ICM Administration Guide for Cisco ICM/IPCC Enterprise & Hosted Editions

Cisco ICM/IPCC Enterprise & Hosted Editions Release 7.0(0)
July 2007

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: 408 527-0883

ICM Administration Guide for Cisco ICM/IPCC Enterprise & Hosted Editions, Release 7.0(0)

CONTENTS

About This Guide

Purpose

This manual describes how to administer and manage Intelligent Contact Management (ICM) software. It includes information about database administration, event management, support services, and ICM software's fault tolerant architecture.

This manual includes some discussions of how ICM software functions in integrated environments with the Cisco E-Mail Manager (Cisco E-Mail Manager Option) and Cisco Collaboration Server (Cisco Web Collaboration Option) components, but it does *not* provide administration information for the E-Mail Manager and Collaboration Server components. Please refer to the approriate E-Mail Manager and Collaboration Server documents for instructions on how to administer these components.

Audience

This manual is intended for personnel responsible for administering ICM software. As an ICM administrator, you should be familiar with Microsoft SQL Server database administration and Windows 2003. This manual also assumes that you have a general understanding of the ICM system components and how they work together as a complete call routing system. Administrators who are responsible for an ICM system that is part of an integrated environment should also have a general understanding of Cisco Collaboration Server and Cisco E-Mail Manager system components.

Organization

The manual is divided into the following chapters.

Chapter	Description
Chapter 1, "Administration Overview"	Describes aspects of the system that are of interest to the administrator.
Chapter 2, "Fault Tolerance"	Describes the main features of the ICM fault tolerant architecture, with special emphasis on how fault tolerance affects the administration of the system.

Chapter	Description
Chapter 3, “The ICM Databases”	Introduces the local and central ICM databases and explains how they are used.
Chapter 4, “Database Administration”	Describes the ICMDBA tool, used for various database administration tasks.
Chapter 5, “General Administration”	Describes the administration that ICM software performs automatically. This chapter also includes several optional administration features that you can use.
Chapter 6, “Event Management”	Describes how ICM software reports events from components and processes throughout the system. This chapter also describes the different tools that you can use to view event data.
Chapter 7, “Support Facilities”	Explains the ICM Distributed Diagnostics and Services Network (DDSN) and several other support provider services.
Chapter 8, “ICM Partitioning Overview”	Discusses the ICM)\ Partitioning feature, which controls what data individuals are allowed to access within an ICM database.

Conventions

This manual uses the following conventions.

Format	Example
Boldface type is used for user entries, keys, buttons, and folder and submenu names.	Choose **Edit > Find** from the ICM Configure menu bar.
Italic type indicates one of the following: • A newly introduced term • For emphasis • A generic syntax item that you must replace with a specific value • A title of a publication	• A *skill group* is a collection of agents who share similar skills. • *Do not* use the numerical naming convention that is used in the predefined templates (for example, **persvc01**). • IF *(condition, true-value, false-value)* • For more information, see the *Cisco ICM Enterprise Edition Database Schema Handbook.*
An arrow (>) indicates an item from a pull-down menu.	The Save command from the File menu is referenced as **File > Save**.

Other Publications

For additional information about Cisco Intelligent Contact Management (ICM) software, see the Cisco web site listing ICM documentation.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. This section explains the product documentation resources that Cisco offers.

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

The Product Documentation DVD is a library of technical product documentation on a portable medium. The DVD enables you to access installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the HTML documentation and some of the PDF files found on the Cisco website at this URL:

http://www.cisco.com/univercd/home/home.htm

The Product Documentation DVD is created and released regularly. DVDs are available singly or by subscription. Registered Cisco.com users can order a Product Documentation DVD (product number DOC-DOCDVD= or DOC-DOCDVD=SUB) from Cisco Marketplace at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

Ordering Documentation

You must be a registered Cisco.com user to access Cisco Marketplace. Registered users may order Cisco documentation at the Product Documentation Store at this URL:

http://www.cisco.com/go/marketplace/docstore

If you do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Documentation Feedback

You can provide feedback about Cisco technical documentation on the Cisco Support site area by entering your comments in the feedback form available in every online document.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you will find information about how to do the following:

- Report security vulnerabilities in Cisco products
- Obtain assistance with security incidents that involve Cisco products
- Register to receive security information from Cisco

A current list of security advisories, security notices, and security responses for Cisco products is available at this URL:

http://www.cisco.com/go/psirt

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you have identified a vulnerability in a Cisco product, contact PSIRT:

- For emergencies only—security-alert@cisco.com

  An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- For nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Tip We encourage you to use Pretty Good Privacy (PGP) or a compatible product (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.*x* through 9.*x*.

Never use a revoked encryption key or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

If you do not have or use PGP, contact PSIRT to find other means of encrypting the data before sending any sensitive material.

Product Alerts and Field Notices

Modifications to or updates about Cisco products are announced in Cisco Product Alerts and Cisco Field Notices. You can receive these announcements by using the Product Alert Tool on Cisco.com. This tool enables you to create a profile and choose those products for which you want to receive information.

To access the Product Alert Tool, you must be a registered Cisco.com user. Registered users can access the tool at this URL:

http://tools.cisco.com/Support/PAT/do/ViewMyProfiles.do?local=en

To register as a Cisco.com user, go to this URL:

http://tools.cisco.com/RPF/register/register.do

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Support website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Support Website

The Cisco Support website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day at this URL:

http://www.cisco.com/en/US/support/index.html

Access to all tools on the Cisco Support website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Note Before you submit a request for service online or by phone, use the **Cisco Product Identification Tool** to locate your product serial number. You can access this tool from the Cisco Support website by clicking the **Get Tools & Resources** link, clicking the **All Tools (A-Z)** tab, and then choosing **Cisco Product Identification Tool** from the alphabetical list. This tool offers three search options: by product ID or model name; by tree view; or, for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Tip **Displaying and Searching on Cisco.com**

If you suspect that the browser is not refreshing a web page, force the browser to update the web page by holding down the Ctrl key while pressing **F5**.

To find technical information, narrow your search to look in technical documentation, not the entire Cisco.com website. After using the Search box on the Cisco.com home page, click the **Advanced Search** link next to the Search box on the resulting page and then click the **Technical Support & Documentation** radio button.

To provide feedback about the Cisco.com website or a particular technical document, click **Contacts & Feedback** at the top of any Cisco.com web page.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests, or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411
Australia: 1 800 805 227
EMEA: +32 2 704 55 55
USA: 1 800 553 2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—An existing network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of the network is impaired while most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Online Subscription Center is the website where you can sign up for a variety of Cisco e-mail newsletters and other communications. Create a profile and then select the subscriptions that you would like to receive. To visit the Cisco Online Subscription Center, go to this URL:

  http://www.cisco.com/offer/subscribe

- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco channel product offerings. To order and find out more about the *Cisco Product Quick Reference Guide*, go to this URL:

  http://www.cisco.com/go/guide

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- Cisco Press publishes a wide range of general networking, training, and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Internet Protocol Journal* is a quarterly journal published by Cisco for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the *Internet Protocol Journal* at this URL:

  http://www.cisco.com/ipj

- Networking products offered by Cisco, as well as customer support services, can be obtained at this URL:

  http://www.cisco.com/en/US/products/index.html

- Networking Professionals Connection is an interactive website where networking professionals share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

  http://www.cisco.com/discuss/networking

- “What’s New in Cisco Documentation” is an online publication that provides information about the latest documentation releases for Cisco products. Updated monthly, this online publication is organized by product category to direct you quickly to the documentation for your products. You can view the latest release of “What’s New in Cisco Documentation” at this URL:

  http://www.cisco.com/univercd/cc/td/doc/abtunicd/136957.htm

- World-class networking training is available from Cisco. You can view current offerings at this URL:

  http://www.cisco.com/en/US/learning/index.html

CHAPTER 1

Administration Overview

Because of its fault tolerant design, ICM software requires little ongoing administration. However, there are some aspects of the ICM system that you should understand:

- **Fault Tolerant Architecture**. The fault tolerant architecture of the ICM system ensures continuous operation in the event of hardware or software failures. Certain system administration tasks may not be necessary depending on the level of fault tolerance present in your ICM system. You should review the ICM's fault tolerant features in order to gain a better understanding of overall system administration.
- **ICM Databases**. The central database resides on the Central Controller and is used for persistent storage of data. In addition, each Distributor Admin Workstation has its own local database. The local database is used for real-time reporting and storing configuration data and scripts. You should understand how these databases are used in the system. You should also become familiar with the tools that manage the data in these databases.
- **Database Storage**. The ICM databases are sized and set up at installation to suit your particular contact center enterprise requirements. However, you may want to become familiar with the aspects of system usage that affect database storage capacity. You might also want to review the criteria for sizing the ICM central database.
- **General Administration**. Although most administration is taken care of automatically by the system, there are several optional administration features you should be aware of (especially if configuration uses a simplexed Central Controller). These include backing up the central database, performing manual integrity checks on the Distributor AW local database, and examining the Logger's event log files.
- **Event Management**. You may want to become familiar with the ICM's event management system. ICM software provides several tools for reviewing event data in the system. Event data can aid you in identifying potential system performance problems.
- **Support Facilities**. ICM software includes several support provider and remote maintenance facilities. You might want to know more about the Distributed Diagnostic and Services Network (DDSN), which is a facility that allows your ICM support provider to remotely diagnose and fix problems in an ICM system. You might also be interested in the DDSN's optional serial alarm and SNMP feeds. For more information on SNMP Feeds see the S*NMP Guide for Cisco ICM/IPCC Enterprise & Hosted Editions.*
- **Partitioning**. ICM Partitioning provides a mechanism for controlling what data individuals are allowed to access within an ICM database.

The following chapters describe these topics in more detail.

For information on registering users and setting up security for the ICM system, see the *ICM Installation Guide for Cisco Enterprise Edition*. The installation guide also contains information on networking requirements and configuration options for the ICM system components.

CHAPTER 2

Fault Tolerance

Intelligent Contact Management (ICM) software is a fault tolerant call routing system that continues to operate without interruption in the case of hardware, software, or communications failures. The main goals of the ICM's fault tolerant architecture are to:

- Minimize time periods during which the system is non-responsive to call routing requests (for example, while the system is being reconfigured due to a component failure or recovery).
- Eliminate all single points of failure that would cause the system to stop.
- Provide disaster protection by allowing the major system components to be geographically separated.

The ICM's fault tolerant mechanisms operate in the background and are not visible from within ICM applications. However, it is still important that you have a general understanding of the fault tolerant architecture and the implications it has for system administration.

In some cases, the level of fault tolerance in the ICM system can affect which administration tasks you need to perform. For example, in duplexed database configurations many typical database administration tasks such as database backups become unnecessary because exact copies of the central database are kept on each side of the system on separate computers.

This chapter provides an overview of ICM fault tolerance with a special emphasis on the fault tolerance of the Central Controller and the central database.

Architecture

The architecture of ICM software allows the system to continue to function if one component fails. This ability is called *fault tolerance*. To ensure that ICM software continues to operate in the case of a computer failure, all critical parts of the system can be physically duplicated. There can be two or more physical Network Interface Controllers (NICs), two physical Peripheral Gateways (PGs) at each call center, and two Central Controllers. The communication paths between critical components can also be duplicated.

The critical components of ICM software include the Central Controller (CallRouter and Logger), PGs, and NICs. Normal Admin Workstations (AWs) are not considered to be critical to the operation of the system since they play no active role in routing calls or storing historical data.

When both instances of a component are available to the system, that component is said to be *duplexed*; when only one of the pair is available, the component is running *simplexed*. You might have some components in your ICM system that are duplexed and others that are simplexed. For example, you might have a duplexed Central Controller (two CallRouters and two Loggers) and simplexed Peripheral Gateways at call center sites.

It takes more than duplicate hardware to achieve fault tolerance. The ICM system can quickly detect that a component has failed, bypass that component, and use its duplicate instead. ICM software can also initiate diagnostics and service so that the failed component can be fixed or replaced and the system returned to duplexed operation.

Approaches to Fault Tolerance

ICM software uses two approaches to fault tolerance: hot standby and synchronized execution. 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. Peripheral Gateways optionally use the hot standby approach to fault tolerance.

ICM software uses *synchronized execution* in the Central Controller. In the synchronized execution approach, all critical processes (CallRouter, Logger, and Database Manager) are duplicated on separate computers. There is no concept of primary or backup. Both process sets run in a synchronized fashion, processing duplicate input and producing duplicate output. Each synchronized system is an equal peer. Each set of peers is a *synchronized process pair*.

In the event that one of the synchronized processes fails (for example, a CallRouter goes off-line), its peer continues to run. There is no loss of data and calls continue to be routed. When the failed member of the pair returns to operation, it is resynchronized with its peer and begins to run again as a synchronized process. Figure 2-1 shows how synchronized execution and hot standby are applied in ICM software.

Figure 2-1 Duplexed ICM Fault Tolerance

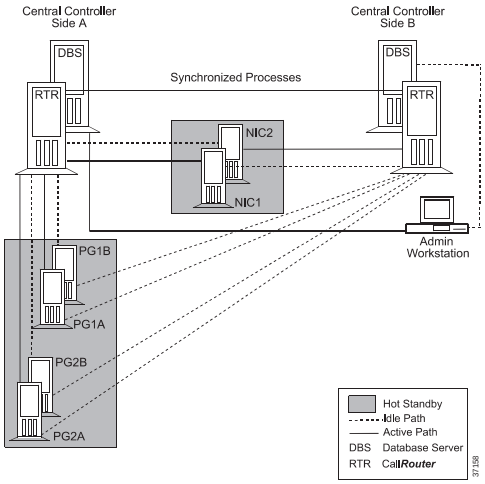

PGs and NICs use the hot standby approach to fault tolerance. Note that the duplexed NIC in Figure 2-1 is implemented on two separate computers. Each computer has active and idle connections to the sides of the Central Controller. NIC fault tolerance is described in more detail later in this chapter.

Duplicated Communication Paths

Each NIC, Peripheral Gateway, and Admin Workstation has two communication paths to the Central Controller (see Figure 2-1). The two paths connect the device (for example, a PG) to a Central Controller Agent process on each side of the Central Controller. The *Central Controller Agent* is a software process that manages communications between the Central Controller and nodes in the ICM system.

At any one time, one of the two communications paths is active and the other is idle. All communication traffic between the Central Controller and the device is sent on the active path. If the active path fails for any reason, the second path is activated and all traffic is switched to the newly active path. The previously active path becomes the idle path.

The communication protocols use buffering and acknowledgments to ensure that no messages are lost during the path failure and switch-over. After a communication path failure, the device periodically attempts to re-establish communication along the failed path.

A different mechanism is used for the real-time data feed to Admin Workstations. See Real-Time Distributors, page 2-13, later in this chapter, for more information.

Node Manager

Each ICM component (except the client-only Admin Workstation) includes a Node Manager process. The Node Manager is in charge of restarting Intelligent Contact Management processes that have failed.

For example, each Logger and each CallRouter has its own Node Manager. If a Logger and CallRouter are installed on the same machine, two separate Node Managers run on that machine. If Loggers for multiple customers run on a single machine, a separate Node Manager runs for each customer.

When a failure occurs in a single-customer ICM system, the Node Manager may shut down the machine to initiate a reboot. However, in a network service provider (NSP) environment when a Logger or CallRouter fails, components for other customers might still be active on the machine. In such a case, the Node Manager for an NSP component *does not* shut down and reboot the machine, and manual intervention is required to restore the failed component.

If the Node Manager does initiate a reboot, the Node Manager itself restarts when the machine reboots. The Node Manager then starts the other processes for the component. On a Distributor Admin Workstation, you can choose whether to have the Node Manager automatically restart when the computer reboots.

For more information on Node Manager start-up options, see the *ICM Installation Guide for Cisco Enterprise Edition*.

Central Controller

The Central Controller includes the CallRouter, Logger, and the Database Manager. The CallRouter and Logger processes are typically on separate computers. However, in smaller call center configurations the CallRouter and Logger processes can be on the same computer. The Database Manager works very closely with the Logger. The Logger and Database Manager processes are always on the same computer.

Note Beginning with ICM 7.0(0), the Logger is changing from a single process to a double process—one process handling configuration data and the other handling historical data. This allows parallel processing of the two kinds of data and, thus, a more efficient Logger. However, these two processes are still part of a single Logger node; that is, the functionality of the Logger remains essentially unchanged.

Therefore, throughout this manual, reference will generally continue to be made to *the* Logger, without distinguishing between the separate processes.

However, you should be aware that the new split in the Logger does affect failure and failover behavior. For example, if the historical Logger on side A fails, the system fails over to the historical Logger on side B; however, the still functioning configuration Logger on side A continues to be used.

A duplexed Central Controller uses the synchronized execution approach to fault tolerance. The Central Controller processes are duplicated and run as synchronized process pairs. In synchronized execution, if one component fails its peer continues running and the system runs without interruption. The Database Manager is also duplicated, but technically it does not run synchronized. Since all modifications to the database come through the Logger, the databases automatically remain synchronized.

Two Sides

All components of the Central Controller, with their duplicates, form one logical duplexed system. The system can be divided into two *sides*, each of which contains one instance of a component. Each side of the Central Controller has a Database Manager, Logger, CallRouter, Synchronizer, and an Agent. By convention, the two sides are referred to as Side A and Side B.

All components within a side are collocated; that is, located on the same local area network (LAN). However, Side A might be geographically separated from Side B. Figure 2-2 shows the two sides of a duplexed Central Controller.

Figure 2-2 Duplexed Central Controller

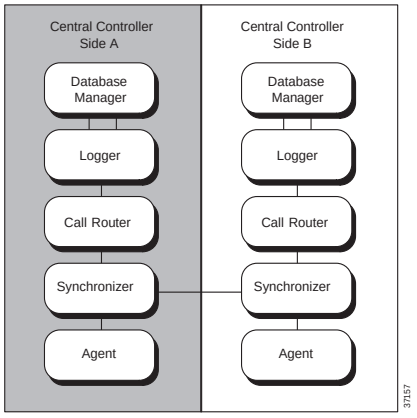

During normal operation, the two sides run in parallel. For example, information about each incoming call is processed by both CallRouters. Both CallRouters, using the same call routing scripts and identical information about the call centers, determine the same destination for the call. Both sides of the Central Controller receive the same information from the Peripheral Gateways and Admin Workstations.

A duplexed Central Controller can tolerate a single device or system failure (for example, the loss of one CallRouter) without losing functions. A double failure, while extremely rare, typically results in some loss of functions. An example of a double failure would be if both CallRouters in a duplexed system were to go off-line.

Single failures are typically caused by system crashes, operating system failures, or disk failures. However, LAN outages and IP router failures can also cause single failures. Figure 2-3 shows five possible Central Controller failure scenarios.

Figure 2-3 Central Controller Failure Scenarios

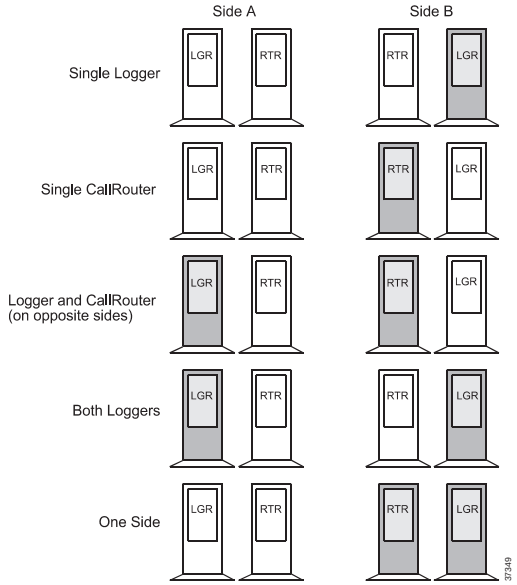

Each of these failures affects system functions differently:

- **Single Logger**. If a single Logger (whether historical or configuration Logger) goes off-line, ICM software runs without interruption. All call routing and reporting functions remain available. The CallRouters continue to operate as a synchronized pair. The remaining Logger runs simplexed. When the failed Logger returns to service, the Loggers return to synchronized execution.
- **Single CallRouter**. When a CallRouter on one side of the Central Controller fails, that entire side of the Central Controller is removed from service. This is because the CallRouter plays an integral part in forwarding historical data to the historical Logger on its side of the system. The on-line side of the Central Controller runs as a simplexed system. Call processing continues uninterrupted and reporting functions are still available. When the failed CallRouter returns to service, both CallRouters and both Loggers return to synchronized execution.
- **Logger and CallRouter (opposite sides)**. In this failure scenario, side B of the Central Controller is removed from service due to the CallRouter failure. Call routing continues uninterrupted with the remaining Side A CallRouter; however, because neither Logger is available, data in both databases slowly becomes out of date. Some reporting functions are not available until the nodes are returned to service and synchronized execution is restored.
- **Both Loggers**. In a double Logger failure, call routing continues uninterrupted. If it is the historical Loggers that failed, all reporting functions are lost until at least one of the historical Loggers returns. If it is the configuration Loggers that failed, you cannot make any configuration changes until at least one configuration Logger is operational. Such a double failure is extremely rare. The CallRouter continues to route calls because it has a copy of the call center enterprise configuration data in its program memory. (The CallRouter loads the configuration data into memory when it is started and keeps it up-to-date as configuration changes occur.)
- **One Side**. If one side of the Central Controller goes off-line, processing continues uninterrupted. ICM software continues to function as a simplexed system until the failed side of the Central Controller returns to service. All functions remain, but the system is running simplexed (without protection against an additional failure). When the off-line side of the Central Controller returns, normal duplexed operation is restored.

A double CallRouter failure would temporarily disrupt call routing and reporting functions. This type of failure is extremely rare (especially in geographically distributed Central Controller configurations).

Geographic Distribution

To provide maximum protection against disasters such as fires, floods, and earthquakes, the two sides of the Central Controller can be in separate locations—even in separate cities. The two Synchronizers communicate with each other via a private wide area network (WAN) to ensure that they remain synchronized. This WAN, called the *private WAN*, is used for no other purpose than to ensure synchronization between the sides of the Central Controller.

For details on collocated and distributed Central Controller configurations, see the *Cisco ICM Enterprise Edition Installation Guide*.

Role of the Synchronizers

The Synchronizers play the key role in maintaining synchronized execution across the two sides of the Central Controller. All input for the CallRouter and any changes to the Logger must pass through the Synchronizers.

Each time a Synchronizer receives input, it passes that input to its duplicate on the other side. The two Synchronizers cooperate to ensure that they are both sending the same input to the Central Controllers on both sides of the system.

Figure 2-4 shows how the Synchronizers combine input messages and send the messages in the same order to each side of the Central Controller.

Figure 2-4 Role of the Synchronizers

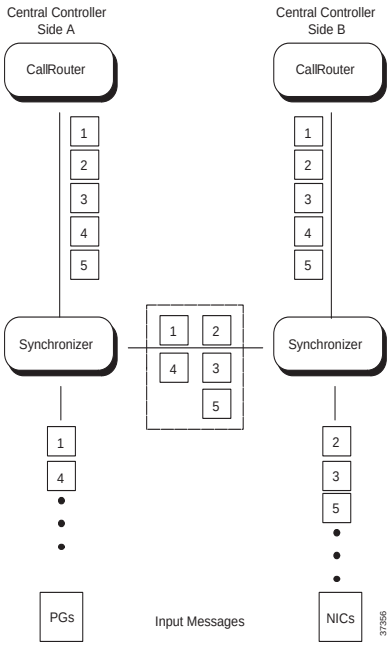

Both CallRouters receive the same input and generate the same output. The Synchronizers ensure that both sides of the Central Controller return identical destinations for the same call and write identical data to the databases.

Figure 2-5 further illustrates the Central Controller and its device connections.

Figure 2-5 ICM Fault-Tolerant Architecture

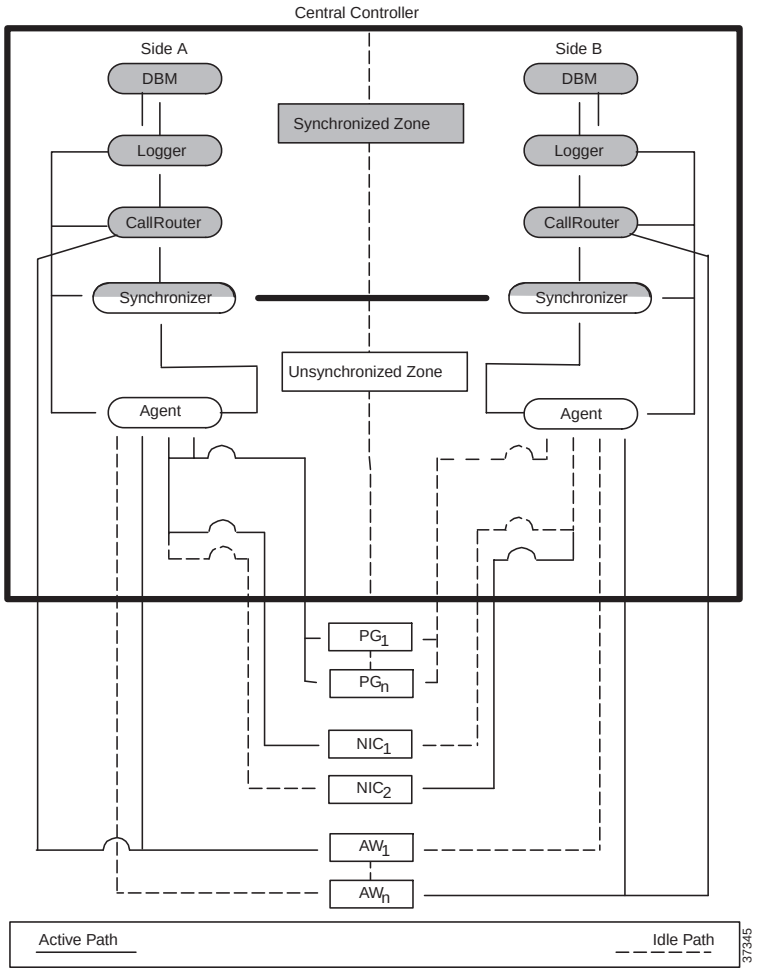

Each PG, NIC, and Admin Workstation has duplicate communication paths to the Central Controller. If there is a failure on one side of the Central Controller, the PGs, NICs, and Admin Workstations can switch their communication paths to the active side of the Central Controller. As shown in Figure 2-5, only one communication path is active at a time. The other communication path is idle (indicated by a dotted line). ICM software sends heartbeats (brief periodic messages) over the idle path to ensure that it can still be used in the event that the active path fails.

