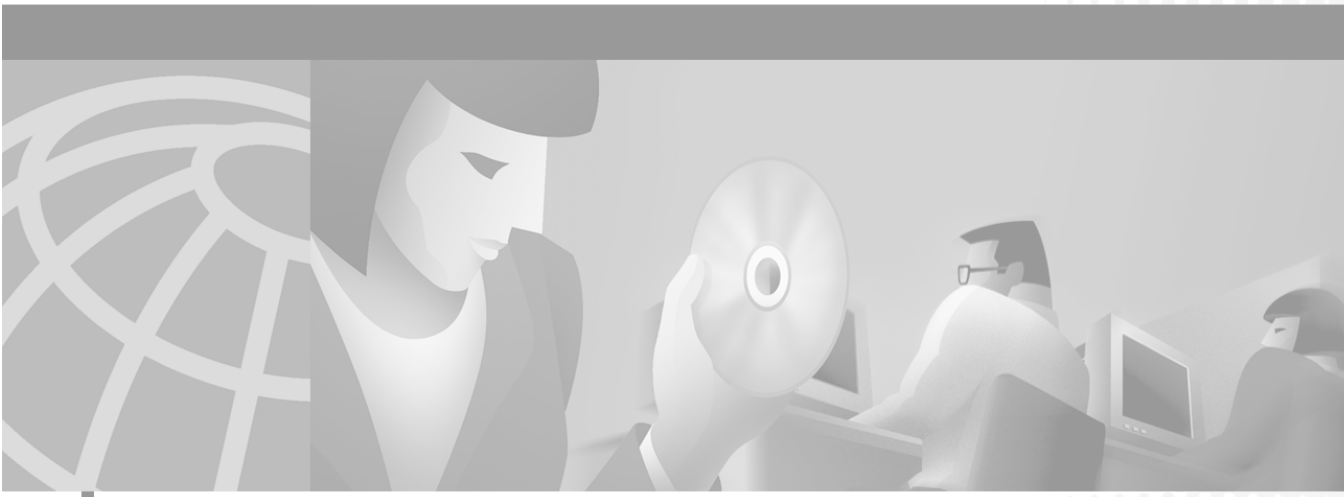

Cisco Resource Policy Management System Solutions Guide Release 2.0.1

Release 2.0.1
Updated September 2002

Corporate 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: 408 526-4100

Customer Order Number:
Text Part Number: OL-3328-01

Cisco Resource Policy Management System Solutions Guide Release 2.0.1

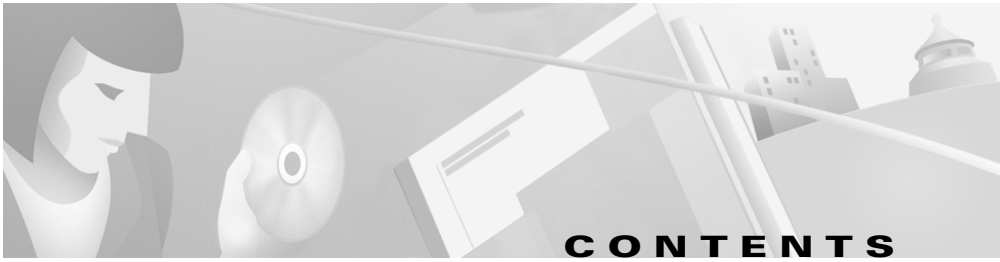

CONTENTS

Preface

This guide describes concepts you must understand before configuring Cisco Resource Policy Management System (Cisco RPMS) and provides examples of common configurations.

Note With Release 2.0, Cisco Resource Pool Manager Server is now referred to as Cisco Resource Policy Management System to emphasize its increased functionality.

This guide was written for system administrators who use Cisco RPMS for universal gateway (UG) resource management and assumes you have experience with UGs, Oracle databases, and general networking technology.

Conventions Used in This Guide

Convention	Description
bold	Command or keyword that you must enter.
italic	Argument for which you supply a value.
[x]	Optional keyword or argument that you enter.
{x | y | z}	Required keyword or argument that you must enter.
[x {y | z}]	Optional keyword or argument that you enter with a required keyword or argument.
string	Set of characters that you enter. Do not use quotation marks around the character string, or the string will include the quotation marks.
screen	Information that appears on the screen.
^ or Ctrl	Control key—for example, ^D means press the Control and the D keys simultaneously.
< >	Nonprinting characters, such as passwords.
!	Comment line at the beginning of a line of code.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss.

Note Means *reader take note*. Notes contain helpful suggestions or reference to materials not contained in this manual.

Tip Means the information *might help the reader solve a problem*.

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

http://www.cisco.com

Translated documentation is available at the following URL:

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

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

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

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on Cisco.com, you can submit technical comments electronically. Click **Leave Feedback** at the bottom of the Cisco Documentation home page. After you complete the form, print it out and fax it to Cisco at 408 527-0730.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.

- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

http://www.cisco.com/register/

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

Defining Cisco Resource Policy Management System

Topics in this chapter include:

Overview: What Is Cisco Resource Policy Management System?

Cisco Resource Policy Management System (Cisco RPMS) is used by wholesale network providers to guarantee service agreements among multiple retail service providers and corporate virtual private networks vying for simultaneous access of the same network resources. These services are applicable not only to wholesale dial networks, but also include support for Any Service, Any Port (ASAP), Session Initiation Protocol (SIP) and H.323 voice networks.

Note With Release 2.0, Cisco Resource Pool Manager Server is now referred to as Cisco Resource Policy Management System to emphasize its increased functionality.

Built with the wholesale customer in mind, Cisco RPMS provides a rich set of tools and call control parameters. It allows you to create and enforce policies shared across multiple customers, as well as to correlate individual calls with a specific customer and enforce service policies specific to that customer.

Cisco RPMS works to enforce and record resource usage and service level agreement (SLA) contracts between a wholesale provider and an Internet Service Providers (ISP). With shared overflow pooling and dedicated customer SLA management, Cisco RPMS provides the flexibility to share pooled resources while still maintaining and enforcing SLA contracts with multiple customers, without necessarily dedicating fixed hardware ports to a customer. This provides cost savings to the wholesaler and broader port and trunk access to customers.

Working as a service level call controller for the telco, Cisco RPMS tracks ISP sessions and service use, and port usage. As a call controller, Cisco RPMS participates in call set up requests before universal gateway (UG) resources are allocated to handle the call. This is commonly known as “preauthentication” call control.

By leveraging a rich set of Remote Authentication Dial-In User Service (RADIUS) attributes, Cisco RPMS can associate a call with a particular SLA and customer policy by using parameters such as:

- Dialed Number Identification Service (DNIS)
- Call type

- Trunk group
- Domain
- A combination of those fields and default groups

With the introduction of Cisco RPMS ASAP, VOIP calls can additionally be associated to policies by using Cisco H.323 Gatekeeper Clear Tokens, originating IP network masks, and various Cisco H.323 and SIP vendor specific attributes (VSAs).

Cisco RPMS is written in Java and leverages a dynamic class loader. This optional module enables the product to easily extend beyond managing Cisco Powered and RFC compliant equipment, to handle even non-standard vendor specific messaging.

A distributed software application, Cisco RPMS supports deployments in which physical policies (such as first come, first serve access to shared POP or community resources) and customer-specific policies (guaranteed port policy or trunk group access for an ISP or an Internet Telephony Service Provider, or ITSP) can be managed on the same host, or on separate hosts.

