution Issues

If you are going to be installing the CTIOS or CTI Desktop software components on multiple desktops, you need to create a distribution strategy. For example, you may decide to place the software on a centralized server and allow certain desktops within the enterprise to download the software. In addition, if you use this strategy, and will be installing software across distributed sites, you must ensure that all sites have access to the centralized server.

Select a Well-known Port for CTI Server

A well-known port number identifies CTI Server as an application running in your intranet. All CTI clients, as well as the system administrator, need to be aware of this well-known port number. If you do not want to use CTI Server's default port numbering scheme, you can choose a well-known port number that fits into your overall network environment. Peripheral Gateway Setup allows you to override the default port settings used to install the CTI Server PGs.

Plan a Fail-over Strategy for CTI Clients

Cisco CTI includes automatic fail-over and recovery mechanisms. Ensure that each CTI client has a clear and established network path to a CTI Server in case of a fail-over. For example, you might plan for each CTI client to have access to local and remote CTI Servers.

Develop a Database Strategy

You might have CTI applications that perform database queries to retrieve customer information for use in call processing. Some CTI applications might acquire database records "pre-call" (that is, before the call arrives at an agent's desktop). Other applications might query a database immediately after the call arrives at the agent's deskset. Plan a strategy for executing database queries in the most efficient and timely manner possible.

CTI Server Message Traffic

The CTI Server makes call data available to applications in real time. To accomplish this task, the CTI Server process responds to requests from clients and originates unsolicited messages. All messages share a common message header and use the same set of data types.

The table titled **CTI Server Message Categories** groups the messages into broad categories based on the nature of the message data.

Table 2: CTI Server Message Categories

Category	Description
Session Management	Messages related to the establishment and maintenance of a client connection to the CTI Server. These messages

Category	Description
	typically happen at client startup, shutdown, and during auto-recovery.
Miscellaneous	Messages related to system-level events on the PG (for example, peripheral offline, loss of PG-to-central controller communications).
Call Events	Messages related to call state changes.
Agent Events	Messages related to agent state changes.
Call Data Update	Messages related to CTI client modification of call data.
Client Control	Messages related to the direct control of agent state (for example, login, logout) as well as control of inbound and outbound calls.

CTI Server imposes varying degrees of message traffic against the PG based on the specific call center and CTI application environment in which it is deployed. Document a typical call scenario in your CTI application environment, prepare for adequate bandwidth, and order the proper server platform.

Note: For a description of the session management messages, see the latest version of the *Cisco Unified ICM Software CTI Server Message Reference Guide*.

Documenting a Typical Call Scenario

To estimate CTI Server message traffic, document a typical call scenario in your CTI application environment. The goal of this exercise is to account for all types of potential message traffic in the link between the CTI Server and the CTI application environment.

For example, a typical call might be handled as follows:

- The call is pre-routed.
- The call receives a call treatment such as a request to set call data.
- Next, maybe a simple call release, hold, transfer, or post-route request takes place.
- During this time an agent state may have changed (for example, from ready to work ready).

Estimating Required Bandwidth

You need to ensure that you have enough bandwidth in the datacom connection to handle the message traffic between the CTI Server and the CTI application environment. For example, are you sure that a 56-Kbps connection will be adequate for your environment?

The call scenario process helps you to estimate the message load and calculate how much bandwidth is required in the link between the CTI Server and the CTI application environment (for example, 56K, 256K, or more).

Choosing the CTI Server Platform

You also need to ensure that the CTI Server platform has adequate CPU processing speed and RAM to handle the message activity. You may require a high-end Cisco CTI Server/PG platform for the CTI Server.

Third-Party Call Control

The term call control refers to the ability of an application that is external to the ACD to programmatically control a telephone call. For example, a CTI application might put a call on hold, transfer the call, or hang up the call.

With first-party call control, the CTI application can control only the teleset that is physically connected to the computer running the CTI application. First-party call control requires a physical connection between the computer and the telephone and other add-on hardware (see the following figure).

Figure 15: Desktop First-Party Call Control

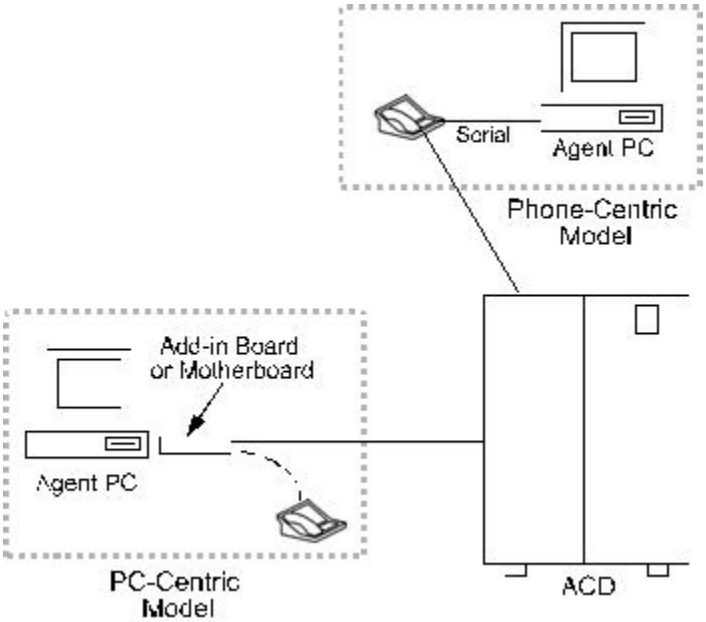

CTI Server products support third-party call control. Any call control initiated outside the ACD/teleset domain is referred to as third-party. With third-party call control, there is no physical connection between the computer and the teleset (see the following figure).

Figure 16: Desktop Third-Party Call Control

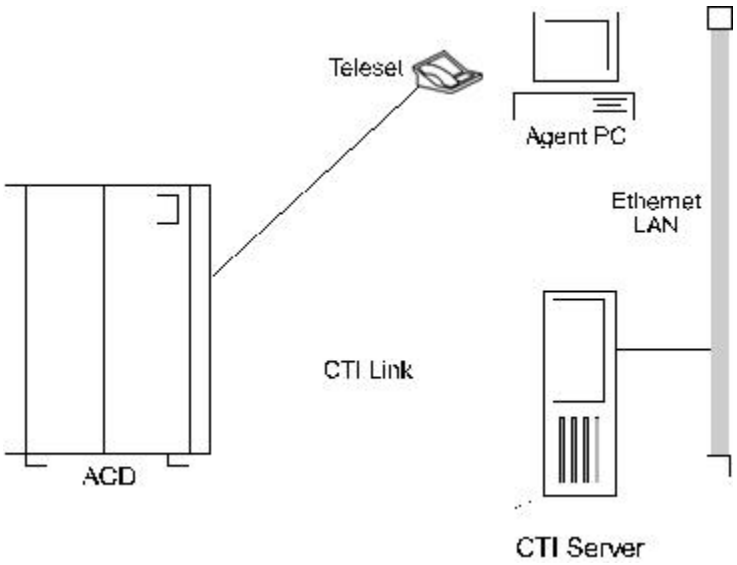

The desktop CTI application communicates with the Cisco CTI Server over a LAN. The CTI Server in turn communicates with the ACD to send call control requests. In this model, the CTI application is not bound to any particular teleset. The CTI application can control any teleset connected to the ACD and CTI Server.

Note: Most, but not all, ACDs support third-party call control.

Depending on the specific ACD, the client application can perform all or most of the following telephony functions:

- Answer/Hang up
- Agent Login and Wrap-up data
- Consultative/Blind Conference
- Consultative/Blind Transfer
- Generate DTMF tones
- Get/Set Agent states
- Get/Set ICM call data (ANI, DNIS, CED, UUI, call vars)
- Hold/Unhold/Swap Hold
- Make a call
- Redirect

ACD Support for Client and Third-Party Call Control

Different peripheral types implement and support varying levels of CTI functionality. For example, a different set of client control requests and call event types may be available depending on the peripheral type. In addition, there may be other CTI-related restrictions and implementation differences based on the type of peripheral. You need to take these differences into account when you write a CTI client application that will interface with third-party switches and devices.

For example,

- The Rockwell Galaxy does not have CTI Server support.
- The Siemens Rolm 9751 CBX does not have CTI Server support, but does support screen-pop applications.

As part of CTI pre-installation planning, you need to review ACD support for client control and third-party call control.

Chapter 7

IVR Planning

Cisco provides an option for running an interface to Interactive Voice Response (IVR) systems. The IVR interface software allows IVRs to take advantage of Unified ICM call routing features. For example, an IVR can use post-routing capabilities to select targets for calls it needs to transfer. The IVR interface software runs as a process on a standard PG hardware platform. It allows the Unified ICM to route calls to targets on an IVR and collect data from an IVR for use in call routing, real-time monitoring, and historical reporting. The IVR interface is not specific to a particular IVR system or manufacturer. It is based on an open IVR model. Many IVR systems support Cisco's Open IVR Interface Specification, including Unified CVP. For a list of IVRs that support this interface, contact your Cisco representative.

To plan for this IVR option:

- Review the options for integrating IVRs into the Unified ICM system.
- Determine if any IVR programming or application development is necessary.
- Review the Peripheral Gateway platform requirements.

Reviewing IVR Configuration Options

IVRs can be located at the customer's call center site or in the IXC network. At the call center, the IVR might be connected on the network side of the ACD or "behind" the ACD. In the IXC network, the IVR may be offered as a service by the network provider.

In an Unified ICM configuration that includes an IVR, the ACD is configured so that it can transfer calls to the IVR. The following figure shows some of the capabilities of the IVR in an Unified ICM system.

Figure 17: IVR/ICM Integration Overview

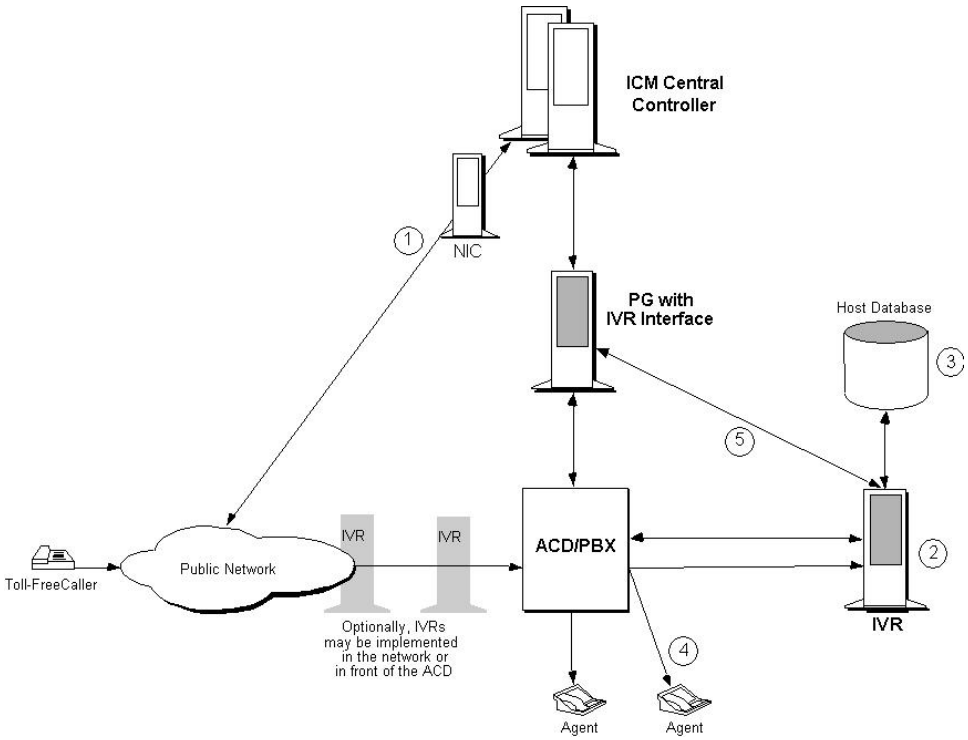

Capabilities of IVR:

1. In most Unified ICM /IVR configurations, calls continued to be Pre-Routed by the Unified ICM system.

2. When a call is routed to an IVR, the IVR answers the call and interacts with the caller.

3. The IVR may access a host system (for example, a customer profile database) to retrieve more information to help process the call.

4. Often, the caller can get all the information he or she needs through simple interaction with the IVR. In some cases, however, the IVR needs to transfer the caller to an agent or another call resource.

5. In some configurations, the IVR can invoke post-routing to select an agent from anywhere in the call center enterprise. To do this, the IVR sends a route request to the PG. The PG forwards the request to the Unified ICM system, which responds with a new destination for the call. The PG returns the new destination to the IVR. The IVR then signals the ACD or network to transfer the call to the specified destination.

The way in which an IVR is integrated into the Unified ICM system affects the flow of call processing and determines the types of data the Unified ICM can collect from the IVR. For example, an IVR that has a direct interface to an IVR PG provides the Unified ICM system with data that can be used in call routing, monitoring, and reporting. A configuration in which the IVR has an interface only to the ACD has more limited capabilities.

You can integrate IVRs into the Unified ICM system in several different ways. Each integration option provides a different set of Unified ICM functionality.

Configuration with an ACD PG Only

In this option, the IVR is attached only to the ACD. The ACD, in turn, is attached to a PG. The PG is running the Cisco peripheral interface software (PG software process) required to communicate with the specific type of ACD. There is no direct interface between the IVR and the Unified ICM system (in other words, an IVR process is not implemented).

Figure 18: Configuration With an ACD PG Only

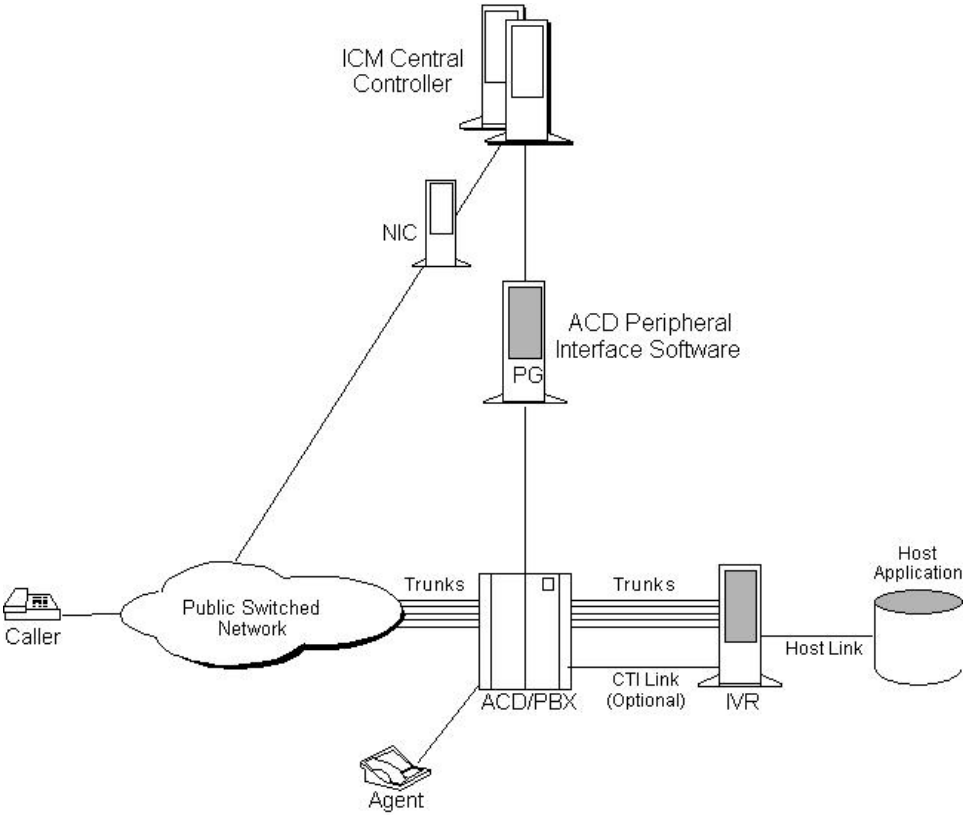

In this configuration, the IVR must be connected to an ACD that supports post-routing. The IVR and ACD cooperate so that calls can be transferred from the IVR to the ACD, and then post-routed by the ACD via the PG. The PG in this configuration has only the ACD peripheral interface software. It does not have the IVR interface software. Therefore, it does not provide the IVR with full access to post-routing.

In the preceding figure, the IVR can handle a call in two different ways:

- The IVR can handle the call to completion (for example, if the caller wanted current billing information and needed no further assistance, the IVR could complete the call.)
- The IVR can transfer the call to the ACD. The ACD could then use the PG to post-route the call.

Configuration with IVR and ACD PGs

This configuration option is similar to the previous option except that an IVR process and host link to the IVR are implemented. In addition to monitoring the ACD for real-time agent and call event data, the PG can monitor the IVR for call and application data and control the movement of calls into and out of the IVR. The IVR data is also forwarded to the CallRouter for use in call routing and reporting.

As shown in the following figure, the IVR and ACD interface software can be installed on the same PG hardware platform.