Synchronization and State Transfer

In synchronized execution, duplicated processes are always processing identical input and generating identical output. If one process fails, the other continues to operate without interrupting system operation. Once the failed process returns, it is immediately updated with the current state of ICM processes running on its peer.

In order to synchronize one peer with another after a failure, the system performs a *state transfer.* The state transfer facility allows a synchronized process (for example, a CallRouter) to copy the variables in its memory to its peer. The recovering system receives the variables from the currently executing system and is able to restart with a copy of the current state of ICM processes. For example, as soon as a failure is detected on the Side A CallRouter, ICM software uses only Side B. When the Side A CallRouter is restarted, ICM software invokes a state transfer to immediately update the Central Controller Side A components with the current state of their counterparts on Side B.

In order to better understand synchronization and state transfer, it might help to take a closer look at CallRouter and Logger recovery.

CallRouter Recovery

When a single CallRouter process fails for any reason, ICM software continues to operate without any loss of functions by using the other side of the Central Controller. All attached devices (PGs, NICs, and Admin Workstations) switch their active communications paths to the remaining side. This ensures that devices such as PGs continue to receive CallRouter output through the active CallRouter on the other side of the system.

As a consequence of the CallRouter failure, the entire side of the Central Controller is removed from service. The Logger and Database Manager associated with the failed CallRouter see no further inputs (and will not until the failed CallRouter is restored to full operation). All components on the failed side of the Central Controller lose synchronization with the other side. The CallRouter, Logger, and Database Manager must all be resynchronized before normal duplexed operation can resume.

For a single-customer ICM, the recovery process begins when the Node Manager notices the failure of a CallRouter process and automatically restarts it. Other processes are not impacted. In a network service provider (NSP) environment, if several ICM instances are running on the same machine, the Node Manager cannot restart the machine. In such NSP environments, manual intervention is required to restart the failed CallRouter process.

The restarted CallRouter immediately initiates a state transfer from its currently executing peer. Each CallRouter sends a message to its Logger. The Loggers then perform their own state transfer.

When the state transfer is completed, all processes are now synchronized. The newly on-line Central Controller sends an in-service status to all local Agents. It then begins processing input messages. After the state transfer, both sides of the Central Controller see exactly the same sequence of input messages. At this point the ICM system is returned to full duplexed operation.

Logger and Database Manager Recovery

Logger recovery is closely linked with central database recovery. In central database recovery, the SQL Server component of the central database is accessed directly through its client interface rather than through proprietary ICM interfaces. Therefore, in addition to synchronization and state transfer, the following database recovery procedures must be performed before the Logger can return to full duplexed operation:

- The Database Manager must run SQL Server automatic recovery.

- The state transfer process may need to update configuration data in the central database (if it was the configuration Logger that failed).
- The Database Manager may need to run historical recovery to recover historical data gathered during the off-line period (if it was the historical Logger that failed).

When a single Logger process fails, ICM software continues to operate with the Logger on the other side. The remaining Logger ensures that output messages continue to reach PGs and Admin Workstations. The ICM's Message Delivery Service detects the failure of the Logger and notifies the PGs and Admin Workstations, which switch their active communication paths to the on-line Logger. At this point, both CallRouters are in service, but only one Logger is available.

For a single-customer ICM, when the Node Manager detects that the Logger has gone off-line, it initiates a shutdown and reboot of the machine. In an NSP environment, the Node Manager does not restart the machine. In this case, manual intervention is needed to restart the failed Logger.

The Logger's Node Manager automatically restarts when the machine reboots. Next, the SQL Server service starts automatically as part of the reboot. *SQL Server automatic recovery* runs to ensure that the returning database is consistent and that all transactions committed before the failure are recorded on disk. Once automatic recovery is completed, the Logger can then go through the application synchronization and state transfer process. If configuration data in the on-line database has changed, the state transfer also updates the configuration data in the returning database. However, in most cases configuration data will not have changed during the failure.

Once the two Loggers are returned to synchronized execution, ICM software may need to recover historical data that was accumulated during the off-line period. This process, referred to as Recovery, is described in the next section, "Database Fault Tolerance".

In a double Logger failure (both Loggers are off-line), the CallRouter continues to route calls. This is possible, even if it is the configuration Loggers that have failed, because the CallRouter loads configuration data in its program memory at system initialization. In a double Logger failure scenario, all messages and data that the CallRouter sends to an off-line Logger are discarded until a Logger is completely recovered.

Database Fault Tolerance

The Central Controller database provides two major ICM functions:

- Permanent storage of the data that describes a call routing configuration.
- Permanent storage for the historical data that is gathered by the ICM system.

Each time a CallRouter starts, it loads configuration data from the central database into its program memory. Once the configuration data is loaded, the CallRouter can begin to route calls (even when the central database is not available). Therefore, when a CallRouter fails and restarts, at least one configuration Logger and central database must be available so that the CallRouter can load the configuration data into memory.

In addition to configuration data, Peripheral Gateways, NICs, and the CallRouter itself all produce historical data. The system components gather historical data and pass it to the CallRouter, which then delivers it to the historical Logger and the central database. The historical Logger passes the historical data on to an Historical Data Server (HDS) facility on a Distributor Admin Workstation.

The ability of the CallRouter to deliver data to the historical Logger and the central database is not necessary for call routing. However, the ICM's monitoring and reporting facilities require both real-time data and historical data from the central database. Database fault tolerance and data recovery, therefore, are extremely important to the reporting functions of ICM software.

ICM Database Recovery

Database recovery is the process of bringing an off-line database up to the same state as an on-line database. In a database device failure (for example, in a disk failure), some manual intervention is required to restore duplexed operation and bring the off-line database up to date. The following scenarios describe what happens in a system failure, a disk failure, and a software failure.

System Failure

When a single Logger, CallRouter, or Database Manager fails (for example, due to a power outage), the associated central database will go off-line. The process of bringing the off-line database back to full synchronization is completely automatic. If the Logger machine reboots, SQL Server automatic recovery runs to ensure that the database is consistent and that all transactions committed before the failure are recorded on disk.

Note If the Logger machine does not reboot, SQL Server automatic recovery is not required.

After SQL Server automatic recovery is completed, the off-line Logger synchronizes its state with the state of the on-line Logger. After the state transfer process takes place, both members of the Logger pair can execute as a synchronized process pair.

During the time that one database is off-line, configuration data may have been added to the contents of the on-line database. If any configuration data changed while one database was off-line, the configuration changes are applied to the database as part of the configuration Logger's state transfer process. This configuration update happens as part of the state transfer *before* synchronized execution begins.

Any historical data that accumulated in the on-line database is recovered *after* synchronized execution begins. Rather than attempting to recover the historical data immediately, ICM software first restores system fault tolerance (that is, duplexed database capability and synchronized execution).

ICM software recovers historical data from the on-line database using a special process called *Recovery.* In Recovery, the historical Logger communicates with its counterpart on the other side of the Central Controller and requests any historical data that was inserted during the off-line period. The counterpart delivers the data over the private network that connects both sides of a duplexed Central Controller.

Disk Failure

A disk failure requires additional steps. If a disk failure disables one side of the Central Controller database, the disk must be repaired or replaced.

Note Contact your ICM support provider if a disk failure occurs.

Your support provider may repair or replace the disk and perform the following steps:

Step 1 Rebuild the database structure from scratch on the new disks.

Step 2 Restore the configuration data, either from:

- A snapshot of the on-line database.
- The most recent backup tape.

- A backup tape taken from the on-line side of the Central Controller database.

At the time of the state transfer, any missing configuration data will be restored. Historical data is restored by the Recovery process, which is run automatically each time the Node Manager process starts on the Logger, or by loading the data from a backup tape.

Software Failure

Cases of software failure that leave a Central Controller database unavailable are handled in the same way as a disk failure (if the failure cannot be repaired by existing software tools). Contact your ICM support provider if such a failure occurs.

Network Interface Controllers

The *NIC* has four physical controllers on each side of the Central Controller. Each of these controllers can simultaneously handle calls from the signaling network. Typically, each physical NIC handles part of the total call routing load for ICM software.

The NIC processes are implemented as non-synchronized process pairs (one process on each side of the Central Controller). The NIC runs as a process on the CallRouter machine.

As a non-synchronized process pair, the NICs operate without knowledge of each other. They are managed by the Node Manager and communicate with other CallRouter processes via the Message Delivery Service (MDS). Figure 2-6 shows how fault tolerance is implemented for various NICs.

Figure 2-6 NIC Fault Tolerance

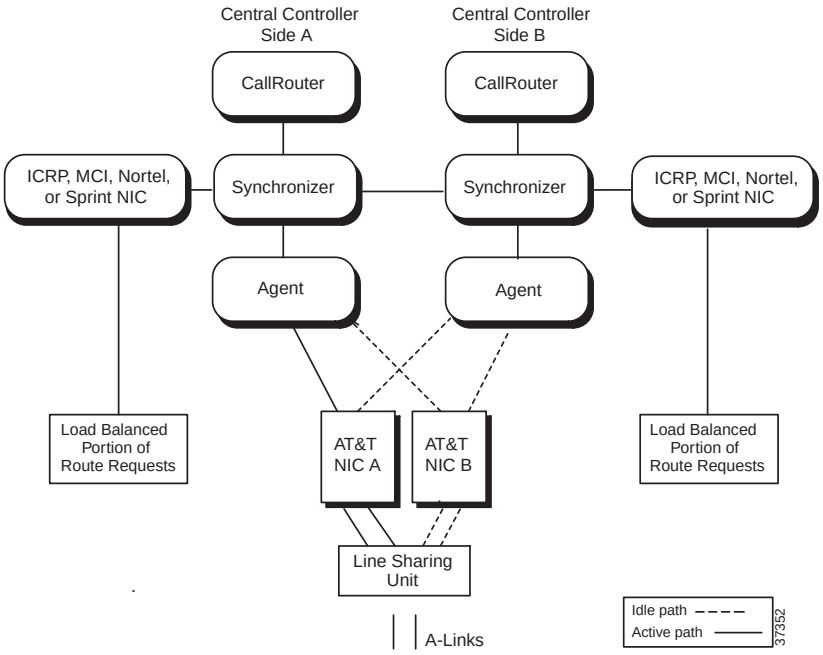

In a duplexed environment, two NICs are on-line and handling routing requests simultaneously. Typically, each NIC handles part of the total call routing load for ICM software. The Synchronizers combine the two input streams to ensure that both sides of the Central Controller receive the same routing requests. If one of the NIC processes fails, or one side of the Central Controller is removed from service, the signaling network detects that communication is no longer possible to that NIC and automatically sends all routing requests to the remaining NIC process.

Peripheral Gateways

Peripheral Gateways use a combination of the synchronization and hot standby approaches to fault tolerance. The Open Peripheral Controller (OPC) operates as a synchronized process pair on a duplexed PG system. The Peripheral Interface Managers (PIMs) typically use the hot standby approach. Figure 2-7 shows how synchronization and hot standby are employed in a duplexed Peripheral Gateway (PG).

Figure 2-7 PG Fault Tolerance

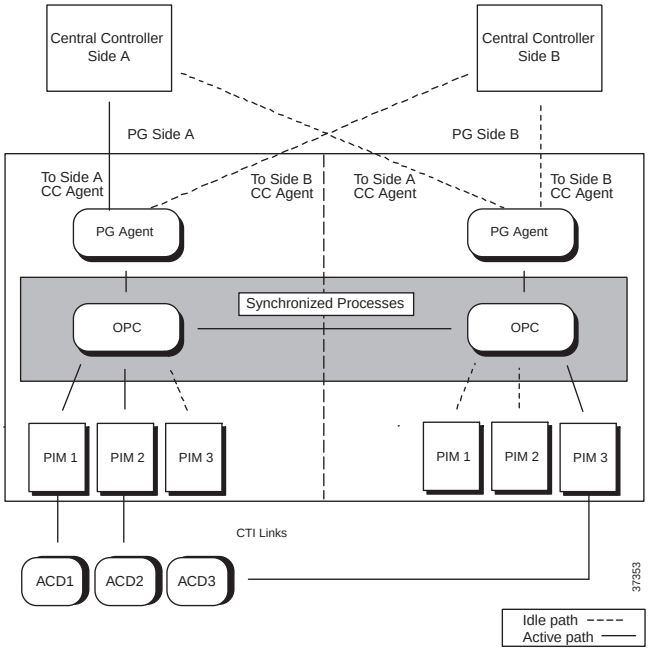

The OPC processes communicate with each other via a private network connection and the Message Delivery Service (MDS). The MDS provides a synchronizer service which combines the input streams from the PIMs and PG Agents on both sides of the PG to ensure that both OPC processes see exactly the same input.

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 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 Figure 2-7 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 Figure 2-7, 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.

Real-Time Distributors

To allow users to monitor current call center activity, ICM software must send up-to-date data to the Distributor Admin Workstation in a reliable and timely manner. The real-time data arrives at the Central Controller from the Peripheral Gateways, which are monitoring activity at each call center. The CallRouter acts as the *real-time server*. The CallRouter for the other side of the Central Controller is the *back-up real-time server*.

Admin Workstations can be located with one or both sides of the Central Controller, at a call center, or at another site. Any site that contains AWs is referred to as an *admin site*.

The CallRouter is responsible for providing real-time data to a Distributor AW at each admin site. Each site has at least one, and usually two, Distributor AWs that serve as real-time data distributors for the site. The primary Distributor AW maintains an active connection to the real-time server through which it receives real-time data.

Client AWs at the site receive their real-time data through a connection to a Distributor AW. These AWs are called Client AWs because they do not have the local database and distributor processes required to receive real-time data directly from the Central Controller real-time server.

If the site has two Distributor AWs, Client AWs are configured to automatically switch to a secondary Distributor AW if the first distributor becomes non-functional for any reason. The secondary Distributor AW also maintains connections to the real-time server; however, these connections remain idle until needed.

You specify whether to install Distributor or Client AWs through the ICM Setup tool.

Historical Data Servers

Historical data is forwarded to the Distributor AW where they are stored in a special database. The distributor then acts as an Historical Data Server (HDS) for the admin site. Admin Workstations at the site query historical data from the HDS rather than directly from the historical Logger.

Two Distributor AWs at a site are set up as HDS machines. Each has its own HDS database. The same fault-tolerant strategy that applies to the real-time Distributor AW also applies to the HDS. That is, when the primary HDS fails, other Admin Workstations at the site automatically switch over to use the backup HDS.

The "Historical Data Server" section in Chapter 4, "Database Administration" provides more information on setting up an HDS AW.

Simplexed Operation

If you have a simplexed Central Controller configuration, you are vulnerable to a single device failure (for example, a system failure, process failure, or a disk crash). Have a strategy in place for keeping daily backups of the central database. Your backup strategy might involve regularly scheduled backups, mirrored disk configurations, or Redundant Array of Inexpensive Disk (RAID) configurations. Always have the central database backed up on removable media.

If the central database becomes unavailable due to disk failure, contact your ICM support provider. A support representative can assist you in replacing the disk, rebuilding the database, and restoring the configuration and historical data.

For more information on database backup and restore procedures for simplexed Central Controllers, see Chapter 5, "General Administration."

Fault Tolerance in Integrated Deployments

Some components in the ICM implement synchronized fault-tolerance, meaning that communication paths to redundant components are utilized simultaneously. This reduces the probability of message loss during transition periods to a very low rate. Cisco Collaboration Server (CS) and Cisco Media Blender (MB) nodes do not implement synchronized fault-tolerance; however, duplexed implementations of these nodes provides a vast improvement over a single path setup. Table 2-1 describes general recovery behavior and its effects when a single node failure occurs in an integrated deployment. Although multiple failures are not considered, an instance of multiple node failure will manifest itself as a superposition of single failure cases. However, a catastrophic failure can occur if all redundant components in an array fail (for example, if all routing CSs fail).

Table 2-1 Integrated Deployment Failure Recovery

Point of Failure	Recovery Action	State Lost	State Recovered	Effect on Web Caller
Routing CS Node	LocalDirector routes to other CS	Current routing sessions	None	Connection fails, must reestablish
Routing MB Node	CSs activate gateway connections to other MBs	None	Current routing sessions requeued by CS	Time delay
Web PG Node	MBs await connection from other PG web side which becomes active	None	Current routing sessions requeued by CS	Time delay
An Agent CS	ICM software routes around this site to another agent site	Current agent sessions at this site, some new sessions may be lost	Other sites not affected	Connection fails, must reestablish CS sessions

Table 2-1 Integrated Deployment Failure Recovery (continued)

Point of Failure	Recovery Action	State Lost	State Recovered	Effect on Web Caller
An Agent CMS[1]	ICM software routes around this site to another agent site	CTI blending and agent reporting at this site	Current sessions continue, unblended	None
An Agent CTI Server	MB connects to the other CTI Server side, if duplexed. Site routing disabled during transition.	Some loss of CTI blending and reporting during transition	Current sessions continue, some loss of blending	Possible loss of web callback.

1. CMS is discussed in Configuration Management Service (CMS), page 3-5.

CHAPTER 3

The ICM Databases

ICM software stores configuration information and call routing scripts in a central database that is part of the Central Controller. You cannot directly alter data in the central database. Instead, you work with a copy of the configuration and script data that resides in the local database of the Distributor AW. When you make changes to these data, (for example, by using the ICM Configuration Manager or Script Editor tools,) the changes are applied to the Distributor AW's local database and automatically to the central database on the Logger.

In addition to safeguarding the integrity of data in the central database, the Distributor AW's local database stores real-time performance statistics. For example, it stores data such as the current Average Delay in Queue, Longest Call in Queue, and the number of Available Agents for each service. This information is updated approximately every ten seconds.

ICM software stores historical performance data in the central database and in a special Historical Data Server (HDS) database on the Distributor AW.

This chapter provides an overview of the ICM databases. In particular, it focuses on the types of data contained in each database and how these data are kept current.

Overview

Figure 3-1 shows the ICM databases and how changes made to the Distributor AW local data are automatically applied to the central database.

Figure 3-1 The ICM Databases

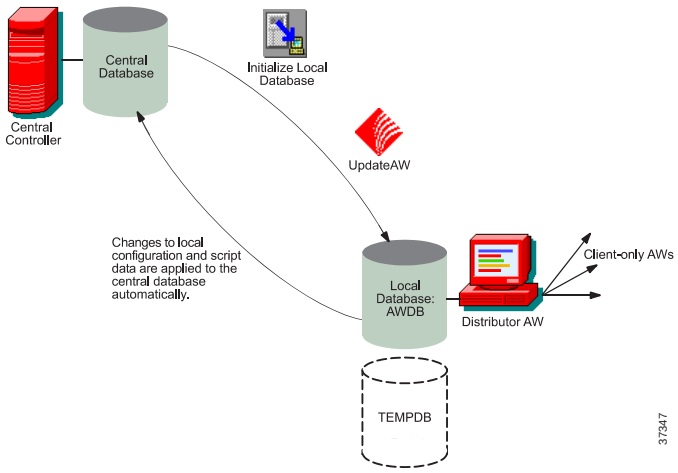

The *Initialize Local Database tool* copies data from the central database to the local database on the Distributor AW. This tool is used to update the local database on the Distributor AW when the AW is first installed. Optionally, you can use this tool to update the local database at any time.

Note Until the Initialize Local Database tool operation is completed, configuration changes cannot be made.

Subsequently, the *Update AW process* forwards to the local database any changes made to the central database. Changes made to configuration data and scripts are automatically copied from the central database back to the local database.

To make access to the real-time data as efficient as possible, the data are stored in memory in temporary tables. These temporary tables actually reside in the TEMPDB database; however, you can access them as if they were tables within the AWDB database.

A Distributor AW that serves as an Historical Database Server (HDS) has a special database to store historical data it receives from the central database. Client AWs can then access historical data from the HDS rather than from the central database.

See “Historical Data Server,” later in this chapter, for more information on the HDS option.

In a network service provider (NSP) environment, a single machine might serve multiple customers. In this case, each customer has its own database. For example, if a Logger machine runs separate instances for each of 10 customers, then it contains 10 central databases. If a Distributor AW needs access to those 10 customers, it needs 10 local databases.

Regardless of the number of databases, only one instance of SQL Server runs on each machine. One set of SQL Server processes maintain all the databases on the machine.

The *Cisco Network Applications Manager Product Description* and the *Cisco Network Applications Manager User Guide* provide more information on multiple customer environments.

Types of Data

ICM software handles three types of data:

- **Configuration data** is stored in the central, HDS, and local databases.

- **Historical data** is stored in the central database and the HDS database.
- **Real-time data** is stored in the local database.

Configuration and Script Data

Configuration data describe your call center enterprise. For example, all of your peripherals, services, dialed numbers, routes, and peripheral targets are part of the configuration data. Configuration data can also include data that has been imported from other systems, such as workforce scheduling data.

In duplexed Central Controller systems, configuration data is kept duplexed on both Loggers. It is always resynchronized when a Logger is restarted.

Script data is also kept on both Loggers. Script data include all call routing and administrative scripts that ICM software uses in call routing (both current and previous versions).

Historical and Real-time Data

Historical data and Real-time data provide information about certain objects in the system such as service, skill groups, and routes.

Real-time data provide current information on these objects.

Historical data fall into four categories: five-minute snapshots, half-hour summaries, call detail records, and events.

The five-minute tables contain *snapshot data*, which are values that are derived from real-time data. Snapshot data provides a view of contact center activity at a particular instant. Since the five-minute values change frequently, they are not synchronized across the central databases of a duplexed Central Controller.

For the tables in which the Historical data and Real-time data are stored, see the *Database Schema Handbook for Cisco ICM /IPCC Enterprise & Hosted Editions*.

Central Database

The central database on the Logger contains the following types of data:

- Full configuration information for the enterprise
- All routing scripts—current and, if you choose to save them, past versions
- Event data
- Call detail data
- Five-minute summary data
- Half-hour historical data

The ICM central database maintains 5-minute summary and half-hour historical data for each:

- Route
- Service
- Skill group
- Trunk group

It also maintains five-minute summary data for Routing Clients and Scripts and half-hour data for Application Gateways and Network Trunk Groups. Although you can view these data in reports, you cannot modify them directly or indirectly.

For specific information on the tables of the ICM databases, see the *Database Schema Handbook for Cisco ICM Enterprise & Hosted Editions* or the on-line *ICM Schema Help*.

The central ICM database resides with the Logger itself. If the Logger is duplexed, each physical Logger has its own copy of the database. If the Logger services multiple customers, each machine has a separate database for each customer.

The name of the central database has the form *inst_sideX*, where *inst* is an up-to-five-character instance name and *X* indicates the side of the central controller (A or B). For example, the central database for cus01 on the Side A Logger is named cus01_sideA.

Database Sizing and Creation

When you first install the Logger software on a machine, you must also create the central database. If you install multiple instances of the Logger software (for multiple customers), you must create a central database for each instance.

Before creating the database, you must determine how much space it requires. The size of the database depends on a number of factors, including the size of your configuration, the expected call load, and how long you want to retain historical data.

To prevent the database from growing indefinitely, old records are automatically purged from the historical tables on a regular basis. You can work with your ICM support provider to decide how long you want to retain historical data in the central database.

For specific information on sizing and creating the database, see Chapter 4, “Database Administration.”

Optimization tips for SQL Server 2000.

SQL Server database files can be configured to grow automatically, using the Autogrow feature. This eliminates many problems that occur when logs or databases run out of space.

In the event that the ICM database in undersized or there is an unplanned increase in the data stored in the database and the ICM processes cannot purge the data to maintain the maximum database size, the autogrow option automatically expands the database when additional space is required.

For more information about autogrow, see the *Microsoft SQL Server System Administrator’s Guide*.

Database Updates and Changes

Each Distributor AW has a copy of the central database configuration data and scripts in its local database. All AWs use the Distributor AW’s local database copy to make changes to configuration and script data. When you change the ICM configuration by using the Configuration Manager, or you create scripts with the Script Editor, you are actually modifying the data in the Distributor AW’s local database. Any changes you make are then automatically applied to the central database.

In the Configuration Manager tool and in the Script Editor, the central database is updated when you perform a Save or Save All operation. Every time you perform a Save or Save All operation, those changes are applied immediately to the central database. Changes to call organization data (call types and schedules) are applied to the central database only when you perform a Save All operation.

As changes are made to the data in the central database, the UpdateAW process copies the changes to all Distributor AWs. This ensures that the local database copy is up-to-date with the central database. The Logger forwards any changes in historical data to the HDS machine.

Distributor AW Local Database

The Distributor AW local database contains the following information:

- Configuration information (copied from the central database)
- Scripts (copied from the central database)
- Real-time data

For specific information on the tables of the ICM databases, see the *Database Schema Handbook for Cisco ICM Enterprise & Hosted Editions* or the on-line *ICM Schema Help*.

Database Creation and Initialization

The local database is created automatically and initialized when you install the Distributor AW software. Its name has the form *inst*_awdb, where *inst* is an up to five-character instance name. For example, the local database for instance ins01 is named ins01_awdb.

Real-Time Data

The real-time client process on the Distributor AW keeps the real-time data in the local database up-to-date. It receives real-time data from the real-time server approximately every ten seconds. Old real-time data is constantly overwritten by new real-time data.

For information on how real-time data is delivered, see Chapter 2, “Fault Tolerance.”

Configuration Management Service (CMS)

In integrated environments, the Configuration Management Service (CMS) coordinates the configuration of objects common to both ICM software and to application instances such as Cisco E-Mail Manager and Cisco Collaboration Server. CMS also authenticates agents when they log in. Common objects are stored both in the application instance local database and the ICM central database. The copy of the data in the ICM central database is used for integrated reporting and for authentication of agents. The copies of the data in the application instance databases is used for application specific reporting and for application operation.

CMS includes the following components:

- A set of client libraries that reside on the application instance system and are called by the application.
- A service process that runs on an ICM Distributor.

The server process reads configuration data from the Distributor database and writes configuration data to the ICM via the Central Controller API (CCAPI).