The architecture allows customers to tailor and expand Cisco RPMS deployments to scale from an entry level single server to highly redundant clusters of interoperating policy servers deployed across wide-area or large-scale networks. This flexibility and expansion potential works primarily behind the scenes of Cisco RPMS' easy-to-follow InstallShield GUI. Even multi-host distributed deployments are typically only a few default clicks away as part of an upgrade. As a service management system, you can deploy Cisco RPMS across several servers to manage services for every call within large networks at call processing rates that could cripple a single server.

The core application components of a Cisco RPMS system are:

- A RASER
- Policy processors
- A Web and SNMP server
- An Oracle database

You can run policy processing on a single host or it can be distributed. Because policy processing must remain stateful throughout the life cycle of a call, Cisco RPMS provides an optional high availability option so that two servers can maintain the same policy state. This also provides additional reliability.

Cisco RPMS also supports policy distribution—so shared overflow pool and POP policies can run on independent servers (POP servers) from customer and SLA policies (SLA servers). A single server can host all policy processing necessary for moderate sized networks; with Cisco RPMS you can also segregate policies that enforce service agreements across multiple POP servers and multiple SLA servers. This allows Cisco RPMS to scale to meet large scale network and high system call rate service enforcement. The role of each of these core components is summarized below with details provided in other parts of the Cisco RPMS documentation.

Remote Access Service Router (RASER)

RASERs (or Cisco RPMS proxies) are completely stateless high-speed RADIUS software switching or routing components used to route RADIUS call control messages between various call and service controllers.

RASERs are configured to receive RADIUS messages directly from universal gateways, and proxy these requests to policy servers, RADIUS AAA authentication and accounting servers. Additionally, with the Cisco RPMS ASAP release, you can deploy RASERs to forward voice control messages from Cisco SIP proxy servers and Cisco H.323 gatekeepers to policy servers.

POP and Shared Overflow Pool Manager

This component manages shared overflow pools. The ability to separate virtual POP (VPOP) policy from ISP customer policy enables the wholesaler to create and enforce customer SLAs that span multiple POPs. That is, one customer could share multiple overflow pools across a region; for example, a wholesaler can deploy Cisco RPMS to handle ISP traffic from Los Angeles and Seattle. In this

instance, the wholesaler could have one SLA server managing system-wide policy limits for an ISP, but two POP managers—one for Los Angeles, and one for Seattle, each enforcing POP specific shared overflow pool limits.

Additionally, with Cisco RPMS ASAP deployments, the POP manager can enforce a ratio of voice and dial service mix within the shared pool of universal port resources.

Service Level Agreement Server

The service level agreement (SLA) server enforces policy as per the agreement between the wholesaler and its customers. For a single or dual server deployments, the SLA resides on the Cisco RPMS host.

An SLA server can handle SLA enforcement for multiple customers, or for a single customer policy. By dedicating an SLA server to a subset of all customers, or even to a single customer, you can distribute SLA enforcement across multiple hosts to scale out to manage very high-speed ASAP traffic or very large-scale pooled resources. The only restriction is that each customer profile limit is completely managed by one SLA server—one customer profile cannot span multiple SLA servers except for High Availability purposes.

SLA servers and POP and shared overflow pool managers are types of policy processors. As such, they are stateful components that maintain an active snapshot of every call within the network that affects or is affected by policies enforced by the processor. Examples include ISP guaranteed session limits across trunks or POPs, or first come, first serve access limits for single POP overflow resources. Policy processors currently run on Solaris servers.

In a single host deployment, policy processing runs on the RASER and all policies are enforced by the single host. With multi-host deployments, you can separate policy processing from RASER functionality, so that a host is dedicated solely to processing all policies, or distribute processing so that each individual policy runs on a server dedicated to enforce that unique policy.

Because separation of POP or VPOP policy from individual customer SLA policy is common, Cisco RPMS introduces the concept of a POP and Overflow Pool server, as well as customer policy or SLA servers.

Secure Web Server and SNMP Access

The Cisco RPMS system embeds Web servers with secure socket connections and an SNMP server during installation. Any server hosting policy processors can be accessed by using the Web or by using SNMP. The Web interface allows for

provisioning and configuring any component in the Cisco RPMS system. The SNMP interface is for monitoring, and generating alarms, and not used for configuration.

Oracle Database

Cisco RPMS uses JDBC drivers that have been verified with Oracle. All persistent administration and operational data for a Cisco RPMS system is stored within a database. Only one Oracle database is required for the whole system; all Cisco RPMS components can share this database, which provides an easy means for centralized provisioning of the entire system. The system performs call processing independent from the database, so a centralized database is not a centralized single point of failure for operations.

The details for using all of the service policy modeling parameters, constructing service policies and configuring Cisco RPMS to apply those to incoming calls are described in the *Cisco Resource Policy Management System Configuration Guide*. The remainder of this guide describes an overview of Cisco RPMS use across solutions with an emphasis on the attributes of Cisco RPMS that are applicable to all solutions.

Overview: Where is Cisco RPMS Used?

You can deploy Cisco RPMS to manage SLA policies in various networks, including wholesale dial networks, originating and terminating H.323 and SIP Voice Over IP (VOIP) networks, and ASAP networks. Following is a brief description for Cisco RPMS' role in each of these network solutions.

Wholesale Dial Networks

Figure 1-1 depicts an abstraction for a simple wholesale dial network. Elements essential for call control in such a network include the universal gateway, Cisco RPMS, and optionally, a AAA proxy server managed by the wholesaler, as well as remote call control systems, such as an ISP's AAA server.

Figure 1-1 Basic Wholesale Dial Call Control Components

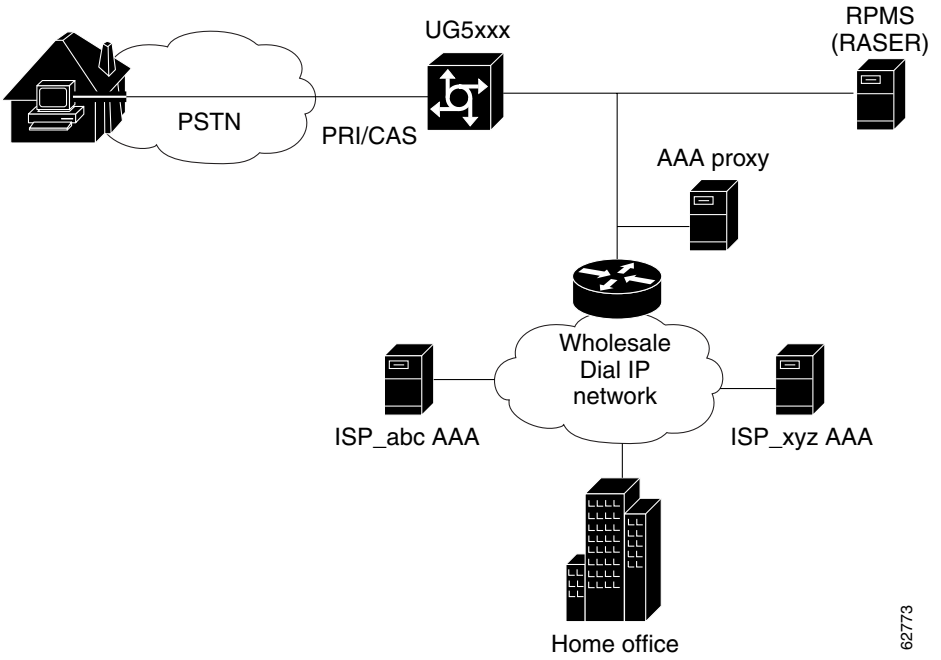

Figure 1-2 depicts a very basic call flow sequence diagram and rudimentary Cisco IOS commands for provisioning a Cisco UG to enable preauthentication with Cisco RPMS. Note that the universal gateway directs all RADIUS messages to the RASER.