Figure 19: Configuration with IVR and ACD PGs

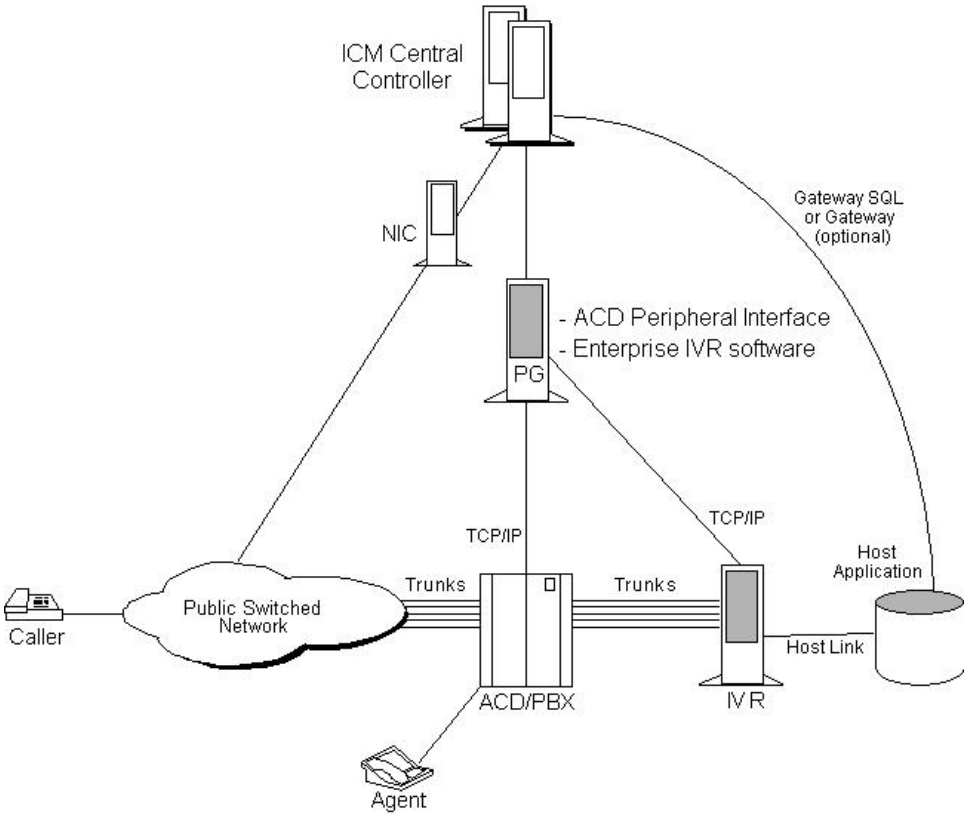

Network-Side IVR with IVR and ACD PGs

The next configuration option places the IVR on the network side of the ACD. In this configuration, the IVR is connected to the network and potentially to the ACD. The IVR can receive calls directly from the network without ACD involvement. These calls might be pre-routed by the Unified ICM , but this is not a requirement.

The IVR may also receive calls from the ACD (for example, when an agent transfers a call to the IVR). Again, these calls may or may not have been routed by the Unified ICM . The following figure shows an example.

Figure 20: Network-Side IVR with IVR and ACD PGs

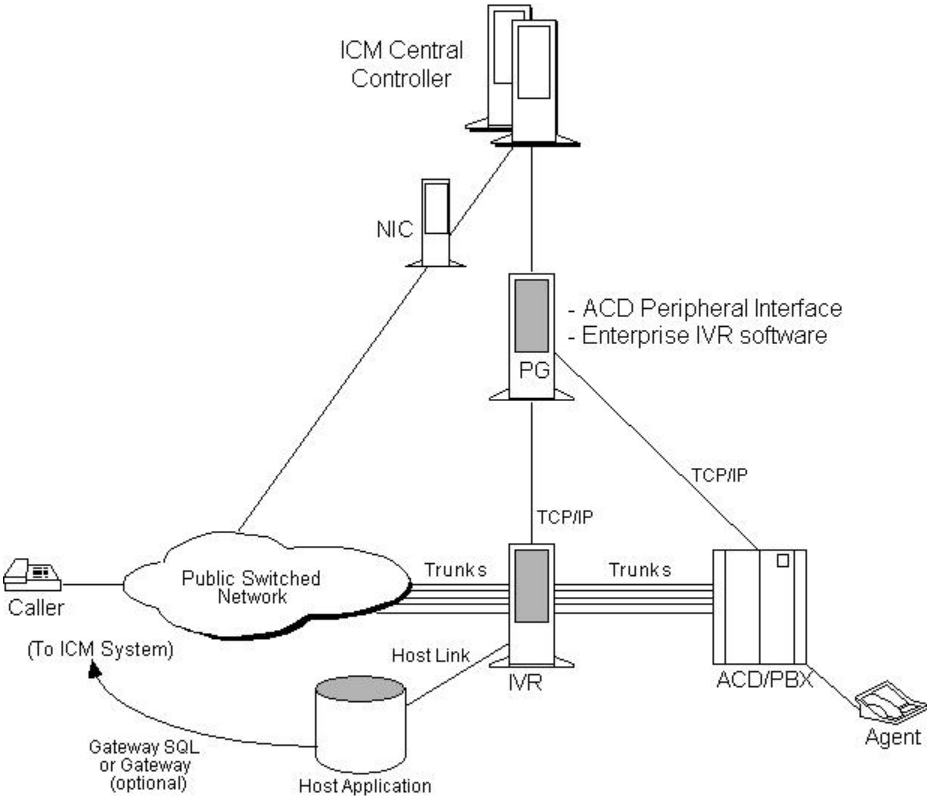

Once the IVR receives a call, it may handle the call to completion or transfer the call off-IVR for subsequent handling. The IVR may also use post-routing to select a target for the transfer. If the IVR transfers the call to an ACD, the IVR may or may not request routing instructions from the Unified ICM .

This configuration is different from the earlier options in several ways:

- The IVR is connected to both the network and the ACD.
- A call that originated in the network can be transferred to the local ACD by tandem connecting a second trunk with the original trunk. A network call can be transferred to a remote ACD either by connecting a second trunk in tandem with the original trunk, or by invoking a “call take-back” feature in the network.
- A call that originated at the local ACD can be transferred to any target using post-routing.

In-Network IVR with an ACD PG Only

In this configuration, the IVR is provided as a service by the network service provider. The PG monitors the ACD and forwards data to the Unified ICM system for use in call routing and reporting.

Figure 21: In-Network IVR with ACD PG Only

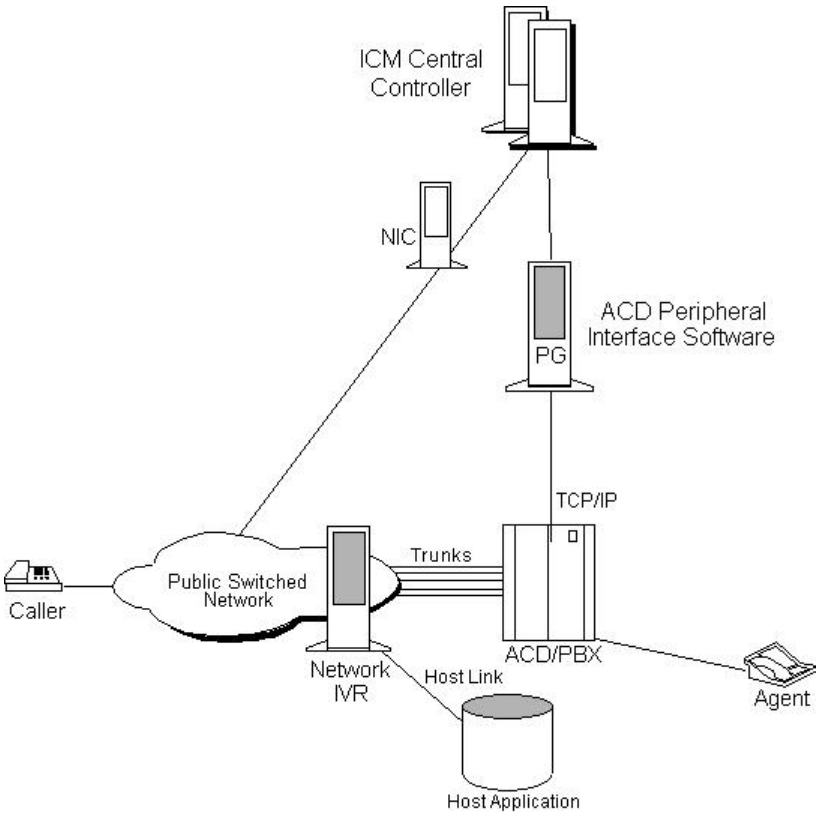

When the caller dials the toll-free number, the Unified ICM instructs the network to transfer the call to the network-based IVR. The network IVR then prompts the caller for input. If the caller requires additional information (such as speaking to an agent), the IVR dials a “hidden” toll-free number. The network then queries the Unified ICM system for a routing destination. The Unified ICM system returns a routing label and the network transfers the call to the specified ACD and DNIS. An agent at the ACD may handle the call to completion or transfer the call for subsequent handling.

In-Network IVR with IVR and ACD PGs

In this configuration, the IVR is provided as a service by the network provider. The network transfers all calls to a destination IVR. The IVR is responsible for handling a call to completion or transferring the call to another resource (for example, an agent at an ACD).

Figure 22: In-Network IVR with IVR and ACD PGs

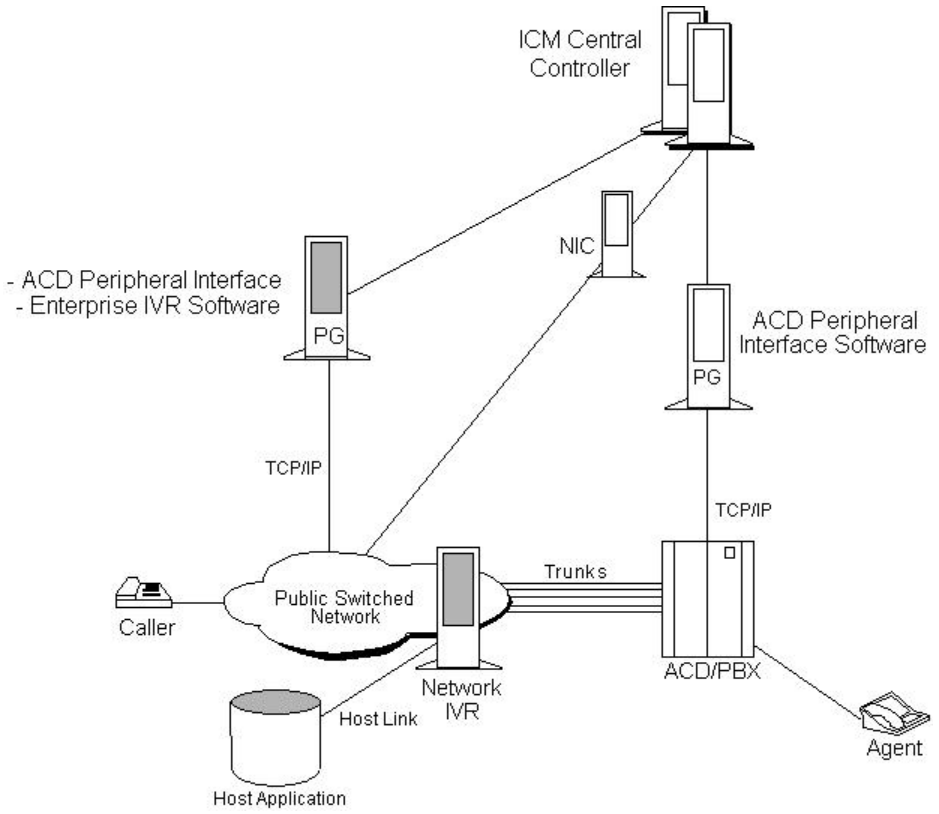

IVR Transfer Routing Using Third-Party Call Control

In this configuration, the IVR invokes a transfer request to transfer a call to the ACD. The IVR uses a CTI link to the ACD which allows it to set variables in the transfer request (for example, CED, DNIS, CLID, Social Security number, or account number). This configuration is viable only if the IVR is attached to an ACD that supports post-routing. The following figure provides an example of this configuration.

Figure 23: IVR Transfer Routing with Third-Party Call Control

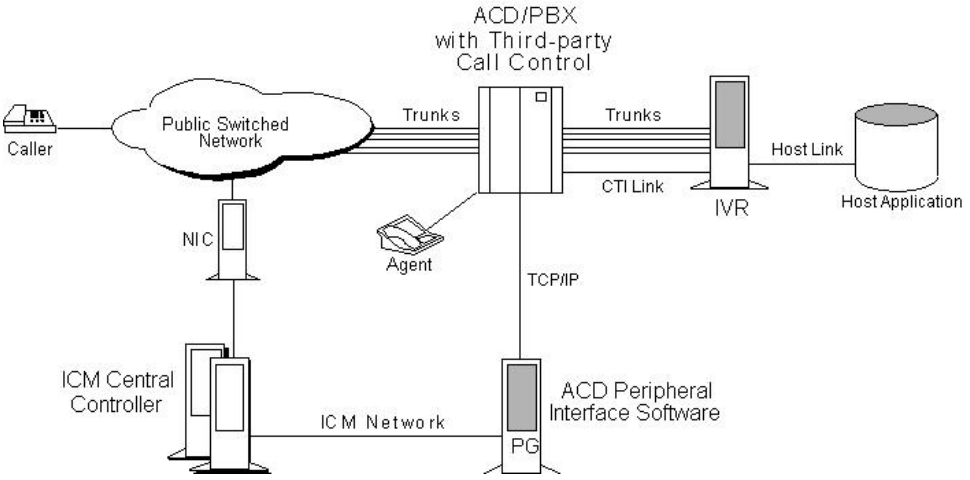

When the ACD receives the transfer from the IVR, it makes a route request to the PG in order to conduct an enterprise-wide agent selection. The PG routing client sends a route request to the CallRouter. The CallRouter passes a response to the PG and on to the ACD. The ACD then transfers the call to the specified destination.

IVR Programming and Application Development

The Open IVR Interface allows the Unified ICM to see some level of IVR application-specific data (for example, menu selections). An IVR application developer can use the Open IVR Interface to implement call routing (routing client) and monitoring capabilities.

The IVR routing client allows the IVR to send route requests to the Unified ICM via the PG. These requests can include data variables such as Customer ID and Menu Selections. The Unified ICM system can use this data to instruct the IVR where to transfer the call. The IVR monitoring interface allows the application developer to send IVR port and application activity data to the Unified ICM system for use in call routing and reporting.

IVR Peripheral Gateway

The Cisco IVR interface software runs as a logical PG on a standard Peripheral Gateway hardware platform. A single PG hardware platform can support a maximum of two logical PGs. A single PG platform might run one or two IVR PGs. Or it might run an IVR PG and an ACD PG. For example, you could have a PG hardware platform that runs an Aspect CallCenter PG and an IVR PG. A logical PG can have PIMs for one type of ACD, plus an IVR PIM. The hardware platform must have sufficient capacity to handle the aggregate load from all attached peripherals.

Note: The multi-instance CTIOS configuration supports up to ten logical PGs on a single PG platform. These PGs are configured as separate customer instances.

In the following figure, a duplexed set of PGs serve both an IVR system and an ACD system. These PGs would be equipped with both ACD and IVR interface software.

Note: The IVR can also be on a System IPCC PG or a IPCC Generic PG.

The IVR Peripheral Gateway can run in simplexed or duplexed configurations. In a duplexed configuration, only one side of the PG has an active connection to the IVR at a time.

Note: When multiple IVRs are connected to a PG, IVRs that use poll-based monitoring may not be mixed with IVRs using any other kind of monitoring.

Figure 24: IVR-to-PG Interface

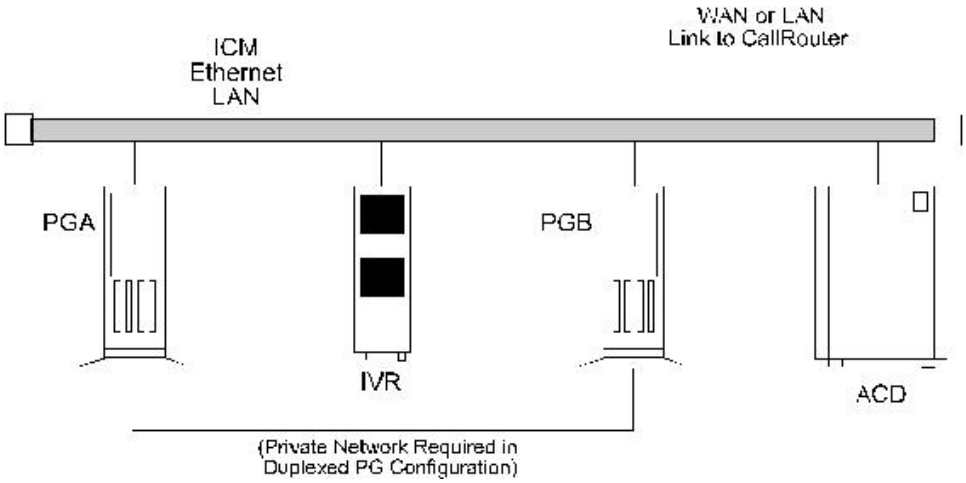

For information on how IVR systems fit into the Unified ICM data communications networks, see Chapter 10, “Determining the Datacom Requirements” (page 73).

Chapter 8

ICM Application Gateway and ICM Gateway SQL Planning

The Unified ICM Application Gateway and Unified ICM Gateway SQL options allow Unified ICM to integrate external contact center applications into the Unified ICME . Each of these options involves some pre-installation planning. For example, you may need to prepare host systems and databases; review fault tolerance issues; and, in the case of Unified ICM Gateway SQL, plan for data transfer.

This chapter contains the following topics:

ICM Application Gateway Planning

The Unified ICM Application Gateway option allows the Unified ICM system to interface with any external call center application. Within the Unified ICM software, the Unified ICM Application Gateway feature is implemented as a node in a call routing script. You add a Gateway node to a script to instruct the Unified ICM to execute an external application. This allows the script to evaluate responses from an external application. The Unified ICM can then base subsequent routing decisions on the results produced by the application.

A typical Unified ICM Application Gateway application might return a variable to the CallRouter that identifies the caller as having a certain type of account. The script could then use this information to control where and how the call is routed. Optionally, the Unified ICM software can pass the retrieved information to the site that is receiving the call. In this case, certain data such as account numbers, dates, billing phone numbers, and addresses are passed along with the call to an answering resource.