The CMS client library communicates with the CMS server process using a message bus. Figure 3-2 shows the CMS integrated architecture.

Figure 3-2 CMS Integrated Architecture

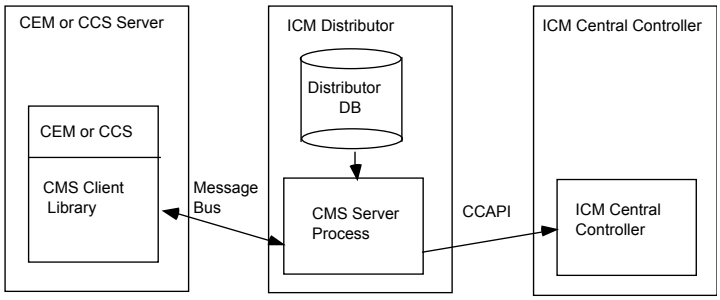

There may be multiple CMS services in a global ICM system, just as there may be multiple Distributors in an ICM system.

Note See the *Cisco Collaboration Server Administration Guide for Cisco ICM/IPCC Enterprise & Hosted Editions*, for information about Cisco Collaboration Server databases. See the *Cisco E-Mail Manager Administration Guide for Cisco ICM/IPCC Enterprise & Hosted Editions*, for information about Cisco E-Mail Manager databases.

Temporary Database

Because real-time data are written and read frequently, the real-time tables are stored in memory as temporary tables. Although these tables physically reside in the temporary database, TEMPDB, you can access them as if they were in AWDB.

Historical Data Server

Admin Workstations need to access historical data (half hour data, call detail, etc.). ICM software normally stores historical data in the central database on the Logger, as well as on the Distributor Admin Workstation that acts as the Historical Data Server (HDS).

One Distributor AW at each admin site is an HDS machine. The Central Controller forwards historical records to the HDS machine for storage in a special local database. Other Admin Workstations at the local site can retrieve historical data from the HDS machine without having to access the central site (see Figure 3-3).

Figure 3-3 Historical Data Server Architecture

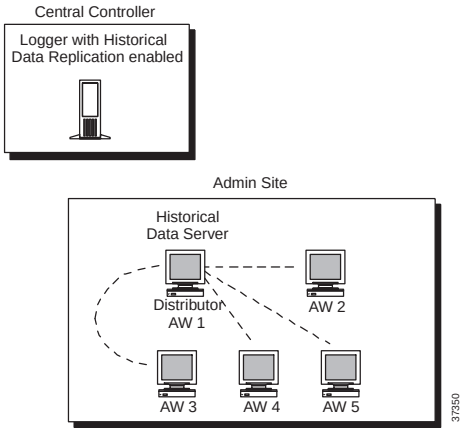

To set up an Historical Data Server, you must configure the Logger to perform historical data replication. You must also configure a Distributor Admin Workstation to be an Historical Data Server. You can use the ICM Setup tool to create an HDS database on a real-time distributor.

For specific information about setting up an HDS database, see Chapter 4, "Database Administration."

Locks

To prevent users from changing the same data simultaneously, ICM configuration data records contain a ChangeStamp field. The ChangeStamp field is incremented when the record is changed in the central database. When an administrator makes changes to configuration data and sends the changed data to the Logger, the value of the ChangeStamp field is compared to the value in the Logger's copy. If the values differ, this means that another administrator has changed the data in the interim. The Logger rejects the change, and the administrator is notified to refresh the data view and try again.

ICM software also provides the Master lock, solely for compatibility with previous ICM releases. The Master lock provides exclusive access to all configuration data and scripts. If a user holds the Master lock, no one can acquire any other locks. You must explicitly acquire and release the Master lock through Lock Admin.Use one of the more specific lock types instead.

The Lock Admin tool (accessible through the Admin Workstations) displays the status of all locks:

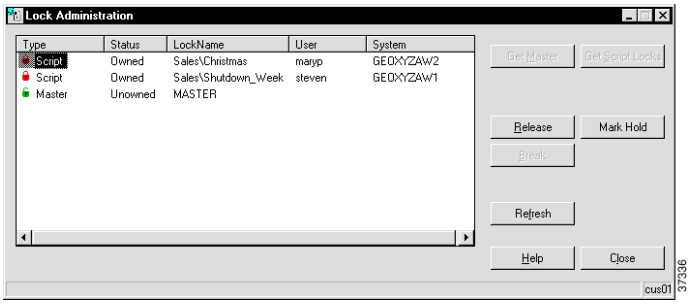

To release a lock:

Step 1 Select the lock by clicking on its name in the Type column.

Step 2 To release the lock, click the **Release** button.

CHAPTER 4

Database Administration

When you install a new Logger, you must create its central database. When you create a database, you must specify the size of its data or log file (or files). The data file (or files) size should be sufficient for all the data you expect the database to hold.

The size of the central database depends on the size of your call center enterprise, your expected call load, and your specific data retention requirements. Based on your expectations and requirements, you can create a central database of the appropriate size.

You must create an HDS database on a real-time Distributor Admin Workstation. The same considerations that affect the size of the central database also affect the size of the HDS database.

Over time, the size of your enterprise or your call volumes may change significantly. Therefore, you may need to resize the central and HDS databases to meet new requirements.

This chapter presents the ICM Database Administration (ICMDBA) tool, that gives you the capability to manage your Logger and Distributor databases.

Overview

When you install a Distributor Admin Workstation, ICM Setup automatically sizes and creates a local database on the machine. Because the real-time data in the local database are constantly overwritten by new data, the database size remains fairly constant. You normally do not need to resize the Distributor AW real-time database. If you do need to resize the Distributor AW database, you can do so using the ICM Database Administration (ICMDBA) tool. The procedures for using the ICMDBA tool are described later in this chapter.

The data in the central database and HDS database grow as they accumulate historical data and call detail records. The growth is directly related to the following factors:

- Size of the ICM configuration; for example, how many services, skill groups, routes, and trunk groups are configured.
- Call rate; that is, how many calls per day ICM software is handling.
- How long historical data is kept in the database.

The amount of configuration data directly affects the amount of historical data generated. ICM software generates a new historical record every half hour for each service, skill group, route, trunk group, etc., configured in the ICM system.

You must size and create the central and HDS databases after installing the ICM software. Use the ICM Database Administration (ICMDBA) tool for estimating the size of these databases, based on the expected usage.

The database size remains adequate as long as your usage is consistent with the values used to size the database. If your configuration expands significantly or if you change the retention times for historical data, you may have to increase the size of the database. This may involve adding additional disks to the system.

Normally, the central database transaction logs are sized to handle the processing of historical data at a call rate of 35 calls per second.

Retaining Historical Data

ICM software initiates a purge process on the Logger and on each HDS AW machine once every day. By default, the purge process runs each night at 12:30 A.M. The purge process deletes records that are older than a specified number of days. When you configure the ICM databases you can specify the number of days to keep data for each historical table.

Table 4-1 lists the *default* settings for retaining historical data.

Table 4-1 Historical Tables

Historical Tables	Default Retention Time
Application_Event, Config_Message_Log, Event	14 days
Logger_Admin, Recovery	30 days
All other historical tables	14 days in Logger, 1095 days in HDS

Database Administration Tool

The ICMDBA tool (icmdba.exe) is included with the ICM software and is located in the \icm\bin directory. This tool provides a central utility that allows you to manage ICM database administration. Use this tool to:

- Estimate size requirements for databases
- Create, edit and delete central databases, local databases, and historical database for installed ICM customers
- Resize database files
- Recreate a database
- Import/export ICM configuration data to/from databases
- View database properties

In addition to these tasks, you can use the ICMDBA tool to start or stop a server, and to do some limited SQL Server configuration.

Note Before using the ICMDBA tool, the ICM software for a customer must be installed. See the *ICM Installation Guide for Cisco ICM Enterprise Edition*, for information on the ICM installation.

Note The ICMDBA Import/Export feature pertians to ICM configuration data only. To import or export ICM historical data, us Microsoft's SQL Server Database Backup and Database Restore utilities.

Starting ICMDBA

Start the ICMDBA by entering the following command in the Windows Run dialog box or command window:

ICMDBA

The ICMDBA main window appears.

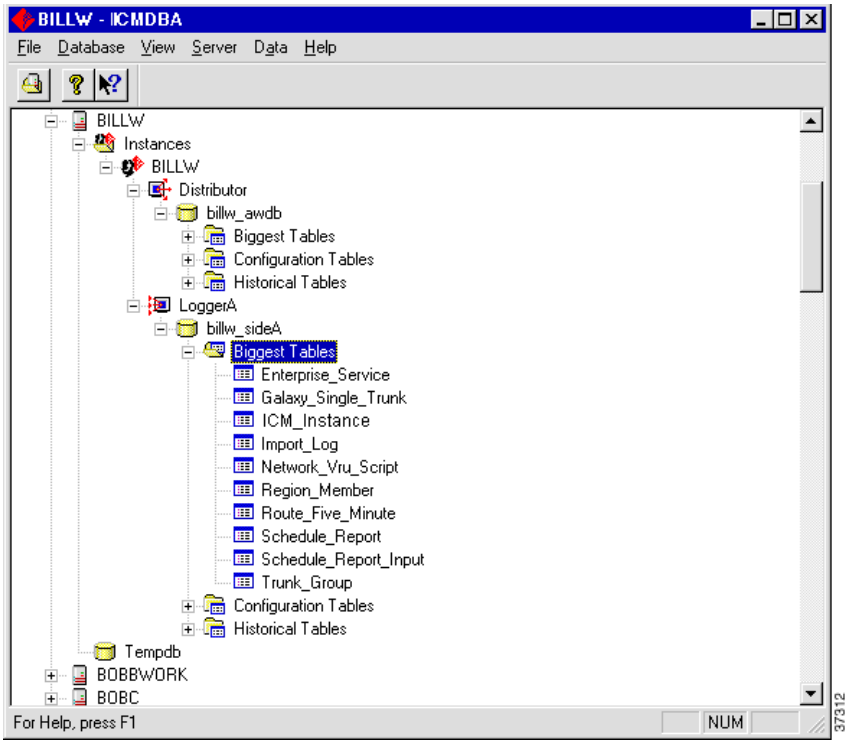

The main window is a tree hierarchy displaying the ICM database servers in the current domain.

Note If you cannot find the server you want in the main window, you can select any computer on your local network by choosing **File** > **Add Computer** from the menu bar.

You can expand the sever by clicking on the plus sign (+) next to its name. This displays the ICM instances that have databases on the server. Expanding the ICM instance displays a specific ICM node or nodes (Distributor and Logger) on machines that have databases for that instance. Expanding the node displays the databases associated with the node. Expanding the node database displays a list of the individual tables in the node database. Under databases are the table groups, and the final level lists the tables in the group.

To view the properties of a table, right-click on the desired table in the list and select Properties from the pop-up menu, or double-click on the table in the list.

There are two ways to access the ICMDBA tool functions. From the main window, choose a node or database from the tree and then select a function from the menu bar menu, or right-click a node or database to display a pop-up menu.

Estimating the Size of a Database

Use the Estimate Database function for the following

- Estimate required database size
- Control the amount of time that data is retained
- Save a database estimate to a file
- Open a previously saved database estimate file

To estimate a database:

Step 1 For the server, instance, and node (Distributor or Logger), select the database that you want to estimate.

Step 2 Choose **Database** > **Estimate** from the menu bar (or click the right mouse button and choose **Estimate**). The Estimate window displays.

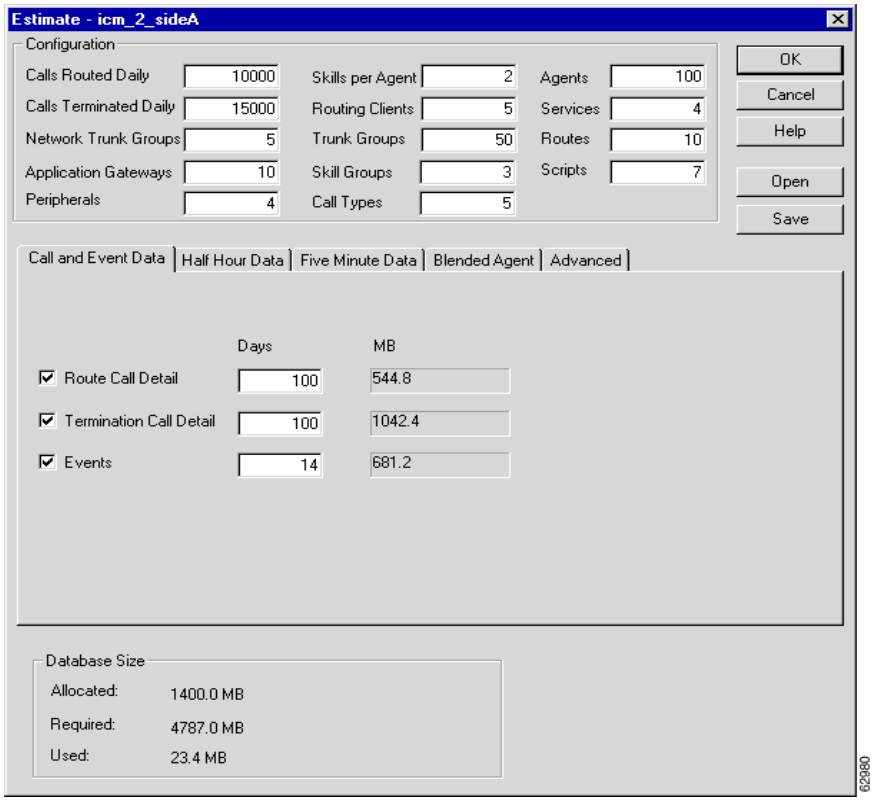

Step 3 Use this window to estimate the size of the database. Use the tabbed sections of the window to configure ICM settings that control the amount of time that data is retained in the database. Use the **Save** button to save the estimate to a file, or the **Open** button to open a previously saved estimate file.

Enter the following information to estimate the size of the database:

Configuration. Use this section of the screen to estimate the ICM variables. Enter your best estimate for each of the fields. The ICMDBA uses these values to estimate the size of the database.

Data. Use the tabbed sections to configure the ICM settings that control the amount of time that data is retained in the database. Changes that you make in this section are applied to the Registry and affect data purging in the database.

- Click on the **Call and Event Data** tab to configure call and event data.
- Click on the **Half Hour** tab to configure half hour data.
- Click on the **Five Minute** tab to configure five minute data.
- Click on the **Outbound Option** (Blended Agent) tab to configure the Outbound Option data, if you are using Outbound Option.
- Click on the **Advanced** tab to configure the Overhead Factor, Average Events Per Day, and Variable Percent Used. By default, Overhead Factor is 2, Average Events Per Day is 10000, and Variable Percent Used is 25.

Overhead Factor. Enter a Value. The value entered in this field affects the database size. This Value is used as a mulitplier for the sum of table sizes. The result is the required database size. Incase the values provided for other fields are not accurate, the Overhead Factor provides the margins.

Variable Percent Used. Use this filed to specify the percentage of maximum length to be used while estimating the required database size.

Database Size. This section displays the actual database size and the required database size, based on the current values entered on the screen.

Step 4 Click the **OK** button to save setting changes to the Registry and exit the screen. A message displays indicating the action was successful.

Warning **The database size estimation algorithm does not factor in customer defined fields in the ECC tables.**

Creating a Database

Use the Create function to create a database for an Admin Workstation or Logger. You can only create one Logger database per side.

To create a new ICM database:

Step 1 With the ICM software running, for the server and instance, select the node (Distributor or Logger) where you want to create the database

Step 2 Choose **Database** > **Create** from the menu bar (or click the right mouse button and choose **Create**). The Create Database window displays.

Figure 4-1 Create Database Window

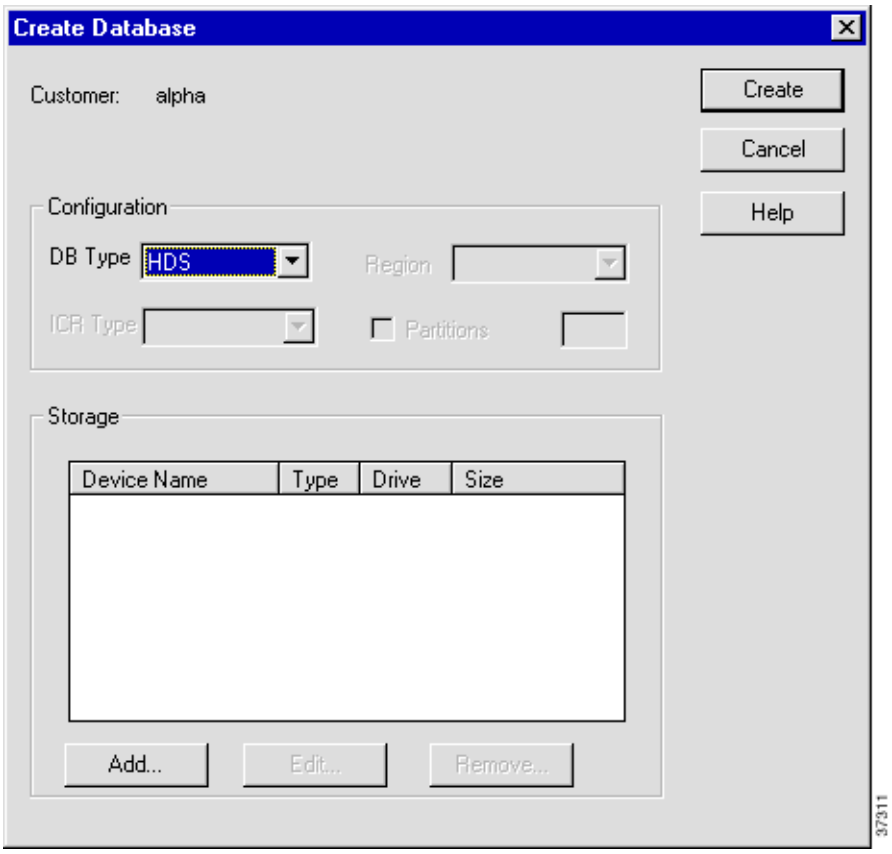

Step 3 Enter the following information for the database:

Database Type. Specify the type of database: HDS (Historical Data Server) for distributor machines, AW for an Admin Workstation local database. For a Logger device, the default database type displays.

ICM Type. Specify whether this is a Standard, NAM, or CICM (Customer ICM) system.

Region. Specify regional information where applicable.

Partitions. If partitioning is enabled, check this box and specify the maximum number of partitions allowed for the customer (1 through 5).

Step 4 Click on the **Add** button. This button invokes the Add Device window.

Use this window to create a new data file and a new log file for the selected database. Specify the disk drive letter and size in megabytes for each new file.

When finished adding the file, click the **OK** button to return to the Create Database screen.

Note By default, the newly created data file will be set to “Automatically Grow”, should it exceed the initially specified size. This setting, as well as the maximum file size, can be modified using SQL Server Enterprise Manager.

Step 5 When you have completed entering information in the Create Database window, click the **Create** button to close the window and create the database.

Note Partitioning is only supported for customers using ICM Enterprise Edition. It is not supported in ICM Hosted Edition, IPCC Enterprise Edition, or IPCC Hosted Edition.

Deleting a Database

Use the Delete function to delete a Distributor or Logger database.

To delete a database:

Step 1 With ICM software running, for the server, instance, and node (Distributor or Logger), select the database that you want to delete.

Step 2 Choose **Database** > **Delete** from the menu bar.

Step 3 The Delete Database prompt displays. Select Yes to delete the database (or No to return to the main window). Another message displays to verify that you want to delete the database. Indicate whether or not to delete the database.

Step 4 Click the **Close** button to exit. Check the main window to verify that the database was deleted.

Expanding a Database

Use this function to add a new storage file.

Note ICMDBA allows a database to be expanded a maximum of 49 times (resulting in 50 segments). In the event that you reach this limit, you must either recreate the database or use SQL Enterprise Manager to modify the database.

To expand database storage on a storage device:

Step 1 For the server, instance, and node (Distributor or Logger), select the database that you want to expand.

Step 2 Choose **Database** > **Expand** from the menu bar (or click the right mouse button and choose Expand).

Step 3 The Expand Database screen displays.

Use the screen to adjust the size allocation on the database's storage device, by completing the following fields:

Component. Specifies whether the file is a data file or log file. Each database must have a file for each type of service.

Available Drives. Specify the drive on which to create the database.

Size. Specifies the size (in MB) of the storage. Field displays a default size. This field may be edited to adjust size as necessary.

Step 4 Click the **OK** button to expand the file and exit the screen.

Recreating a Database

Use this function to recreate a database. The procedure for recreating a database is the same as when you create a database.

Note When you recreate a database, the information currently stored in the database is deleted.

To recreate a database:

Step 1 For the server, instance, and node (Distributor or Logger), select the database that you want to recreate.

Step 2 Choose **Database** > **Recreate** from the menu. The Recreate Database window displays.

Step 3 Enter the database information. Refer to the Creating a Database section of this document for a description of the fields.

Step 4 Click the **Create** button to continue. A message displays asking if you are sure you want to recreate the database. Click **Yes** to continue the operation.

Step 5 The next Recreate Database window displays. Click the **Start** button to recreate the database. When the process is completed, a message displays indicating the action was successful. Click **OK** and then click the **Close** button to exit.

Viewing Database Properties

The ICMDBA tool allows you to view the properties of specified databases.

To view the properties of a database:

Step 1 For the server, instance, and node (Distributor or Logger), select the database that you want to view.

Step 2 Choose **Database** > **Properties** from the menu bar (or click the right mouse button and choose Properties). The Properties window displays.

The screen display includes the following information:

- Customer and database name
- The database configuration
- The size of the data and log files
- The size and percentage full of the combined files

Step 3 When you are finished viewing the database properties, click the **Close** button to exit the screen.

Viewing Table Properties

ICMDBA also allows you to view the properties of each table in the database.

To view the properties of a table:

Step 1 Select and expand the database to display the tables of a database.

Step 2 Double-click on the table you want to view. The properties screen for the table displays.

Step 3 When you are finished viewing the table properties, click on the **Close** button to exit the screen.

Importing/Exporting Data

You can use Import/Export functions to import/export ICM configuation data from one database to another.

Note The ICMDBA Import/Export feature pertians to ICM configuration data only. To import or export ICM historical data, us Microsoft's SQL Server Database Backup and Database Restore utilities.

To import/export data:

Step 1 For the server, instance, and node (Distributor or Logger), select the database from which you want to import/export data.

Step 2 Choose **Data** > **Import** (or **Export**) from the menu bar. The Import (or Export) window displays.

Step 3 Check the **Lockout Changes** box if you want to ensure that changes cannot be made to the database during the import or export operation.

Step 4 Check the **Truncate Config Message Log** box if you want to truncate the Config_Message_Log table in the Logger database.

Step 5 Indicate the path for the source/destination of the data.

Step 6 Click the **Import/Export** button to display the next Import/Export screen.

Step 7 Click the **Start** button to import/export the data. When the process is completed, a message displays indicating the action was successful. Click **OK** and then click the **Close** button to exit. You can click the **Cancel** button at any time to end the process.

Synchronizing Database Data

Use the Synchronize function to synchronize the data of two Logger databases.

Note Whenever an ICM database is restored from a previous backup or the Logger databases are synchronized using ICM config tools, the Verify Sync utility must be run on all connected Cisco E-Mail Manager (EM) and Cisco Collaboration Server (CS) instances prior to performing any configuration via any of the EM or CS instances. Failure to do so may result in an unrecoverable condition for EM and CS instances that may require complete reinstall of EM/CS instances. The Sync utility may not be able to fix any errors but at least manual recovery may be attempted.

To synchronize databases:

Step 1 For the server and instance, select the Logger database to synchronize.

Step 2 Choose **Data** > **Synchronize** from the menu bar. The Synchronize window displays.

Step 3 Check the **Lockout Changes** box if you want to ensure that changes cannot be made to the database during the synchronize operation.

Step 4 Check the **Truncate Config Message** Log box if you want to truncate the Config_Message_Log table in the Logger database.

Step 5 Select the server name and database for both source and target from the drop down lists. To select a server that is not on the drop down list, click the **Add** button and enter the server name in the Add Server box.

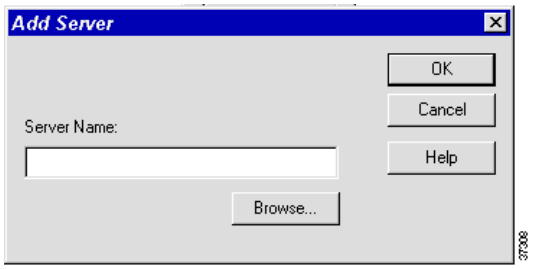

Step 6 Click the **Synchronize** button.

Step 7 A message box appears asking for confirmation. Click **OK** to continue.

Step 8 The next Synchronize window displays. Click the **Start** button to import/export the data. When the process is completed, a message displays indicating the action was successful. Click **OK** and then click the **Close** button to exit. You can click the **Cancel** button at any time to end the process.

Configuring the Server

ICMDBA allows you to start or stop a server and to do some limited server configuration.

To start or stop a server select the node from the list and choose **Server** > **Start/Stop** from the menu bar.

To configure a server:

Step 1 Select the server and choose **Server** > **Configure** from the menu bar. The Configure window displays.

Step 2 Use this window to modify the following SQL Server parameters:

User Connections. Indicate the maximum number of users that may connect to SQL Server at one time.

Locks. Indicate the maximum number of available locks.

Open Objects. Indicate the maximum number of available open objects.

Note User Connections, Locks, and Open Objects are "dynamically allocated" by SQL Server. ICM does not allow you to change these options, so they are grayed out.

Open Databases. Indicate the maximum number of available open databases.

Note Open Databases is not available in SQL 7.0 or SQL 2000.

Memory. Indicates the amount of memory (in megabytes) allocated to SQL Server processing.

Note Memory can be configured to be a specific value instead of the SQL Server default of "Dynamic". Specifying a value of 0 can set the Memory setting to "Dynamic".

Recovery Interval. This setting controls checkpoint frequency.

Max Async ID. Indicates the maximum number of outstanding asynchronous disk input/output (I/O) requests that the entire server can issue against a file.

Note Max Async ID is not available in SQL 2000.