The Cisco RPMS RASER can proxy authentication and accounting messages directly to remote service sites, or route them to an existing AAA proxy (or list of proxies) in the wholesale network as depicted.

Figure 1-2 Single RPMS Host Wholesale Dial Call Flow

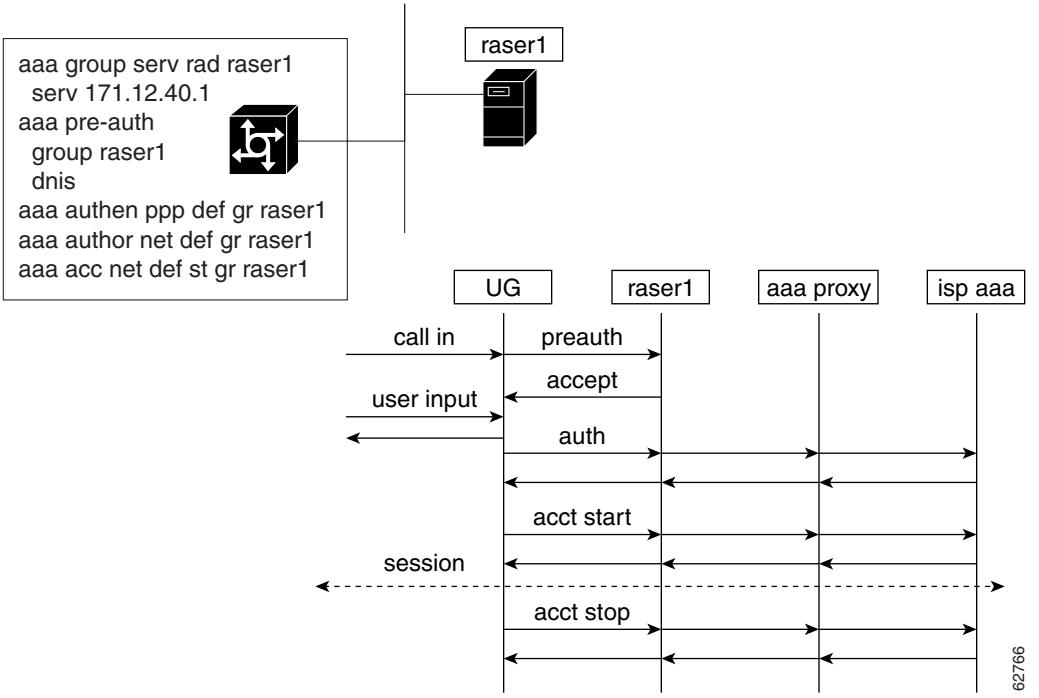

Cisco RPMS deployment options for wholesale dial solutions are described in detail in Chapter 5, "Cisco RPMS Deployment Scenarios."

Incoming H.323 or SIP VOIP Networks

Figure 1-3 depicts an abstraction for a simple wholesale H.323 VOIP network. Elements essential for call control in such a network include the universal gateway, Cisco RPMS, gatekeepers and, optionally, a voice billing server managed by the wholesaler.

Figure 1-3 Call Control for Incoming H.323 VOIP

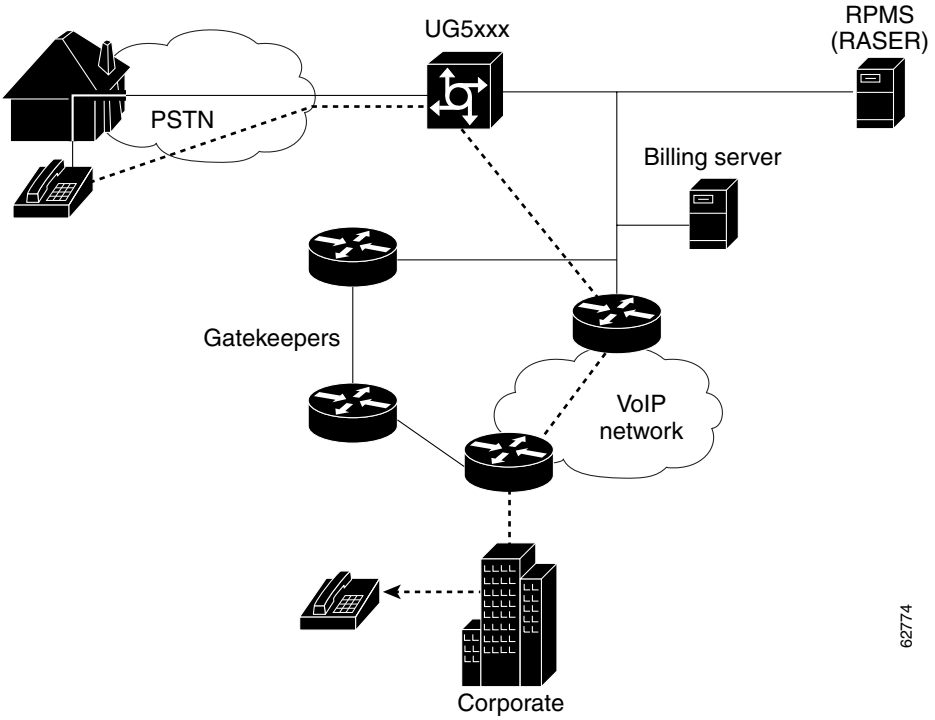

Call control is applied to the call that enters the VOIP wholesale network from the public switched telephone network (PSTN) or to the call arriving from a VOIP IP network. The call control for Incoming PSTN Voice call is about the same as for incoming dial. The call control sequence for VOIP includes interaction with the gatekeeper for H.323 VOIP or SIP proxy server for SIP.

Figure 1-4 depicts a very basic call flow sequence diagram and rudimentary Cisco IOS commands for provisioning a Cisco UG in a VOIP network to enable preauthentication with Cisco RPMS. Note that the universal gateway directs all RADIUS messages to the RASER.