Preparing the Host System

To prepare for the Unified ICM Application Gateway option, you must set up the host system to communicate with the Unified ICM system. This involves configuring the host application to listen to a socket on the target Unified ICM machine. You also need to configure a name and port number to be used to connect the host system to the Unified ICM central database. These steps are performed at system installation. However, you can begin preparing the host applications ahead of time.

During system installation, when connectivity between the Unified ICM system and host system is established, you need to identify the host system to be queried by entering data in the Application_Gateway table.

Fault Tolerance

You can configure access to a single host application or duplicate host applications. In a single host configuration, configure the same host for both CallRouters (Side A and Side B). The single host method provides no protection against host failures; however, it does protect against connection failures.

In order to achieve a higher level of fault tolerance in an Unified ICM Application Gateway application, you can connect duplicate host applications to the CallRouter. For example, the Side A and Side B CallRouters can each manage a connection to one of the duplicated host applications. Each time a script initiates a request, both CallRouters query their corresponding host. The CallRouters use the response from the host that responds first. This method is highly reliable. Even if a host or a connection fails, all query requests are satisfied.

ICM Gateway SQL Planning

The Unified ICM Gateway SQL option allows the CallRouter to query an external SQL Server database and use the data in call routing.

If you are going to use the Gateway SQL option, you need to review several pre-installation planning issues:

- First, Unified ICM Gateway SQL requires an additional Database Server hardware platform.
- You also need to be aware of the tasks involved in setting up the external host database and populating it with the data you want to use in call routing.

Database Server Platform

The Unified ICM Gateway SQL option requires a host database server. The host database server can be duplexed in order to maintain Unified ICM fault tolerance. A duplexed Unified ICM Gateway SQL system requires two identical host database server platforms. Each host database

server resides on the same LAN segment as its corresponding Unified ICM CallRouter. The following figure shows a duplexed Unified ICM system that has a duplexed Unified ICM Gateway SQL host database server.

Figure 25: ICM Gateway SQL Duplexed Configuration

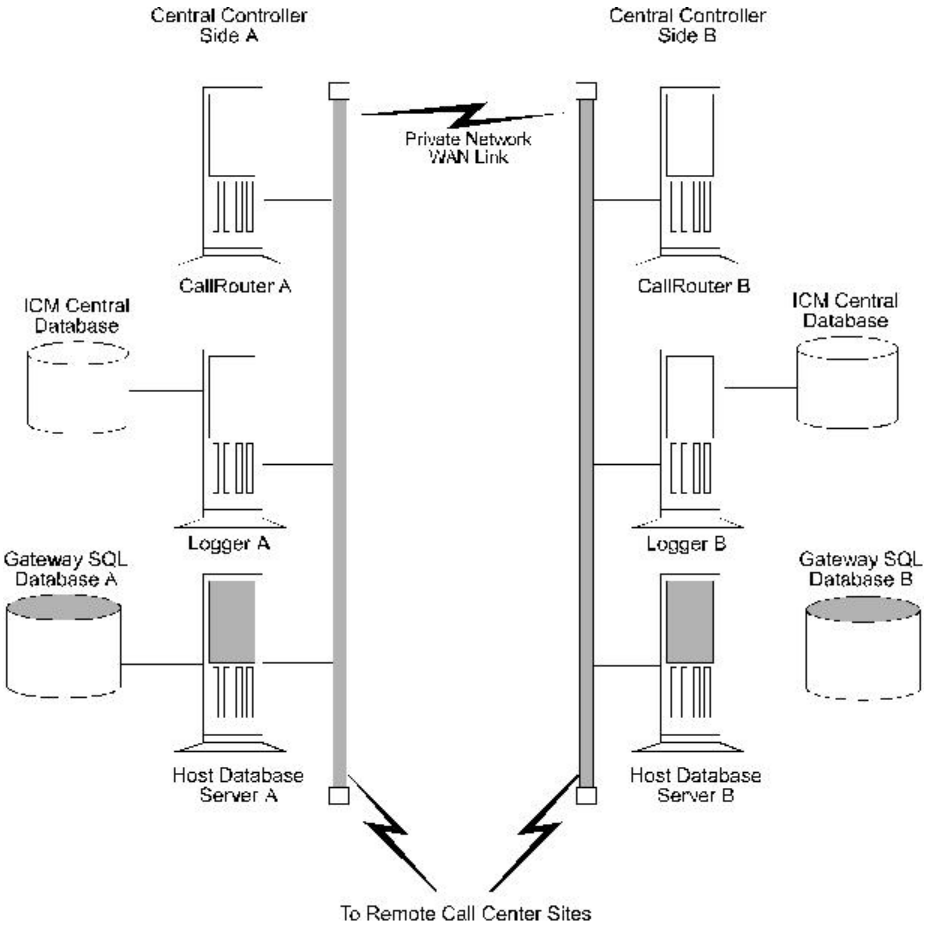

Planning for Data Transfer

To prepare an Unified ICM system for Unified ICM Gateway SQL, you need to make several decisions:

Decide which data you want to use in the external database. For example, will you be using:

- Customer records?
- Account information?
- Other types of data?

Decide where the data is coming from:

- Another database?

- A flat file?
- Other sources?

Make a plan to transfer the data to the external database:

- What type of media will you use to transfer the data (tape, disk, network)?
- Will the transferred data be in a certain format (comma-separated values, text file, SQL Server syntax)?

Configuration Overview

For Unified ICM Gateway SQL, you must set up and configure one or more host databases to function with the Unified ICM system.

Configuration considerations:

- Choosing a host database server platform. The host database server must have adequate processing power and disk space. Cisco can provide you with specifications for basic and high-end host database server platforms.
- Setting up the host database. This includes:

  Installing SQL Server

  – Creating a database on the host database server platform

  – Defining fields and indexes

  – Setting up permissions and replication
- Transferring data from a data source. This task involves transferring data to populate the database with the data to be used in call routing (for example, you might want to transfer customer records to the database).
- Configuring the Unified ICM system to access the host database. This task involves setting up user names and passwords that the Unified ICM system can use to access the data in the host database.
- Writing test scripts to test the Unified ICM Gateway SQL option. This task involves monitoring test scripts that use the Script Editor DB Lookup node. The monitoring results are captured and stored in the Route_Call_Detail table to validate that the Unified ICM Gateway SQL feature is functioning.

Chapter 9

Planning for ICM Platforms

Once you have the system sizing recommendations, you can begin to order the appropriate hardware configuration. First, however, you must determine how many Unified ICM nodes you will need.

The number of servers required in an Unified ICM system depends on the configuration of the central controller, PGs, NICs, and other nodes. For example, a duplexed central controller configuration requires additional servers because the CallRouter and Logger are duplicated.

This chapter contains the following topics:

Determining the Number of Servers Required

The table titled **Sample Server Requirements** shows how to determine the number of servers required in your system.

The counts of servers in this example are based on an Unified ICM configuration that has the following characteristics:

- The Unified ICM system has a duplexed, geographically distributed central controller (in other words, each central site has a CallRouter and a Logger).
- One side of the central controller (Central Site 1) is located at a call center and consequently has a PG to serve one or more ACDs. The PG is duplexed (two servers) for fault tolerance.
- This Unified ICM installation has three remote call center sites and two Admin sites.

Table 3: Sample Server Requirements

Sites	**Node Types**					
	CallRtr	Lgr	Call/Lgr	DB Server[1]	PG[2]	Administration & Data Server with HDS
Central Site 1	1	1	-	-	2	1
Central Site 2[3]	1	1	-	-	-	1
Remote Call Center 1	-----	-----	-----	-----	2	-
Remote Call Center 2	-----	-----	-----	-----	2	-
Remote Call Center 3	-----	-----	-----	-----	2	-
Admin Site 1	-----	-----	-----	-----	-----	1
Admin Site 2	-----	-----	-----	-----	-----	1
Total Nodes:	2	2			8	4
Key:		--------- These servers are not installed at this type of site. – Not selected as an option in this particular configuration.				

ICM Platform Considerations

Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1) contains information on server configurations and provides examples of supported server platforms.

Processor Utilization

As a general rule for all Unified ICM nodes, processor utilization should be kept below 60 percent at the maximum expected call load on the system. This is needed in order to smooth out call request “spikes” as well as to allow enough reserve capacity to perform activities such as re-synchronization and background cleanup. Non-ICM software can make up a part of the 60 percent maximum load. The processor utilization figure (60 percent) covers all software running on the platform.

In addition to the utilization requirement, it is necessary that no software on the system run at a priority equal to or higher than the Unified ICM software for more than 100 milliseconds in uninterrupted bursts. In other words, the Unified ICM software needs to run on the system at least as frequently as once every 100 milliseconds. This is usually not a problem unless device

2) Only installed at the central site if that site also serves as a call center or you are using the remote ACD option

1) Required only in ICM Gateway SQL configurations.

3) A second central site is not required in duplexed-collocated Central Controller configurations.

drivers or other kernel-level software is installed, or process/thread priorities have changed incorrectly.

Paging Requirements

The most time-critical component of the Unified ICM system, the CallRouter node, must not be delayed due to disk I/O (that is, paging). The only disk I/O that should be occurring on Unified ICM machines is for log file writes and database I/O. The database I/Os occur on Logger and Administration & Data Server machines. The simple rule is to provide enough main memory so that the entire working sets of critical processes remain in memory.

For complete and current information about RAM and other platform requirements, see the Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1).

The database platforms (Loggers, Administration & Data Servers, and Unified ICM Gateway SQL machines) should have enough main memory so that all first level index pages are kept in main memory cache.

Logger Expansion

The Logger platform you order may include a combination of internal and external SCSI hard drives. As your call center enterprise grows, your database requirements will typically grow as well. You might have more services, skill groups, and routes in your configuration, and you might be routing more calls each day. This will result in more historical data being stored in the central database.

When your database requirements change, contact your Unified ICM software support provider to have the storage capacity of the central database increased.

Note: Refer to Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1) for more information on data storage specifications.

They can allocate more database space after your system is installed by:

- Remotely adding database space (if current disk space allows).
- Installing “hot-plugable” disk drives and configuring the disks while the system is running.

Note: *Administration Guide for Cisco Unified ICM Enterprise* provides information for managing database space once the Unified ICM system is installed and running.

Planning for Administration & Data Servers

To allow users to monitor current call center activity, the Unified ICM system forwards real-time data to Administration & Data Servers at selected sites throughout the call center enterprise. The following figure illustrates the real-time architecture of the Unified ICM system.

Figure 26: Real-Time Data Architecture

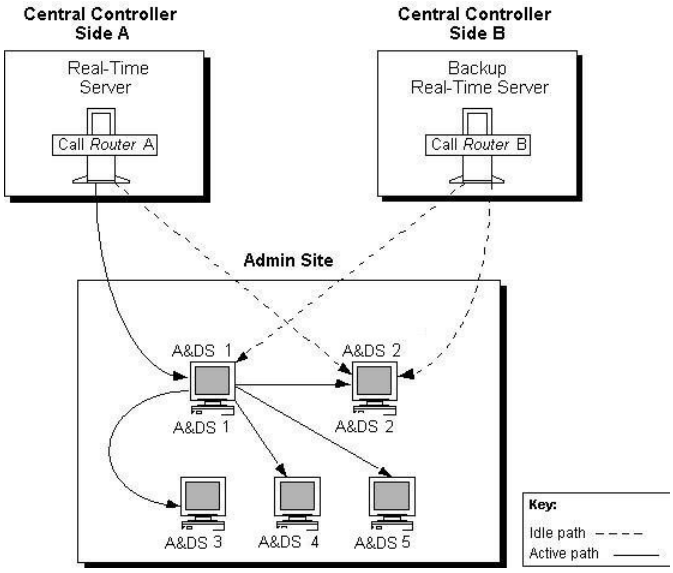

Real-time call and agent group status data arrives at the central controller from the Peripheral Gateways, which are constantly monitoring activity at each call center. The CallRouter acts as the real-time server. The CallRouter for the other side of the central controller acts as a back-up real-time server.

The CallRouter is responsible for providing real-time data to one or more Administration & Data Servers at each administrator site. Administration Clients at the site receive their real-time data through a connection to a Administration & Data Server. Administration Clients do not have the local database and Administration & Data Server processes are required to receive real-time data directly from the CallRouter.

Administration & Data Servers and Admin Sites

Administration & Data Servers can be located with one or both sides of the central controller, at a call center, or at another site. Any site that contains Administration & Data Servers is referred to as an administrator site. Each administrator site requires at least one Administration & Data Server. Two Administration & Data Servers should be used (as shown in Figure 9-1) (page 70) to provide fault tolerance in the real-time data distribution architecture.

The primary Administration & Data Server maintains an active connection to the real-time server through which it receives real-time data. The secondary Administration & Data Server also maintains connections to the real-time server; however, these connections remain idle until needed (for example, in cases where the primary Administration & Data Server is unavailable for some reason). In sites that have two Administration & Data Servers, the Administration Clients are configured to automatically switch to a secondary Administration & Data Server if the first distributor becomes non-functional for any reason.

Administration & Data Servers and Administration Client Requirements

There is no set limit to how many Administration Clients can be served by a Administration & Data Server. Refer to the Hardware & System Software Specification (Bill of Materials) for

Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1) for information about requirements for Distributor and Administration Clients.

Planning for Historical Data Servers

Historical data is stored both as individual call detail records and also rolled up and stored as interval records. A Administration & Data Server with a Historical Data Server (HDS) stores historical data that supports reporting queries. Administration & Data Servers at the site query historical data from the HDS rather than directly from the Logger.

Figure 27: Historical Data Server Architecture

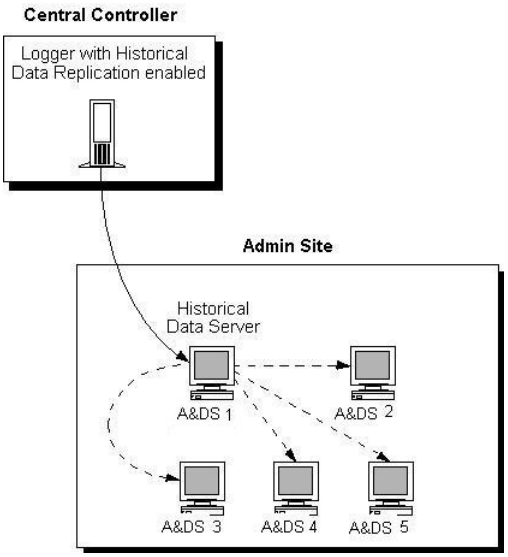

To set up a Historical Data Server, you must configure the Logger to perform historical data replication. You must also configure the real-time Administration & Data Server to be an HDS. You can then create an HDS database on the real-time distributor.

Information in the real-time feed tells each Administration Client where to obtain historical data. If the real-time distributor is a Historical Data Server, then it instructs its clients to get historical data from it. Otherwise, it instructs its clients to get historical data from the Logger.

Each Logger can support up to two HDS's. These servers can be configured either as two primary distributors or with one as a secondary distributor. If these systems are needed to support reporting requirements, review the "Primary/Secondary Administration & Data Server Deployment" section of the WebView Installation and Administration Guide to understand all of the factors for consideration before deciding the best deployment strategy for your organization. The same fault-tolerant strategy that applies to the real-time Administration & Data Server also applies to the HDS. That is, when the primary HDS fails, other Administration Clients at the site automatically switch over to use the backup HDS.

HDS Features

The HDS eliminates the performance impact on the central database caused when multiple Administration & Data Servers need to access the central database to generate reports. In systems

that have multiple remote Administration & Data Servers, the HDS brings Unified ICM historical reporting data closer to the end user. Each HDS provides a set of database tables. You can set specific times for retaining data in these tables. These capabilities give you flexibility in setting up reporting capabilities on a site-by-site basis.

The Historical Data Server also provides:

- Greater flexibility in leveraging Internet applications.
- An open interface for data mining and data warehousing applications.
- The ability to host other database tables and have them work with the HDS.
- Improved security and data access capabilities.

  The HDS Administration & Data Server requires a high-end platform with a more powerful CPU, greater disk capacity, and more RAM. For complete and current information on HDS requirements, see the Hardware & System Software Specification (Bill of Materials) for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.0(1).

Chapter 10

Determining the Datacom Requirements

The Unified ICM system needs highly reliable networks to ensure sufficient real-time responsiveness and fault tolerance. Because the Unified ICM system is a mission-critical, fault-tolerant system, it must be able to respond quickly when a node goes offline for any reason. Depending on the situation, the Unified ICM system might need to switch communication paths or activate other nodes to keep the system running without interruption.

In addition to responding to node failures, the Unified ICM system needs to perform diagnostics on failed nodes so they can be returned to service as soon as possible. Often, the Unified ICM diagnostic procedures take place over a Wide Area Network (WAN).

The Unified ICM system must also be able to respond to route requests from the Interexchange Carriers (IXCs) within a certain minimum time-out period. For example, the AT&T intelligent call processing network requires a response from the Unified ICM system within 200 milliseconds of receiving a route request. In a geographically distributed Unified ICM configuration, this means that the Unified ICM system must perform communications between the NICs and CallRouters on both sides of the central controller and return a route response all within the 200 millisecond time-out period.

This chapter helps you to prepare network facilities for an Unified ICM system installation. In this chapter, complete the following tasks:

- **Determine requirements for visible and private networking.** The Unified ICM networks must meet certain minimum bandwidth and latency requirements.
- **Allocate IP addresses.** Assess the IP address requirements for Unified ICM nodes at each site in the system.
- **Fill out IP address worksheets.** Use the worksheets in Chapter 12, “IP Address Worksheets” (page 111)to assign IP addresses.
- **Order any additional network hardware.** To prepare the network facilities, you may need to order routers, bridges, or cabling.