Step 3 When you are finished configuring the server, click on the **OK** button to complete the operation or click on **Cancel** to end the operation without making any changes.

Note When you use the Configure option, the SQL Server, Distributor and Logger restart automatically. However, when you use the Stop option from the Server menu, the Logger and Distributor must be manually restarted from ICM Service Control.

Historical Data Server

To set up an Historical Data Server machine:

Step 1 Run ICM Setup and install the standard Admin Workstation software on the machine. Select the Historical Data Server option.

Step 2 Create the HDS database on the machine.

Note For information about running Setup for an Admin Workstation, see the *Cisco ICM Enterprise Edition Installation Guide*.

Use the ICMDBA tool to determine the size of the database and to create the database. (See "Estimating the Size of a Database" and "Creating a Database" earlier in this chapter.)

When a Database Nears Capacity

ICM software has automatic checks to prevent the central database from becoming full:

Warning message. When the central database begins to approach its capacity, ICM software issues a warning message. By default this occurs when the database is 85% full, but this value can be configured.

Automatic purge. If you select the Automatic Purge option when you install the Logger software, the ICM software automatically deletes the oldest historical data, if it exceeds the retention period, when the central or HDS database nears its capacity. If the data has not exceeded the retention period, it does not get deleted. By default, automatic purge occurs when the database is 90% full, but you can set the percentage when you set up the Logger. You can also set the retention period for data when you set up the Logger.

Note Refer to the *Cisco ICM Installation Guide for Cisco Enterprise Edition* for more on purging information from databases.

The automatic purge ensures that the database can never become completely full. The worst that can happen is you begin to lose older historical data.

Monitoring the Database Size

You should regularly monitor the space used by the central database and transaction logs. You can monitor database size by viewing the Logger's per-process log files. The per-process log files contain specific information on Logger and database activity. The following example shows a per-process event log file for a side A Logger.

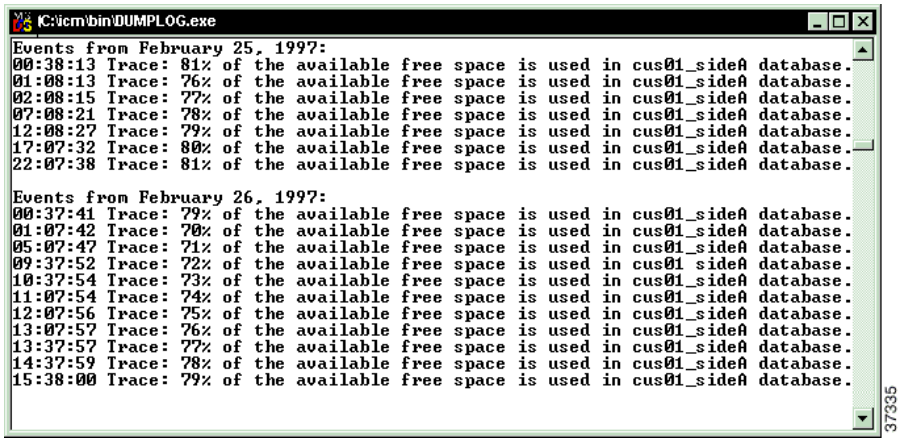

The Logger logs events and trace messages that show the percentage of space used in the database. These files are stored in a \logfiles subdirectory in the Logger's directory (la or lb). You can view the Logger's per-process log files by using the ICM dumplog utility.

When the database becomes 85 percent full, the Logger logs an EMS warning message to the central database. The "85 percent full" warning message might also immediately be sent to your ICM support provider where the appropriate customer support engineer would be notified.

Note For more information on using the dumplog utility, see Viewing Per-Process Log Files, page 6-11.

If you decide that you need additional database space, contact your ICM support provider.

Allocating Additional Space

If the central database is growing too large, you might have to allocate additional space. Your ICM support provider may have options for allocating more space, including:

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

If you require additional space in the central database, you must back up the master database before more space is added.

Note For more information on backing up the database, see Chapter 5, "General Administration."

Initializing the Local Database

It should not be necessary to initialize the local database, since it is done automatically when the AWDB is created. However, if you should ever need to initialize it again after it is first installed, you can do so.

To initialize the local database:

Step 1 Double-click on the Initialize Local Database icon within the Admin Workstation program group of the Program Manager. The Initialize Local Database main window appears.

Step 2 Click the **Start** button to transfer the data. As data are copied, the screen displays the number of rows processed for each table.

Step 3 When the transfer is completed, click the **Close** button to exit.

Troubleshooting

Problem:

Viewing historical data from an AW database does not return the expected data.

Possible Causes:

1. The AW distributor was installed with the HDS option enabled, and the HDS database was then created. This ends up creating historical data views without including the HDS database name.
2. The Logger is partitioned and the AW is not partitioned, or vice versa.

Possible Solutions:

1. In the first case: Delete the AW database. Run the local AW distributor setup.
2. In the second case: Make the Logger and the AW both either partitioned or not partitioned.

Problem:

The "Select into/bulkcopy" option is missing on the AW or Logger database.

Possible Cause:

The database was dropped at some point and not recreated properly. (The proper method would have been to use ICMDBA, which would have set the following default database options during the database creation: Trunc. log on chkpt.; Select into/bulkcopy.)

Possible Solutions:

- In either case, you could recreate this database option — on the AW or the Logger, as appropriate.

or

- If this database option is missing on the AW database:
  a. Delete the AW database.
  b. Run the AW local setup (which recreates the AW database).

or

- If this database option is missing on the Logger database:

  Recreate the database using the "sp_db option" in order to add the "Select into/bulkcopy" option.

CHAPTER 5

General Administration

Because Intelligent Contact Management is a mission-critical application that runs 24 hours a day, ICM software takes care of many routine administration tasks automatically. In general, the ICM software retains control of most of the database administration functions in order to keep external interference to a minimum.

This chapter describes the data integrity checks that ICM software performs on configuration data. It also describes the scheduled database maintenance jobs that run on automatically.

As the ICM administrator, you might be responsible for performing several optional ICM administration tasks:

- Setting networking options
- Monitoring Logger activity
- Backing up the central database
- Restoring the central database from backup
- Comparing databases
- Resynchronizing databases

This chapter describes each of these administration tasks.

Note In order to conserve system resources, minimize all ICM process windows prior to configuring your system.

Built-In Administration

ICM software maintains a database on each side of the Central Controller and a local database on each Distributor Admin Workstation. Each database consists of a group of interrelated tables. As you add or update data in the database, you must ensure that logical relationships are maintained. For example, if you delete a trunk group, you must not leave trunks in the database that reference that trunk group. If you do, the integrity of the database is broken.

Configuration Manager prevents you from making certain changes that disrupt the integrity of the data in the database. However, it cannot prevent all such changes. Usually, if data integrity in the local database is temporarily disrupted, no major problems occur. However, integrity problems in the central ICM database could cause errors in system processing.

Note To protect the integrity of the ICM databases, do not use third-party tools to modify them. These tools do not protect against disruptions of database integrity. (You may use third-party tools to view ICM data.)

When your ICM support provider installs the ICM system, they perform integrity checks to make sure that the database is configured correctly. After that, the integrity of the central database is maintained by the ICM software. You do not need to manually check the integrity of the ICM central database. If you ever have a problem with data integrity in the central database, the problem is most likely a software problem that needs to be addressed by your ICM support provider.

Caution Any manual integrity checks of the central database must involve your ICM support provider. *Do not* run the DBCC CHECKDB procedure on the central database while the ICM system is running. This procedure will stop the Logger.

Optional Administration

You can perform optional administration functions for ICM software such as manually checking data integrity in the local database, monitoring central database space, and viewing a Logger's event logs. These tasks are not required, but you may find them useful in situations when you need to check the system immediately.

Checking Data Integrity in the Local Database

You can manually check the integrity of data in the local database. Configuration Manager provides a Check Integrity option under the Administer menu. Configuration Manager allows you to choose which checks you want to execute.

To check data integrity at any time:

Step 1 Invoke Configuration Manager by clicking on its icon in the Admin Workstation program group.

Step 2 Choose **Configure ICM**> **Administration** > **Integrity Check** from the menu bar. The following dialog box appears:

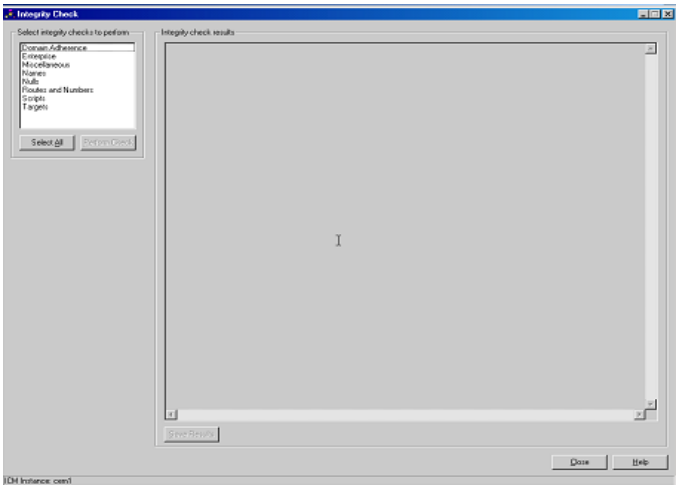

Step 3 Choose specific checks to execute, or choose All to perform all the checks.

Step 4 Click the **Start** button to perform the checks. If any integrity problems are found, Configuration Manager displays a message describing the problems.

Step 5 When you have performed all the checks you want, click the **Done** button to dismiss the Integrity Check dialog box.

The specific data integrity check procedures are listed in Table 5-1.

Table 5-1 Local Database Data Integrity Check Procedures

Procedure	Description
Nulls	Checks for the value NULL in specific fields in the database that must not be null. It also checks that the value of the RoutingClient.PeripheralID is NULL for routing clients associated with a NIC.
Targets	Checks for appropriate relationships among peripherals, targets at peripherals (services, skill groups, agents, and translation routes), trunk groups, network targets, announcements and peripheral targets.
Routes and Numbers	Checks that ID fields cross-referenced from several tables correspond to existing records.
Scripts	Checks for valid cross-references among scripts, call types, and dialed numbers.

Table 5-1 Local Database Data Integrity Check Procedures

Enterprise	Checks for valid cross-references among enterprise services and services, and between enterprise skill groups and skill groups. Also performs several other checks on skill groups, trunks, etc.
Domain Adherence	Checks for valid relationships between agents and skill groups, between skill groups and services, between labels and routing clients, between dialed numbers and routes, and between peripherals and routing clients.
Names	Checks for invalid characters in enterprise names (EnterpriseName field) in various database tables. Enterprise names provide unique character-string names for objects in the ICM configuration.
Miscelleaneuos	Checks rules for Outbound Option Configuration.

For more information on the specific fields checked by these procedures, see the on-line help for the Configuration Manager tool.

Viewing Logger Events

You can view recent Logger activity by viewing the Logger's per-process log files. Per-process log files document events for the specific processes running on a computer. These files are useful in diagnosing problems with processes on the Logger (and on other nodes in the ICM system).

You can also view Logger event data in the central database. The Event Management System (EMS) logs events to the central database. You should be especially aware of Error and Warning events generated by the Logger. For example, ICM software logs a Warning event when the central database becomes 85 percent full.

For more information on viewing the per-process log files and central database event data, see Chapter 6, "Event Management."

Database Networking Support

You can use the SQL Server Setup program to specify which network protocols the database manager supports. Named pipes are the default for Windows; you need not change this default.

For more information about database networking, see the *Microsoft SQL Server System Administrator's Guide*.

Performance Monitoring

When you install ICM software, ICM Setup installs a DLL and sets registry values that enable you to monitor ICM software through the Performance Monitor (perfmon.exe) utility. You can use this tool to monitor ICM software from the local machine or from a remote computer.

The Performance Monitor utility is a standard Windows administrative tool. It graphically tracks one or more variables that you select. You can track variables related to the processor, memory, or various processes and services running on a machine. You can monitor the local machine or choose a remote machine to monitor.

Start the Performance Monitor by choosing **Programs** > **Administrative Tools** > **Performance** from the Windows Start menu. A blank performance chart appears.

Click the Add button ("plus sign" icon) above the blank chart. The Add Counters dialog box appears:

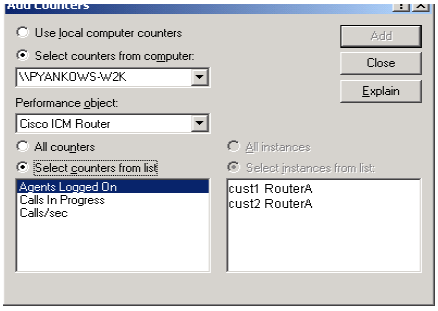

If running from a remote computer, choose the ICM machine in the Computer field.

ICM Router

To chart Cisco router values, choose Cisco Router from the Performance Object drop-down list. The Instance field then lists all the instances running on the machine.

The Counter field displays the values that you can track:

- **Agents Logged On**. This represents the number of agents currently logged on.
- **Calls In Progress**. This represents the number of active calls in progress.
- **Calls/sec**. This represents the number of calls per second.

Choose the instance and value to track and click the **Add** button to add it to the current chart.

ICM QoS PerfMon Objects and Counters

The following is an example Performance Monitor Add ICM QoS Counter window on a router machine.

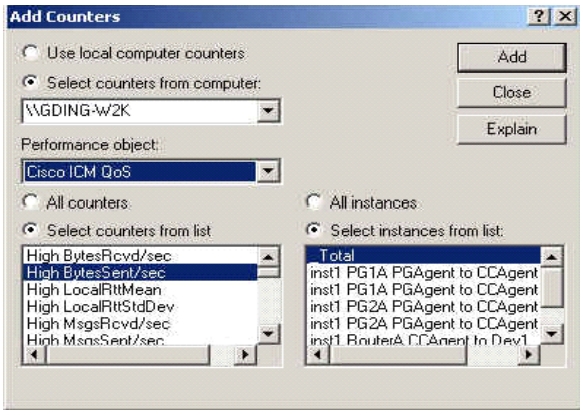

The following is an example Performance Monitor Add ICM QoS Counter window on a PG machine.

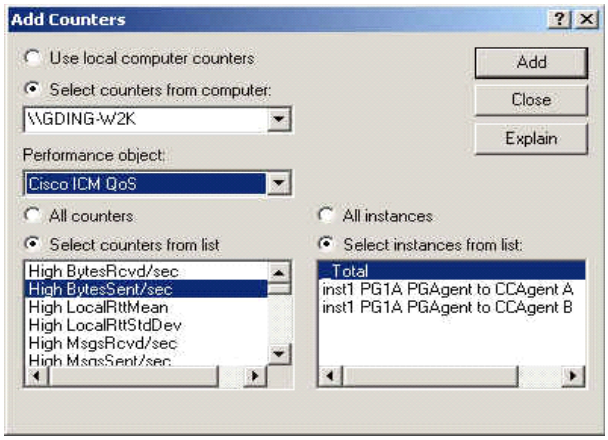

In both windows there is a single performance object, Cisco ICM QoS, which contains all ICM QoS performance counters defined. Copies of the same objects are differentiated in the Instance list. Since performance data is centralized to the Router, PG instances are visible in the Instance list of the Router's Add Counter window. Notice that the Instance list includes a pseudo-instance called _Total. If _Total is selected, each selected counter will contain the sum of the values for all the instances.

Registry Settings and Risks

There are overheads in maintaining ICM QoS counters. The application needs to have a block of memory that stores current counter data, and periodically it must update these values. Furthermore, synchronizing access to the counter values adds serious burdens to the system. For these reasons, the performance monitoring feature is turned off by default. To turn on the feature, change the following registry key value to 1 and cycle the application process.

```
HKEY_LOCAL_MACHINE\Software\Cisco Systems, Inc.\ICM\<instance>\<node>\DMP\
CurrentVersion\EnablePerformanceMonitor
```

System Performance Monitor introduces overheads itself and the overheads depend on the periodic update interval, which is set as the minimum 1 second by default. This interval should be set reasonably high to minimize the impact on the system.

Charting QoS Values

To chart performance values associated with links between the Peripheral Gateway Agent (PG Agent) and the Central Controller Agent (CC Agent), perform the following steps from the Add Counters dialog box:

1. From the Performance Object drop-down list, select **Cisco ICM QoS**.
2. From the Instance list, select the link that you want to chart.

   On the PG agent side, instance names are listed in the following format:

   ```
   <from node name> PGAgent to CCAgent <A or B>
   ```

   On the Central Controller side, instance names are listed in the following format:

   ```
   <from node name> CCAgent to Dev <device id>
   ```

3. From the Counter list, select the counter values that you wish to chart (Table 5-2).
4. Click the **Add** button to add the instances and values that you selected to the current chart.

Table 5-2 Cisco ICM QoS Counters

Counter	Description
High BytesSent/sec	The number of bytes per second sent to the other side over high priority connection.
High MsgsSent/sec	The number of messages per second sent to the other side over high priority connection.
High BytesRcvd/sec	The number of bytes received from the other side over high priority connection.
High MsgsRcvd/sec	The number of messages received from the other side over high priority connection.
High LocalRttMean	The mean round trip time in milliseconds of high priority messages as measured by local node.
High LocalRttStdDev	The standard deviation of round trip time of high priority messages as measured by local node.
High RemoteRttMean	The mean round trip time in milliseconds of high priority messages as measured by remote node.
High RemoteRttStdDev	The standard deviation of round trip time of high priority messages as measured by remote node.
High Xmit NowQueueDepth	The current number of messages in the transmit queue for high priority traffic.
High Xmit MaxQueueDepth	The maximum number of messages observed in the transmit queue for high priority traffic.
High Xmit NowBytesQueued	The current number of bytes in the retransmit queue for high priority traffic.

Table 5-2 Cisco ICM QoS Counters (continued)

Counter	Description
High Xmit MaxBytesQueued	The maximum number of bytes observed in the retransmit queue for high priority traffic.
High TotalQoSReallocations	The total number of times QoS resources had to be reallocated for high priority connection because actual usage has exceeded previous allocation over defined threshold levels.
Med BytesSent/sec	The number of bytes per second sent to the other side over medium priority connection.
Med MsgsSent/sec	The number of messages sent to the other side over medium priority connection.
Med BytesRcvd/sec	The number of bytes received from the other side over medium priority connection.
Med MsgsRcvd/sec	The number of messages received from the other side over medium priority connection.
Med LocalRttMean	The mean round trip time in milliseconds of medium priority messages as measured by local node.
Med LocalRttStdDev	The standard deviation of round trip time of medium priority messages as measured by local node.
Med RemoteRttMean	The mean round trip time in milliseconds of medium priority messages as measured by remote node.
Med RemoteRttStdDev	The standard deviation of round trip time of medium priority messages as measured by remote node.
Med Xmit NowQueueDepth	The current number of messages in the transmit queue for medium priority traffic.
Med Xmit MaxQueueDepth	The maximum number of messages observed in the transmit queue for medium priority traffic.
Med Xmit NowBytesQueued	The current number of bytes in the retransmit queue for medium priority traffic.
Med Xmit MaxBytesQueued	The maximum number of bytes observed in the retransmit queue for medium priority traffic.
Med TotalQoSReallocations	The total number of times QoS resources had to be reallocated for medium priority connection because actual usage has exceeded previous allocation over defined threshold levels.
Low BytesSent/sec	The number of bytes per second sent to the other side over low priority connection.
Low MsgsSent/sec	The number of messages sent to the other side over low priority connection.
Low BytesRcvd/sec	The number of bytes received from the other side over low priority connection.
Low MsgsRcvd/sec	The number of messages received from the other side over low priority connection.
Low LocalRttMean	The mean round trip time in milliseconds of low priority messages as measured by local node.

Table 5-2 Cisco ICM QoS Counters (continued)

Counter	Description
Low LocalRttStdDev	The standard deviation of round trip time of low priority messages as measured by local node.
Low RemoteRttMean	The mean round trip time in milliseconds of low priority messages as measured by remote node.
Low RemoteRttStdDev	The standard deviation of round trip time of low priority messages as measured by remote node.
Low Xmit NowQueueDepth	The current number of messages in the transmit queue for low priority traffic.
Low Xmit MaxQueueDepth	The maximum number of messages observed in the transmit queue for low priority traffic.
Low Xmit NowBytesQueued	The current number of bytes in the retransmit queue for low priority traffic.
Low Xmit MaxBytesQueued	The maximum number of bytes observed in the retransmit queue for low priority traffic.
Low TotalQoSReallocations	The total number of times QoS resources had to be reallocated for low priority connection because actual usage has exceeded previous allocation over defined threshold levels.

ICM Message Delivery Service (MDS) Performance Meters

There are overheads in maintaining MDS meters. The application needs to have a block of memory that stores current counter data, and periodically it must update these values. Furthermore, synchronizing access to the counter values adds serious burdens to the system. For these reasons, the performance monitoring feature is turned off by default. To turn on the feature, change one or both of the following registry key values to 1.

- For the MDS process (required for the **Cisco ICM MDSPROC** and **Cisco ICM MDSPROCCLIENT** performance objects mentioned below):

```
HKEY_LOCAL_MACHINE\Software\Cisco Systems, Inc.\ICM\<instance>\<node>\MDS\
CurrentVersion\Process\EnablePerformanceMonitor
```

- For the MDS client (required for the **Cisco ICM MDSCLIENT** performance object mentioned below):

```
HKEY_LOCAL_MACHINE\Software\Cisco Systems, Inc.\\ICM\<instance>\<node>\MDS\
CurrentVersion\Clients\<client>\EnablePerformanceMonitor
```

where rtr (router) is an example of a <client>.

System Performance Monitor introduces overheads itself and the overheads depend on the periodic update interval, which is set as the minimum 1 second by default. This interval should be set reasonably high to minimize the impact on the system.

The assumption is that extensive metering will not be performed while the system is performing as expected. Only in exceptional cases would close monitoring of the system be desirable.

Charting MDS Values

To chart MDS, perform the following steps from the Add Counters dialog box:

1. Click the radio button: **Select counters from computer**.
2. From the **Performance object** drop-down list, select **Cisco ICM MDSCLIENT**, **Cisco ICM MDSPROC**, or **Cisco ICM MDSPROCCLIENT**.
3. From the Instance list, select that instances that you want to chart:
<cust>, <nodes>, <process>
(_Total is not actually an instance of the object, but a pseudo-instance. If _Total is selected, each selected counter will contain the sum of the values for all the instances.)
4. Click the radio button: **Select counters from list**, and then select the counter values that you wish to chart; or click the radio button: **All counters**, to select all the counter values for charting.
5. Click the **Add** button to add the values that you selected to the current chart.

Table 5-3 and Table 5-4 list the meters that are provided if you select Cisco ICM MDSPROC.

Table 5-5 lists the meters that are provided if you select Cisco ICM MDSPROCCLIENT.

Table 5-6 lists the meters that are provided if you select Cisco ICM MDSCLIENT.

MDS maintains a number of different queues, among which are:

Local Incoming Queue
Remote Output Queue
Local Ordering Queue
Remote Ordering Queue
Timed Delivery Queue

A high, medium, and low priority counter is provided for each of the above queues.

Table 5-3 below is presented in a condensed form to reduce repetition. For the string QQQ you can substitute any of the following values:

LocalHighInQ
LocalMedInQ
LocalLowInQ
RemoteHighOutQ
RemoteMedOutQ
RemoteLowOutQ
LocalHighOrderQ
LocalMedOrderQ
LocalLowOrderQ
RemoteHighOrderQ
RemoteMedOrderQ
RemoteLowOrderQ
TDHighQ
TDMedQ
TDLowQ

where:

High stands for "high priority"
Med stands for "medium priority"
Low stands for "low priority"
In stands for "incoming"
Out stands for "output"

Order stands for "ordering"
Q stands for "queue"
TD stands for "timed delivery"

Table 5-3 Cisco ICM MDS Meters

Counter	Description
QQQ Current Depth	Current number of messages in the queue.
QQQ Now Messages In/sec	Total number of messages added to the queue during last second.
QQQ Now Messages Out/sec	Total number of messages removed from the queue during last second.
QQQ Now Bytes In/sec	Total number of bytes added for all the messages to the queue during last second.
QQQ Now Bytes Out/sec	Total number of bytes removed for all the messages from the queue during last second.
QQQ Now Traffic Intensity	Ratio (x 100) of the **number of messages added** to the **number of messages removed** from the queue during last second.
QQQ Avg. Queue Response Time [ms]	Average time in milliseconds a message waits in the queue.
QQQ 90% Queue Response Time [ms]	The response time in milliseconds that 90% of all messages passing through the queue will experience.

The following meters are also provided for the MDS process.

Table 5-4 Additional Cisco ICM MDS Meters

Counter	Description
Current Buffers Memory Allocated	Total number of bytes used by all currently allocated buffers.
Current Buffers Allocated	Total number of buffers currently allocated from buffer pool.
Buffers Allocation Requests/sec	Number of buffers allocated during last second.
Buffers Free Requests/sec	Number of buffers freed during last second.
Current Buffers Memory Limit	Maximum amount of memory (in bytes) allowed to be allocated for buffers for this process.
Initial Buffers Memory Limit	Amount of memory limit (in bytes) reserved for buffers for this process.
Synch Messages Ordered/sec	Number of messages ordered by the MDS synchronizer during last second.
Synch MDS Duplicates/sec	Number of duplicate MDS messages detected by the synchronizer during last second.
Synch DMP Duplicates/sec	Number of duplicate DMP messages detected by the synchronizer during last second.

Table 5-4 Additional Cisco ICM MDS Meters (continued)

Counter	Description
Output Waits/sec	Number of times output from critical client (Route or OPC) waited for ACK from MDS peer during last second.
Average Output Wait Time	Average number of milliseconds MDS output waits to receive an ACK message from MDS peer.
Private Net Min RTT	Minimum time (in milliseconds) it took MDS to send a message over the private network and receive an ACK response from MDS peer.
Private Net Avg RTT	Average time (in milliseconds) it took MDS to send a message over the private network and receive an ACK response from MDS peer.
Private Net Max RTT	Maximum time (in milliseconds) it took MDS to send a message over the private network and receive an ACK response from MDS peer.

The following meters are provided by the MDS process for each MDS client (statistics as seen by the MDS process, not by the MDS client itself).