Figure 1-4 Pre-Auth Incoming H.323 VOIP

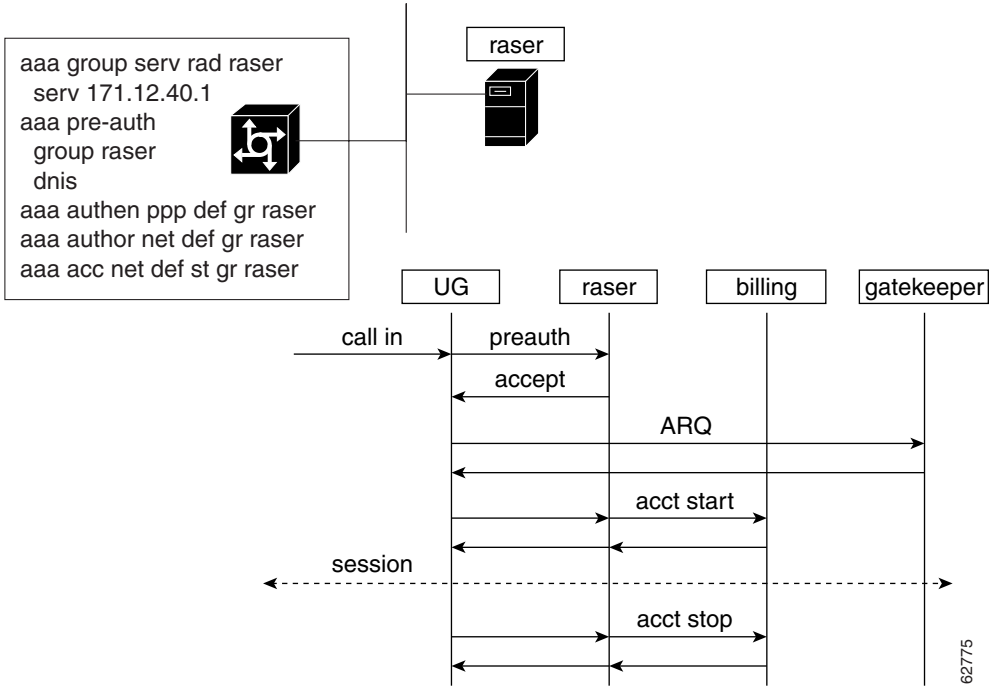

Terminating H.323 or SIP Networks

You can also use Cisco RPMS to manage policies affected by terminating VOIP calls. Cisco RPMS can interoperate with Cisco H.323 gatekeepers by using Gatekeeper Transaction Message Protocol (GKTMP) and the universal gateways with preauthentication control. Figure 1-5 depicts an arrangement of call control components controlling H.323 VOIP call termination.

Figure 1-5 Call Control for Terminating H.323 VOIP

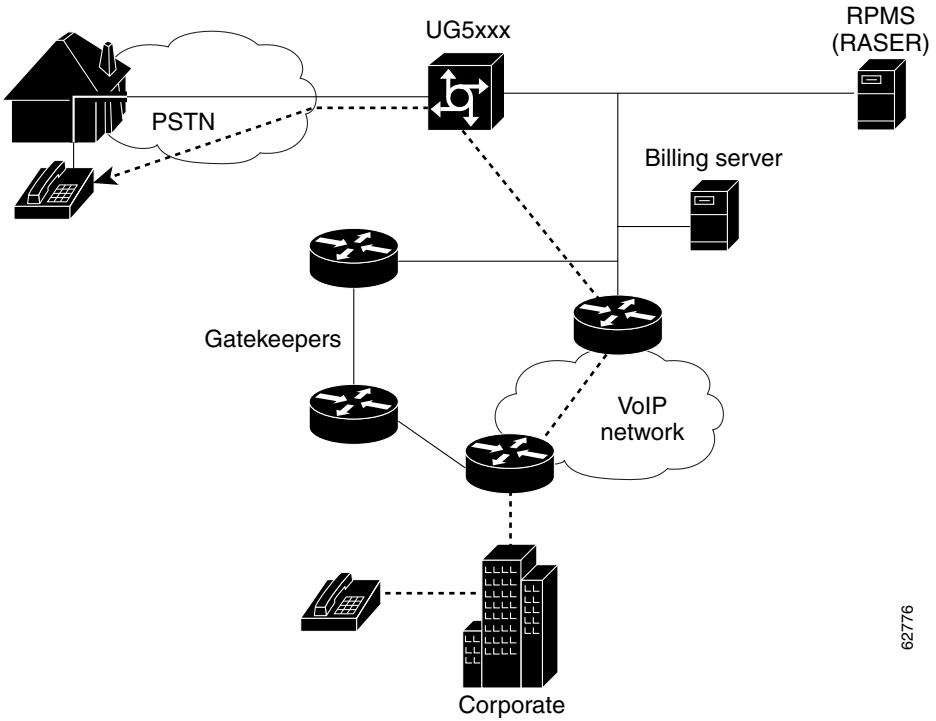

The basic call control sequence for enforcing policies impacted by terminating VOIP calls is depicted in Figure 1-6. The sequence is initiated by a Location Request (LRQ) received from the originating gatekeeper. You can provision Cisco gatekeepers and RPMS to communicate preauthentication queries by using GKTMP. This query service means that Cisco RPMS provides a policy

response—to accept or reject the call based on policy—without a policy lock or reservation. Later, when the universal gateway receives the H.225 set up request, the UG sends a preauth reservation request to Cisco RPMS.

Figure 1-6 Pre-Auth Terminating VOIP Call Termination

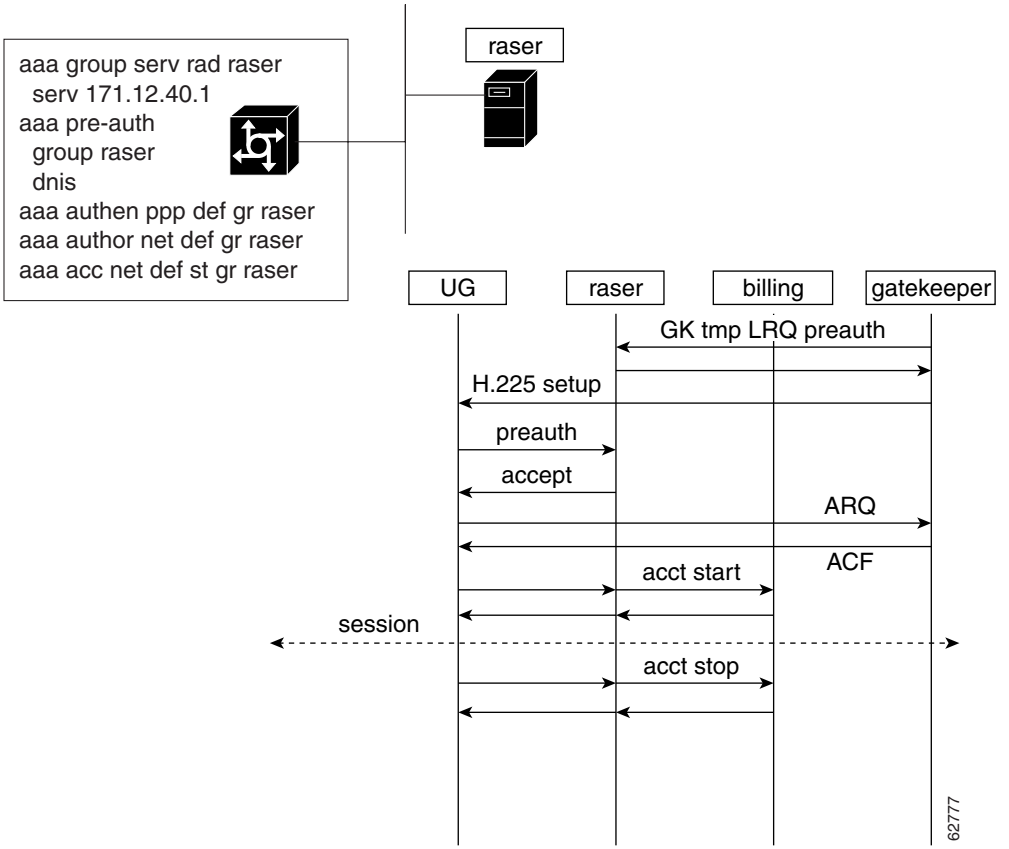

ASAP Networks

Cisco RPMS can manage voice and dial access to universal ports from the Cisco universal gateways.

The Cisco RPMS system was designed to be integrated into the Cisco Any Service, Any Port (ASAP) network architecture solution. Refer to the ASAP Solution documentation or your Cisco representative for the latest status on solution integration.

The Cisco ASAP solution is a unified network architecture that delivers integrated data, voice, fax and wireless services on a single network. The Cisco AS5000 series universal gateways are the foundation of the ASAP network infrastructure.

Cisco RPMS is an integral component of the ASAP solution with its capability to manage voice and dial access to the universal ports on the Cisco UGs on a call-by-call basis. Cisco RPMS will support the Cisco ASAP network solution in mid-2002, and requires an upgrade for both Cisco RPMS and the Cisco IOS release.