This chapter also covers some of the options for configuring the Unified ICM networks and integrating them with your existing networks.

This chapter contains the following topics:

ICM Sites

The Unified ICM system consists of a number of computers, or nodes, which are typically located at more than one site. An Unified ICM system can be distributed among anywhere from three to fifty sites or more. Each site might contain one or more nodes. The Unified ICM system requires several networks to interconnect nodes within and among the sites.

Unified ICM sites are of three basic types:

- **Central sites.** Contain one or both sides of the central controller (that is, the CallRouter and Logger) and possibly a separate Network Interface Controller. Central sites can also contain Administration & Data Servers and Peripheral Gateways.

- **Contact center sites.** Contain one or more Peripheral Gateways (PGs) and possibly Administration & Data Servers. Sites also support Agents, phone applications and CTI applications.

- **Admin sites.** Contain one or more Administration & Data Servers.

An Unified ICM site might be a combination of any two or more of these. For example, a single location might be both a central site and a contact center site.

The ICM Networks

The Unified ICM system uses three independent communications networks

Following are the ill three independent communications networks:

- **Private network.** This is a dedicated network that allows specific nodes to communicate with each other without outside interference. This network carries the data that is necessary to maintain and restore synchronization between the systems. The private network is not to be used for any other purpose.

- **Visible network.** This is a shared network that allows the central controller to communicate with local and remote nodes. It carries traffic between each side of the synchronized system

and foreign systems. The visible network may also be used by the fault tolerance software as an alternate network to distinguish between node failures and network failures.

- **Signaling Access Network.** This network connects the Unified ICM system to a carrier network or client network. When a SAN is implemented, the Unified ICM system uses the SAN (not the private network) to communicate with the carrier network.

The following figure shows the two sides of the central controller, a contact center site, and an administrator site. A private WAN links both sides of the duplexed central controller. A visible WAN links the contact center and administrator sites to each side of the central controller. Nodes within each site are linked by a local area network (LAN).

Figure 28: ICM System Network Overview

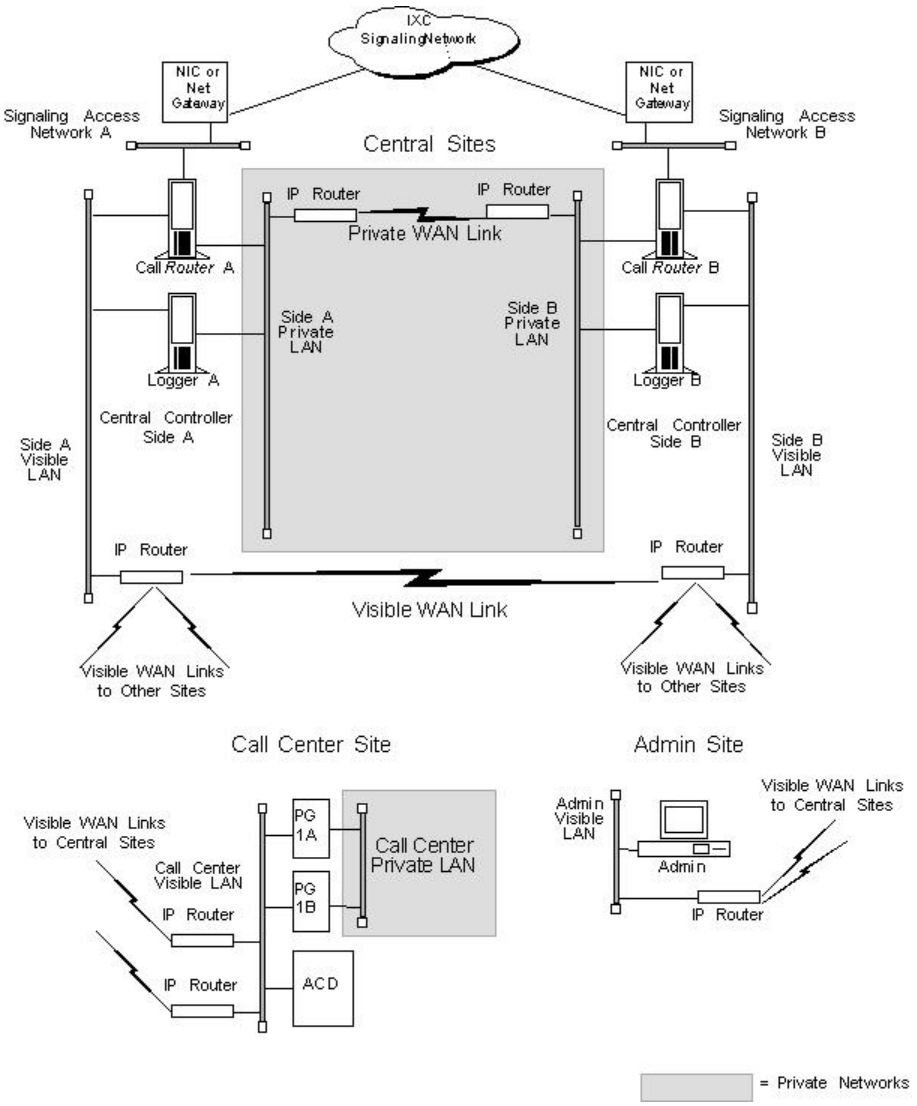

In the preceding figure, the two sides of the central controller are geographically separated. The wide area network connections in both the private and visible networks are referred to as WAN links. WAN links in the Unified ICM system are typically high-availability provisioned circuits. These links must possess extremely low and extremely predictable latency characteristics.

Therefore, some types of WAN service cannot be used for WAN links within the Unified ICM system (for example, packet routing).

Private and Visible WAN Links

The two sides of the duplexed Unified ICM central controller share a single private network and are linked via a private WAN link. They also share a visible network which connects the two sides via a visible WAN link. To ensure a high level of fault tolerance, the private WAN link and visible WAN links must be independent (that is, they must use different trunks and possibly even different service providers).

When the two sides of the central controller are co-located, the visible WAN link between the sites is not needed. The standard visible WAN links to remote contact center sites provide adequate connectivity between the two sides. In a co-located central controller configuration, the private network is implemented locally by using Ethernet switches.

Remote contact centers connect to each side of the central controller via the visible network. Each visible WAN link to a contact center must have adequate bandwidth to support PGs and Administration & Data Servers at the contact center (the bandwidth requirement varies greatly as the configuration changes, that is, the call load, the number of agents, and so on).

When a contact center is co-located with a side of the central controller, the PGs and Administration & Data Servers connect to the visible LAN on that side. The PGs and Administration & Data Servers connect to the other side of the central controller via a visible WAN link. In such a configuration, a direct visible WAN link between the sides of the central controller is required to ensure adequate connectivity between the two sides. LAN bridges may optionally be deployed to isolate PGs from the Administration & Data Server LAN segment and to enhance protection against LAN outages.

Note: See the section titled Central Sites (page 89), for some examples of co-located central controller configurations.

Signaling Access Networking

The CallRouter machine connects to the IXC signaling network via the Signaling Access Network (SAN). A separate LAN interface card in the CallRouter is dedicated for use just by the SAN. The SAN connects the NICs on each side of the duplexed system to the IXC signaling network. In most cases, the NIC software runs on the CallRouter computer. For clarity, in Figure 10-1 (page 75), the NIC is shown as a separate computer installed on the SAN.

A node called the Unified ICM Network Gateway may also be installed on the SAN and used to interface to some SS7-based networks. The Unified ICM Network Gateway is a dedicated machine that provides SS7 protocol handling services.

In Sigtran SS7 networks, Sigtran Gateways may be co-located on the CallRouter machine or on a separate machine. Sigtran Gateways connect to the Sigtran SS7 network using the SCCP User Adaptation layer (SUA) with Stream Control Transmission Protocol (SCTP) as the network transport. Sigtran Gateways can serve the following functions, depending on customer requirements:

- Communicate with the Service Switching Point (SSP)
- Communicate with the Media Gateway Controller
- Communicate directly to a Signaling Gateway. In this deployment Sigtran connections are established using a Client / Server message exchange, in which the Sigtran Client requests connections with the Sigtran Server. The Signaling Gateway (such as Cisco's Internet Transfer Point) is a server in this model and accepts incoming connection requests. The Sigtran Gateways act as the Client when connected to a Signaling Gateway.

Local Area Networks

The Unified ICM system uses Ethernet for local area network connectivity. The particular Ethernet topology used is immaterial from an architectural standpoint. However, the topology used may be relevant from a network or systems management perspective. Typically, UTP is used in the private, visible, and signaling access LANs.

The three networks (visible, private, and signaling) should be on separate LAN segments. This requires the use of three Ethernet cards in the CallRouter machine.

Network Bandwidth Requirements

The visible network bandwidth requirements for a typical Unified ICM system are about 1,000 bytes of data per call over the networks that carry call data. For example, if a remote PG is managing 15 calls per second at a contact center site, it needs to transfer 15,000 bytes of data over the visible WAN to the central site every second (a total of 120,000 bits per second, ignoring packet overhead).

The bandwidth for the private WAN between the two sides of a duplexed central controller must support the total sustained call load for all ACD sites. In addition, bandwidth on this private WAN must provide some degree of burst resilience and enough reserve capacity to perform fault tolerant messaging and synchronization. The table titled **Network Circuit Requirements** summarizes the network circuit requirements for visible and private networks within the Unified ICM system.

Table 4: Network Circuit Requirements

Network	Purpose	Facilities	Min. Bandwidth
Private WAN	Dedicated path that connects both sides of a duplexed, distributed ICM central controller.	T1	T1 dedicated
Visible WAN	Circuits that connect PGs and Administration & Data Servers at remote sites to	Typically, a T1 or a fractional T1.	128-Kbps dedicated.[4]

4) Variable, depending on load. See the section Calculating QoS Bandwidth Requirements, page 11-11, for a means of calculating the minimum required bandwidth for a Quality of Service (QoS) compliant network.

Network	Purpose	Facilities	Min. Bandwidth
	each side of the ICM central controller.		
Signaling Access Network	Local area network that connects the NIC to the IXC carrier network or client network.[5]	Ethernet Unshielded Twisted Pair (UTP)	100 Mbps
Visible and private LANs	Local area networks that connect ICM nodes at a central site and PGs and Administration & Data Servers at remote contact center sites.	Ethernet Unshielded Twisted Pair (UTP). Cisco requires using manageable hubs.	100 Mbps

You may require additional bandwidth on the visible WAN. The actual requirement depends on a number of factors, including call load, the number of ACDs, the number of agents, and the number of Admin sites.

Note: If your network will be utilizing the Cisco Unified ICM Quality of Service (QoS) feature, see also the section Cisco ICM QoS (page 81), for additional latency considerations.

Network Latency Requirements

The Unified ICM system is a real-time, fault-tolerant distributed system.

To guarantee the real-time nature of the system and to support the methods used in fault tolerance, the WAN links in the Unified ICM system must have extremely low and predictable message latency characteristics, especially in these critical areas:

- Route requests and route responses between the CallRouter/NIC and IXC. This communication must meet the strict message latency requirements of the carrier networks.
- Communications involving post-routing requests from PGs and route responses from the CallRouter. This communication must also be extremely fast since callers are online expecting the call to be answered by an appropriate agent.
- Communications from the PGs to the CallRouter concerning the real-time status of the contact center. The CallRouter needs this information to base its routing decisions on the latest data available from the contact center.

Three fault tolerance mechanisms of the Unified ICM system require reliable, low latency communications. These are heartbeat detection, synchronization, and state transfer.

Note: If your network will be utilizing the Cisco Unified ICM Quality of Service (QoS) feature, see also the section Cisco ICM QoS (page 81), for additional latency considerations.

5) For the Sprint NIC, the local Ethernet Signaling Access Network is not implemented. Instead, X.25 WAN cards in the CallRouter platform serve as the Signaling Access Network and allow the CallRouter-NIC machine to connect to the IXC signaling network.

Heartbeat Detection

As part of its fault-tolerant design, the Unified ICM system must be able to quickly respond when a component goes offline for any reason (typically, because of a failure in the node or in a network link). Each critical component in the system periodically sends short messages through the network to indicate that it is still online. These messages are called heartbeats.

Communicating Unified ICM components send heartbeats to each other at regular intervals. If a component fails to receive a heartbeat for five consecutive intervals, it assumes that the component or a network link has failed and it initiates recovery actions. The table titled **Heartbeat Configuration** lists some of the nodes that send heartbeats, the network on which they are sent, and how often they are sent.

Table 5: Heartbeat Configuration

Node	Medium	Interval
AT&T NIC (or Network Gateway) to CallRouter	Signaling Access Network	200 milliseconds
CallRouter to CallRouter	Private network	100 milliseconds
PG to CallRouter	Visible network	400 milliseconds
PG to PG (if duplexed)	PG to PG (if duplexed) Private network	100 milliseconds

The two sides of a duplexed Unified ICM central controller periodically test each other to see if the other side is operating correctly. As shown in the table titled **Heartbeat Configuration**, network latency from CallRouter-to-CallRouter over the private network must support round trip messaging of 100 milliseconds. If the bandwidth of the private network is not adequate, packets may need to be fragmented by IP routers in order to prevent long messages (greater than 1,500 bytes). Such long messages can delay transmission of User Diagram Protocol (UDP) packets, which indicate that the other side of the central controller is still operating.

Note: A consistent heartbeat or keep-alive mechanism is enforced for both the public and private network interface. When QoS is enabled on the network interface, a TCP keep-alive message is sent; otherwise UDP heartbeats are retained.

Another requirement of fault tolerance is that messages cannot be released back to a NIC or PG until the other side of the central controller has acknowledged receipt of a copy of the message. Therefore, in order to meet the 200 millisecond response times established by the carrier networks, and to leave some margin for queuing, a 100 millisecond round trip requirement is established.

Heartbeats from a remote PG to the CallRouter must compete with other network traffic on the visible WAN.

Synchronization

In a duplexed central controller configuration, the private network allows the CallRouters and Loggers on each side to run in a synchronized fashion. This means that the CallRouter and

Logger processes on each side of the system receive the same input and generate the same output.

To ensure synchronization, each message intended for the CallRouter or Logger is received by a Synchronizer process that runs on the CallRouter node. The Synchronizer forwards the message across the private network to the Synchronizer on the other side. The Synchronizers then remove any duplicates before passing the messages on to the CallRouter processes. If a message is intended for the Logger, the CallRouter passes it along.

Figure 29: Role of Synchronizers

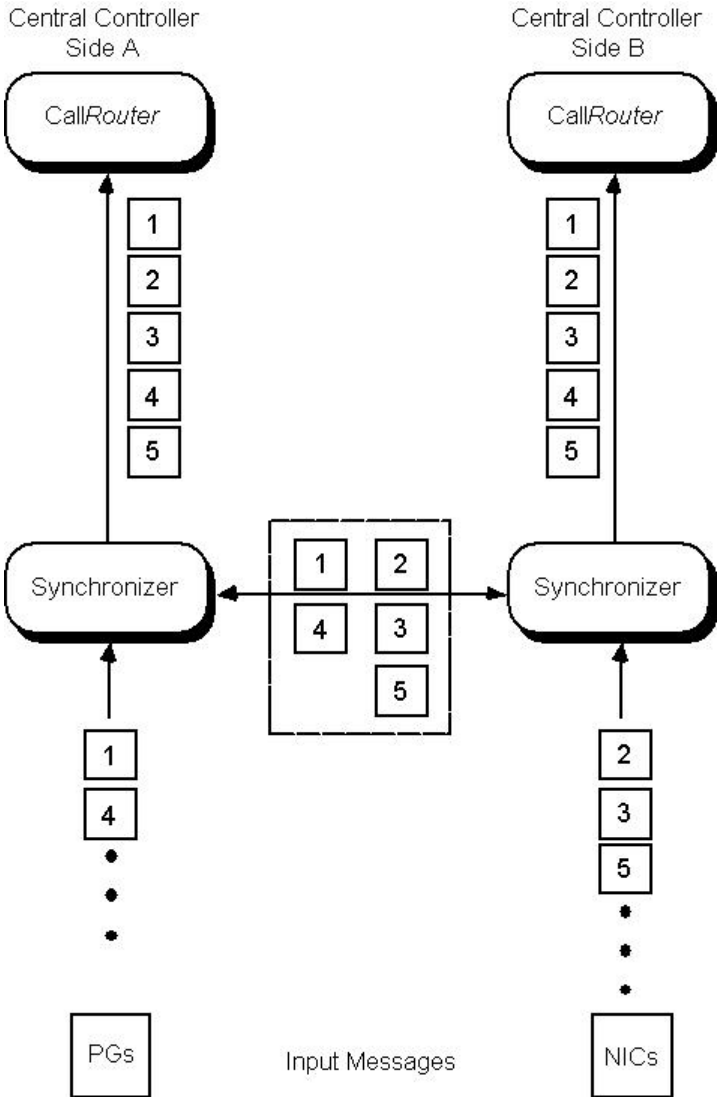

The preceding figure shows how the Synchronizers combine input messages and send the messages in the same order to each side of the central controller. Both CallRouters receive the same input and generate the same output. The Synchronizers ensure that both sides of the central controller return identical destinations for the same call and write identical data to the databases.

State Transfer

The fault tolerance of the Unified ICM system enables nodes to restart after a failure. However, when a failed node restarts, the values of variables in its memory are out-of-date. Before returning the node to service, the Unified ICM system must copy the values from its peer on the other side to the recovering node. That is, it must transfer the state of the running machine to the recovering machine. This transfer takes place over the private network.

Note that such state transfers occur after the failure and restart of any synchronized MDS client: PG, Logger, CallRouter, and so on.

Diverse Facilities

The private WAN between central controllers (when the central controllers are geographically separated) and the visible WAN must be on separate facilities. They must use different circuits and different IP routers. As added protection, you might also want to use diverse routes or even different service providers for the private and visible WAN links. Otherwise, you run the risk of having a single network failure disable both the Unified ICM private and visible WANs.