Table 5-5 Cisco ICM MDS Meters - client statistics as seen by process

Counter	Description
Client Handle ID	Handle for this MDS client. It is used to uniquely identify the MDS client connected to the MDS process.
Total MDS Client Connects	Total number of times the MDS client has connected to the MDS process
Total MDS Client Disconnects	Total number of times the MDS client has disconnected from the MDS process.
Now Message Received from Client	Number of messages received from the MDS client during last second.
Now Message Sent to Client	Number of messages sent to the MDS client during last second.
Now Bytes Received from Client	Number of bytes received from the MDS client during last second.
Now Bytes Sent to Client	Number of messages sent to the MDS client during last second.
ToClientQ Current Depth	Current number of messages in the Send to MDS Client Queue.
ToClientQ Now Messages In/sec	Total number of messages added to the Send to MDS Client Queue during last second.

Table 5-5 Cisco ICM MDS Meters - client statistics as seen by process (continued)

Counter	Description
ToClientQ Now Messages Out/sec	Total number of messages removed from the Send to MDS Client Queue during last second.
ToClientQ Now Bytes In/sec	Total number of bytes added for all the messages to the Send to MDS Client Queue during last second.
ToClientQ Now Bytes Out/sec	Total number of bytes removed for all the messages from the Send to MDS Client Queue during last second.
ToClientQ Now Traffic Intensity	Ratio (x 100) of the **number of messages added** to the **number of messages removed** from the Send to MDS Client Queue during last second.
ToClientQ Avg. Queue Response Time [ms]	Average time in milliseconds a message waits in the Send to MDS Client Queue.
ToClientQ 90% Queue Response Time [ms]	The response time in milliseconds that 90% of all messages passing through the Send to MDS Client Queue will experience.

The following meters are provided for each MDS client (statistics as seen by the MDS client itself, not by the MDS process).

Table 5-6 Cisco ICM MDS Meters - client statistics as seen by client

Counter	Description
Client Handle ID	Handle for this MDS client. It is used to uniquely identify the MDS client connected to the MDS process.
Now Message Received	Number of messages received by the MDS client during last second.
Now Message Sent	Number of messages sent by the MDS client during last second.
Now Bytes Received	Number of bytes received by the MDS client during last second.
Now Bytes Sent	Number of bytes sent by the MDS client during last second.
Current Buffers Memory Allocated	Total number of bytes used by all currently allocated buffers.
Current Buffers Allocated	Total number of buffers currently allocated from buffer pool.
Buffers Allocation Requests/sec	Number of buffers allocated during last second.

Table 5-6 Cisco ICM MDS Meters - client statistics as seen by client (continued)

Counter	Description
Buffers Free Requests/sec	Number of buffers freed during last second.
Current Buffers Memory Limit	Maximum amount of memory (in bytes) allowed to be allocated for buffers for this process.
Initial Buffers Memory Limit	Amount of memory limit (in bytes) reserved for buffers for this process.
SendClientQ Current Depth	Current number of messages in the Send by MDS Client Queue.
SendClientQ Now Messages In/sec	Total number of messages added to the Send by MDS Client Queue during last second.
SendClientQ Now Messages Out/sec	Total number of messages removed from the Send by MDS Client Queue during last second.
SendClientQ Now Bytes In/sec	Total number of bytes added for all the messages to the Send by MDS Client Queue during last second.
SendClientQ Now Bytes Out/sec	Total number of bytes removed for all the messages from the Send by MDS Client Queue during last second.
SendClientQ Now Traffic Intensity	Ratio (x 100) of the **number of messages added** to the **number of messages removed** from the Send by MDS Client Queue during last second.
SendClientQ Avg. Queue Response Time [ms]	Average time in milliseconds a message waits in the Send by MDS Client Queue.
SendClientQ 90% Queue Response Time [ms]	The response time in milliseconds that 90% of all messages passing through the Send by MDS Client Queue will experience.

Backup and Restore

A database can be lost or corrupted for several reasons, such as:

- Disk drive failure
- Bad media
- Software error (in the Database Manager or elsewhere in the system)

Because you cannot protect against all these conditions, you must have a backup strategy in place. This is especially important if you have a simplexed central database configuration. However, even for a duplexed system, you still need to perform backups to protect against software problems that corrupt both sides of the system.

Database

The following database backup strategies are commonly used:

- Regularly scheduled database backups
- Mirrored disk configurations
- Redundant Array of Inexpensive Disks (RAID) configurations

Although the last two strategies might decrease system performance, they have the advantage of not requiring manual intervention. However, while these configurations protect against disk drive failure and bad media, they might not protect against some software errors.

In a simplexed database configuration, you need to ensure protection against all types of errors. To protect your data, regularly back up the central database to Digital Audio Tape (DAT).

To perform a database backup, use the SQL Administrator tool provided with SQL Server.

Note The SQL Monitor service must be running during a backup. If SQL Server is not configured to start SQL Monitor automatically, you must start the service manually before beginning the backup.

When you restore a database, you can only restore up to the last backup. Any transactions after that backup are lost. Therefore daily backups are recommended for simplexed systems.

Note You must backup the entire database at each backup interval. ICM software *does not* support the use of transaction log dumps as incremental backups.

For general information about developing a backup strategy, including the use of mirrored disks, see Microsoft's *SQL Server System Administrator's Guide*.

For specific information about backing up a database using SQL Administrator, see Microsoft's *SQL Administrator User's Guide*.

Best practices for performing a backup

To backup an ICM database:

- Stop ICM services for the processes that are using the database being backed up.
- Run the backup.
- Restart ICM services after the backup completes.

Note On a duplexed system, calls will continue to be processed if sideA is backed up at a different time than sideB.

Comparing Databases

For diagnostic purposes, you might want to check that two databases have the same data in a specific table. For example, you might want to check that the ICM_Locks table contains the same data on both sides of a Central Controller. The tool dbdiff.exe performs this type of check. Its syntax is as follows:

dbdiff ***database1.table@host1 database2.table@host2***

For example:

dbdiff cust1_sideA.ICM_Locks@geoxyzlgra cust1_sideB.ICM_Locks@geoxyzlgrb

The batch script diffconfig.bat invokes dbdiff for various tables to automatically compare two ICM databases. Its syntax is as follows:

diffconfig ***database1 host1 database2 host2***

For example:

diffconfig cust1_sideA geoxyzlgra cust1_sideB geoxyzlgrb

Resynchronizing Databases

It may occasionally be necessary to repair a corrupt Logger database on one side of a duplexed ICM by copying the Logger database from the other side. You can synchronize the databases using either the DOS Command window or the ICM Database Administration (ICMDBA) tool.

Synchronizing Databases from the Command Window

The following directions explain how to perform this copy from side A to side B for a customer named CustX.

To copy CustX database from Side A to Side B:

Step 1 Stop the Side B Logger, if it is running.

Step 2 In a DOS Command window on the Side B Logger, change to the \icm directory.

Step 3 Run the following command:

install\syncloggers geoCustXlgrA CustX_SideA geoCustlgrB CustX_SideB

Step 4 When prompted, verify that the configuration will be deleted from the correct database and type Y to continue.

When the command completes, you can restart the side B Logger.

Synchronizing Databases with ICMDBA

To synchronize the data of two Logger databases:

Step 1 Start the ICMDBA tool from the Admin Workstation program group.

Step 2 Select a database from the main window and choose **Synchronize** from the Data menu. The Synchronize window appears:

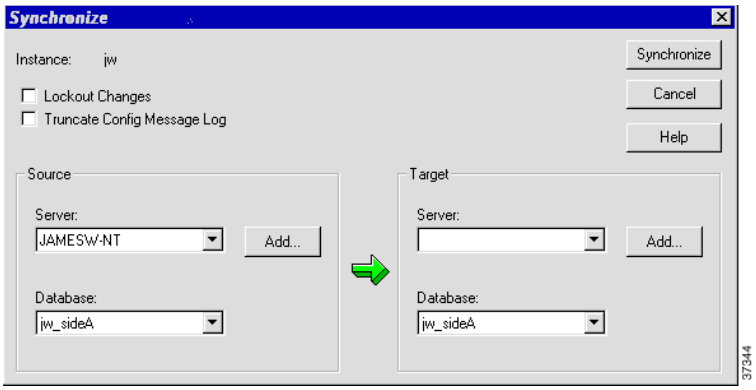

Step 3 Enter the server name and database for both source and target and click the **Synchronize** button.

Note The ICMDBA synchronize process involves dropping the targeted side data and copying the data from the source. For example, if you are synchronizing side B data to side A data, the side B data will be replaced with the data stored in side A.

ICM Time Synchronization

This section describes the components that are involved in to keeping the time-of-day clocks synchronized across all machines that comprise an ICM system.

MDS

The Message Delivery Service (MDS) Synchronizer attempts to keep the system clocks of both sides of a duplexed system synchronized. The enabled Synchronizer is the time master, and the disabled Synchronizer is the time slave. The enabled Synchronizer supplies time messages every half-second to the synchronized application processes as well as to the disabled Synchronizer. To insulate applications against time discontinuities, the time supplied by the enabled Synchronizer is smoothed. If the system clock on the enabled side is changed, the enabled Synchronizer will supply time messages that appear to run 10% faster or slower (as needed) until the MDS time has converged with the system time.

The disabled Synchronizer receives time messages from the enabled Synchronizer every half second, and periodically compares the received time to the system time. In the event of a discrepancy greater than 100 milliseconds, the disabled Synchronizer uses system calls to run the system clock 10% faster or

slower (as needed) until the discrepancy falls within tolerance. The net effect is that the systems clock on the disabled side smoothly tracks the MDS time, which in turn smoothly tracks the system clock on the enabled side.

MDS provides a "Set System Time" message API for setting the time. When MDS receives this message, the enabled Synchronizer sets the system time and the disabled Synchronizer does nothing (since it will track the clock of the enabled side). The "Set System Time" message should be sent by a synchronous process (on both sides of the system), since it cannot be determined which Synchronizer is the enabled one.

VRU PIM

The VRU PIM supports a mechanism for keeping the VRU time synchronized with ICM time. The PIM periodically compares the time reported by the VRU to its own time, and sends a time adjustment message to the VRU if the times differ by more than about 15 seconds.

NAM/Customer ICM

In a Service Bureau installation, there can be multiple INCRP NICs running on the same machine (one per Customer ICM instance). However, only one of these NICs will be running on the ICM instance that is in control of the system clock, and only this NIC is able to affect the time on the CICM machine. This NIC will participate in time synchronization with the Network Applications Managers (NAMs).

While initializing, each INCRP NIC will query MDS to determine whether it can control the clock. Only one NIC on each machine will receive a positive response. This NIC will then periodically (every few hours) send an unsolicited "Time Indication" message to one member of each duplexed CIC pair. The time indication will include the NIC's machine name and current time.

Each duplexed CIC pair may be configured to be a "Time Server", a "Time Client", or "Passive". A CIC that is passive with regard to time synchronization forwards all received time indications to the Routers for reporting purposes. A "Time Server" CIC additionally responds to each "Time Indication" message from an INCRP NIC with a "Set Time Request" message containing its own (MDS) time. (If the NIC and CIC reside on the same machine, as determined by the machine name in the "Time Indication" message, the CIC will not send a "Set Time Request" – this prevents the NAM and CICM from conflicting over control over the system clock.)

A "Time Client" CIC receives "Time Indication" messages and forwards them to the Routers for reporting purposes. In addition, a time client CIC may also receive "Set Time Request" messages. Upon receipt of a "Set Time Request" message, the CIC sends a "Set System Time" message to the Routers, which in turn send a "Set System Time" message to MDS. Participation of the Routers is required to ensure that the "Set System Time" message is delivered to MDS on both sides of the system.

Similar processing occurs when an INCRP NIC receives a "Set Time Indication" from a (time server) CIC: the NIC sends a "Set System Time" message to the Routers, and the Routers in turn send a "Set System Time" message to MDS on both sides of the system. In addition, the INCRP NIC forwards the "Set Time Indication" message to one member of each duplexed CIC pair other than the one from which the original message was received. This serves to keep multiple duplexed NAM pairs synchronized with each other.

CIC

The CIC can be configured with respect to time synchronization. A CIC configured to be a "Time Server" responds to a "Time Indication" message from an INCRP NIC with a "Set Time Request" message. A CIC configured to be a "Time Client" accepts "Set Time Request" messages from INCRP NICs and sends a "Set System Time" message to the Router. All CICs forward CICM time indications to the Routers for reporting.