Managing an ASAP network involves taking the concepts introduced in this chapter, and concatenating them. With an ASAP network, these steps happen simultaneously. An ASAP network also allows for running multiple services from the same UG.

Voice call provisioning and solutions will be described more thoroughly in the Cisco RPMS Voice and ASAP Solution Guides.

CHAPTER 2

Cisco RPMS Features

Topics in this chapter include:

Overview: Major Features of Cisco RPMS Release 2.0.1

Cisco RPMS Release 2.0.1 ports the Cisco RPMS Release 1.*x* feature set into a world managed by RADIUS. The bundled RFC compliant RADIUS drivers permit Cisco RPMS 2.0.1 to manage call setup requests (preauth) and accounting from compliant non-Cisco gateways and from all Cisco universal gateways.

The RPMS Dynamic Class Loader further extends Cisco RPMS by providing a means for loading customized RADIUS drivers with vendor specific messaging into an operational RPMS system. This optional module permits Cisco RPMS to be extended to manage even non-RFC compliant gateways.

Cisco RPMS enhancements provide more services to wholesale dial networks providers. The Cisco RPMS architecture introduces reliable, scalable, distributed-policy management to wholesale Voice and ASAP networks, where call hold times are an order of magnitude smaller than comparable dial networks.

Cisco RPMS ASAP is built to enable service agreement criteria as part of VOIP call setup requests—even for call setup requests not triggered from port connection—but from terminating VOIP.

This chapter describes these and other Cisco RPMS features.

Overview: RADIUS-Based Support

Cisco RPMS uses the RADIUS protocol to communicate between the universal gateway and AAA servers. The transition from TACACS+ to RADIUS allows Cisco RPMS to communicate with both Cisco UGs and non-Cisco network access servers (NASs) that support RADIUS preauthentication.

As such, UGs are no longer provisioned with the “resource pooling” directives, but instead use the RADIUS “preauth” and RADIUS server features. Each UG client uses RADIUS preauthentication and accounting to communicate resource management requests and resource changes to the Cisco RPMS server. For more information about UG configuration, refer to Appendix A, “Configuring the Universal Gateway” of the *Cisco Resource Pool Manager Server Configuration Guide*.

For general information on RADIUS, refer to the following Web pages:

http://www.cisco.com/warp/public/cc/techno/protocol/ipsecur/prodlit/555_pp.htm

http://www.cisco.com/warp/public/707/32.html

Figure 2-1 shows a network in which the UG, Cisco RPMS RASER and AAA proxy all use the RADIUS protocol to communicate.

Figure 2-1 Cisco RPMS and the RADIUS Protocol for Wholesale Dial

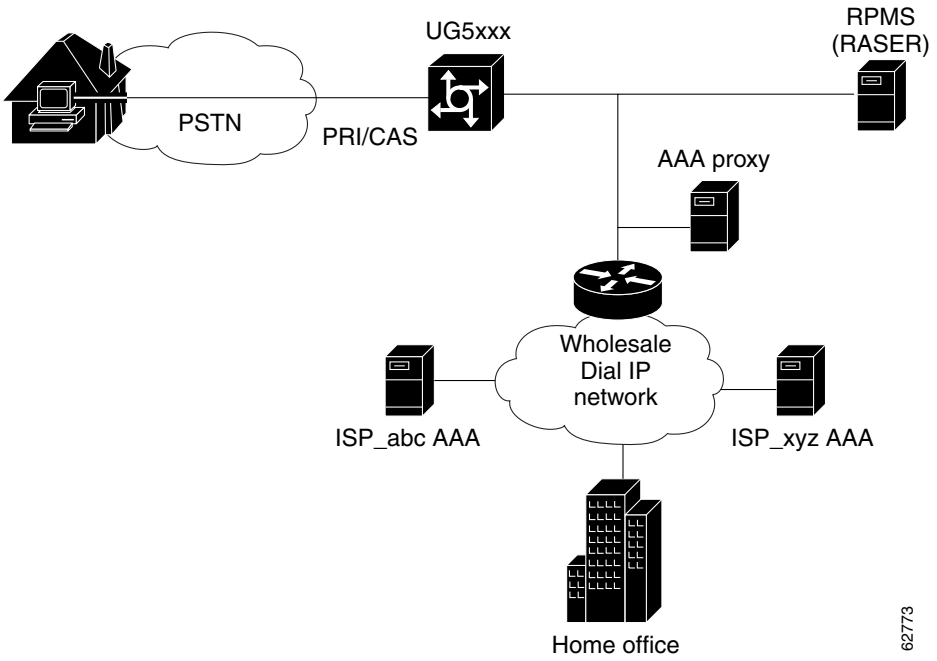

Note You must configure RADIUS preauthentication on all UGs communicating with Cisco RPMS 2.0.1.

Figure 2-2 shows a single host wholesale dial call flow and partial provisioning for a Cisco UG to enable preauthentication with Cisco RPMS. In the call flow, the universal gateway directs all RADIUS messages to the RPMS RASER that can be configured to proxy AAA messages to other AAA servers.

Figure 2-2 Single Cisco RPMS Host Wholesale Dial Call Flow

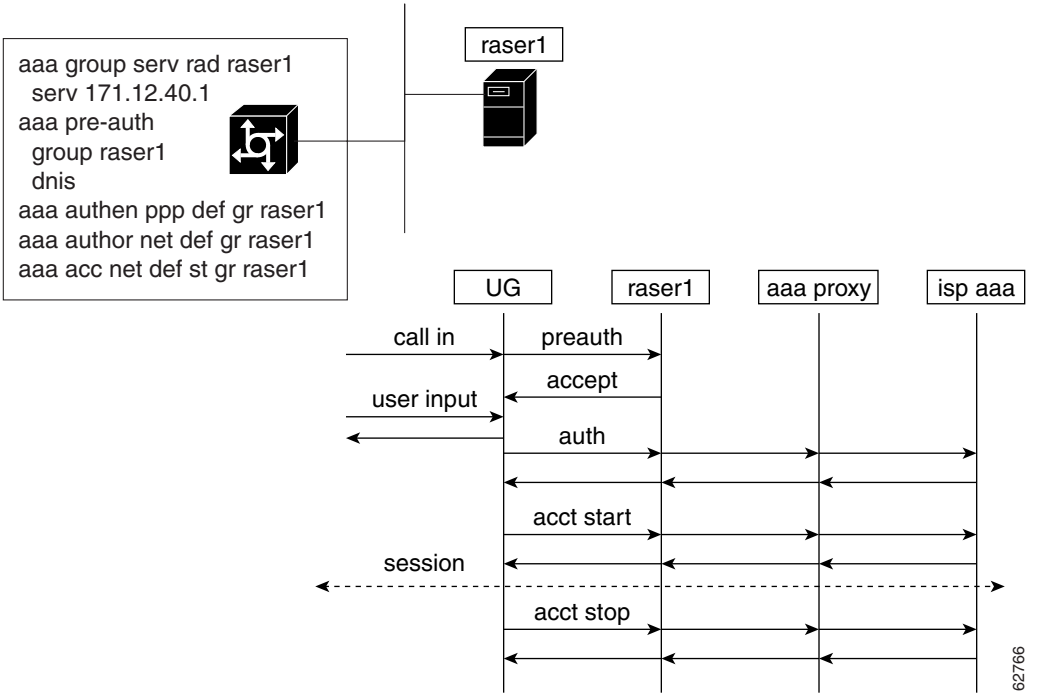

Some of the core RADIUS features of Cisco RPMS include:

- RADIUS Proxy (RASER)
- Accounting and Billing Support
- Specifying User-Based Authentication
- Managing Non-Cisco Equipment

The following sections describe these features.