For example, if the private WAN failed, or a visible WAN link to one side of the central controller failed, the Unified ICM system would continue to route calls and function normally. However, if the private WAN and the visible WAN were on the same facilities and failed simultaneously, the fault tolerance of the system would be compromised. In such a scenario, the failure of any single node on either side of the central controller would interrupt system processing. By provisioning the private WAN and visible WAN on separate facilities, you can eliminate this potential point of failure.

Cisco ICM Quality Of Service (QoS)

This section describes the Cisco Unified ICM Quality of Service (QoS) feature, and discusses considerations to take into account when planning for and deploying Unified ICM networks that will utilize QoS.

What Is Quality of Service?

Quality of Service is a set of capabilities that enables you to define a level of performance in a data communications network. QoS allows you to create differentiated services for network traffic, thereby providing better service for selected network traffic. For example, with QoS, you can increase bandwidth for critical traffic, limit bandwidth for non-critical traffic, and provide consistent network response. This enables you to use expensive network connections more efficiently, lets you establish service level agreements with customers of the network, and eliminates the need of having dedicated leased lines for connection with Unified ICM components.

QoS capabilities enable Unified ICM software to overcome the following architectural limitations:

- Unified ICM software requires dedicated leased lines. This means that Unified ICM cannot be deployed in a converged network, which is more cost-effective and has more transmission capacity.

- Lack of a congestion control mechanism over the LAN segment. This is often considered not important because LAN resources tend to be less costly than WAN resources. However, with the increasing usage of multimedia applications on LANs, delays through LAN switches do become problematic. The QoS technology 802.1p tackles these delays.

- Lack of support for Cisco's AVVID (Architecture for Voice, Video and Integrated Data) enterprise network architecture. AVVID defines the network design principles to optimize the integration of mission-critical applications in a convergent network environment. QoS is a key technology for AVVID. Unified ICM should be AVVID compliant to be better deployed in a Cisco AVVID network.

- Problematic UDP heartbeats. The use of UDP heartbeats creates unnecessary complexity for Unified ICM deployment in firewall and NAT (Network Address Translation) environments. For this reason, UDP heartbeats are replaced by TCP keep-alive messages in Unified ICM QoS implementation.

To implement QoS, you define QoS policies on network devices (routers and switches), and apply the policies to traffic based on DSCP markings, IP precedence, IP address, port, and so on.

QoS primarily comes into play when the amount of traffic through an interface is greater than the interface's bandwidth. When the traffic through an interface exceeds the bandwidth, packets form one or more queues from which the device selects the next packet to send. By setting the queuing property on a device or interface, you can control how the queues are serviced, thus determining the priority of the traffic.

Unified ICM supports DSCP and 802.1p markings for both the public network link (connecting the PG to the CC) and the private network link (connecting the duplexed sides of PG or CC).

Deploying Cisco ICM QoS

The process of deploying and implementing QoS is a combined effort supported by Cisco System Engineers, Unified ICM Deployment Groups, and Cisco Partners.

These Cisco representatives provide customers who plan to deploy QoS with the following assistance:

- Defining customer requirements. Cisco Professional Services and Cisco Partners utilize their historical knowledge of the customer's Unified ICM deployment, and QoS bandwidth calculation tools (see the section Calculating QoS Bandwidth Requirements (page 84)), to assess these requirements.

- Reviewing the Unified ICM portion of the customer's QoS migration plan.

- Meeting with the customer to prepare a statement of work that defines the level of support Cisco will provide.

Meeting with the customer to prepare a statement of work that defines the level of support Cisco will provide.

Alongside these steps, there are the following tasks to consider when planning to implement a QoS-compliant network in your Unified ICM environment.

- Where to mark traffic
- Determining QoS markings.
- Projecting bandwidth requirements.
- Installing Microsoft Packet Scheduler (optional).
- Installing and configuring 802.1p-capable network components (optional).
- Configuring QoS on IP routers.

Where to Mark Traffic

In planning QoS, a question often arises about where to mark traffic, in the application or at the network edge. Marking traffic in the application saves the access lists for classifying traffic in IP routers/switches, and it may be the only option if traffic flows can not be differentiated by IP address, port and/or other TCP/IP header fields. As mentioned earlier, Unified ICM currently supports DSCP markings on the visible network connection between the central controller and the PG, as well as on the private network connection between duplexed sides of the Router or PG. Additionally, when deployed with Windows Packet Scheduler, it offers shaping and 802.1p.

Traffic can also be marked or remarked in edge IP routers/switches if it is not marked at Unified ICM servers, or if the QoS trust is disabled. QoS trust may be disabled in an attempt to prevent nonpriority users in the network from falsely setting the DSCP or 802.1p values of their packets to inflated levels so that they receive priority service. For classification criteria definitions on edge routers and switches, refer to the tables titled **Public Network Traffic Marking (default) and Latency Requirements** and **Private Network Traffic Marking (default) and Latency Requirements** in the next section.

Determining QoS Markings

The default Unified ICM QoS markings are set in compliance with Cisco AVVID recommendations (but can be overwritten if necessary). See Cisco AVVID Solution IP Telephony QoS Classification for details about Cisco AVVID packet classifications.

Before QoS implementation, IP-based prioritization is used to provide two externally visible priority levels: high and non-high. Internally, however, there are three different priorities for application messages: high, medium, and low. In the public network, medium priority messages are sent through the same high IP connection as the high priority messages; yet in the private network, they are sent through the non-high IP connection.

The tables titled **Public Network Traffic Marking (default) and Latency Requirements** and **Private Network Traffic Marking (default) and Latency Requirements** list the IP address and port, latency requirement, default marking under each priority for the public network connection and the private network connection respectively.

Table 6: Public Network Traffic Marking (default) and Latency Requirements

Priority	IP Address and Port	Latency Requirement	DSCP / 802.1p Marking
High	Public high IP and high priority connection port	200 ms	AF31 / 3
Medium	Public high IP and medium priority connection port	1,000 ms	AF31 / 3
Low	Public non-high IP and low priority connection port	5 seconds	AF11 / 1

Table 7: Private Network Traffic Marking (default) and Latency Requirements

Priority	IP Address and Port	Latency Requirement	DSCP / 802.1p Marking
High	Private high IP and high priority connection port	100 ms (50ms preferred)	AF31 / 3
Medium	Private non-high IP and medium priority connection port	1,000 ms	AF11 / 1
Low	Private non-high IP and low priority connection port	1,000 ms	AF11 / 1

Note:

- Microsoft Packet Scheduler supports at most two marking levels excluding best effort, and therefore the medium priority traffic is either marked same as the high priority traffic (in public network) or marked same as the low priority (in private network). This is similar to the IP-based prioritization approach and no priority level is lost from the network perspective. When Packet Scheduler is bypassed, however, three marking levels are available and the medium priority messages can be marked differently.

- Cisco makes the QoS marking recommendation for Call-Signaling traffic to DSCP CS3 because Class-Selector code points, defined in RFC 2474, are not subject to markdown and aggressive dropping as Assured Forwarding Per-Hop Behaviors are. Some Cisco IP Telephony products already have begun transitioning to DSCP CS3 for Call-Signaling marking. During this interim period, both code points (CS3 and AF31) should be reserved for Call-Signaling marking until the transition is complete. The Unified ICM QoS markings are configurable through the Peripheral Gateway or Web setup and the default Assured Forwarding code points can be replaced with the Class-Selector code points to fit into the existing QoS infrastructure.

Calculating QoS Bandwidth Requirements

Although QoS alleviates bandwidth usage and increases network throughput, network congestion is still unavoidable unless sufficient physical bandwidth is available along the path. For Unified ICM , the bandwidth requirement at each priority is a function of traffic volume and latency

requirement. It varies greatly for Unified ICM systems depending on factors such as call load, traffic composition, call context information, and configuration settings.

Cisco provides the following QoS bandwidth calculators and sizing worksheets to help Cisco System Engineers, Unified ICM Deployment Groups, and Cisco Partners project traffic volume as well as bandwidth requirement.

- ACD/CallManager PG to CC Bandwidth Calculator.
- VRU PG to CC Bandwidth Calculator.
- Router Private Link Sizing Worksheet.
- PG Private Link Sizing Worksheet.

Note:

- The network administrator should clearly understand the bandwidth requirement of Unified ICM flows under each priority and factor it in the bandwidth definition of QoS policies configured on network routers/switches.
- Unified ICM applications are not RSVP (Resource Reservation Protocol) aware and therefore IntServ (Integrated Service) is not supported. If Packet Scheduler is used, the QoS bandwidth reservations are only made within the local box for the purpose of shaping; no reservations are made in the network.

Installing Microsoft Packet Scheduler

Note: The Unified ICM DSCP markings can be done with or without Packet Scheduler. Cisco recommends that you do not use Packet Scheduler unless: 1. Bandwidth requirements are clearly understood and correctly configured, and 2. the convergent network link is occasionally congested and shaping Unified ICM traffic at the source can be helpful.

Warning: While using the Microsoft Packet Scheduler does provide shaping and 802.1p features, there are significant risks when using this option, as follows: 1. Multiple defects have been submitted to Microsoft. Currently, some fixes are have been released by Microsoft, but some have not. 2. If the shaping bandwidth is configured too low, Packet Scheduler may introduce excessive delay and as a result it may cause timed-out calls, queue overflows and buffer exhaustion. 3. Shaping at the Unified ICM server may neither be necessary or helpful given that the LAN is rarely the bottleneck of communications over WAN and a QoS-enabled network is more capable of shaping/queuing/policing traffic based on the resource usage.

Microsoft Packet Scheduler is a key component in creating the Windows Server 2003 QoS solution. It regulates how much data a given flow is allowed, when those packets are put onto the network, and in which order such packets (those ready for transmission) are sent.

The Packet Scheduler installation is not required for Unified ICM . However, some benefit may be gained from the following:

- The Packet Scheduler's shaping functionality mitigates the burst nature of Unified ICM transmissions by smoothing transmission peaks over a given period of time, and thereby smoothing out network usage to affect a more steady use of the network.
- 802.1p tagging on Windows Server 2003 is available only if the Packet Scheduler is installed. Without the use of 802.1p, there is no physical guarantee that any prioritized data transmissions will receive a better service than best-effort transmissions receive in the LAN segment.

To install Microsoft Packet Scheduler, perform the following steps on both the CallRouter machines and the PG machines.

Note: All current TCP connections are terminated when the Packet Scheduler is installed. Therefore, even though a reboot of the machine is not required when the Packet Scheduler is installed, current TCP connections are terminated. You should not install the Packet Scheduler if important connections are in progress.

To install Microsoft Packet Scheduler

Perform the following steps:

Step 1 Open Network and Dial-up Connections.

Step 2 Right click the Network Connection (Public Visible) on which you want to install the QoS Packet Scheduler. Select Properties.

Step 3 Click Install. The Select Network Component Type dialog box appears.

Step 4 Select Service and click Add. The Select Network Service dialog box appears.

Step 5 Select QoS Packet Scheduler. Click OK to begin the installation process.

Installing and Configuring 802.1p-Capable Components

Note: 802.1p is optional. However, see the section Installing Microsoft Packet Scheduler, page 11-17, for reasons why you might want to use it.

802.1p expresses priority class by setting three bits in the Layer 2 MAC header, whose binary values 0 through 7 represent eight distinct priority classes (named as Class of Service). For Unified ICM , the default 802.1p settings are compliant with Cisco AVVID recommendations. Specifically, the high and medium traffic uses the value of 3, and the low priority traffic uses the value of 1. See Cisco AVVID Solution IP Telephony QoS Classification for details about Cisco AVVID packet classifications.

If you wish to enable 802.1p marking capabilities as part of your QoS implementation, you must perform the following tasks:

- Install and enable Microsoft Packet Scheduler, as discussed in the section Installing Microsoft Packet Scheduler (page 85).

- Install 802.1p-capable NICs in the QoS-enabled Unified ICM computers (the Router and the PG)
- Enable 802.1p on the NICs, through the Advanced tab on the NIC Properties screen. This is done by enabling a selection most often referred to as QoS Packet Tagging.
- Install 802.1p-capable switches on the LAN segment.
- Configure 802.1p-capable switches and coordinate their configuration with the settings on the Router and/or the PG.

  Note: You should install NIC cards before you install Unified ICM software. If you add a NIC card after you install Unified ICM software, you will need to reinstall Unified ICM software.

  Refer to Cisco AVVID Network Infrastructure Enterprise Quality of Service Design for details about AVVID-Enabled campus network design, switch selection, and QoS configuration commands.

Configuring QoS on IP Routes

See Cisco AVVID Network Infrastructure Enterprise Quality of Service Design for details about AVVID-Enabled WAN design, router selection, and QoS configuration commands.

Additional Tasks

This section briefly discusses a few additional tasks that you need to perform, after the deployment tasks listed in the previous sections, to ensure that your QoS-enabled network runs correctly and efficiently.

ICM QoS Setup

Refer to the Installation Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted for details about Unified ICM QoS setup.

Performance Monitoring

You can use the Windows Performance Monitor to track the performance counters associated with QoS-enabled connections. Refer to the Administration Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted for information on using the Windows Performance Monitor.

Note: Depending on your operating system version, this tool may be named System Monitor

For More Information on QoS

The following are Cisco documents that contain additional information on QoS.

You can access most Cisco documentation from the Cisco corporate website at **http://www.cisco.com**

- Cisco IP Contact Center Enterprise Edition Solution Reference Network Design (SRND)
- Cisco AVVID Network Infrastructure Overview
- Cisco AVVID Network Infrastructure Enterprise Quality of Service Design
- Cisco AVVID Solution: IP Telephony QoS Classification
- Planning for Quality of Service
- Quality of Service Networking
- Cisco IP Telephony QoS Design Guide

Active Directory Model

Microsoft Windows Active Directory provides a central repository for managing network resources. Unified ICM software uses Active Directory services to control users' access rights to perform setup, configuration, and reporting tasks. Active Directory services also grant permissions for different components of Unified ICM software to interact; for example, it grants permissions for a Distributor to read the Logger database.

Unified ICME supports both Windows 2003 and Windows 2008 Active Directory domains. Native mode is required. Unified ICM user configuration data is stored in Active Directory Organizational Units (OU).

For more information see the Staging Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted, Release 8.x(y).

TCP/IP Configuration

To set up IP addresses for Windows Server 2003 nodes, use the TCP/IP Properties dialog box. To display this dialog box, go to the Windows Server 2003 Start menu and choose Settings > Network and Dialup Connections > Local Area Connection. In the Local Area Connection Status window, click on Properties. Select Internet Protocol (TCP/IP) and click on Properties.

Select “Use the following IP address.” Enter the IP address and click OK. To enter additional IP addresses, open the TCP/IP Properties window again and click Advanced. The Advanced TCP/IP Settings window allows you to enter additional IP addresses.

Central Sites

Each side of the central controller includes the CallRouter, Logger, and Network Interface Controller (NIC). These may be on three separate nodes, two nodes, or a single node. Although the NICs are indicated as separate nodes for clarity, in fact a NIC is implemented as a process within the CallRouter node. The two sides of the central controller may be at two different central sites as shown in the following figure.

Figure 30: Geographically Distributed Central Controller

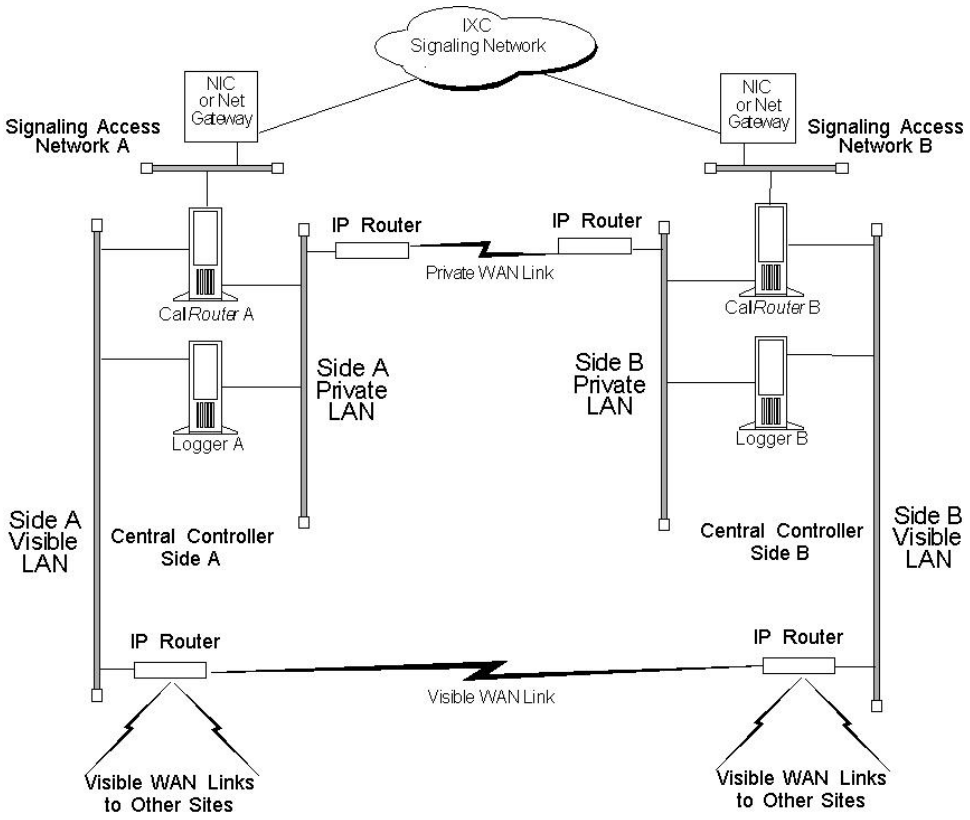

The private network carries Unified ICM system traffic between the nodes on one side of the central controller and between the nodes on both sides of the system. The traffic between the two sides of the central controller consists of synchronization and state transfer messaging between the CallRouters and Loggers. Most communications between the CallRouter and Logger on one side take place over the private network.