The CIC Time Manager Type key in the registry is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Cisco Systems, Inc.\ICM\
instanceName\RouterA(B)\CIC\CurrentVersion\Configuration\CIC Time Manager Type
```

To configure a CIC to be a Time Server, set the type to 2. To configure a CIC to be a Time Client, set the type to 1. Setting the Type to 0 will only forward Time Indications to the Router.

Router

The Router attempts to keep the clocks of all controllers (NICs and PGs) synchronized with its own MDS time. It periodically queries each controller for its time. If the time discrepancy between the Router and controller is sufficiently large (15 seconds or more), the Router sends a time adjustment message to the controller instructing it to adjust its time by a delta value. The Router uses the round-trip delay of the query-response to account for transmission delay when computing the time adjustment.

Different controllers handle the time adjustment message in different ways. On a PG, OPC uses the MDS API to adjust the time of the (possibly duplexed) PG. A NIC ignores the time adjustment message, since adjusting the time on the Router machine could have unwanted feedback effects.

The Router records the time skew of all controllers and peripherals and can report these statistics via rttest.

In addition, the Router can optionally be configured via the Registry to designate one peripheral (usually an ACD) as a reference time source. When the Router receives a time update from the named peripheral, it invokes the MDS "Set System Time" API to set the Router time. This effectively synchronizes the Routers and all controllers to the reference time provided by the ACD.

Logger

The Router attempts to keep Loggers synchronized with its own Message Deliver Service (MDS) time. It periodically queries each Logger for its time. If the time discrepancy between the Router and Logger is 15 seconds or more, the Router sends a time adjustment message to the Logger instructing it to adjust its time by a delta value. The Logger then uses this delta value to adjust its time. The Router uses the round-trip delay of the query-response to account for transmission delay when computing the time adjustment.

CHAPTER 6

Event Management

Intelligent Contact Management software tracks events for processes and applications running in the system. An event is any significant occurrence within the ICM system that you might want to know about. Events are recorded on a local and system-wide basis to aid you in maintaining the ICM system.

This chapter provides an overview of event logging and management in the ICM system. It also describes how to use the ICM's event viewing tool.

Overview

Intelligent Contact Management software is a distributed call routing system with components that span several networks. The major components of ICM software generate event data that can be useful in troubleshooting and maintaining the system.

The ICM Event Management System (EMS) logs events from processes throughout the system and stores the event data in the central database. For example, a typical EMS event might record that a system component has been disconnected.

The EMS also saves events from individual processes in per-process log files on the local computer. These files document events for a specific process running on a specific computer.

Several components and processes log events through the EMS:

- Peripheral Gateways
- Network Interface Controllers
- CallRouters
- Loggers

These ICM components are critical to the effective routing of calls in the ICM system. As a system administrator, you need to be informed almost immediately when significant events occur on these components. Admin Workstations also log EMS events, but only to the Application event log. This is because Admin Workstations are not as critical to call routing as the other components of ICM software.

Figure 6-1 summarizes how EMS logs events.

Figure 6-1 Event Logging Overview

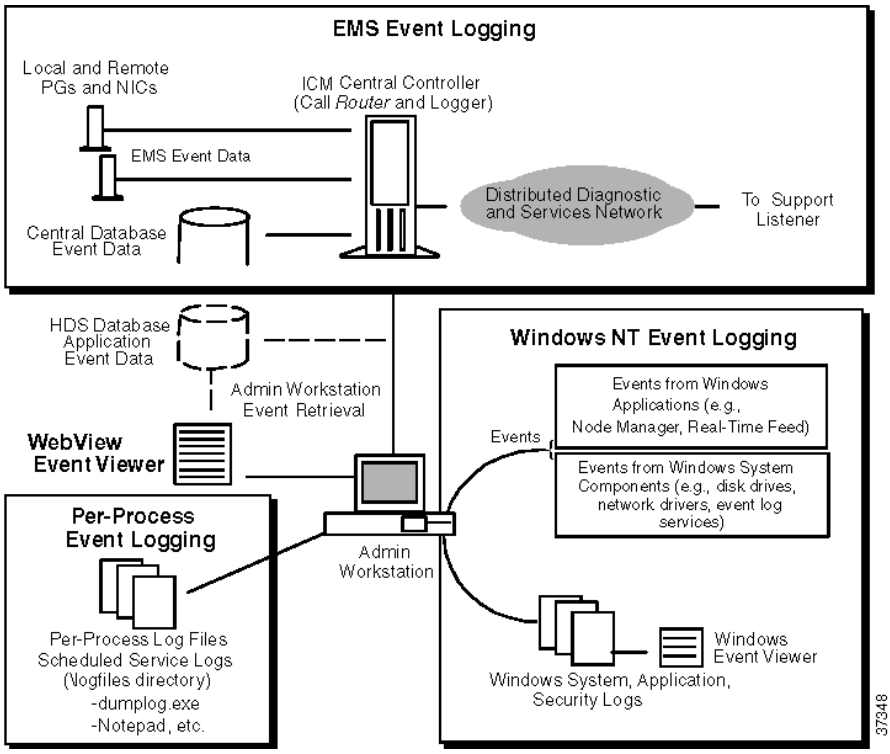

As show in Figure 6-1, event logging in the ICM system involves central and remote system components. The Event Management System (EMS) enables ICM components and the processes that run on them to report events back to the Central Controller. The Central Controller then forwards the events to the Logger for storage in the central database. Events are also forwarded to the HDS database on the Distributor AW. Some of these EMS events may also be forwarded to your ICM support provider's Listener process by the Distributed Diagnostic and Services Network (DDSN).

The DDSN and Listener are described in Chapter 7, "Support Facilities."

ICM software classifies events based on their severity. Table 6-1 lists the severity levels for ICM events.

Table 6-1 Event Severity

Severity	Description
Error	Indicates a significant problem such as a loss of data, incorrect configuration data, or a loss of function. For example, an error would be logged if a Peripheral Gateway were to become disconnected.
Warning	Indicates a potential problem in the ICM system. For example, a warning event might be logged if a user attempted to add a duplicate record to the configuration. Although an event such as this does not cause a loss of function, it is something that you should note.

Table 6-1 Event Severity

Severity	Description
Informational	Documents a successful operation for a major process or application in the system. For example, an informational event might indicate that the peripheral data service was activated on a specific Peripheral Gateway in the system.
Trace	Used for internal testing and diagnostics only.

Trace events are stored in log files, but not in the ICM database.

Event Data Storage

Table 6-2 summarizes the types of events stored in different locations.

Table 6-2 Event Logging Locations

Location	Events	Viewer
Windows event logs	Event data from the local computer. This event data includes EMS Warning and Error events that were generated by ICM processes on the computer.	Windows Event Viewer
ICM per-process log files (.ems)	All EMS events and trace messages logged by processes on the individual computer. The log files are saved in the ICM component \logfiles directory on each computer. For example, on an Admin Workstation the log files are stored in the aw\logfiles directory.	ICM Dumplog utility
ICM command log files (.log)	Status information reported by scheduled jobs. These files are saved in the \logfiles directory along with the per-process log files.	Notepad or WordPad

All computers that have SQL Server also contain SQL Server transaction log files. These files are found under the SQL installation directory on individual computers. You can examine the transaction logs using a standard text editing tool such as Notepad.

For more information on SQL Server log files, see the *Microsoft SQL Server System Administrator's Guide*.

Event Viewing Tools

Viewing event data in the ICM system requires that you use different tools to view the event data that reside in different parts of the system.

You can use the following tools to view event data:

- **Windows Event Viewer**. This tool is part of Windows. Use the Windows Event Viewer to manage event logs for Windows systems.
- **Dumplog.exe utility**. This is a utility for displaying per-process log files at individual ICM computers. You can view the log files on the screen or export them to text files.
- **Notepad or WordPad**. These Microsoft tools can be used to view command log files and any other event log files that have been saved in a text file format.

When to View Events and Log Files

The following guidelines apply to examining the different types of event data collected by the ICM system:

- **Component check**. Use the Windows Event Viewer as needed to examine the Application and System event logs on systems you have identified as having problems. For example, if you notice EMS error events being generated by the CallRouter, you can use the Windows Event Viewer on an Admin Workstation to examine the event data on the CallRouter computer.
- **Process check**. Use the per-process log files as needed to evaluate the specific processes that may be responsible for generating errors. To view these logs, use the Dumplog utility provided with ICM software.

Windows Event Logs

Each Windows computer logs events to its own local System, Application, and Security event logs. You can view event data through the Windows Event Viewer (on the local computer or from a remote computer). Windows computers include CallRouters, Loggers, PGs, and Admin Workstations.

All EMS events that are logged to the central database with an Error or Warning severity level are also logged to the local computer's Windows Application Event Log. This ensures that ICM events are logged at the source and can be viewed locally through the Event Viewer.

Event Log Settings

ICM software requires the Event Log settings shown in Table 6-3.

Table 6-3 Event Log Settings

Log	Size	Wrapping
Application	16384K	Overwrite as Needed
System	16384K	Overwrite Events Older than 7 days
Security	81920K	Overwrite Events Older than 7 days

These values ensure that none of the logs become full. The 1024K setting ensures that large log files can be accommodated in any of the logs. The Application log must overwrite events as needed because it logs EMS Errors and Warnings, application events, and SQL Server events. If it could not overwrite events, the Application log could quickly become full.

Viewing the Event Logs

The Microsoft Windows Event Viewer allows you to view and manage events on a system-by-system basis. You can use the Event Viewer to isolate problems on specific computers. Once you identify an individual computer as generating errors, you can use the Windows Event Viewer to view the computer's local event data. All EMS-generated Error and Warning events are logged to the local computer's Windows Application Event Log.

The Windows event logging process starts automatically each time a Windows system is started. At an Admin Workstation, you can use the Event Viewer to view event data for that computer or for other locally connected computers. For example, you might select a PG or a Logger and view the event data for those computers.

To start the Windows Event Viewer:

In the Administrative Tools group in the Windows Program Manager, double-click the Event Viewer icon. The Event Viewer window is displayed:

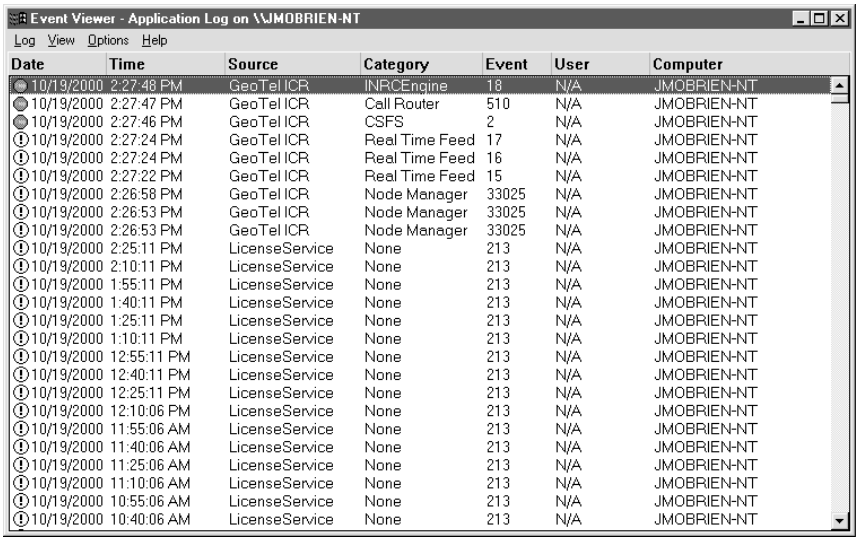

You can change to different logs (for example, the Application, System, or Security logs) by choosing Options from the Log menu.

Windows Logs and Event Types

You can choose between three different logs, depending on the type of event data you want to view. You can view these log files for any Windows computer. The Application log is typically the most useful log since it contains ICM-related events.

- **Application log**. Records events logged by Windows applications (including all ICM applications and processes running on the local computer). For example, when the Node Manager restarts on the local computer, an informational event is entered in the Application log.
- **System log**. Records events logged by the Windows local computer system components (for example, disk drives, network drivers, event log services). For example, the failure of a driver or other local component to load during startup is recorded in the System log.
- **Security log**. Records security events. This log keeps track of changes to system security. For example, attempts to log on might be recorded in the security log.

The event types in the Windows Event Viewer (Error, Warning, and Informational) have similar meanings to those listed earlier in Table 6-1. The Event Viewer provides two additional event types related to system security:

- **Success Audit**. Audited security access attempts that were successful. For example, a user's successful attempt to log onto the system might be logged as a Success Audit event.
- **Failure Audit**. Audited security access attempts that failed. For example, if a user tried to access a network drive and failed, the attempt might be logged as a Failure Audit event.

Viewing Event Data from Other Systems

When you first start the Windows Event Viewer, event data for the local computer is displayed. However, you can connect to another computer in the local network (for example, a Peripheral Gateway) in order to examine its event data.

Note To view events for other computers, you must be logged in as an Administrator.

To connect to another computer:

Step 1 From the Log menu of the Event Viewer, choose Select Computer.

Step 2 In the Computer field, type the computer name of the computer to view. You can also select a Computer name from the Select Computer list.

Step 3 Click the **OK** button. The Event Viewer displays event data for the selected computer.

Per-Process Log Files

The per-process EMS log files are stored in the ICM component \logfiles directory on the local computer as well as forwarded to the central database. For example, per-process log files on Admin Workstations are stored in the aw\logfiles directory. EMS log files have the suffix .ems.

The \logfiles directory also contains the command log file **purgeold.log**. Unlike the per-process log files, you can view purgeold.log directly with a text editor such as Notepad or WordPad.

ICM automatically schedules the command **purgeold** to run nightly. This command removes records over 30 days old from ICM per-process (.ems) log files. Typical **purgeold.log** entries include how many .ems files were found and how many were deleted. purgeold.log is updated each time that purgeold is run.

Naming Conventions

Each per-process log file has a prefix that indicates the process within ICM that generated the event. Each file name includes the date and time the log was created. All log files end with an .ems file extension.

Table 6-4 lists the process names and prefixes and provides brief descriptions of each process. The following example shows the format of a log file name:

PPP_YYMMDD_HHMMSS.ems

The PPP is a prefix that indicates the process. For example, the following log file is for the real-time distributor process. It was created on February 8, 1996 at 9:48:39 A.M.

rtd_960208_094839.ems

The timestamp on a log file is in 24-hour format. For example, 3:00 P.M. is indicated as 15:00; 9:00 A.M is indicated as 09:00.

Table 6-4 Process Prefixes and Descriptions

Prefix	Process	Description	Node(s)
acdsim	ACDSIM	An ICM software process that simulates the functions of an ACD. Used for testing purposes.	AW, Logger, CallRouter, PG
agi	APPGW	The Application Gateway process, which allows ICM software to interact with external host applications.	CallRouter
ccag	CCAGENT	Central Controller DMP Agent. Device Management Protocol Agent that manages session layer communications with ICM nodes.	CallRouter
cic	CIC	The Customer Interface Controller. A process that maintains communication between the NAM, on which it runs, and one or more CICMs.	CallRouter
clgr	CONFIGLOGGER	Configuration Database Logger. Process that stores configuration data in the central database.	Logger

Table 6-4 Process Prefixes and Descriptions (continued)

Prefix	Process	Description	Node(s)
csfs	CSFS	Customer Support Forwarding Service. Receives, filters, and saves appropriate events for delivery to your ICM support provider.	Logger
ctisvr	CTILINK	Computer Telephony Integration server. A PG process that serves as an interface between ICM software and client CTI applications.	PG
dba	DBAGENT	Central Controller Database Agent. Communications process that validates access to the central database.	CallRouter
dbw	DBWORKER	Host Database Lookup. Process that queries external databases and uses that data in call routing.	CallRouter
dcserver	DCSERVER	Rockwell Demand Command Server. Admin Workstation process that provides access to Demand Commands on attached Galaxy ACDs.	Admin Workstation
dtp	DTP	Customer Support Data Transfer Process. Transfers events from the Logger to your ICM support provider.	Logger
edt	SCRIPTED	ICM Script Editor. Tool used to create and schedule call routing scripts.	Admin Workstation
ftp	FTPPROC	File Transfer Protocol. Transfers Rockwell Resource Management Center (RMC) reports to the Admin Workstation.	Logger
hlgr	HISTLOGGER	Historical Database Logger. Process that stores historical data in the central database.	Logger
hsl	HSLTRACE	Northern Telecom High-Speed Link diagnostic tool.	PG

Table 6-4 Process Prefixes and Descriptions (continued)

Prefix	Process	Description	Node(s)
hsltomei	HSLtoMEI	Northern Telecom High-Speed Link and Meridian Event Interface diagnostic tool.	PG
mci	MCI	NIC for ICM communication with the MCI signaling network.	CallRouter
mds	MDS	Message Delivery Service. Process that provides reliable message delivery between ICM processes.	CallRouter, PG
nm	NODEMAN	Node Manager. Process that manages, restarts, and initializes processes on ICM nodes.	Dist. AW, CallRouter, Logger, PG
nmm	NMM	Node Manager Manager. Process that manages, restarts, and initializes the Node Manager process on each ICM node.	Dist. AW, CallRouter, Logger, PG
nic	nic	A special Generic Network Interface Controller (NIC) used in testing. The Generic NIC receives route requests from the ICM's call generator (CallGen).	CallRouter
nortelnic	NTNIC	NIC for ICM communication with the Nortel signaling network.	CallRouter
opc	OPC	Open Peripheral Controller. Interface between the PIM and the CallRouter. Supplies the CallRouter with uniform message sets from different PG types.	PG
pgag	PGAGENT	Peripheral Gateway DMP Agent. The Device Management Protocol Agent that manages session layer communications between the PG and CallRouter.	PG

Table 6-4 Process Prefixes and Descriptions (continued)

Prefix	Process	Description	Node(s)
pim1, pim2, pim3, etc.	varies	Peripheral Interface Manager. The proprietary interface between a peripheral and the PG.	PG
rcv	RECOVERY	Central Database Recovery. Recovers central database historical data.	Logger
rmc	RMCPROC	Rockwell Resource Management Center process. Periodically generates a Rockwell RMC report and places it in a file.	Logger
rtc	RTCLIENT	Real Time Feed Client. A Distributor AW process that receives real-time data from the Real-Time Distributor.	Distributor AW
rtd	RTDIST	Real Time Feed Distributor. A Distributor AW process that distributes real-time data to client-only Admin Workstations.	Distributor AW
rtr	ROUTER	CallRouter. Process that receives call routing requests, determines call destinations, and collects information about the entire system.	CallRouter
rts	RTSERVER	Real Time Server. Process that takes real-time data retrieved from PGs and forwards it to Admin Workstations.	CallRouter
sef	SERIALFD	Serial Event Feed. Provides an alarm feed to an external management station.	Logger
spr	SPR	NIC for ICM communication with the Sprint signaling network.	CallRouter
stentornic	STENTORNIC	NIC for ICM communication with the Stentor signaling network.	CallRouter
tsyp	TESTSYNC	Diagnostic tool.	PG
tsyl	TESTSYNC	Diagnostic tool.	Logger

Table 6-4 Process Prefixes and Descriptions (continued)

Prefix	Process	Description	Node(s)
tsyr	TESTSYNC	Diagnostic tool.	CallRouter
upcc	UPDATECC	Update ICM Central Database tool. Copies data from the local database to the central database.	Admin Workstation

Sample File

The EMS creates a new log file each time a process initializes. This means that messages documenting the end of a process can always be found at the end of a log file; messages documenting the initialization of a process can always be found at the beginning of the log file.

The following is an example of a typical per-process log file:

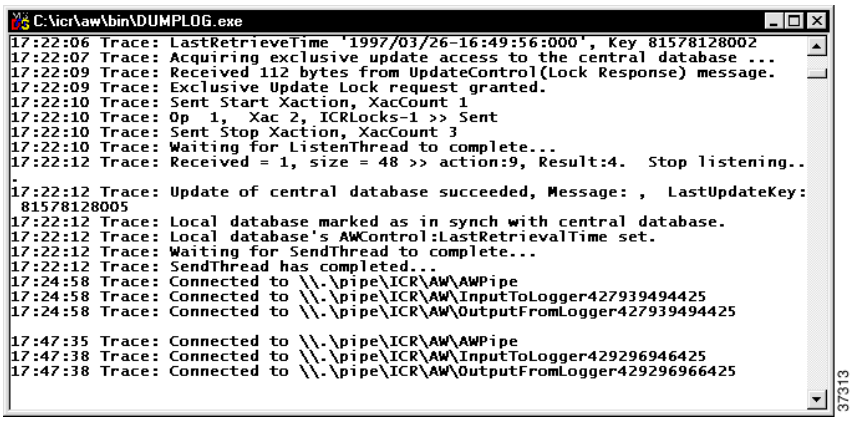

Viewing Per-Process Log Files

You can view per-process log files by using the dumplog.exe command. The dumplog.exe command reads the file, formats the event data, and writes the formatted data to the workstation screen. You can also redirect output to a file using either the /o or /of arguments.

To view per-process log files:

Step 1 Open a DOS Command Prompt window.

Step 2 Change to the \logfiles directory. For example, at an Admin Workstation the directory is icm*instance*\aw\logfiles.

You have several options for viewing log files. The most common option is to display the most recent events for a process on the screen.

To display today's events on the screen, type:

dumplog rtr

This command displays all of today's CallRouter (rtr) events. You can specify any process prefix. You can build on this basic dumplog command by adding date and time arguments.

To dump events for a specific day:

dumplog rtr /bd 1/15/97

This command displays all rtr information that was logged on January 15, 1997 (the begin date, /bd). To see more than one day's log, use the end date (/ed) argument.

The complete syntax for the dumplog command is as follows:

dumplog [ProcessName(s)] [/dir Dirs] [/if InputFile] [/o] [/of OutputFile] [/c] [/bd BeginDate(mm/dd/yyyy)] [/bt BeginTime(hh:mm:ss)] [/ed EndDate(mm/dd/yyyy)] [/et EndTime(hh:mm:ss)] [/hr HoursBack] [/all] [/last] [/prev] [/bin] [/m MatchString] [/x ExcludeString] [/ms] [/mc] [/debug] [/help] [/?]

The specific parameters are shown in Table 6-5.

Table 6-5 Dumplog Parameters

Parameter	Description
ProcessName(s)	Specifies a process prefix from Table 6-4. The command dumps the current day's log for that process, unless you specify different dates or times with other arguments.
/dir Dirs	Specifies the location (directory) of the log files for any processes listed on the command line after the /dir switch. If no /dir switch is used, the current directory is used by default.
/if InputFile	Specifies a specific .ems file to dump. The /if token is optional. If you specify an input file, the /bd, /bt, /ed, /et, /hr, and /all arguments are ignored.
/o	Writes output to a text file in the \logfiles directory. The filename is formed by adding the .txt suffix to the specified process prefix or input file name (without the .ems suffix). The file is written to the current directory.
/of OutputFile	Specifies an output text file; for example, c:\temp\mylog.txt.
/c	Specifies continuous output. The command does not exit after reaching the end of the log. Instead, it waits and writes any further entries that appear in the log.
/bd BeginDate(mm/dd/yyyy)	Specifies the begin date. If specified with /bt, this specifies a range of dates. Otherwise, dumplog dumps events for only the specified date.

Table 6-5 Dumplog Parameters (continued)

Parameter	Description
/bt BeginTime(hh:mm:ss)	Specifies the begin time. Use with /et to specify a range of time.
/ed EndDate(mm/dd/yyyy)	Specifies the end date. Use with /bd to specify a range of days.
/et EndTime(hh:mm:ss)	Specifies the end time. Use with /bt to specify a range of time.
/hr HoursBack	Specifies a number of hours back from the current time.
/all	Displays all information from the specified process's log files.
/last	Displays information from the most recent log file for the process.
/prev	Displays information from the next to last log file for the process.
/m MatchString	Displays only events that contain a match for the specified string.
/x ExcludeString	Displays only events that do not contain a match for the specified string.
[/ms]	Displays milliseconds in time stamps.
[/mc]	Use multiple colors when dumping merged logs. Each process is given a different color.

You must specify either a ProcessPrefix or an InputFile. If you give only a ProcessPrefix value (for example: rtr, nm), dumplog displays the current day's log for that process by default.

To view redirected log files through Notepad:

If you save the log file to a text file (using the dumplog /of argument), you can open the text file from the command prompt by typing:

notepad *filename*

You can also print the file or include it in an e-mail message. To deliver a log file to the your ICM support provider, it may be sufficient to save it as a text file and place it in the Logger's export directory. If used, the Distributed Diagnostic and Service Network (DDSN) would automatically deliver the file to your ICM support provider.

For more information on the DDSN, see Chapter 7, "Support Facilities".

CHAPTER 7

Support Facilities

The ICM's Logger collects events and messages from all components of the system. The Logger can pass this information to a process called the Listener, which can reside at your ICM support provider's facility. Depending on the installation, the Logger may connect to the Listener via a dial-up connection or via a normal network connection.

The facilities that allow the Logger to transfer events and messages to the Listener are collectively called the Distributed Diagnostics and Services Network (DDSN). The DDSN allows Support representatives to remotely diagnose, and in some cases remotely fix, problems in your system.

This chapter also provides an overview of the DDSN and Cisco Support Tools.

The DDSN

For customer sites not connected to the monitoring site via a VPN / LAN or WAN, each computer running the ICM Logger at a customer site is equipped with a modem in order to support the DDSN. The Logger sends data to the Listener through a dial-up connection using the Windows Remote Access Service (RAS) or through a direct network connection. Loggers located at customer premises also allow dial-in or direct network connections. Figure 7-1 shows the basic parts of the DDSN.

Figure 7-1 DDSN Overview

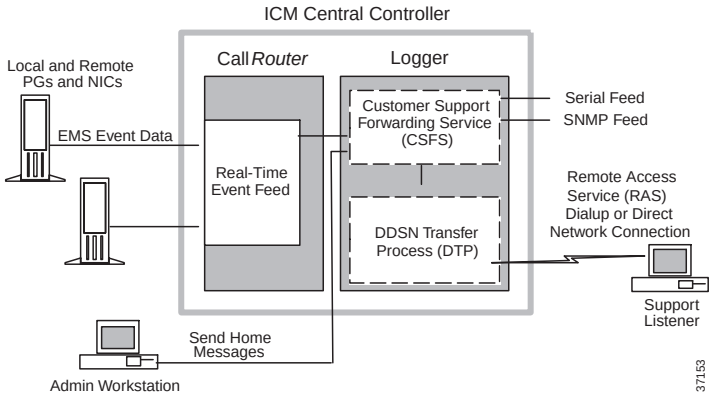

The DDSN Transfer Process (DTP) keeps EMS events in memory until delivering them to the Listener. To minimize the traffic to the Listener (and particularly the number of dial-up connections that may needed over time), messages are batched and sent periodically. However, if the DTP receives a high

priority event, it immediately sends the event to the Listener. If an attempt to establish a RAS connection fails because of a busy phone or no answer, the DTP process periodically tries to re-establish the RAS connection.

You can place exported log files (for example, .txt files) in the export directory on the local machine.

Every 30 minutes, the DTP checks to see if there are EMS events in memory or any new files in the Logger's export directory. When there are new events and files, the DTP sends the events and files to the Listener, establishing a RAS connection, if necessary. Any files sent to the Listener are then deleted from the Logger's export directory.

Error Reporting

The ICM Logger immediately informs the Listener of any significant errors or unexpected conditions it encounters. Error reporting is handled by two processes on the Logger:

- **Customer Support Forwarding Service (CSFS)**. Receives events, filters them, and holds then in memory.
- **DDSN Transfer Process (DTP)**. Transfers the events and export files to the machine running the Listener. It uses either a dial-up connection and the Remote Access Service (RAS) or a direct network connection. The Listener stores the events in a customer-specific directory on its machine.

ICM software sends two types of data to the Listener:

- Event information generated by any process within ICM software.
- Export files placed in the Logger's export directory.

The event messages received by the Listener include information about when and where the error occurred and the full message as reported on the event feed.

File Transfer

You can transfer any file to the Listener by copying it to the Logger's export directory. For example, you might transfer a per-process log file that you exported to a text file (.txt). DTP automatically transfers the file to the Listener during the next transfer cycle. At the Listener machine, the file is held in a customer-specific directory.

Support Processing

When your messages arrive at the Listener, they are stored in a customer-specific directory. For error messages, appropriate Support representatives receive automatic and immediate notification. Representatives assigned to a specific customer are notified of all error messages from that customer. Representatives assigned to specific areas of the ICM product are notified of all error messages related to their areas.

Serial Alarm Feed

ICM software provides an optional serial alarm feed that allows you to establish your own alarm/event links to the DDSN. The Serial Alarm Feed process (SERIALFD) uses the Customer Support Forwarding Service (CSFS) to communicate alarm information to an external system. The Serial Alarm Feed process receives events and sends alarms in ASCII format to a communications port on the Logger. Once the SERIALFD process is started, alarm messages are sent to the communications port as they occur.

The Serial Alarm Feed consists of a series of alarm messages that are sent out over a 9600 baud serial connection. The Alarm Messages are formatted as shown in Table 7-1.

Table 7-1 Alarm Message Format

Meaning	Example
Trap Number	6
System Name	CSOXYZRTRB
System Type	2
Process Name	rtr
Trap Severity	6
Date (format: YYYYMMDD)	19961219
Time (format: hh:mm:ss)	16:08:51
Number of Optional Arguments Following	1
1st Optional Argument	pim1
Description	Restarting process pim1 after having delayed restart for 60 seconds
End of message sequence (0xD, 0xA)	[CR][LF]

Note that all the fields in Table 7-1 are delimited by a single SPACE character. All fields are variable length.

You can find information about specific traps in the ICM Management Information Base (MIB). The MIB correlates to the driving table used by the Customer Support Forwarding Service (CSFS). You can look up each trap number in the MIB to see the descriptions and appropriate ASN.1 syntax used to generate the SNMP traps.

Note For more information on SNMP Feeds and ICM MIB, see the *SNMP Guide for Cisco ICM/IPCC Enterprise & Hosted Editions.*

You typically see alarms from the following sources:

- Nodes
- Processes
- Connections
- Peripherals
- Sessions/Links

Since the Serial Alarm Feed is an alarm process, only events that have triggered a state change in an object are forwarded to the communications port. All other events are discarded. For example, if a process stops, an alarm is generated and forwarded to the communications port. All subsequent alarms indicating that the process has stopped are discarded. When the process restarts, another alarm is generated. The latest alarm indicates a state change, so it is forwarded to the communications port.

Syslog Compatible Feed

The ICM system supports the Syslog event reporting mechanism for CiscoWorks 2000. If you are using CiscoWorks 2000 for monitoring other Cisco products, you can optionally add the ICM system by configuring the ICM Logger for CiscoWorks 2000 support. Please refer to the CiscoWorks 2000 documentation for details on how to add the ICM system as a managed device.

Figure 7-2 shows an example CiscoWorks 2000 Syslog ICM report.

CiscoWorks 2000 Syslog event reports show the EMS event data in a web browser.

Figure 7-2 CiscoWorks 2000 Syslog Display for ICM

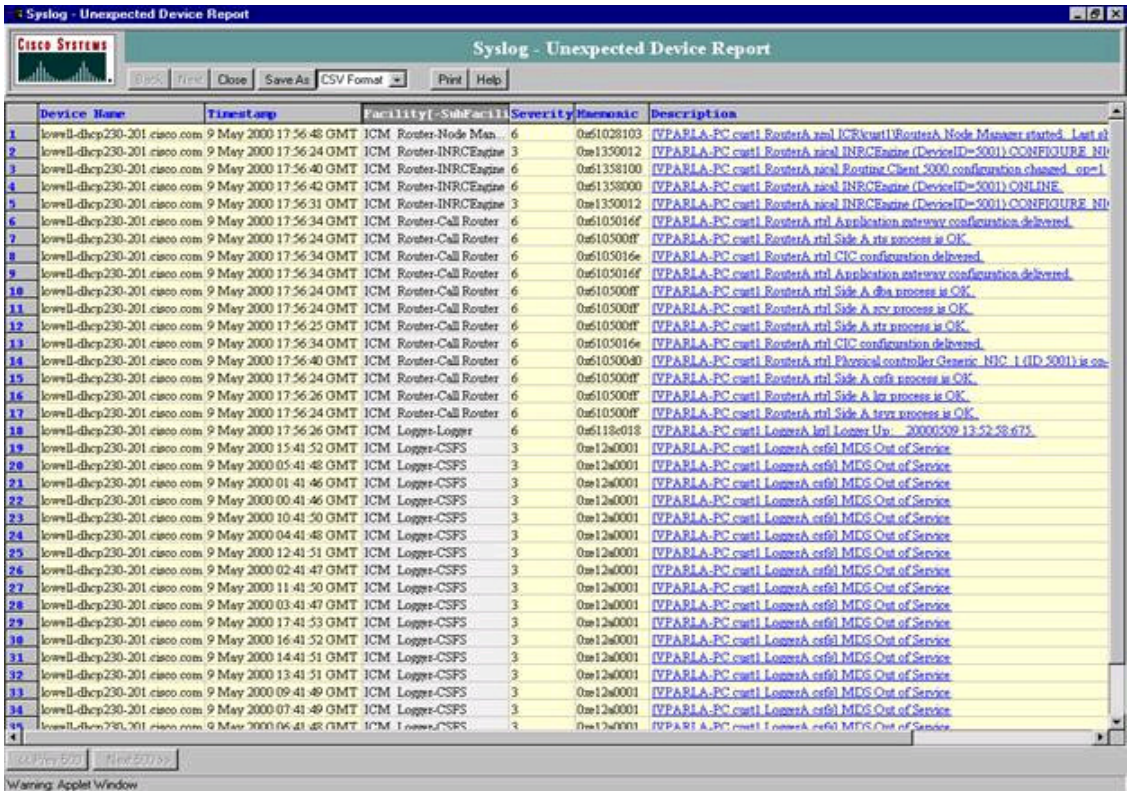

Cisco Support Tools

The Support Tools suite includes the full set of standard diagnostic tools delivered with earlier ICM versions. It also provides key new functionality including:

- The ability to interrogate individual Support Tools nodes for their hardware/OS, Cisco component, and third party product information.
- The ability to view, stop, and start services running on Support Tools nodes.
- The ability to view and terminate processes running on Support Tools nodes.
- The ability to compare and synchronize registry settings from different Support Tools nodes.
- The ability to pull logs from most Support Tools nodes including ICM CallRouters, Loggers, Peripheral Gateways (PGs), Admin Workstations (AWs), CTI Object Server (CTI OS), Cisco Collaboration Server, Cisco E-Mail Manager, and Cisco Media Blender, as well as from Cisco CallManager.
- The ability to create enhanced time-synchronized merged logs across servers.

CHAPTER 8

ICM Partitioning

This chapter discusses the Intelligent Contact Management (ICM) Partitioning feature, which controls what data individuals are allowed to access within an ICM database.

Note Partitioning is only supported for customers using ICM Enterprise Edition. It is not supported in ICM Hosted Edition, IPCC Enterprise Edition, or IPCC Hosted Edition.

Note Partitioning has some limitations. Read the section on Problems with Partitioning, page 8-47,- before you enable partitioning.

ICM Partitioning Overview

People often equate the word *Partitioning* with the computer management utility that logically divides a hard drive into sections to improve data storage. The ICM Partition feature **does not** split the database into sections; it only **controls access** to the ICM database.

The ICM Partitioning feature is used to control the data that individual ICM administrative users are allowed to see within the ICM database. This control is enforced in the standard ICM Administrative Workstation tools, as well as WebView and Script Editor.

Other methods of protecting data—such as through firewalls, encryption, and so on—are not described in this chapter.

Why Use ICM Partitioning?

The data stored in the ICM database describes the ICM enterprise. In the simplest case, all users of ICM software have access to all the data in the enterprise. However, there are several reasons why you might want to restrict access to specific data:

- To limit the users who can make changes that affect call handling or monitoring.
- To restrict the users who can see sensitive data.
- To allow separate divisions to act independently without interference.

Depending on the data, you might want to limit a user's access to scripts, routes, peripherals, services, enterprise services, skill groups, and so on. For example, one administrator might have access to only the sales services while another administrator might have access to only scripts. You may also need one or more business entities, depending on what data you need to segregate.

The optional ICM Partition feature allows you to apply these types of security measures to ICM software.

Classes and Objects

You can grant access to broad classes of data or to specific objects within the ICM enterprise. A *class* represents a group of *objects*; an *object* represents a specific *element* and its *related data*.

An object might control other objects. In turn, each object and controlled object has a group of ICM database configuration tables it is associated with. Figure 8-1 illustrates this hierarchy.

Figure 8-1 Class and Object Hierarchy

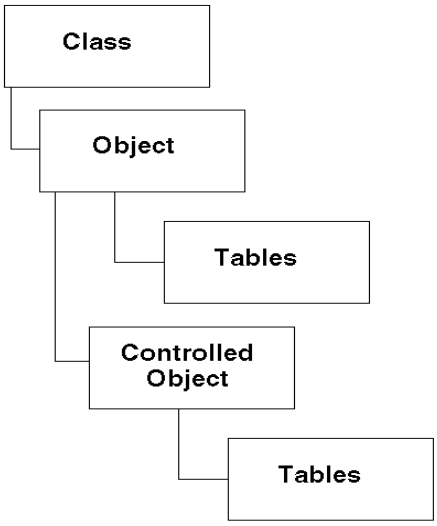

For example, if you grant a user access to the Peripheral *class*, that user can access the configuration data for *all* peripherals and all the data associated with each of those peripherals in the enterprise. On the other hand, if you grant a user access to a peripheral *object*—the Scranton ACD, for example—that user can access only the configuration data for that *specific* peripheral and its related data (trunks, services, skill groups, agents, etc.).