RADIUS Proxy (RASER)

Cisco RPMS 2.0.1 is a AAA proxy for the UG (see the note below). As such, the preferred deployment is for UG to direct all RADIUS traffic into the Cisco RPMS system.

Cisco RPMS processes all RADIUS preauthentication and VPDN auth requests and can proxy all auth and accounting messages. As a completely stateless RADIUS proxy, Cisco RPMS leverages the standards-based RADIUS Proxy State attribute when proxying requests. Cisco RPMS also provides means for multi-host deployment, so you can configure each UG to use a list of Cisco RPMS servers (RASERs) as the RADIUS proxy devices. Similarly, the Cisco RPMS system itself provides means for communicating with lists of other AAA servers.

Fault tolerance, including the ability to proxy to lists of AAA systems, and leverage of UG AAA server list, is described in Chapter 4, "Fault Tolerance." For information on configuring fault tolerance, refer to Chapter 5, "Configuring Cisco RPMS Fault Tolerance" of the *Cisco Resource Pool Manager Server Configuration Guide*.

Note Cisco RPMS 2.0.1 alone does not support IP Pooling or full AAA filtering. Consult your Cisco representative to determine whether a AAA proxy, such as Cisco Access Registrar, should be provisioned in your network to complement Cisco RPMS 2.0.1 proxy and SLA management.

Accounting and Billing Support

AAA accounting start and stop records are created in the universal gateway for every call and are forwarded to Cisco RPMS. You can configure Cisco RPMS to create and manage its own customer-based call detail records (CDRs), and to forward accounting records to other AAA billing or accounting systems. Certain Cisco IOS releases permit routing accounting records directly from the UG to both Cisco RPMS systems and independent accounting or billing systems.

When configured for creating CDRs, Cisco RPMS provides additional information pertaining to preauthentication rejection reasons (before accounting), and to policy limit information. Every customer profile has a dedicated CDR, which makes it easy to use the CDR as data for billing or other customer specific purposes.

You can provision Cisco RPMS ASAP to proxy Voice accounting records to one billing system and Dial records to another AAA server.

For more information on CDRs, see the "Overview: Reports and Monitoring" section on page 2-16.

Specifying User-Based Authentication

Another new feature in Cisco RPMS is specifying which authentication mechanism to use for customer calls. This information is sent to the UG from Cisco RPMS during the RADIUS preauthentication phase as a vendor specific attribute (VSA). It is possible to specify any of the following values for authentication method:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)
- MS-CHAP

Managing Non-Cisco Equipment

As mentioned in the "Overview: Port Policies and Service Level Agreement Enforcement" section on page 2-8, Cisco RPMS allows you to manage Cisco universal gateways or non-Cisco equipment.

By implementing the RADIUS protocol, Cisco RPMS allows you to manage non-Cisco NAS types, providing they conform to the RADIUS RFC standards. The RPMS Dynamic Class Loader further provides a means to associate calls from a UG or NAS with a specific call-processing model, allowing the product to extend to manage non-RFC compliant NASs. For more information regarding customized RADIUS drivers, contact your Cisco representative.

Overview: Port Policies and Service Level Agreement Enforcement

Cisco RPMS manages port policies and enforces service level agreements within a network. Service policies allow for partitioning the network or networks of various sizes to help manage policy, while port policies allow you to guarantee port limits for each of your ISPs.

Customer and Policy Modeling

With Cisco RPMS, you can enforce policies as either customer-based (limits per customer) or port-based (limits per overflow pool or trunk group).

Cisco RPMS uses a *customer profile* to organize elements to constitute a policy that impacts a wholesale customer. The customer profile provides an association between calls and specific policy enforcement. It also allows you to readily organize and quickly monitor specific customers and customer SLA policies by policy name. You can abstract and associate a call request for a customer by associating combinations of multiple attributes, such as DNIS numbers, call type, domain name and trunk groups with that customer.

Shared overflow pool profiles allow you to readily organize and quickly monitor pooled resources that are shared among customers that have exceeded their guaranteed limits (overflow) and are contending for first come, first served access to a shared virtual pool.

The building blocks for customer and policy modeling and call association are described next.

Session Limits

Cisco RPMS provides three limits for a customer profile: session count limit, oversubscription limit and shared overflow limit.

Session Count Limit

The session count limit determines the maximum number of standard guaranteed sessions allowed for the customer.

Oversubscription Limit

The oversubscription limit determines the maximum number of guaranteed sessions allowed above the session count limit. For example, when a wholesale service provider, say Wholesaler1, defines a port policy with a retail ISP, say BigISP, the wholesaler guarantees a specific number of ports that the ISP can use. Let's say that Wholesaler1 guarantees BigISP 100 ports, charging a standard rate for the first 75 ports that BigISP uses and a higher, premium rate for each call beyond the 75th. This agreement wherein a premium rate is charged for a percentage of a retail ISP's guaranteed ports is called an oversubscription policy.

Shared Overflow Limit

The shared overflow limit determines the maximum number of overflow sessions allowed above the total guaranteed limit (session count limit plus oversubscription limit). Overflow, like oversubscription may also be used to charge a premium rate, but it is different from oversubscription in that it is not guaranteed: multiple customers compete on a first come, first serve basis for a virtual pool of shared resources. Any customer who exceeds the guaranteed limit and overflows to this pool competes for access to the remaining shared resources in the pool. As a result, any customer session accepted in an overflow pool counts against both that individual customer's overflow limit, and the shared session limit associated with the shared overflow pool. With guaranteed customer limits, sessions for that customer count against the guaranteed session limit for that customer alone.

Session Limit Enforcement

When the session count limit is reached and if both the oversubscription limit and the shared overflow limit are set to zero, new incoming calls will be rejected until the session count falls below the barrier. If the oversubscription limit is greater

than zero, all new sessions are accepted and marked as oversubscription for call handling and accounting, until the oversubscription limit is reached. When the oversubscription limit is reached, new calls will be accepted only if the shared overflow limit is greater than zero. New calls are then marked as shared overflow. When the shared overflow limit is reached, any new calls are rejected.

Note When a call is identified as a shared overflow call, it maintains overflow status throughout its duration, even if the number of current sessions returns below the total guaranteed limit.

The maximum number of allowed sessions for the customer profile equals the session count limit plus the oversubscription limit plus the shared overflow limit.

Enabling oversubscription or shared overflow sessions help isolate excess sessions for preferred customers or premium rates. Use oversubscription or shared overflow sessions to encourage customers to adequately forecast bandwidth usage or for special events when normal session usage has been exceeded. For example, if a customer has a corporate-wide program and expects many people to request remote access, enable a large number of oversubscription or shared overflow sessions and charge a premium rate for the excess bandwidth requirements.

The following table shows some scenarios:

Table 2-1 Session Count, Oversubscription and Session Overflow Limit Settings Combinations

Session Count Limit	Oversubscription Limit	Shared Overflow Limit	Call Handling
1. 0	0	0	Reject all calls.
2. 10	0	0	Accept up to 10 sessions.
3. 10	10	0	Accept up to 20 sessions. Mark sessions 11 – 20 as oversubscription sessions.
4. 10	10	10	Accept up to 30 sessions. Mark sessions 11 - 20 as oversubscription sessions and mark sessions 21-30 as shared overflow sessions.
5. 0	10	0	Accept up to 10 sessions and mark all sessions as oversubscription sessions.
6. 0	0	10	Accept up to 10 sessions and mark all sessions as shared overflow sessions.
7. All	0	0	Accept all calls.
8. 0	All	0	Accept all calls and mark all calls as oversubscription.
9. 0	0	All	Accept all calls and mark all calls as shared overflow.