The private WAN link (see the preceding figure) is critical to the overall responsiveness of the Unified ICM system. First, it must provide sufficient bandwidth to handle simultaneous synchronizer and state transfer traffic. It must also have enough bandwidth left over in case additional data must be transferred as part of a recovery operation. Since the private WAN link is the only link that carries central controller synchronization and state transfer traffic, you may want to provision backup service of some kind as a contingency for network outages.

The IP routers in the private network always use traffic prioritization, and frequently use IP fragmentation, to ensure that high priority Unified ICM system traffic does not experience

excessive queuing delay. Alternately, both sides of the central controller may be co-located at a single site as shown in the following figure.

Figure 31: Collocated Central Controller

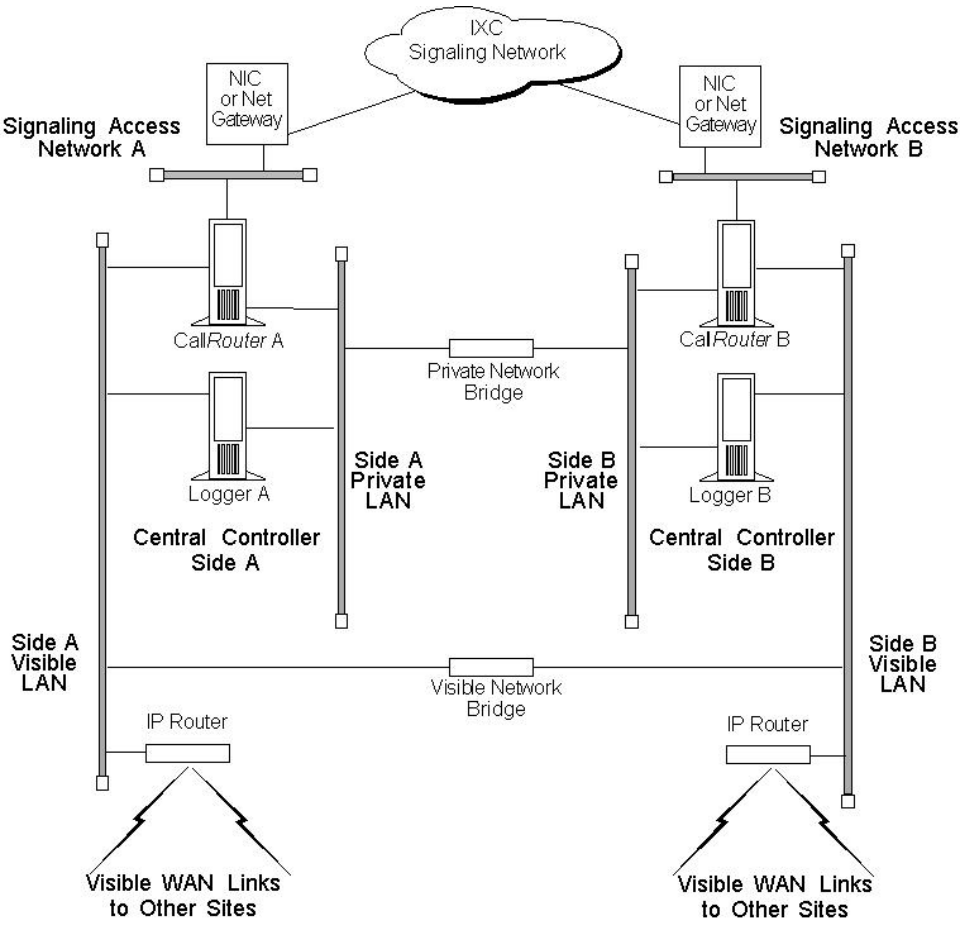

In a co-located central controller configuration, Ethernet switches separate the Side A and Side B private Ethernet LANs for fault tolerance. This private network bridge replaces the private WAN link. A visible network bridge also connects the Side A and Side B visible networks.

The Visible Network

Each central site has a visible network that connects nodes within that site. To allow communication between sites, each side of the central controller must have one IP router on its visible LAN.

Note: When a Peripheral Gateway is co-located with one side of a duplexed, geographically distributed central controller, you must have a direct connection between the visible WAN IP routers at the two central sites. This ensures that there is adequate visible network connectivity between the sides of the central controller.

The IP router requires a single address on the LAN. It also requires that you define a static route on the IP router for each contact center's visible LAN and for each administrator site's visible LAN.

Visible IP Router Configuration

To allow optimal tuning of the network, Cisco requires using IP routers that allow you to prioritize packets based on a range of source or destination port numbers. Typically, you need to set up the IP router to give higher priority to certain outgoing network packets. Also, depending on the bandwidth available on the visible WAN, you may need to set up IP fragmentation. The table titled **Central Site Visible IP Router Configuration** summarizes the configuration for the visible network IP router.

Table 8: Central Site Visible IP Router Configuration

Attribute	Requirements
IP Addresses	One address required.
Default Gateway	The network bridge (or the IP router used as bridge), if any. Otherwise, the IP router does not have a default gateway.
Static Routes	Define one static route for the visible LAN at each remote contact center site and each administrator site. If the central sites are geographically separated, add a static route for the other central site.
Other	Turn off any preset routing protocols. Give higher priority to specific network packets. Use fragmentation if necessary to limit the queuing delay.

You may need to prioritize packets as described in the table titled **Visible Network Packet Priorities from Central Site**.

Table 9: Visible Network Packet Priorities from Central Site

Packet Type	High Priority	Low Priority
TCP	If received from the CallRouter's high priority address (as derived from the packet's source address).	If received from any other address.
UDP	If source or destination port number is in the range 39000–39999	All other UDP packets.

The maximum queuing delay is 50 milliseconds to contact center sites that use post-routing or translation routes and 200 milliseconds to other contact center sites. You may have to implement fragmentation to meet these limits.

The Private Network

Each central site must also have its own private LAN. If the sides of the central controller are geographically separated, each private LAN has one IP router to communicate with the private WAN that connects the two sides.

If the two sides of the central controller are co-located, the IP router on the private LAN is not needed. If two central sites are geographically separated, each side requires an IP router on the private network.

The table titled **Central Site Private IP Router Configuration** summarizes the configuration for the private network IP router.

Table 10: Central Site Private IP Router Configuration

Setting	Requirements
IP Addresses	None.
Default Gateway	Define one static route for the private LAN at the other central site.
Static Routes	Define one static route for the private LAN at the other central site.
Other	Turn off any preset routing protocols. Give higher priority to specific network packets.

The table titled **Private Network Packet Priorities from Central Site** describes how private network packets must be prioritized.

Table 11: Private Network Packet Priorities from Central Site

Packet Type	High Priority	Low Priority
TCP	If the source address is the local CallRouter's high priority address or the destination address is the other CallRouter's high priority address.	All other TCP packets.
UDP	If source or destination port number is in the range 39000–39999.	All other UDP packets.

The Signaling Access Network

Each central site must have its own Signaling Access Network (SAN). The Unified ICM system uses the Signaling Access Network to communicate with the IXC signaling network. The Signaling Access Network for the MCI, AT&T, Nortel, and Stentor NICs is implemented as an Ethernet LAN. This LAN is separate from the Unified ICM private LAN.

In Sprint NIC configurations, the Signaling Access Network is implemented via Eicon X.25 WAN cards on the CallRouter platform. These cards allow the Unified ICM system to join the IXC signaling network. The X.25 links to the IXC signaling network are considered the Signaling Access Network. In these configurations, the separate Ethernet Signaling Access Network is not required.

The following figure shows a typical Signaling Access Network for a single central site. It assumes that the two sides are geographically separated.

Note: The IP addresses shown in this and subsequent figures are examples only. Use addresses specific to your networks.

Figure 32: Central Site Signaling Access Network

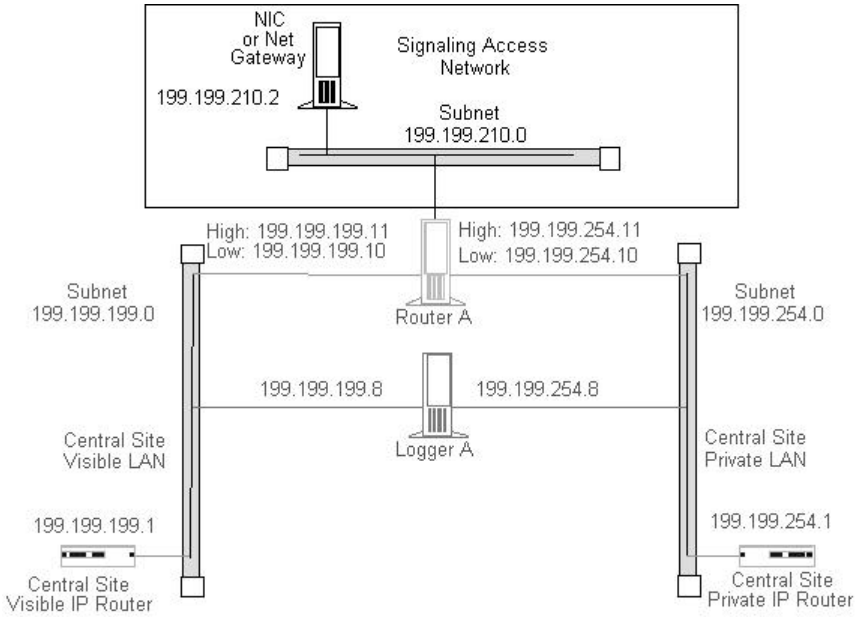

The CallRouter Node

The CallRouter connects to the visible network through the visible LAN; and to the private network through the private LAN. The CallRouter also has a connection to the Signaling Access Network.

Figure 33: CallRouter Network Connections

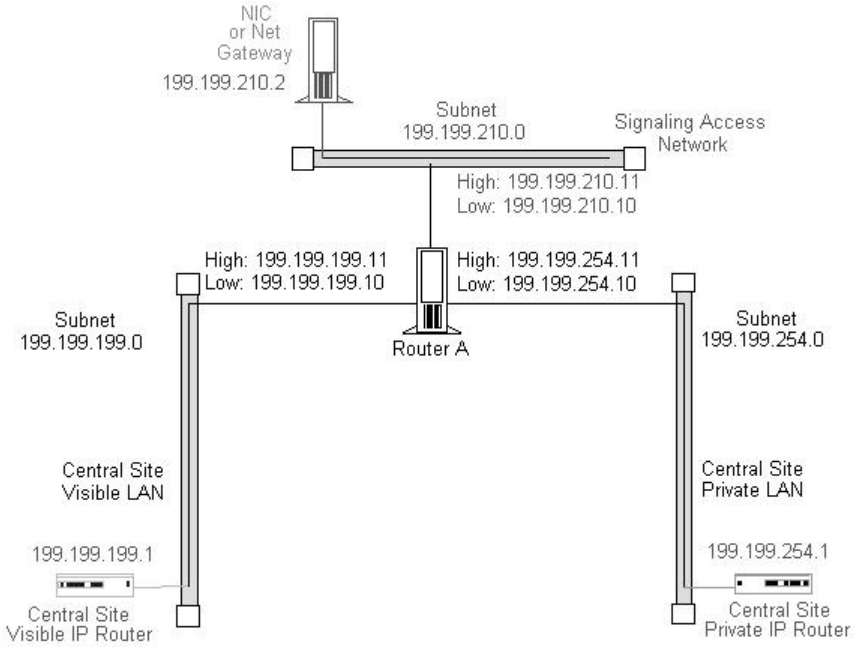

As shown in the preceding figure, the CallRouter requires two addresses on the visible LAN; two addresses on the private LAN; and two addresses on the signaling access LAN. This allows the Unified ICM system to separate high-priority network traffic from low-priority traffic.

The table titled **CallRouter Visible Network Configuration** summarizes the visible network configuration for the CallRouter.

Table 12: CallRouter Visible Network Configuration

Setting	Requirements
IP Addresses	Two required: one for high priority data; one for low (normal) priority data. Note that only one address is required if you are using QoS.
Default Gateway	Visible network IP router.
Static Routes	None.
Other	Preferred and alternate DNS server. See Active Directory Model, page 11-21.

The table titled **CallRouter Private Network Configuration**summarizes the private network configuration for the CallRouter.

Table 13: CallRouter Private Network Configuration

Setting	Requirements
IP Addresses	Two required: one for high priority data; one for low (normal) priority data.
Default Gateway	None. (The default gateway is on the visible LAN.)
Static Routes	If the sides of the central controller are geographically separated, define one static route for the subnet address of the private LAN on the other side of the central controller.
Other	Disable Windows Server 2003 networking on the private LAN.

Note: Instructions on disabling Windows Server 2003 networking on the private LAN appear later in this section.

The table titled **CallRouter Signaling Access LAN Configuration** summarizes the Signaling Access Network configuration for the CallRouter.

Table 14: CallRouter Signaling Access LAN Configuration

Setting	Requirements
IP Addresses	Two may be required, the second functioning as a serviceability interface for your Unified ICM service provider.
Default Gateway	None.
Static Routes	None.

Setting	Requirements
Other	Disable Windows Server 2003 networking on the Signaling Access Network.

Disabling Windows 2000 Server and Windows Server 2003 Networking

You need to disable network bindings for the private LAN adaptor on machines that connect to the Unified ICM private network.

You can disable Windows 2000 Server and Windows Server 2003 networking on the private LAN interface through the Network and Dial-up Connections window. Right click on the My Network Places icon on the Windows 2000 Server or 2003 desktop. The Network and Dial-up Connections window appears. (Optionally, you can right-click on the My Computer icon, select Explore, then right click on My Network Places and select Properties.)

Choose Advanced > Advanced Settings to display the Advanced Settings window:

Figure 34: Advanced Settings

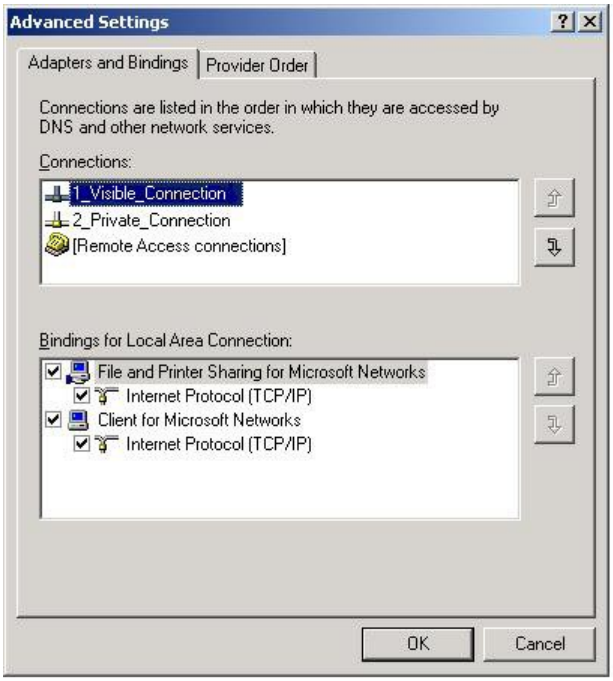

Make sure that the visible network connection appears first in the list, followed by the private network connection. You can change the order in which the network connections appear by using the arrows on the right side of the window. Select the private network connection and disable both “File and Printer Sharing for Microsoft Networks” and “Client for Microsoft Networks.”

The Logger Node

The Logger can be on the same node as the CallRouter, or it can be a separate node.

Figure 35: CallRouter and Logger Combination

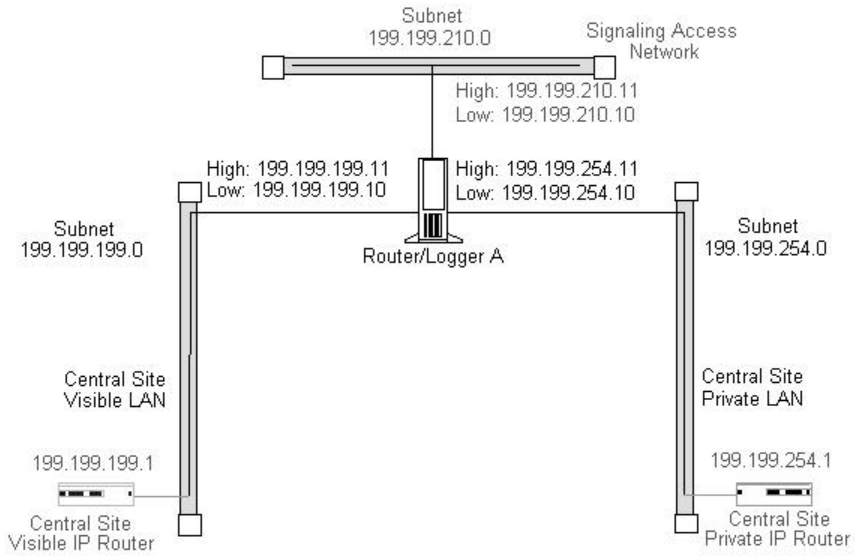

If the CallRouter and Logger are on the same node, then the Logger has no specific requirements; it uses low priority addresses defined for the node on the visible and private networks. If the two are on separate nodes, then the Logger requires its own connections to both the visible and private LANs.