By selectively granting access to specific classes and objects, you can ensure that each user has the full access he or she needs without allowing unnecessary or unwanted access to other data.

Mapping Objects

Most objects have a direct mapping between the database security object and a configuration item. For example, the Agent database security object directly maps to Agents created through the Agent Explorer or the Agent Bulk Configuration Tool. However, some objects do not have a configuration object, but rather, serve only as a mapping device between the Class and Object levels.

Note This mapping is required because Classes do not have a direct association with tables.

Table 8-1 lists these mapping objects:

Table 8-1 Special Mapping Objects

Class	Objects
Call	Call
Network Interface	Network Interface, Network/Peripheral
Peripheral	Peripheral Global, Network/Peripheral
System	System

Some objects intersect the Network Interface Class and the Peripheral Class, where access levels can be assigned to the object either through the Network Interface class or the Peripheral Class. The Network/Peripheral object exists for such objects.

For example, a Dialed Number object requires an association to a Routing Client. If the Routing Client is associated with:

- A *Network Interface Controller*, then access to that the Dialed Number comes from the Network Interface Class.
- A *Peripheral*, then access to the Dialed Number comes from the Peripheral class.

For specific information about the classes and objects recognized by ICM software, see Class and Object Security, page 8-8.

Business Entities

A business entity is a subset of the ICM software enterprise and is an object in the ICM database. Once you create business entities, then you can define your own set of objects that belong to the business entity objects, such as:

- Routing and administrative scripts
- Enterprise services
- Enterprise skill groups
- Enterprise agent groups
- Enterprise routes

By default, the ICM software enterprise consists of only one business entity. However, if the ICM Partition feature is enabled, you have the option to logically divide the ICM enterprise into several business entities. For example, in a large corporation, you might create business entities to represent specific divisions.

Note The number of business entities on an ICM system must be less-than-or-equal-to the maximum number of Partitions (five). The number of Partitions is defined using the ICM Database Administration (ICMDBA) tool. For more information, see Creating a Database, page 4-5.

You can limit the access of individual users and user groups to specific business entities. For example, you might grant a system manager within one business entity the privileges to create and modify routing and administrative scripts for that business entity. However, this same manager might not have any access to the scripts of another business entity

For more information about setting the security access for business entities, see Installing and Configuring ICM Partitioning, page 8-22.

Access Privilege Levels

For each class and object in the ICM software database, you can grant users or groups a specific access privilege level. The access level determines what rights the user has to the associated data, as described in Table 8-2.

Table 8-2 Access Levels

Access Level	Description	Example: Peripheral Access
Maintenance	Permits the User or Group to read, update, and delete the object.	Allows the user to create, modify, or delete the services, skill groups, etc., for a peripheral.
Reference	Permits the User or Group to read the object and use it in a script.	Allows the user to reference peripheral-level variables in a script.
Read	Permits the User or Group to see the object, but not change it or use it in a script.	Allows the user to see the peripheral and the associated peripheral real-time and historical data
No access	Restricts access to the object. (This is the default if a User or Group is not explicitly assigned an access level.)	Allows no access.

Not all access levels can be applied to all classes and objects. For example, a user can only have *Read* or *Reference* access to call detail data; ICM software does not permit *Maintenance* access to call detail data.

You can intermix different levels of security. For example, you might choose to give some users Read access to a wide range of data, but grant them Maintenance access to only a subset of that data.

Note The highest—that is, the most *permissive*—access level to a particular piece of data “wins.”

ICM Partitioning Security

The security provided by Partitioning involves checking the user's access privileges.

User Privileges

Security settings can be assigned directly to a user. A user account is created on the Windows domain.The system administrator then defines the access rights of the user to different objects in the ICM database tables.

Note Non-system administrators can assign security settings if they have maintenance access to the ICM system class and maintenance access to the object for which they wish to assign access.

Users can be granted access to data by their access level to the:

- Class to which the data belongs.
- Data object

A user's access level to data is determined both by the user's access privileges and the access privileges of the group(s) to which the user belongs.

User Groups

To simplify security administration, you can define user groups and assign users to these groups. A user group is a collection of ICM users that exists only in ICM, not in the domain. You can grant to each group the appropriate set of rights within the system for the tasks that they will need to perform.

For example, you might want to create groups for the following:

- Users who can make changes to the carrier interfaces
- Users who can add and remove peripherals
- For each peripheral, users who can change the configuration within that specific peripheral

You can define any number of groups with broader or narrower rights than in these examples. In addition to granting rights to user groups, you can also grant specific rights to individual users. However, it is usually simpler and easier to use groups as much as possible.

For more information about defining the security access for user groups, see Installing and Configuring ICM Partitioning, page 8-22.

How Partitioning Works

The ICM Partitioning feature restricts access to objects in the ICM system by restricting users from directly accessing the raw data in the ICM system. Instead, users are granted access to database defined views. These views are defined by SQL queries which only select the data to which the given user is granted access.

In order to populate these views from the raw data in the database tables, the ICM must keep track of which users are granted access to which data, and make the data easily accessible from a SQL query. This is done by generating a table (User_Security_Control) from the security definitions that has a single

row (on the order of 100 bytes) for each object to which each user has access. The data in this table is then joined with the raw data in the database tables to generate views that are specific to each user and only contain the data that user is permitted to see.

The User_Security_Control table could get very large. If a given database has 20 users and 10,000 objects the resulting User_Secuirty_Control table could contain as many as 200,000 rows. If a given user does not have access to a given object, no row appears in the table, therefore (Number of Users * Number of Objects) is an upper bound on the table size. Rows in most (but not all) ICM configuration tables count as objects. See Table 8-6 for a comprehensive list of configuration object types.

It is important to note that the number of objects to which a given user has access is not just the number of objects to which they are granted explicit access; this number also includes those objects to which the user implicitly has access. In other words, if a user is granted access to an entire Business Entity object, that user will not only have a single row in the User_Security_Control table for the Business Entity object, but will also have one row for each object that falls under that Business Entity (each Enterprise Service, Enterprise Skill Group, Schedule, Schedule Report, Schedule Source and Script).

The User_Security_Control table is kept up to date through the use of triggers which execute stored procedures. The tables are updated when database object changes affect object access. The triggers attempt only to regenerate sections of the User_Security_Control table instead of the entire table when they can do so. Table 8-3 shows a list of conditions that lead to partial updates, and the updates that occur:

Table 8-3 User_Secuirty_Control table partial recalculation conditions

What Changed	Modified Records
Object Added	Create all records for the object
Object Deleted	Delete all records for the object
User added to group	Recreate all records for the user
User deleted from group	Recreate all records for the user
New group or user given access to object	Recreate all records for the object
Group or user access to object deleted	Recreate all records for the object
Group or user access to object changed	Recreate all records for the object
New group or user given access to class	Recreate all records for the class
Group or user access to class deleted	Recreate all records for the class
Group or user access to class changed	Recreate all records for the class
New group or user given access to global	Recreate all records for all members of the group or user
Group or user access to global changed	Recreate all records for all members of the group or user
Group or user access to global deleted	Recreate all records for all members of the group or user

Even though these are partial updates, they still have the potential of generating large numbers of records. For example, adding a user to a group (that may only give access to a few objects) not only causes new records to be created for that user to grant access to the new object(s), but also recreates all

existing records for objects the user had previously been granted access to. These records include both objects on which rights are explicitly granted (such as a Peripheral) as well as those for which rights are implicitly granted (such as Agents, Skill Groups, Labels, etc).

Getting Started

Before you begin setting up ICM Partitioning, take some time to plan the process. Designing your Partitioning system carefully before beginning to implement it ultimately makes the task easier and less error-prone. In particular, you should:

- Determine tasks that need to be performed.
- Create a group for each task.
- Add users to groups representing the tasks the user needs to perform.

For example, one task might be to add Peripherals. To accommodate this task, you would take the following steps:

1. Create a group name, for instance, AddPeripherals.
2. Grant the group maintenance Peripheral Class Access.
3. Assign users to the group.

See Class and Object Security, page 8-8to work out which specific access rights you need to assign to each group. See Installing and Configuring ICM Partitioning, page 8-22for instructions on setting up Partitioning.

Class and Object Security

This section provides information about ICM class and object security.

Class and Object Security Overview

Class security defines access to a group of ICM configuration objects. *Object security* sets access privileges for specific records within a table or a group of tables.

Note For details regarding the classes and objects that affect the security for a specific database table, see the C*isco ICM Schema On-Line Help*.

Class and object security settings determine a User or Group's access level to ICM data. Certain access levels are valid for each class and each object and can be any combination of the levels described in Table 8-2:

Table 8-4 Access Levels

Access Level	Description
Maintenance	Permits the User or Group to read, update, and delete the object.
Reference	Permits the User or Group to read the object and use it in a script.
Read	Permits the User or Group to see the object, but not change it or use it in a script.
No access	Restricts access to the object. (This is the default if a User or Group is not explicitly assigned an access level.)

Note The highest—that is, the most permissive—access level to a particular piece of data "wins."

A user can belong to multiple groups or be assigned settings at the class level that conflict with settings at the object level.

For example: UserX might have been assigned only read access to PeripheralZ. However, UserX might also belong to Group1 and Group2. Group1 might have reference access to PeripheralZ and Group2 might have maintenance access to PeripheralZ. Consequently, even though UserX as an individual has only read access to PeripheralZ, since he belongs to Group2, he has maintenance access to PeripheralZ.

Another example:UserA might have read-only access to the global class but maintenance access to the peripheral object. Because the peripheral object controls the skill group object, UserA has maintenance access to the skill group object even though his global access gives him only read access.

You assign class security to a User or Group using the Class Security List. You assign object security using the Security Dialog or tab on the Explorer, List, or Bulk Configuration tool.

Class Security

Database *class security* defines access to a group of ICM configuration objects.Table 8-5 describes the classes the ICM software supports.

Table 8-5 Security Classes (Sheet 1 of 2)

Class Name	Description	Access Levels	Objects
Call	Provides security for viewing routing and call history tables.	Reference Read	Call
Global	Provides security to all objects and tables. **Note** A user with Maintenance access to the Global class has full access to all ICM software data. Administrative users automatically have this level of access.	Read, Reference, Maintenance	All
Network Interface	Provides security for setting up the ICM network interface.	Read, Maintenance	Announcement Call Type Device Target Dialed Number Label Network Interface Network Trunk Group Network VRU Network Vru Script Network/Peripheral Scheduled Target

Table 8-5 Security Classes (Sheet 2 of 2)

Class Name	Description	Access Levels	Objects
Peripheral	Provides security for configuring ICM peripherals.	Read, Maintenance	Agent Agent Team Call Type DialedNumber Dialer Label Network TrunkGroup Network/Peripheral Peripheral Peripheral global Route Service Service Array Skill Group Translation Route Trunk Group
System	Provides security for ICM security and configuration objects and tables.	Read, Maintenance	Agent Desk Settings Application Gateway Business Entity Campaign Database Lookup Enterprise Route Enterprise Service Enterprise Skill Group Expanded Call Variable Import RuleQuery Rule Schedule Schedule Report Script System User Formula User Variable

Object Security

Object security sets access privileges for specific records within a table or a group of tables. There are two types of security objects:

- A *controlling object* sets security on the object itself and a set of other objects. For example, the Peripheral object is a controlling object that groups Agents on a particular peripheral.
- A *controlled object* derives its security from a controlling object. For example, Agent is controlled by the Peripheral object.

Table 8-6 lists ICM objects and describes which database table each controls.

Note For complete details on each of the tables listed inTable 8-6, see the *Database Schema Handbook for Cisco ICM/IPCC Enterprise and Hosted Editions.*.

Table 8-6 Security Objects (Sheet 1 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Agent	Provides security on an Agent	Reference Read	Peripheral	Global Peripheral	Agent Agent_Half_ Hour Agent_Logout Agent_Real Time Agent_State_ Trace	
Agent Desk Settings	Provides security to use a set of Agent Desk Settings	Reference Read	System	Global System	Agent_Desk_ Settings Application_ Event ICR_View	
Agent Team	Provides security to use an Agent Team	Reference Read	Peripheral	Global Peripheral	Agent_Team Agent_Team_ Member Agent_Team_ Supervisor	
Announcement	Provides security to use an announce-ment	Reference Read	Network Interface	Global Network Interface	Announcement	
Application Gateway	Provides security to use an Application Gateway	Reference Read	System	Global System	Application_ Gateway Application_ Gateway_ Half_Hour Application_ Gateway_ Connection	

Table 8-6 Security Objects (Sheet 2 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Business Entity	Provides security to create objects within the Business Entity	Maintenance Reference Read	System	Global System	Business_ Entity	Enterprise Route Enterprise Service Enterprise Skill Group Schedule Schedule Report Schedule Source Script
Call	Provides security to read the call related tables		Call Global		Route_Call Detail Route_Call Variable Termination_ Call_Detail Termination_ Call_Variable	
Call Type	Provides security on a call type	Reference Read	Network/ Peripheral	Global Network Interface Peripheral	Call_Type Call_Type_ Half_Hour Call_Type_ Map Call_Type_ Real_Time	

Table 8-6 Security Objects (Sheet 3 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Campaign	Provides security to use a Campaign	Maintenance Reference Read	System	Global System	Campaign Campaign_ Query_Rule Campaign_ Query_Rule_ Half_Hour Campaign_ Query_Rule Real_Time Campaign_ Skill_Group Campaign_ Target_ Sequence	
Database Lookup	Provides security to use a Database Lookup	Reference Read	System	Global System	Script_Table Script_Table_ Column	
Device Target	Provides security to use a Device Target	Reference Read	Network Interface	Global Network Interface	Device_Target	
Dialed Number	Provides security on a dialed number	Reference Read	Network/ Peripheral	Global Network Interface Peripheral	Dialed_ Number Dialed_ Number_ Label Dialed_ Number_Map	
Dialer	Provides security to use a Dialer	Maintenance Reference Read	Peripheral global	Global Peripheral	Dialer Dialer_Half_ Hour Dialer_Port _Map Dialer_Port _Real_Time	
Enterprise Agent Group	Provides security to use an Enterprise Agent Group	Reference Read		Global	Enterprise_ Agent_Group Enterprise_ Agent_Group _Member	

Table 8-6 Security Objects (Sheet 4 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Enterprise Route	Provides security to use an Enterprise Route	Reference Read	Business Entity	Global System	Enterprise_ Route Enterprise_ Route_ Member	
Enterprise Service	Provides security to use an Enterprise Service	Reference Read	Business Entity	Global System	Enterprise_ Service Enterprise_ Service_ Member	
Enterprise Skill Group	Provides security to use an Enterprise Skill Group	Reference Read	Business Entity	Global System	Enterprise_ Skill_Group Enterprise_ Skill_Group_ Member	
Expanded Call Variable	Provides security to use an Expanded Call Variable	Maintenance Reference Read	System	Global System	Expanded_Call _Variable	
Import Rule	Provides security to maintain a Import Rule	Maintenance Reference Read	System	Global System	Import_Rule Import_Rule_ Clause Import_Rule_ History Import_Rule_ Real_Time	
Label	Provides security to use a Label	Reference Read	Network Interface Peripheral	Global Network Interface Peripheral	Label	
Network Interface	Provides security to use the network interface tables	Reference Read		Global Network Interface	Network_ Event_Detail Network_ Target	Announce-ment Device Target Network Vru Script Scheduled Target

Table 8-6 Security Objects (Sheet 5 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Network Trunk Group	Provides security to use a Network Trunk Group	Reference Read	Peripheral global	Global Network Interface Peripheral	Network_ Trunk_Group Network_ Trunk_Group Half_Hour Network_ Trunk_Group Real_Time Peripheral_ Target	
Network VRU	Provides security to use a Network VRU	Maintenance Reference Read	Network/ Peripheral	Global Network Interface	Network_Vru	
Network Vru Script	Provides security to use a Network VRU Script	Maintenance Reference Read	Network Interface	Global Network Interface	Network_Vru_ Script	
Network/ Peripheral	Provides security to read the tables that are used for both the Peripheral and the Network Interface			Global Network Interface Peripheral	Logical_ Interface_ Controller Physical_ Interface_ Controller Physical_ Controller_ Half_Hour Routing_Client Routing_Client Five_Minute Call Type Dialed_ Number Label Network_Vru	

Table 8-6 Security Objects (Sheet 6 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Peripheral	Provides security on a peripheral and the services, skill groups, etc. on it	Maintenance Reference Read	Peripheral global	Global Peripheral	Agent_Team_ Service Peripheral Peripheral_ Default_ Route Peripheral_ Half_Hour Peripheral Monitor Peripheral_ Real_Time Service_Level _Threshold Galaxy_Agent _Call_Count Galaxy_Agent _Igroup Galaxy_Agent _Performance Galaxy_Alarm Galaxy_DNIS Galaxy_PBX Galaxy_ Transaction_ Code	Agent Skill Group Trunk Group
Peripheral Gateway	Provides security to use a Peripheral Gateway	Maintenance	Network/ Peripheral	Global	Default_Call_ Type Dial_Number_ Plan	

Table 8-6 Security Objects (Sheet 7 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Peripheral global	Provides security to use the peripheral related tables			Global Peripheral		Dialer Network Trunk Group Peripheral Person Route Service Array Translat-ion Route
Person	Provides security to use a Person	Reference Read	Peripheral global	Global Peripheral	Person	
Query Rule	Provides security to use a Query Rule	Maintenance Reference Read	System	Global System	Query_Rule Query_Rule_ Clause	
Route	Provides security to use a Route	Reference Read	Peripheral global	Global Peripheral	Route Route_Half_ Hour Route_Five_ Minute Route_Real_ Time	
Schedule	Provides security to use a Schedule	Reference Read	Business Entity	Global System	Import_Log Import_ Schedule Recurring_ Schedule_ Map Schedule Schedule_Map Schedule_ Import Schedule_ Import_Real_ Time	

Table 8-6 Security Objects (Sheet 8 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Schedule Report	Provides security to maintain a Schedule Report	Maintenance Reference Read	Business Entity	Global System	Schedule_ Report Schedule_ Report_Input	
Schedule Source	Provides security to use a Schedule Source	Reference Read	Business Entity	Global	Schedule_ Source	
Scheduled Target	Provides security to maintain a Scheduled Target	Maintenance Reference Read	Network Interface	Global Network Interface	Scheduled_ Target Scheduled_ Target_Real_ Time	
Script	Provides security on a script and the associated real-time and historical tables	Reference Read	Business Entity	Global System	Admin_Script_ Schedule_ Map Master_Script Script Script_Cross_ Reference Script_Data Script_Five_ Minute Script_Print_ Control Script_Real_ Time	

Table 8-6 Security Objects (Sheet 9 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Service	Provides security on a Service and the associated real-time and historical tables	Reference Read	Peripheral	Global Peripheral	Service Service_ Member Service_Five_ Minute Service_Half_ Hour Service_Real_ Time Galaxy_Gate Galaxy_Gate_ Delayed_Call Galaxy_ Overflow	
Service Array	Provides security to use a Service Array	Reference Read	Peripheral global	Global Peripheral	Service_Array Service_Array _Member	
Skill Group	Provides security on a Skill Group and the associated real-time and historical tables	Reference Read	Peripheral	Global Peripheral	Agent_Skill_ Group_Half_ Hour Agent_Skill_ Group_ Logout Agent_Skill_ Group_Real_ Time Skill_Group Skill_Group_ Member Skill_Group_ Five_Minute Skill_Group_ Half_Hour Skill_Group_ Real_Time Skill_Target	

Table 8-6 Security Objects (Sheet 10 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
System	Provides security to read the ICM security and configuration tables			Global System	Agent_ Distribution Application Application_ Gateway_ Globals Application_ Instance Application_ Path Application_ Path_Member Application_ Path_Real_ Time Blended_ Agent_ Options ClassID_To_ ObjectType Class_Access_ Xref Class_List Class_Security Customer_ Definition Customer_ Options ICR_Instance ICR_Node	Agent Desk Settings Application Gateway Business Entity Campaign Database Lookup Expanded Call Variable Import Rule Query Rule User Formula User Variable

Table 8-6 Security Objects (Sheet 11 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
System (continued)					Feature_ Control_Set Media_Class Media_Rout- ing_Domain Object_Access _Xref Object_List Object_ Security Region Region_ Member Region_Prefix Region_View Region_View_ Member User_Group User_Group_ Member User_Superv- isor_Map VRU _Currency VRU_Defaults VRU_Locale	
Translation Route	Provides security to use a Translation Route	Reference Read	Peripheral global	Global Peripheral	Translation_ Route	

Table 8-6 Security Objects (Sheet 12 of 12)

Object	Description	Access Levels	Controlling Object	Classes	ICM Database Tables	Objects Controlled
Trunk Group	Provides security on a Trunk Group and the associated real-time and historical tables	Reference Read	Peripheral	Global Peripheral	Trunk Trunk_Group Trunk_Group_ Five_Minute Trunk_Group_ Half_Hour Trunk_Group_ Real_Time Vru_Port_Map Galaxy_Single _Trunk Galaxy_Trunk _Call_Count Galaxy_Trunk _IGroup	
User Formula	Provides security to maintain a User Formula	Maintenance Reference Read	System	Global System	User_Formula User_Formula _Equation	
User Variable	Provides security to maintain a User Variable	Maintenance Reference Read	System	Global System	Persistent_ Variable User_Variable	

Installing and Configuring ICM Partitioning

This section lists tips for using and installing Partitioning and describes:

- How to install ICM Partitioning
- The tools used to configure Partitioning
- How to configure Partitioning by:
  - Creating and maintaining user groups
  - Granting groups and individual users access to classes
  - Creating and administering individual user accounts
  - Setting access for individual ICM database objects and scripts

Installing ICM Partitioning

You enable the ICM Partition feature during the Admin Workstation Setup phase of the ICM software installation process. To install the ICM Partition feature:

Step 1 When creating the Logger database, using the ICMDBA tool, enable Partitioning and set the maximum number of Partitions. The maximum number of Partitions equals the maximum number of business entities.

Step 2 Run ICM Setup for each distributor AW in your system and enable Partitioning on those AWs.

Note For more information on installing Partitioning, see the *ICM Installation Guide for Cisco ICM Enterprise Edition*. For more information on the ICMDBA tool, see Database Administration Tool, page 4-2.

ICM Security Tools

In a system that **does not** have the ICM Partition feature enabled, the Configuration Manager's Security menu contains two options:

- User List
- Business Entity List

In a system that **does** have the ICM Partition feature enabled, the Security menu has two additional options:

- User Group List
- Class List

You use these Configuration Manager list tools to create ICM users and groups of users.

- Use the Configuration Manager list, explorer, and bulk tools to also set user and group access rights to classes of ICM objects and individual ICM objects.
- Use the Configuration Manager's Bulk Configuration tools to set security access to multiple data records at a time.
- Use the Script Editor to set security access to scripts.

Defining User Groups

Begin setting up security by creating a user group for each set of users who will have the same access rights within the ICM system. For example, you might create separate user groups for:

- Users who can make changes to the network interface
- Users who can add new peripherals
- Users who can change configuration data within each peripheral

Note By carefully defining the user groups you need and assigning the appropriate users to them, you ultimately make ICM security administration easier to maintain.

How to view currently-defined user groups

Step 1 From the Configuration Manager's Configure ICM menu, choose **Security > User Group List**. The User Group List window appears.

Step 2 In the *Select filter data* box, click **Retrieve**. The User Group List window lists the existing security groups.

How to add a new user group and assign members to the group

Step 1 In the User Group List window, click **Add**. A new user group displays in the Attributes tab.

Step 2 Fill in the Attributes tab fields. See the online help if you have any questions.

Step 3 Click **Save**. This saves the new group to the database.

Step 4 Optionally, click the **User Membership** tab and then, in that tab, click **Add** to assign users to the group. This displays the Add Users dialog box.

Note You can perform this step only if user accounts have already been defined. If user accounts have not been defined, you can assign users to the group later as you set up their accounts.

Step 5 In the Name list, select the user or users you want to add to the group and click **OK**. The Add Users dialog box closes.

Step 6 In the User Group List window, click **Save** to save the data in the database.

Repeat this procedure until all new groups have been added. Click **Close** in the User Groups List window to close the window.

How to delete a user group

Step 1 Select a Group Name from the User Groups List window and click **Delete**. The marked for deletion icon appears next to the group's name in the list box.

Step 2 Click **Save** to delete the security group. The User Groups name disappears from the list box.

Defining Users

After defining the security groups and specifying their levels of access, you can assign ICM users to the appropriate user groups.

How to see the users who are currently defined

Step 1 From the Configuration Manager's Configure ICM menu, choose **Security > User List**. The Users List window appears.

Step 2 In the *Select filter data* box, click **Retrieve**. The Users List window lists all users currently defined for the ICM system.

Note If the ICM Partition feature is enabled, the WebView Script Only and Customer properties are not visible in the attributes tab of a selected user.

How to add a new ICM user

Step 1 In the *User List* window, click **Add**. A new Attributes tab appears for the new user.

Step 2 Select the Domain Name.

- **Domain Name**. This is the Windows 2000 domain name, the unique host name on the internet to which the user belongs. Domain names are always in capital letters.

  Note The domain name must start with a letter and contain only letters and numbers. *Domain* refers to a set of servers and workstations grouped together for efficiency and security. A domain is the basic administrative unit in a server running Windows 2000. A network can be divided into domains by any convenient method, such as by department, workgroup, or building floor.

Step 3 Enter the following information:

- **User Name**. Select the name of the windows domain account.

  Note ICM user names must begin with a letter and can contain only letters and numbers. If the Windows user name contains characters other than the preceding, remove those characters from the ICM user name. (For example, the pound characters (#) and dollar characters ($) are not allowed in usernames.) The software appends the user name to the Domain name to form the User group name

- **Description**. Enter additional information about the user, such as the name of the person assigned to this account.
- **Password**. Enter the Windows password for the account. Only asterisks appear in the field as you type.
- **Change Password**. Click to change your password.

  Note Clicking this box enables the Password and Confirm password fields.

- **Confirm Password**. Enter the Windows password again to confirm that you have typed it correctly.

  Note Users created in the NAM domain are, by definition, given access rights to the NAM configuration database. Users created in the Limited AW (Customer) domain are given access to only their customer specific configuration data.

- **Read Only**. Check this box to give the user read-only access to the ICM.
- **Customer**. Not available when the ICM Partition feature is enabled.
- **Feature Control Set.** Select the Feature Set for the selected User.

Step 4 Select the user who is already a member of the following:

- **Configuration**
  - **Read only**. Added when Configuration is checked.
- **Setup**
- **Webview**

Step 5 Optionally, assign the user to one or more groups. Select the Group membership tab, and in that tab click **Add**. Then in the Add Groups dialog box, select the group(s) to which you want to add the user and click **OK**. This closes the Add Groups dialog box.

Step 6 In the User List window, click **Save** to create the user account.

Repeat this procedure until all users are created.

How to delete a user

Step 1 In the User List window, select a User Name from the list box and click **Delete**. A marked for deletion icon appears next to the name in the list box.

Step 2 When prompted to confirm the deletion, click **OK**.

Step 3 Click **Save** to delete the user. The user's name disappears from the list box.

Defining Business Entities Security

Each ICM enterprise consists of between one and five business entities. You can change the names and descriptions of business entities and set the access rights to business entities.

Note You can *only* create business entities on Partitioned systems that have more than one Partition.

How to find out how many Partitions are on the system

In the ICM Configuration Manager, open the System Information tools and look at the Max Partitions field.

How to view the business entities in your enterprise

Step 1 From the Configuration Manager's **Configure ICM** menu, choose **Security > Business Entity List**. The Business Entity List window appears.

Step 2 In the *Select filter data* box, click **Retrieve**. The Business Entity List window lists the defined business entities.

How to change names and descriptions of a business entity

Step 1 In the *Business Entity List* window, select the business entity you want to change.

Step 2 In the Attributes tab of the selected business entity, modify the **Entity Name** and/or **Description**.

Step 3 Click **Save** to submit your changes to the database.

Repeat this procedure to make changes to other business entity names and descriptions.

How to assign business entity security access

Step 1 In the *Business Entity List* window, select the business entity.