Cisco RPMS also lets you set threshold settings and generate e-mail notification when thresholds for session count limit, shared overflow limit, call rejections, and overflow call rejections have been exceeded. For more information, see the "Alerts" section on page 2-24 or the "Overview: Thresholds" section on page 6-14.

Upgrading from Previous Cisco RPMS Releases to Cisco RPMS Release 2.0.1

If you upgrade to Cisco RPMS Release 2.0.1 from Cisco RPMS Releases 1.1 or 2.0, there are rules for reassigning limits to the new limits. The limits are outlined in the following table:

Table 2-2 1.x Limits and the Corresponding 2.x Limits

1.*x* Limits	Corresponding 2.*x* Limits
Session Count Limit	Session Count Limit
Overflow Limit	Oversubscription Limit

Upgrading from Cisco RPMS Release 1.1

For Cisco RPMS Release 1.*x*, the 1.*x* overflow limit is equivalent to the 2.*x* oversubscription limit. In this upgrade scenario, the 2.*x* shared overflow limit is set to 0.

Table 2-3 1x Limits and the Corresponding 2.x Limits

1.*x* Limits	Corresponding 2.*x* Limits
Session Count Limit	Session Count Limit
Overflow Limit	Shared Overflow Limit

Upgrading from Cisco RPMS Release Cisco RPMS Releases 2.0.1

For Cisco RPMS Release 2.0, the 2.0 overflow limit is equivalent to the 2*x* shared overflow limit. In this upgrade scenario, the 2.*x* oversubscription limit is set to 0.

Resource Management

By using Cisco RPMS, you can direct a Cisco UG to handle a call with customer specific modem configuration and authentication methods (PAP versus CHAP). Sample modem configuration options are provided in the "Overview: Configuring RADIUS Vendors" section on page 2-9 of the *Cisco Resource Policy Management System Configuration Guide*.

Cisco vendor specific attributes are sent in the preauthentication accept message from Cisco RPMS to direct the UG how to handle the call. Cisco RPMS provides a Web based interface to administer and associate any VSA with a customer based call accept message providing a means to control even greater resource and service management for each call.

DNIS Groups

With Cisco RPMS, you can use DNIS groups to identify policies to calls. A DNIS group is a configured list of DNIS numbers corresponding to the numbers dialed by particular customers, service offerings, or both.

Cisco RPMS checks the DNIS number of inbound calls against the configured DNIS groups or the default DNIS group. If there is a match, Cisco RPMS enforces the information configured in the customer profile to which the DNIS group is assigned. If there is no match, Cisco RPMS uses the default DNIS group with other matching criteria (such as trunk group or domain) to match a customer profile. The call is rejected if a default customer profile is not configured.

For 800, 888, or similar toll-free numbers, the DNIS number the telco switch delivers to the UG differs from the user-dialed number. Cisco RPMS provides an optional reference number in each DNIS group for numbers requiring public network database lookup, which results in a PSTN-routable DNIS.

For Cisco RPMS ASAP, you can provision DNIS tech prefixes on the UG and voice platforms. These prefixes do not interfere with Cisco RPMS DNIS mapping, because Cisco RPMS filters the tech prefix and only uses DNIS to match customer profiles.

For information on configuring DNIS groups, refer to the "Overview: Configuring DNIS Groups" section on page 3-4 of the *Cisco Resource Pool Manager Server Configuration Guide*.

DNIS Wildcards

With Cisco RPMS, you can specify DNIS wildcards in a DNIS group. The supported wildcard character is * (the star symbol) and can be used only at the beginning or at the end of a DNIS string (for example, *1212 or 555*). You can construct DNIS groups with DNIS wildcard entries and fully qualified DNIS numbers.

Restrictions for using DNIS wildcards:

- You can use the wildcard only to prefix or postfix a DNIS number (not both). No other combination of wildcards is valid.
- You can use only one wildcard character in a DNIS number.
- The wildcard must appear as part of a DNIS number and an “all wildcard” DNIS (for example, “*”) is invalid. Valid DNIS wildcards may include *4455 or 512*. Invalid wildcards might include 44*732, *467*, and *51*6 for example.

Overlapping wildcards are permitted; Cisco RPMS uses the most specific wildcard to match an incoming request. For example, you can add 45* and 455* as wildcards. If an incoming request has the DNIS=4556776, Cisco RPMS matches the more specific wildcard 455*.

Trunk Groups

You can use the trunk group feature in Cisco RPMS 2.0.1 to assign trunk groups to customer profiles with a call type for resource allocation and management.

Cisco RPMS maps incoming call requests to a specific customer profile by DNIS, call type (any, digital, speech, v.110, and v.120), and originating trunk.

Trunks and trunk groups allow you to manage multiple calls originating from different trunks that are calling the same DNIS to map to separate customer profile. For example, there may be one POP that handles calls from multiple geographical regions, and a Virtual Private Dial-up Network (VPDN) customer who wants to enforce counts for those different areas.

Cisco RPMS uniquely identifies a trunk by the following parameters: originating UG IP address, and the port and slot number the trunk terminates on. The information tells you from which area the call came, and this forms a trunk. A

trunk group comprises several trunks. Without the UG IP address, slot and port information, the DNIS and call type map to just one customer profile, and you cannot distinguish from where the call came.

Trunk Wildcards

Cisco RPMS allows for using wildcards with trunks. Cisco RPMS defines trunks as a combination of three things: the originating UG, and the slot and port numbers that the trunk terminates on. You can use wildcards to define the slot and port numbers for a trunk.

The wildcard character is the same as for DNIS groups, an asterisk (*). You can create trunk groups with trunk wildcard entries and fully qualified trunks.

Restrictions for using trunk wildcards:

- You can use the wildcard only to replace an entire slot or port. You cannot use a partial number wildcard. For example, you can use slot=4 or port=23, but not slot=*4 or port=*3
- You can use only one wildcard character in a trunk slot number, and only one wildcard character in a trunk port number. For example, you could use slot=* or port=*, but not slot=** or port=**.

Overlapping wildcards are permitted; Cisco RPMS uses the most specific wildcard to match an incoming request. For example, you can have the following slot and port numbers:

```
trunk(1.1.1.1, slot=*, port=23)
trunk(1.1.1.1, slot=4, port=*)
```

If an incoming request uses trunk(1.1.1.1, slot=4, port=23), Cisco RPMS matches the more specific wildcard, trunk(1.1.1.1, slot=4, port=*).

Note The trunk matching algorithm in Cisco RPMS is as follows:

Exact matches first (from the previous example, this would be slot=4, port=23)

Exact slot and wildcard port (slot=4, port=*)

Wildcard slot and exact port (slot=*, port=23)

Wildcard slot and wildcard port (slot=*, port=*)

For information on configuring trunk wildcards, refer to the Configuring Trunk section of the *Cisco Resource Policy Management System Configuration Guide*.

Voice Over IP Address Groups

Cisco RPMS ASAP introduces IP address groups to help manage terminating Voice calls. You can use these groups to track VOIP calls originating from a remote gateway or gatekeeper.

IP address groups contain a list of IP addresses defining all points of ingress from an ITSP into a wholesale service provider's network. After defining the IP address group, you can associate it with a customer profile.