Figure 36: Logger as a Separate Node

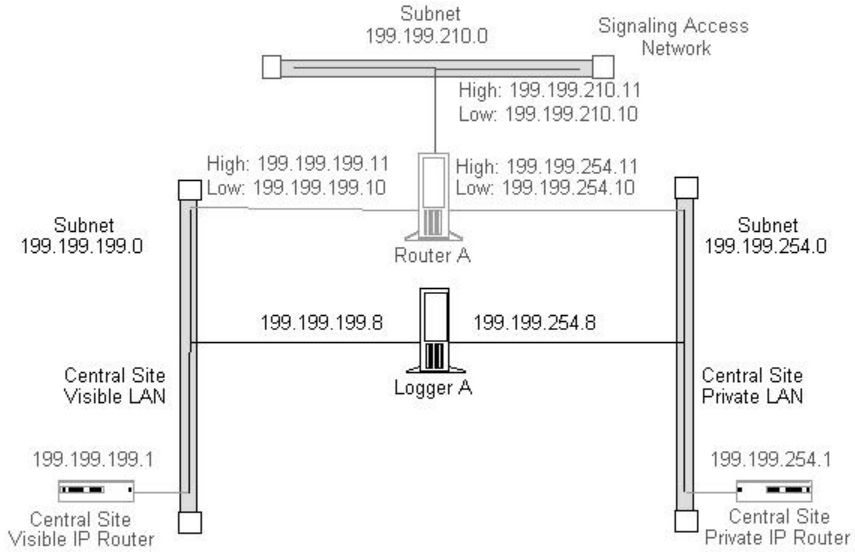

In addition to the IP addresses shown, the Logger node may require two additional addresses on the visible network. These addresses allow for dial-in connections by your Unified ICM support provider's Distributed Diagnostic and Service Network (DDSN). The table titled **Logger Visible Network Configuration** summarizes the visible network connections for the Logger.

Table 15: Logger Visible Network Configuration

Setting	Requirements
IP Addresses	TThree addresses may be required: one for normal data; two more for DDSN dial-up connections.
Default Gateway	Visible network IP router.
Static Routes	None.
Other	Preferred and alternate DNS server. See Active Directory Model (page 88).

The table titled **Logger Private Network Configuration** summarizes the private network configuration for the Logger.

Table 16: Logger Private Network Configuration

Setting	Requirements
IP Addresses	One address required.
Default Gateway	None. (The default gateway is on the visible LAN.)
Static Routes	If the two sides of the central controller are geographically separated, define one static route for the subnet address of the private LAN for the other side of the central controller.
Other	Disable Windows 2000 Server or Windows Server 2003 networking on the private LAN interface. (See Disabling Windows 2000 Server and Windows Server 2003 Networking, (page 95) for more information.)

If the Logger is on the same computer as the CallRouter, then the visible and private network IP configuration for the CallRouter is all that is required.

If the Logger is a separate node, you must disable networking on the private LAN interface (as was required for the CallRouter).

Define a static route in ICMEXEC.BAT, as for the CallRouter.

Optional Database Server Platform

If you order the Cisco Unified ICM Gateway SQL option, you need to set up an additional SQL Server database platform. The database server requires one IP address and one connection to the Unified ICM visible network.

Figure 37: Optional Database Server

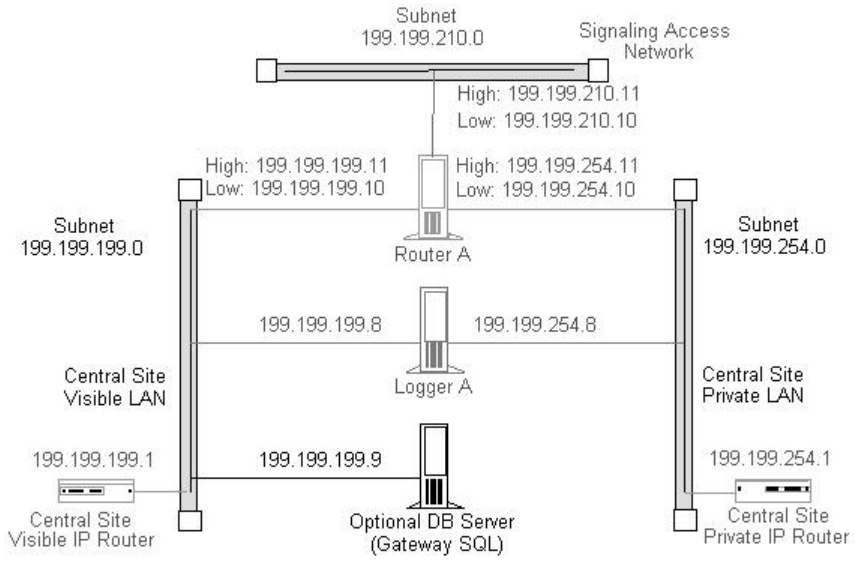

An Unified ICM Network Gateway may be deployed on the Signaling Access Network in SS7 network environments. The Unified ICM Network Gateway is a dedicated Windows Server 2003 machine that provides SS7 protocol handling. When an Unified ICM Network Gateway is used, the NIC software is installed on the CallRouter machine and a separate Gateway machine is used as the interface between the CallRouter and the carrier's SS7 signalling network.

The Network Gateway is installed on a dedicated machine. It connects to both the Signaling Access Network (SAN) and to the Unified ICM visible network. The visible network connection is used strictly for management and maintenance. The Unified ICM Network Gateway does not connect to other nodes at the central site or to nodes at other sites. For example, it does not communicate over the private network with a network gateway on the other side of the system.

The Unified ICM Network Gateway can support up to sixteen signaling links (four PCI cards) to the IXC signaling network. Therefore, the host server must have one free PCI slot for every four signaling links. Each adapter card supports four links with an individual V.35 interface for each link.

In Sigtran SS7 networks you can deploy a Sigtran Gateway on either the CallRouter machine or a separate machine. This Sigtran Gateway can communicate with either a Service Switching Point or a Media Gateway Controller, or it can communicate directly with a Signaling Gateway. In this deployment Sigtran connections are established using a Client / Server message exchange, in which the Sigtran Client requests connections with the Sigtran Server. The Signaling Gateway (such as Cisco's Internet Transfer Point) is a server in this model and accepts incoming connection requests. The Sigtran Gateways act as the Client when connected to a Signaling Gateway.

The table titled **ICM Network Gateway Signaling Access Network Configuration** summarizes the Signaling Access Network requirements for an Unified ICM Network Gateway.

Table 17: ICM Network Gateway Signaling Access Network Configuration

Setting	Requirements
IP Addresses	One address required.

Setting	Requirements
Default Gateway	None.
Static Routes	None.
Other	A HOSTS file is set up and changes are made to the CONFIG.SYS and AUTOEXEC.BAT files. Consult with your ICM support provider before changing these settings.

The table titled **ICM Network Gateway Visible Network Configuration** summarizes the visible network requirements for an Unified ICM Network Gateway.

Table 18: ICM Network Gateway Visible Network Configuration

Setting	Requirements
IP Addresses	One address required.
Default Gateway	Visible network IP router.
Static Routes	None.
Other	A HOSTS file is set up and changes are made to the CONFIG.SYS and AUTOEXEC.BAT files. Consult with your Unified ICM support provider before changing these settings.

In Sigtran SS7 networks you can deploy a Sigtran Gateway on either the CallRouter machine or a separate machine. This Sigtran Gateway can communicate with either a Service Switching Point or a Media Gateway Controller, or it can communicate directly with a Signaling Gateway. In this deployment Sigtran connections are established using a Client / Server message exchange, in which the Sigtran Client requests connections with the Sigtran Server. The Signaling Gateway (such as Cisco's Internet Transfer Point) is a server in this model and accepts incoming connection requests. The Sigtran Gateways act as the Client when connected to a Signaling Gateway.

Administration & Data Servers at a Central Site

Cisco requires that you isolate the CallRouter, Logger, and PGs from the Administration & Data Server LAN segment by using Ethernet switches. This limits the impact of one network's problems on another. By isolating the central controller and PGs from the Administration & Data Server LAN segment, you can protect critical components from network hardware and software failures (for example, an open Ethernet tap or a network error burst).

For further protection against LAN outages, you can use an IP router instead of a bridge. You can then place the Administration & Data Server on a separate LAN with other contact center computers and applications. The IP router is a better option in this situation. LAN bridges tend to forward network error bursts from one side of a LAN to the other. IP routers provide a better fire wall, since they do not forward network errors to other LANs.

The Administration & Data Server must reside on a network visible to the Unified ICM software. The following figure shows how you can use a LAN bridge or an IP router to isolate PGs and the central controller from the Administration & Data Server LAN segment.

Figure 38: Administration & Data Server at a Central Site

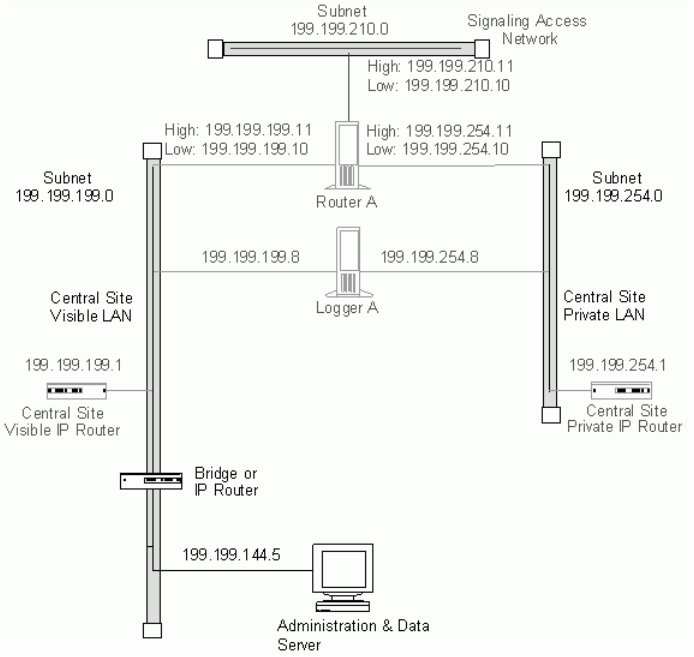

Peripheral Gateways at a Central Site

A Peripheral Gateway (PG) that is co-located with one or both sides of the central controller can share the same visible LAN segment as the CallRouter and Logger nodes. The PG can communicate with the local CallRouter through the visible LAN. If the sides of the central controller are geographically separated, the PG communicates with the other side through the visible IP router and a WAN link. (If both sides of the central controller are co-located with the PG, then the PG communicates with both sides through the visible LAN.)

The following figure shows the network connection for a PG at a central site.

Figure 39: Peripheral Gateway at a Central Site

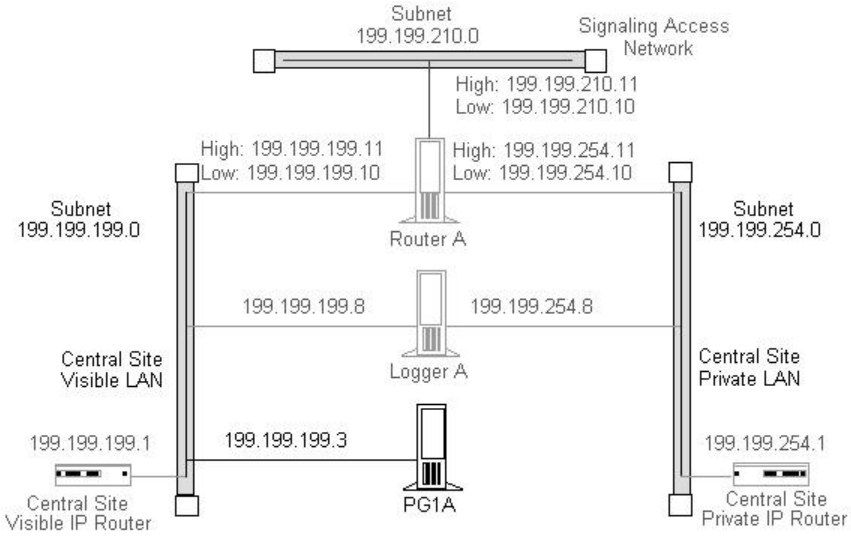

The ACD itself can also be on the visible LAN.

If the PG is duplexed, then the two duplexed PGs must be connected through a separate private network. (They cannot use the same private network as the CallRouter and Logger.) See the following figure.

Figure 40: Duplexed Peripheral Gateways at a Central Site

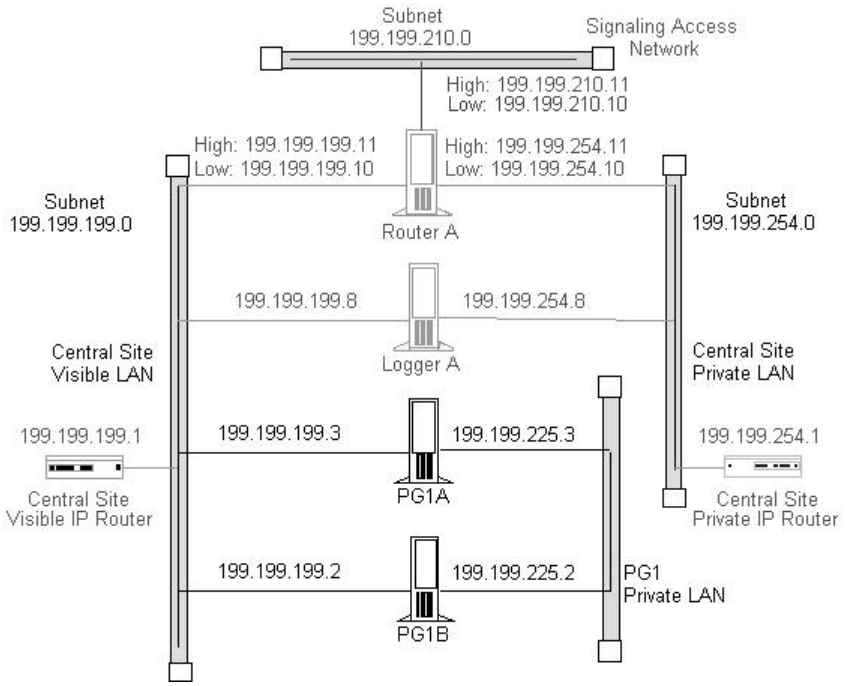

If you have more than one pair of duplexed PGs at a site, each pair requires its own private LAN. The private LAN for the PGs allows for synchronization and state transfer between the PGs. It is not used for any other purpose.

Note: When a Peripheral Gateway is located with one side of a geographically distributed central controller, you must have a WAN link directly connecting the visible WAN IP routers at the two central sites. This ensures that there is adequate visible network connectivity between the sides of the central controller. For more information on PG networking requirements, see the next section, “Contact Center Sites.”

Contact Center Sites

Each contact center site includes at least one ACD, at least one Peripheral Gateway (PG), and optionally, one or more Administration & Data Servers. Contact centers may also have an Interactive Voice Response (IVR) unit. For fault-tolerance, the contact center site must include a duplexed pair of PGs.

A remote contact center complex is reached via the visible network, often with multiple access paths and through multiple IP routers. The contact center site must have at least one IP router on the visible network for communication with the central controller. For maximum fault-tolerance, the site should have two IP routers, each connecting to one side of the central controller.

Note: For information on installing and configuring the Unified ICM Peripheral Gateway software, see the Installation Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted.

Simplexed PG Site

The following figure shows one option for a contact center configuration with a simplexed PG and an Administration & Data Server. This site contains an ACD and an IVR system. The IVR PG software and the ACD PG software may be installed on the same server hardware platform.

Figure 41: Contact Center with Simplexed PG

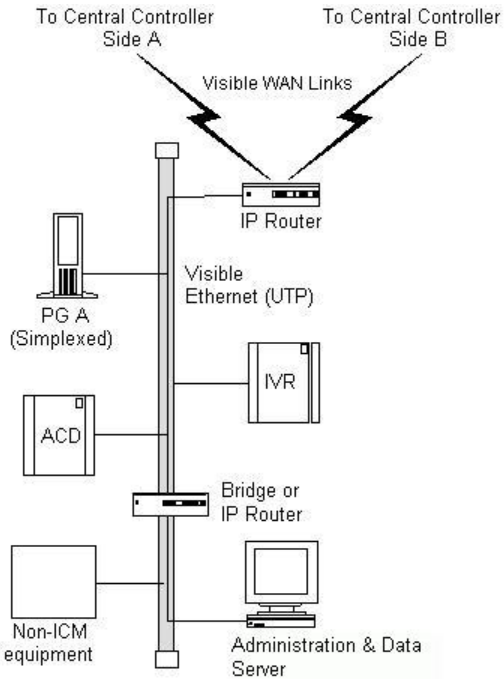

As shown in the preceding figure, the PG and Administration & Data Server share a single Ethernet LAN and an IP router. The IP router uses prioritization and IP fragmentation to minimize queuing delays for high-priority Unified ICM system traffic. Cisco requires that the PG, ACD, IVR, and IP router be separated from other devices by a bridge or IP router. This isolates the critical Unified ICM components from outages that might be caused by other equipment and networks.

The contact center example shown in the preceding figure is a low fault tolerance configuration. It is only for non-fault tolerant sites (for example, for contact center sites with one PG or administrator sites with Administration & Data Servers only). A simplexed PG configuration can represent a single point of failure. Loss of the only PG would stop the flow of real-time data from the contact center to the CallRouter and prevent the use of post-routing and translation routes. You can protect against possible failures by using duplexed PGs.

Duplexed PG Site

A duplexed PG configuration provides enhanced fault-tolerance.