Step 2 Click the **Security** tab.

Step 3 In the Security tab, click **Add**. The Add Users and Groups dialog box appears.

Step 4 In the *Type* box, select User or Group, depending on whether you want to give access rights to a user or a group of users.

Step 5 In the *Names* list, select a User Name or Group Name to which you want to assign access to the business entity.

You can select multiple names if you want to assign access to more than one user or user group.

Step 6 From the *Access type* drop-down list (directly below the User Name and Group Name lists), choose the level of access you want to assign: Read or Maintenance.

Step 7 Click **OK**. The Add Users and Groups dialog box disappears and the user name or group name is displayed in the Security tab User Access list.

Step 8 Repeat steps 4 through 7 to give other Users/Groups access to the business entity.

Step 9 When you have finished assigning access, click **Save** to apply the changes.

Repeat this procedure to set the access rights for other business entities.

Defining Class Security Access

After you have created security groups, you can use the Class Security List tool to grant each group a specific level of access to each ICM security class.

How to assign class security access

Step 1 From the Configuration Manager's Configure ICM menu, choose **Security > Class List**. The Class Security List window appears.

Step 2 In the *Select filter data* box, click **Retrieve**. The Class Security List window lists the existing security classes.

Step 3 Select a class from the list box. (For example, to set access for the Network Interface class, choose Network Interface. See Table 8-5 on page 8-9 for a definition of each class.)

Step 4 Click the **Security** tab and then, in the Security tab, click **Add**. This displays the Add Users and Groups dialog box.

Step 5 In the *Type* box, select User or Group, depending on whether you want to give access rights to a user or a group of users.

Step 6 In the *Names* list, select a User Name or Group Name to which you want to assign access to the class.

You can select multiple names if you want to assign access to more than one user or user group.

Step 7 From the *Access type* drop-down list (directly below the User Name and Group Name lists), choose the level of access you want to assign: Read, Reference, or Maintenance.

Note Not all access levels are available to all classes.

Step 8 Click **OK**. The Add Users and Groups dialog box disappears and the user name or group name is displayed in the Security tab User Access list.

Step 9 Repeat steps 5through 8 to give other Users/Groups access to the class.

Step 10 When you have finished assigning access for the class, click **Save** to apply the changes.

Repeat this procedure to assign access for other classes.

Defining Object Security Access

Many of the elements that you define in Configure ICM are considered to be ICM objects. (For a list of ICM objects, see Class and Object Security, page 8-8.)

If your ICM system has the ICM Partition feature enabled, then whenever you create an ICM object, you have the option of using the security feature to set access rights to it.

How to define access rights for a new object

Step 1 Within the *Configuration Manager*, use the appropriate Configuration tool to specify information about the object.

Step 2 In the configuration tool, click the **Security** button (explorer and bulk tools) or tab (list tools). A Security dialog box appears for the object you are creating.

Step 3 Use the Security dialog box to specify which groups and individual users have access to the object.

Step 4 Click **OK** when done.

How to change the access rights for an existing object

Step 1 Within the *Configuration Manager*, click **Security** in the configuration tool window for the object you want to modify. The Security dialog box appears.

Step 2 Use the Security dialog box to specify which groups and individual users have access to the object.

Step 3 Click **OK** when done.

Defining Security for One or More Records at a Time

Step 1 In the *Configuration Manager* menu, select the *Bulk Configuration Tool* appropriate for the data type records for which you want to set access rights (for example, dialed numbers or labels). You can define security for multiple records in either Edit or Insert mode.

Step 2 In the selected data-type window, select the desired row or rows of records.

Step 3 Click **Security**. The Security dialog box displays. If there are security settings on the selected records and they are mixed (different records having different settings), no security data is displayed. Otherwise, the security settings for the selected record(s) are displayed.

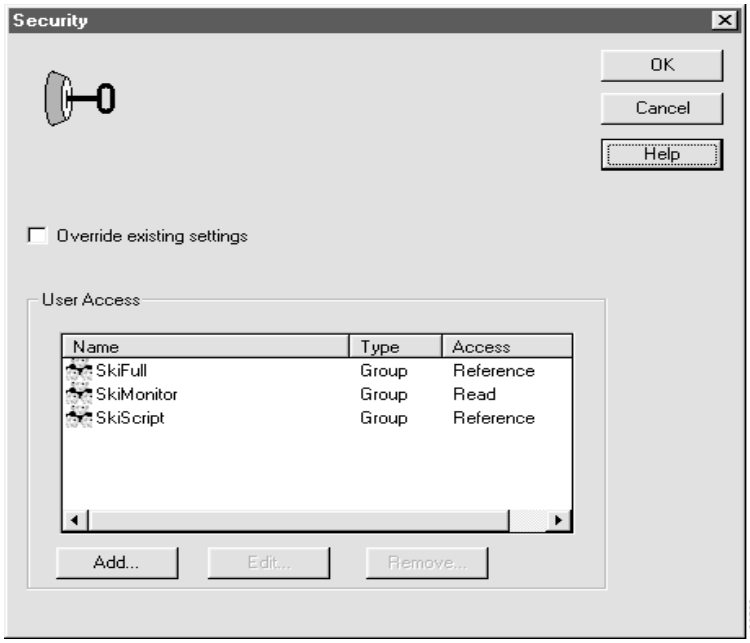

Step 4 If you want to apply one setting to records with mixed settings, select **Override existing settings**.

Note You can set or change security settings on a group of records only if they have the same security settings, if they have no security settings, or if you have selected **Override existing settings**.

Step 5 Make changes to the security settings:

- To add access to the selected records:

  Click **Add** and in the Add Users and Groups dialog box, (a) select user or group, (b) select the user or group name(s), (c) select the Access type (**Read**, **Reference**, or **Maintenance**), and (d) click **OK**.

  Note Not all access levels are available for all objects.

- To remove access to the records:

  In the User Access display box, select the user or group to remove and click **Remove.**

- To edit access to the records:

In the User Access display box, select the user or group to edit and click **Edit** or double-click on the item you want to edit. Then in the Edit Permissions dialog box, select the access type and click **OK**.

Step 6 When done, click **OK**.

Defining Script Security

Scripts that you create with the Script Editor are also ICM objects you can specify security for.

How to assign script security access

Step 1 Within the *Script Editor*, open the script.

Step 2 Right-click in the script to display the pop-up menu.

Step 3 Choose the **Security** option. The Script Security dialog box appears.

Step 4 Choose a User or Group from the lists at the lower right of the dialog box.

Step 5 From the drop-down list above those lists, choose an access level: Read or Reference.

Step 6 Click **Add**. The user or group you selected moves to the list on the left side of the dialog box.

Step 7 Repeat steps 4 through 6 to grant access to other users or groups.

Step 8 Click **OK** to submit your changes. The Script Security dialog box closes.

Repeat this procedure to set the access rights for other scripts.

How to access Script Security from the Script Properties dialog

You can also open the Script Security dialog from the Script Properties dialog box.

Step 1 Chose the **Security** tab in the Script Properties dialog box.

Step 2 Click **Modify Security**.

The Script Security dialog box opens. Continue with Step 4 of the preceding procedure.

Partitioning and Database Access

This section describes database access changes that result from partitioning your ICM Enterprise system.

Concepts to understand when reading this section include:

- Access Level - determines the actions that a user can perform on database objects. Access levels include Read, Reference, and Maintenance. Maintenance provides permission to create, read, update, and delete. These access levels are represented numerically as follows:
  - 10 = Read
  - 20 = Reference
  - 30 = Maintenance (create/read/update/delete)
- Class Security- sets access privileges for a group of objects

- Object Security - sets access privileges for a group of tables
- Controlling Object - sets access privileges for itself and a set of other objects. For example, the Peripheral object is a controlling object that groups Agents on that particular peripheral.

Database tables not included in the table below have no access controls or restrictions if partitioning is enabled.

Database Object Type	Specific Tables	Effect of Partitioning
Agent	• Agent • Agent_Half_Hour • Agent_Log_Out • Agent_Real_Time • Agent_State_Trace	If the ICM partitioning feature is enabled, these tables has the following access controls: • The Agent object provides Reference and Read access. • The Peripheral object is the controlling object. The Agent object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Agent Desk Settings	• Agent_Desk_Settings • Application_Event • ICR_View • Sec_User • Sec_Group • User_Security_Control • View_Column	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Agent_Desk_Settings object provides Reference and Read access. • The System object is the controlling object. • The Agent_Desk_Settings object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Agent Team	• Agent_Team • Agent_Team_Member • Agent_Team_Supervisor	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Agent Team object provides Reference and Read access. • The Peripheral object is the controlling object. The Agent Team object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Announcement	Announcement	If the ICM partitioning feature is enabled, this table has the following access controls: • The Announcement object provides Reference and Read access. • The Network Interface object is the controlling object. The Announcement object is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Application Gateway	• Application_Gateway • Application_Gateway_Connection • Application_Gateway_Half_Hour	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Application Gateway object provides Reference and Read access. • The System object is the controlling object. The Application Gateway object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Business Entity	Business_Entity	If the ICM partitioning feature is enabled, this table has the following access controls: • The Business Entity object provides Maintenance and Read access. It controls the following objects: – Enterprise Route – Enterprise Service – Enterprise Skill Group – Schedule – Schedule Report – Schedule Source – Script • The System object is the controlling object. The Business Entity object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Call Detail	• Route_Call_Detail • Route_Call_Variable • Terminiation_Call_Detail • Termination_Call_Variable	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Call object is a special object that provides a direct mapping between this table and all call objects. The Call object - which provides security to read all call-related tables - is a member of the Global and Call classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Call Type	• Call_Type • Call_Type_Half_Hour • Call_Type_Map • Call_Type_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Call Type object provides Reference and Read access. • The Network/Peripheral object is the controlling object. The Call Type object is a member of the Global, Network Interface, and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Campaign	• Campaign • Campaign_Query_Rule • Campaign_Query_Rule_Half_Hour • Campaign_Query_Rule_Real_Time • Campaign_Skill_Group • Campaign_Target_Sequence	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Campaign object provides Maintenance, Reference, and Read access. • The System object is controlling object. The Campaign object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Configuration Manager	• Cnfg_Manager_Globals • Cnfg_Manager_Snapshot_Stat • Cnfg_Manager_User • Cnfg_Manager_User_Desktop • Cnfg_Manager_User_Settings • Cnfg_Manager_View • ICR_Locks	If the ICM partitioning feature is enabled, these tables have no access restrictions.
Database Lookup	• Script • Script_Table_Column	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Database Lookup object provides Reference and Read access. • The System object is the controlling object. The Database Lookup object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Device Target	Device_Target	If the ICM partitioning feature is enabled, this table has the following access controls: • The Device Target object provides Reference and Read access. • The Network Interface object is the controlling object. The Device Target object is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Dialed Number	• Dialed_Number • Dialed_Number_Label • Dialed_Number_Map	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Dialed Number object provides Reference and Read access. • The Network/Peripheral object is the controlling object. The Dialed Number object is a member of the Global, Network Interface, and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Dialer	• Dialer • Dialer_Half_Hour • Dialer_Port_Map • Dialer_Port_Map_Real_Time • Dialer_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Dialer object provides Maintenance, Reference, and Read access. • The Peripheral Global object is the controlling object. The Dialer object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Enterprise Route	• Enterprise_Route • Enterprise_Route_Member	If the ICM partitioning feature is enabled, these tables have the following access controls: • Enterprise Route object provides Reference and Read access. • The Business Entity object is the controlling object. The Enterprise Skill Group object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Enterprise Service	• Enterprise_Service • Enterprise_Service_Member	If the ICM partitioning feature is enabled, these tables have the following access controls: • Enterprise Service object provides Reference and Read access. • The Business Entity object is the controlling object. The Enterprise Service object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Enterprise Skill Group	• Enterprise_Skill_Group • Enterprise_Skill_Group_Member	If the ICM partitioning feature is enabled, these tables have the following access controls: • Enterprise Skill Group object provides Reference and Read access. • The Business Entity object is the controlling object. The Enterprise Skill Group object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Expanded Call Variables	Expanded_Call_Variable	If the ICM partitioning feature is enabled, this table has the following access controls: • The Expanded Call Variable object provides Maintenance, Reference, and Read access. • The System object is the controlling object. The Expanded Call Variable object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Import Rule	• Import_Rule • Import_Rule_Clause • Import_Rule_History • Import_Rule_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Import Rule object provides Reference and Read access. • The System object is the controlling object. The Import Rule object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Label	Label	If the ICM partitioning feature is enabled, this table has the following access controls: • Label object provides Reference and Read access. • The Network/Peripheral object is the controlling object. The Label object is a member of the Global, Network Interface, and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Network / Peripheral	• Logical_Interface_Controller • Physical_Controller_Half_Hour • Physical_Interface_Controller_Five_Minute • Routing_Client • Routing_Client_Five_Minute	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Network/Peripheral object is a special object that provides a direct mapping between this table and other objects it controls: – Call Type – Dialed Number – Label – Network VRU – Peripheral Gateway The Network/Peripheral object - which provides security to read the tables that are used for both the Peripheral and Network Interface - is a member of the Global, Network Interface, and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Network Interface	• Network_Event_Detail • Network_Target	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Network Interface object is a special object that provides a direct mapping between this table and the other objects it controls: – Announcement – Device Target – Network VRU Script – Scheduled Target The Network Interface object - which provides security to read the network interface tables - is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Network Trunk Group	• Network_Trunk_Group • Network_Trunk_Group_Half_Hour • Network_Trunk_Group_Real_time • Peripheral_Target	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Network Trunk Group object provides Reference and Read access. • The Peripheral Global object is the controlling object. The Network Trunk Group object is a member of the Global, Network Interface, and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Network VRU	• Network_VRU • Network_VRU_Bank	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Network VRU object provides Maintenance, Reference and Read access. • The Network/Peripheral object is the controlling object. The Network VRU object is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Network VRU Script	Network_VRU_Script	If the ICM partitioning feature is enabled, this table has the following access controls: • The Network VRU Script object provides Maintenance, Reference and Read access. • The Network Interface object is the controlling object. The Network VRU object is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Peripheral	• Galaxy_Transaction_Code • Galaxy_PBX • Galaxy_DNIS • Galaxy_Alarm • Galaxy_Agent_Performance • Galaxy_Agent_IGroup • Galaxy_Agent_Call_Count • Peripheral • Peripheral_Default_Route • Peripheral_Half_Hour • Peripheral_Monitor • Peripheral_Real_Time • Service_Level_Threshold	If the ICM partitioning feature is enabled, this table has the following access controls: • The Peripheral object provides Maintenance, Reference, and Read access. It controls the following objects: – Agent – Agent Team – Service – Skill Group – Trunk Group • The Peripheral Global object is the controlling object. The Peripheral object - which provides security on a peripheral and the servcies, skill group, etc. on it - is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Peripheral Gateway	• Default_Call_Type • Dialed_Number_Plan	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Peripheral Gateway object provides Maintenance access. • The Network/Peripheral object is the controlling object. The Peripheral Gateway object is a member of the Global class. The Global class provides Maintenance, Reference, and Read access to all tables.
Person	Person	If the ICM partitioning feature is enabled, this table has the following access controls: • The Person object provides Reference and Read access. • The Peripheral Global object is the controlling object. The Person object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Query Rule	• Query_Rule • Query_Rule_Clause	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Query Rule object provides Reference and Read access. • The System object is the controlling object. The Query Rule object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Response Template	NONE	If the ICM partitioning feature is enabled, this table has the following access controls: • The Response Template object provides Reference and Read access. • The Peripheral object is the controlling object. The Response Template object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Route	• Route • Route_Five_Minute • Route_Half_Hour • Route_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Route object provides Reference and Read access. • The Peripheral Global object is the controlling object. The Route object is a member of Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Schedule	• Import_Log • Import_Schedule • Recurring_Schedule_Map • Schedule • Schedule_Import • Schedule_Import_Real_Time • Schedule_Map	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Schedule object provides Reference and Read access. • The Business Entity object is the controlling object. The Schedule object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Schedule Source	Schedule_Source	If the ICM partitioning feature is enabled, this table has the following access controls: • The Schedule Source object provides Reference and Read access. • The Business Entity object is the controlling object. The Schedule Source object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Schedule Target	• Schedule_Target • Schedule_Target_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Scheduled Target object provides Maintenance, Reference, and Read access. • The Network Interface object is the controlling object. The Scheduled Target object is a member of the Global and Network Interface classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Scheduled Report	• Schedule_Report • Schedule_Report_Input	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Schedule Report object provides Maintenance, Reference, and Read access. • The Business Entity object is the controlling object. The Schedule Report object is a member of the Global classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Script	• Admin_Script_Schedule_Map • Script • Script_Cross_Reference • Script_Data • Script_Five_Minute • Script_Print_Control • Script_Queue_Real_Time • Script_Real_Time • Master_Script • Media_Class • Media_Routing_Domain	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Script object provides Reference and Read access. • The Business Entity object is the controlling object. The Script object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Service	• Galaxy_Gate • Galaxy_Gate_Delayed_Call • Galaxy_Overflow • Service • Service_Five_Minute • Service_Half_Hour • Service_Member • Service_Real_Time	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Service object provides Reference and Read access. • The Peripheral object is the controlling object. The Service object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
Service Array	• Service_Array • Service_Array_Member	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Service Array object provides Reference and Read access. • The Global Peripheral object is the controlling object. The Service Array object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Skill Group	• Agent_Skill_Group_Half_Hour • Agent_Skill_Group_Logout • Agent_Skill_Group_Real_Time • Skill_Group • Skill_Group_Five_Minute • Skill_Group_Half_Hour • Skill_Group_Member • Skill_Group_Real_Time • Skill_Group_Target	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Skill Group object provides Reference and Read access. • The Peripheral object is the controlling object. The Skill Group object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
System	• Agent_Distribution • Application_Gateway_Global • Application_Path • Application_Path_Member • Application_Path_Real_Time • Blended_Agent_Options • Class_Access_Xref • Class_List • Class_Security • ClassID_To_ObjectType • Customer_Definition • Customer_Options • Feature_Control_Set • ICR_Globals • ICR_Instance • ICR_Node • Machine_Info • Object_Access_Xref • Object_List • Object_Security • Reason_Code • Region • Region_Member • Region_Prefix • Region_View • Region_View_Member	If the ICM partitioning feature is enabled, this table has the following access controls: • The System object is a special object that provides a direct mapping between this table and the other objects it controls: – Agent Desk Settings – Application Gateway – Business Entity – Campaign – Database Lookup – Expanded Call Variable – Import Rule – Query Rule – User Formula – User Variable The System object – which provides security to read the ICM security and configuration tables – is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
	• User_Group • User_Group_Member • User_Supervisor_Map • VRU_Currency	

Database Object Type	Specific Tables	Effect of Partitioning
Translation Route	Translation_Route	If the ICM partitioning feature is enabled, this table has the following access controls: • The Translation Route object provides Reference and Read access. • The Peripheral Global object is the controlling object. The Translation Route object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.
Trunk Group	• Galaxy_Single_Trunk_Table • Galaxy_Trunk_Call_Count • Galaxy_Trunk_IGroup • Trunk • Trunk_Group • Trunk_Group_Five_Minute • Trunk_Group_Half_Hour • Trunk_Group_Real_Time • VRU_Port_Map	If the ICM partitioning feature is enabled, these tables have the following access controls: • The Trunk Group object provides Reference and Read access. • The Peripheral object is the controlling object. The Trunk Group object is a member of the Global and Peripheral classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Database Object Type	Specific Tables	Effect of Partitioning
User Formula	User_Formula_Equation	If the ICM partitioning feature is enabled, this table has the following access controls: • The User Formula object provides Maintenance, Reference, and Read access. • The System object is the controlling object. The User Formula object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.
User Variable	• Persistent_Variable • User_Variable	If the ICM partitioning feature is enabled, this table has the following access controls: • The User Variable object provides Reference and Read access. • The System object is the controlling object. The User Variable object is a member of the Global and System classes. The Global class provides Maintenance, Reference, and Read access to all tables.

Problems with Partitioning

Problems with partitioning occur because of the implicit generation of the large User_Security_Control table. Because this table is so large, transaction sizes can become extremely large and cause the logger to detect a timeout and restart.

In order to ensure the consistency of the ICM configuration, every aspect of a configuration change must happen or the entire change must be rejected. The ICM application forces that behavior by wrapping the entire configuration change into a single database transaction.

The ICM logger has been designed so that if any transaction takes longer than a predefined period of time, it considers the transaction a failure, and resets the system (causing the logger process to restart and, depending on configuration settings, potentially rebooting the server). This design was put in place to ensure that if a transaction deadlocked the ICM logger would reboot, and then continue to service requests. When a configuration update gets too large (including the updates to the User_Security_Control table) this limit can be reached, causing the logger process to restart. When this problem is caused by both loggers trying to perform a configuration update it leads to a double logger failure.

Because there is no way to directly correlate transaction size to time, and database behaviors such as triggers can change both the size and time it takes for a transaction to complete, it is not possible to determine when a transaction will fail due to its size before it is executed. It is therefore necessary to understand what the stored procedures are doing in order to keep transaction size to a minimum.

To further complicate the problem, the same transaction can have greatly varying execution times depending on conditions in the database (both existing data and concurrent activity) at the time of execution. As the User_Security_Control table gets larger, transaction times tend to go up significantly. As the number of rows in the User_Security_Control table approaches 800,000, the chance of any configuration change causing a transaction timeout to occur increases significantly. For Cisco support, it is required to keep this table below 800,000 rows.

Best Practices

Because there is no way to be assured that any given transaction will succeed, it is important to follow the best practices and guidelines provided in this document that will allow transactions the highest chance of success. Following these guidelines does not guarantee that a transaction on a system that is approaching or is over the recommended 800,000 row count in the User_Security_Control table will succeed, but not following them will often cause configuration transactions on a partitioned system to fail.

These required practices must be followed to maintain Cisco support.

Only enable partitioning if it is really needed

In order to avoid the problems listed in this document, do not use partitioning unless it is absolutely necessary for your business. If you have systems with partitioning enabled and want to disable it, contact Cisco Advanced Services or your partner for assistance.

ICM software provides two alternatives to using partitioning.

- Feature Control is an alternative to partitioning that provides some level of security on a user by user basis. Feature Control allows users to be associated with Feature Control Sets. Feature Control Sets restrict which configuration tools a user is allowed to use or see, as well as what script nodes they are allowed to use in the Script Editor.

For example, Feature Control could be used to specify that a User can edit Services but not Skill Groups. This would be accomplished by using the Feature Control Set List tool to define a Feature Set that allowed running the Service Explorer but not the Skill Group Explorer. A User would then be assigned to that Feature Set. The User would not be able to run the Skill Group Explorer.

- Another alternative to partitioning is a WebView feature that allows administrators to restrict which call types users can see by associating users and call types with a specific customer object. Both the Call Type list tool and the User list tool have a field for Customer in them. If this is set to a given user, that user is only able to see Call Types in WebView that have the field also set to associated them with the same Customer.

Hardware can help the problem

Because the nature of the problem is timing based, faster hardware can alleviate the problem. While better hardware is not the complete answer ensuring that the Logger system does not have any hardware bottlenecks. Examples of hardware problems that may cause transactions to fail more frequently are:

- Insufficient memory; having too little memory available to the SQL Server database can cause transactions to take longer than normal.
- Poorly performing disks; slow disks will also increase transaction time

Do changes in small batches

Make changes to configuration in small increments. Follow these rules to help ensure that the change is completed successfully in the allotted transaction time:

- Follow the partial update rules. Involving as few objects as possible in a configuration change helps ensure that the change is performed successfully in the allotted transaction time.

- Grant access to objects to a group one at a time, not all at once.
- Add users to groups one at a time.

Note Do not add or remove several users from a group at once; restrict changes to a single change at a time.

Ensure the logger system is idle during changes

Avoid making changes while other activity is happening on the system because this can increase the time required to execute a transaction.

Specific examples are:

- Avoid making updates at ½ hour boundaries. The ICM transfers 30 minute summary data every half hour. Avoid making changes at this time as the database is busy loading that data into the historical tables.
- Avoid updates when other background tasks are executing, such as:
  - Database purge
  - Backups
  - System updates
- Avoid updates during peak call volume hours. When the ICM is actively handling calls data is being written to the database. During high call volume periods the writing of Call Detail Records can significantly increase the time required for transactions to execute.

Minimize the number of objects in the system

Minimizing the number of objects on a partitioned system can significantly increase the performance during configuration changes. Ensuring that the number of objects in the system stays low helps keep the User_Security_Control table size below 1 million rows.

Permanently delete old or unused objects. If a system has 1000 extra objects and 50 users, the User_Security_Control table would have 50,000 extra rows. Specifically look at unused Agents, Labels and old versions of Scripts for objects that can be permanently deleted.

ICM enables you to configure the number of old versions of a script to save. The default is to keep all old versions. Set the script retention to a specific number (based on business requirements) to help keep the total number of objects in the system under control.

Minimize the number of users in the system

Removing old or unused users in the system significantly decreases the size of the User_Security_Control table. This is especially true of administrative users.

Supervisors increase the row count in the User_Security_Control table as well, but not to as great an extent because they do not have access to as many objects as most administrators. Supervisors are granted access to a limited set of objects on the peripheral that the particular supervisor is configured on. These objects are:

- Agent
- Agent Team
- Service
- Skill Group
- Trunk Group

Sharing login information and user definitions between administrators or WebView users helps reduce the number of users configured in the ICM. For example, reducing the number of users from 15 to 10 could reduce the size of the User_Security_Control table by 33%.

Partitioning Tips

The partitioning feature of ICM Enterprise is an optional feature designed to be used to flexibly limit the access of ICM users to configuration objects. The feature is implemented using functionality available in the SQL Server database to limit access directly at the database level.

Note In order to implement this functionality the User_Security_Control table contains a comprehensive list of the access rights for all users in the system. Due to the way this table is constructed its size grows quickly. The size of the User_Security_Control table is bounded by the formula: **Number of ICM Users * Number of ICM Objects**

The size of the User_Security_Control table limits the performance of the database. In severe cases, a transaction (usually initiated because of a configuration change) takes longer than the ICM Logger can accomodate. When this happens, the ICM Logger restarts itself and may reboot the system. This can lead to service outages.

Ensuring that a partitioned ICM system does not run into these problems requires a two-pronged approach:

- Ensuring that the size of the User_Security_Control table stays as small as possible. Problems begin to appear as the size of this table approaches 800,000 rows.
- Following a set of required practices (as outlined in section 4) to ensure that database transactions happen as efficiently as possible on the system.

INDEX

Symbols

Numerics

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W