Figure 42: Fault Tolerant Contact Center

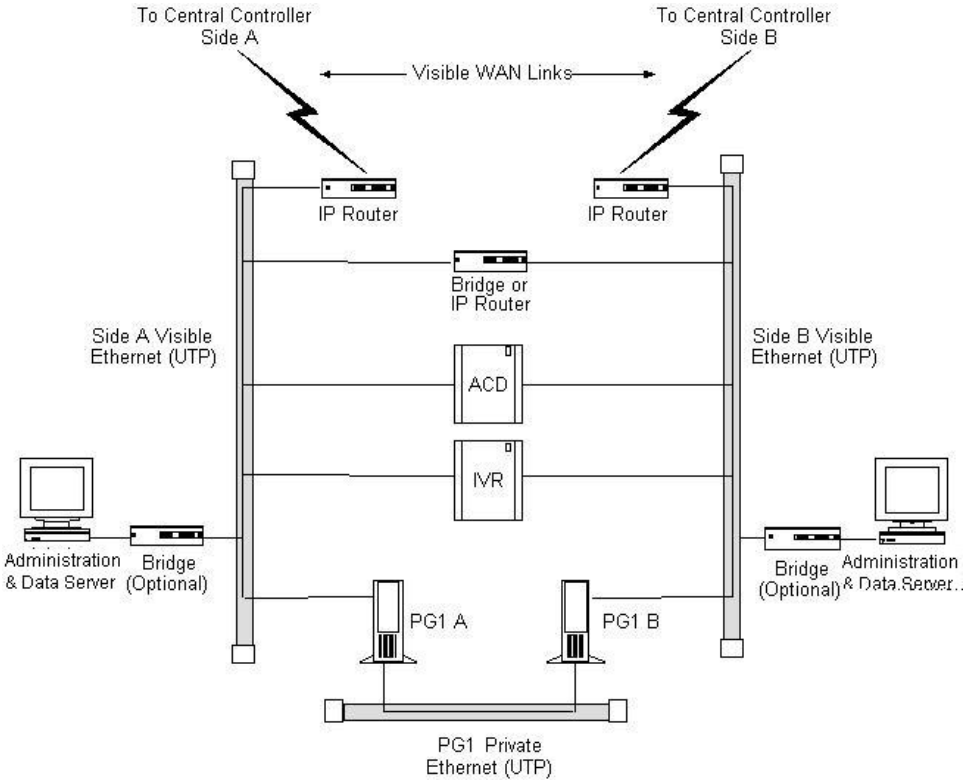

Note that a PG private LAN is added to allow direct communication between the two PGs. If you have more than one duplexed pair of PGs at a site, each PG pair requires its own private LAN.

To further enhance the fault-tolerance of the contact center, you can configure each PG with its own visible LAN and IP router. This eliminates the LAN as a single point of failure. Each PG communicates with one side of the central controller using its own LAN and IP router.

If you used a single IP router instead of two, you introduce a potential single point of failure to the contact center site. Loss of the one IP router would stop the flow of real-time data from the contact center to the CallRouter and stop the flow of monitoring data from the central controller to the Administration & Data Server. It would also prevent the use of post-routing and translation routes for this contact center.

One of the two IP routers shown in the preceding figure serves as the default gateway for the PG. By default, the PG communicates with that side of the central controller. The PG must have a static route defined to the other side of the central controller through the other IP router.

Each PG may contain a modem to allow dial-in access through your Unified ICM support provider's Distributed Diagnostic and Service Network (DDSN). In addition to its normal address on the visible network, the PG would then require two additional visible LAN addresses for this dial-in access.

Duplexed PG Site with Separate IVR LAN

Another contact center configuration may be used in cases where IVR systems need to be separated due to security concerns or when management of the IVRs must be carefully protected. The following figure shows an example of such a fault tolerant contact center site.

Figure 43: Fault Tolerant Contact Center—IVR on Separate LAN

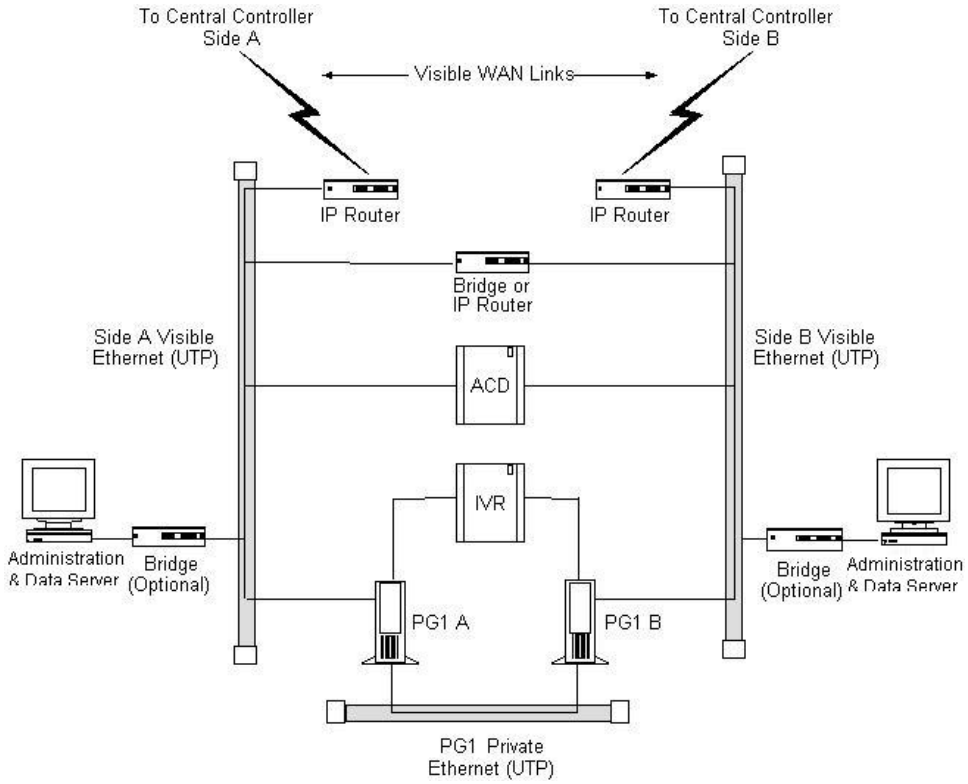

With this option, the ACD is on the visible LAN under the assumption that another CTI application needs to interface to the ACD. An alternative would be to have the ACD on the same LAN as the IVR system.

PG Network Configuration

The table titled **Simplexed PG Network Configuration** summarizes the network configuration for a simplexed PG.

Table 19: Simplexed PG Network Configuration

Setting	Requirements
IP Addresses	Three addresses may be required on the visible LAN: one for normal data and two for use by the DDSN.
Default Gateway	Define one of the visible network IP routers as the default gateway for the PG.

Setting	Requirements
Static Routes	Define one static route to the visible LAN at the central site that is not targeted by the default gateway IP router.
Other	Preferred and alternate DNS server. See Active Directory Model (page 88).

The table titled **Duplexed PG Network Configuration** summarizes the network configuration for a duplexed PG.

Table 20: Duplexed PG Network Configuration

Setting	Requirements
IP Addresses	Each PG may require three addresses on the visible LAN (one for normal traffic plus two addresses for DDSN dial-up connections) and two addresses on the private LAN (one for high priority and one for low priority data).
Default Gateway	Define one of the visible network IP routers as the default gateway for each PG. Do not use the same IP router as the default gateway for both PGs.
Static Routes	Each PG requires a static route to the side of the central controller that is not targeted by its default gateway IP router.
Other	Preferred and alternate DNS server. See Active Directory Model (page 88).

Note: For more information on how Peripheral Gateways connect to ACDs, see Chapter 5, “Peripheral Gateway Configurations” (page 33).

Contact Center IP Routers

The IP router requires a single address on the LAN. It also requires that you define a static route on the IP router to the side of the central controller (central site visible LAN) that is not targeted by the PG’s default gateway IP router.

To allow optimal tuning of the network, Cisco requires that you use IP routers that allow you to prioritize packets based on a range of source or destination port numbers. Typically, you need to set up the IP router to give higher priority to certain outgoing network packets. Also, depending on the bandwidth available on the visible WAN, you may need to set up IP fragmentation.

The table titled **Contact Center IP Router Configuration**summarizes the configuration for the IP routers.

Table 21: Contact Center IP Router Configuration

Setting	Requirements
IP Addresses	Each IP router requires one address on the visible LAN.
Default Gateway	Network bridge or IP router used as bridge, if any. Otherwise, the IP router does not have a default gateway.

Setting	Requirements
Static Routes	Each IP router must have a static route to reach one central site visible LAN.
Other	Turn off any preset routing protocols. Give higher priority to specific network packets. Use fragmentation if necessary to limit the queuing delay.

Note: See the following table for information about packet priorities.

Table 22: Contact Center Packet Priorities

Packet Type	High Priority	Low Priority
TCP	If sending to the CallRouter's high priority address (as derived from the packet's destination address).	If sending to any other address.
UDP	If source or destination port number is in the range 39000–39999.	All other UDP packets.

The maximum queuing delay is 50 milliseconds if the site uses post-routing or translation routes; 200 milliseconds otherwise. You may have to set up fragmentation to meet these limits.

Admin Sites

An administrator site contains one or more Administration & Data Servers. Each administrator site must have a visible LAN and an IP router to communicate with the central sites. An administrator site does not require a private LAN.

Figure 44: Admin Site Configuration

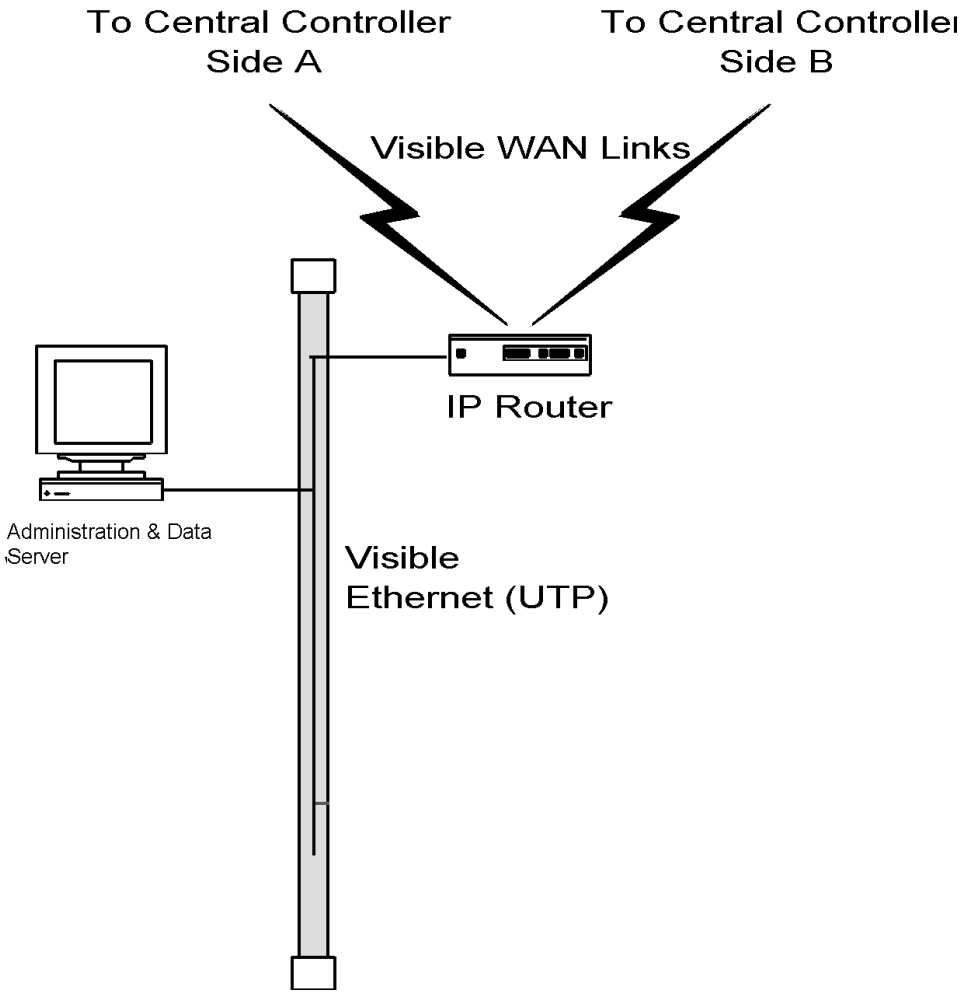

You can have multiple Administration & Data Servers on a single LAN.

Chapter 11

Site Preparation

Once you have provisioned IXC access, ordered the required ACD/PBX options, ordered the server platform, and determined your data communications requirements, you can begin preparing for the arrival of the Unified ICM equipment. You need to prepare each site that is to contain Unified ICM equipment. The sites must have adequate power facilities, security, and space for equipment layout.

Be sure to consider the following site preparation tasks:

- Meet basic site requirements. Prepare for the arrival of equipment; provide a secure staging area; ensure that sites are ready for occupancy; order and assemble equipment racks.
- Design a floor plan for each site. Consider operator workspace, cabling distribution, and maintenance access to Unified ICM nodes.
- Meet the power and environmental requirements at each site. Review the server hardware documentation for specifics on power and environmental requirements.
- Provide adequate security for the Unified ICM system. Allow only authorized access to the Unified ICM system and any backed-up data.
- Determine additional cabling or other equipment required. You may need equipment such as rack-mounting hardware or an uninterruptible power supply (UPS).
- Order any additional cabling or equipment. Order any additional equipment in time for the arrival of the Unified ICM system components.

Chapter 12

IP Address Worksheets

This chapter provides worksheets you can use to record IP addresses for the visible and private networks. You also need to define static routes for some of the nodes in the Unified ICM system.

This chapter contains the following topics:

Visible Network IP Address Requirements

The table titled **Visible Network IP Address Requirements** lists the IP address requirements for Unified ICM node connections to the visible network. The Unified ICM nodes are listed as duplexed pairs (for example, CallRouter A and CallRouter B). You may or may not have duplexed nodes in your configuration. Supply IP addresses only for the nodes you have in your configuration.

Table 23: Visible Network IP Address Requirements

Node	Location	Address Type	IP Address
CallRouter A		High Priority	
		Low Priority	
		Default IP Gateway	
		Netmask	
CallRouter B		High Priority	
		Low Priority	
		Default IP Gateway	

Node	Location	Address Type	IP Address
		Netmask	
Logger A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
Logger B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
Central Site IP Router		Normal data	
Remote Contact Center Site IP Route		Normal data	
PG1A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG1B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG2A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG2B		Normal Data	

Node	Location	Address Type	IP Address
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG3A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG3B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
AW1		Normal Data	
		Default IP Gateway1	
		Netmask	
AW2		Normal Data	
		Default IP Gateway1	
		Netmask	

Private Network IP Address Requirements

The table titled **Private Network IP Address Requirements** lists the IP address requirements for Unified ICM node connections to the private network. The Unified ICM nodes are listed as duplexed pairs (for example, CallRouter A and CallRouter B). You may or may not have duplexed nodes in your configuration. You need to supply IP addresses only for the nodes you have in your configuration.

Table 24: Private Network IP Address Requirements (continued)

Node	Location	Addres Type	IP Address
CallRouter A		High Priority	
		Low Priority	

Node	Location	Addres Type	IP Address
		Default IP Gateway	
		Netmask	
CallRouter B		High Priority	
		Low Priority	
		Default IP Gateway	
		Netmask	
Logger A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
Logger B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
Central Site IP Router		Normal Data	
Remote Contact Center Site IP Router		Normal Data	
PG1A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG1B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG2A		Normal Data	
		RAS 1	

Node	Location	Addres Type	IP Address
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG2B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG3A		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
PG3B		Normal Data	
		RAS 1	
		RAS 2	
		Default IP Gateway1	
		Netmask	
		Modem Tel. Number	
AW 1		Normal Data	
		Default IP Gateway1	
		Netmask	
AW 2		Modem Tel. Number	
		Normal Data	
		Default IP Gateway1	
		Netmask	

Signaling Access Network IP Requirements

The table titled **Signaling Access Network IP Requirements**lists the IP address requirements for Unified ICM node connections to the Signaling Access Network. The Unified ICM nodes are listed as duplexed pairs (for example, CallRouter A and CallRouter B). You may or may

not have duplexed nodes in your configuration. You need to supply IP addresses only for the nodes you have in your configuration

Table 25: Signaling Access Network IP Requirements

Node	Location	Address Type	IP Address
CallRouter A		Normal data	
CallRouter B		Normal data	
Network Gateway 1A		Normal data	
Network Gateway 1B		Normal data	

Static Route Requirements

The IP routers used in the Unified ICM networks must have static routes defined in order to provide the necessary connectivity between the visible LAN at the central site and the visible LANs at remote contact center sites. The static route ensures that the IP router can forward traffic from the central site to the remote site. In addition, CallRouters and Loggers must have a static route defined for the remote private LAN. This static route ensures that private network traffic is segregated from visible network traffic.

All the static routes required in your configuration must be defined. However, these static routes cannot be defined until all Unified ICM nodes have been assigned IP addresses.

Table 26: Static Route Requirements

Node	Network	Static Route
Central Site Visible Network IP Router—Side A and Side B	Visible	Define one static route for the visible LAN at each remote contact center site and each administrator site. If the central sites are geographically separated, add another static route for the other central site.
Central Site Private Network IP Router—Side A and Side B	Private	Define one static route for the private LAN at the other central site.
CallRouter—Side A and Side B	Private	If the sides of the central controller are geographically separated, define one static route for the subnet address of the private LAN on the other side of the central controller.
Logger—Side A and Side B	Private	If the sides of the central controller are geographically separated, define one static route for the subnet address of the private LAN on the other side of the central controller.
PG (all PGs)	Visible	One of the two IP routers at a contact center is targeted as the default gateway

Node	Network	Static Route
		for the PG. However, the PG needs IP connectivity to both sides of the central controller. Therefore, for each PG you must define a static route to the other IP router (that is, to the IP router that is not targeted as the PG's default gateway IP router).
Remote Contact Center IP Routers	Visible	For each IP router, define a static route to one side of the central controller (to the central site visible network IP router).
Admin Site IP Routers	Visible	For each Admin Site IP router, define a static route to one side of the central controller (to the central site visible network IP router).

Index