Customer profiles can contain IP address groups in addition to DNIS groups. An IP address group is similar to a DNIS group and identifies a wholesale customer. However, since the IP address groups are always used for voice calls, do not specify a call type while associating the IP address group with a customer profile.

Cisco RPMS ASAP also supports Cisco Inter Zone Clear Tokens (IZCT). As such, call origin and customer profile can be determined according to IZCT received from gateway.

Overview: Reports and Monitoring

Cisco RPMS records report data so that you can analyze the network dial service and troubleshoot it. The report data is a snapshot of the current activity and appears on the screen, can be printed, or saved to disk in a comma-separated value (.csv) format.

For more information on reports, see Chapter 6, "Reporting and Accounting" of the *Cisco Resource Policy Management System Configuration Guide*.

For more information on accounting, see the "Accounting and Billing Support" section on page 2-6.

Call Detail Records

Cisco RPMS creates CDRs for all calls initiated to and terminated from the UG. CDRs use the Start/Stop and Call ID attributes to identify whether the call is being initiated or terminated, and once initiated, how to map the call detail records to a particular call. Both the start and stop call detail records for a specific call use the same call ID.

The CDRs contain additional information for troubleshooting and planning capacity. For example, Cisco UGs do not generate accounting records for unanswered calls; Cisco RPMS does.

The CDRS also can identify whether a call was accepted against a guaranteed, oversubscription or shared overflow policy.

CDRs store the following information for each call:

- Timestamp
- UG IP address, and slot or port
- Call type
- Customer profile name, active sessions, oversubscription limit and count, and shared overflow
- MLPPP session ID
- Modem connect or disconnect transmit and receive speed
- Reject reason and terminate cause
- VPDN tunnel ID, home gateway IP address, and name

For more information on CDRs, including the CDR fields, refer to the "Generating Call Detail Records" section on page 6-18 of the *Cisco Resource Pool Manager Server Configuration Guide*.

Web-Based Reporting

With Cisco RPMS, you can use a Web browser to generate reports. Cisco RPMS reports show real-time data for network analysis and troubleshooting. You can generate and display different reports onscreen, save the reports to a disk, or print them. You can even set up special filters to see specific parameters in a report.

For more information on running or filtering reports, refer to Chapter 6, "Reporting and Accounting" of the *Cisco Resource Pool Manager Server Configuration Guide*.

Overview: Shared Overflow Pools

Cisco RPMS is used to pool resources across a wholesale network and enforce limits so that no one customer can oversubscribe their resource usage and adversely impact another guaranteed service agreement.

Cisco RPMS also provides shared overflow pools. The wholesaler can actually have more resources within the network than the sum of guaranteed limits. Additional resources are just one reason the wholesaler may contract shared overflow. Customers who exceed their guaranteed session limits can access the additional resources on a first come, first serve basis.

With Cisco RPMS, you can control overflow access through shared overflow pools. To create an overflow pool, you must give it a name, and then associate one or more trunk groups to it. The basic idea behind Cisco RPMS is to allow a pool of resources to be shared across customers; overflow pools extend that idea for overflow calls. Therefore, shared overflow pools allow wholesalers to maintain and enforce SLA limits that span multiple customers.

Overview: VPDN

You can also use Cisco RPMS to configure session count and overflow limits per VPDN group and to manage sessions.

Note The VDPN session count and shared overflow limit are independent of the customer profile limits.

Session and Overflow Limits

The session count limit determines the maximum number of nonoverflow sessions supported for a VPDN group. When the session count limit is reached, all new VPDN calls using the VPDN group sessions are rejected if shared overflow sessions are allowed. But if shared overflow sessions are not allowed, all new sessions are processed and marked as overflow for call handling and accounting, until the overflow limit is reached.

The VPDN session overflow limit determines how many sessions are allowed beyond the session count limit for the VPDN group. If the overflow limit is greater than zero and shared overflow sessions are enabled, the maximum number of allowed sessions equals the session count limit plus the overflow limit. When the overflow limit is reached, any new calls are rejected.

Enabling VPDN shared overflow sessions helps to isolate excess sessions for preferred customers or premium rates. Use shared overflow sessions to encourage customers to adequately forecast bandwidth usage or for special events when exceeding normal session usage.

VPDN Tunnel Limits

For increased VPDN tunnel management, you can set an IP endpoint session count limit for each IP endpoint that is configured for VPDN groups.

VPDN Tunnel Set Up

When setting up a VPDN tunnel, there are two options to receive the information. Cisco RPMS can provide the information, or retrieve it from a AAA server.

Cisco RPMS monitors tunnel connections enforced by Cisco RPMS. You can configure it to load balance tunnel connections by sending sorted VPDN endpoint (or LNS) lists to the UG.

Trunk-Based VPDN Groups

Trunk-based VPDN groups help you manage networks by associating a domain name or DNIS number with a configured trunk group. For example, let's say you had two trunk groups, San Francisco and San Jose. The VPDN group could manage VPDN calls from the domain "cisco.com" originating from the San Francisco trunk group, or VPDN calls originating from the San Jose trunk group.

Overview: SNMP Support

Another feature Cisco RPMS includes is an SNMP agent to monitor customer or call statistics. SNMP traps are also available for events such as a Cisco RPMS start or stop, a database shut down or start up, exceeded subscription or shared overflow limits, and exceeded rejection limits.

Overview: Fault Tolerance and Resiliency

Cisco RPMS allows you to build fault tolerance and resiliency into your dial service by offering the following features.

Hot Standby and High Availability

Hot standby and high availability allow for redundancy in your deployment. When using hot standby, you can deploy two Cisco RPMS servers and configure them as peers. The servers have identical data configurations and must share the same database for configuration purposes.

High availability is a feature of many of the critical Cisco RPMS components. You can deploy any stateful Cisco RPMS component you want to preserve in case of server failure as part of a high availability pair.

Tolerance to Database Failures

Cisco RPMS provides tolerance to database failure by continuing to accept UG requests if the Oracle database fails. Additionally, in case of database failure, the Cisco RPMS server generates an e-mail notifying the system administrator that the database is down.

Multiple Cisco RPMS Servers Using the Same Datastore

To facilitate a single point of provisioning for all Cisco RPMS servers that need identical configurations, you can connect multiple Cisco RPMS servers to the same Oracle database.

Database Replication

You can replicate databases by using the Oracle Database Replication Manager and the Cisco RPMS DBServer component. All configuration is replicated among active and fail-over Cisco RPMS servers.

Note Cisco RPMS does not need Oracle replication; RPMS provides its own database synchronizing and polling with triggers. Rather, replication is provided as an option for customers wanting to replicate the database across several hosts.

Cisco RPMS Autorestart

Cisco RPMS server can automatically start and stop an individual process, independent of all others, if it detects a process failure. Cisco RPMS can also automatically restart any Cisco RPMS processes that failed abnormally.

Detection of Universal Gateway Failures

An SNMP manager is used to detect UG failures. The SNMP manager regularly polls the UG to ensure it is still alive, regardless of the returned value. This method assumes that the UG being checked has an SNMP agent running.

For more detailed information on fault tolerance features, see Chapter 4, “Fault Tolerance.”

Cisco RPMS Fail-Over Detection

When deploying a high availability distributed solution with dual Cisco RPMS servers and multiple RASERs, you can configure the RASER servers to detect AAA and Cisco RPMS server failures.

For more information on fault tolerance, see the "Tolerance to AAA Server Failure" section on page 4-4.

Overview: Cisco RPMS System Administration

There are many administrative tasks that Cisco RPMS can perform to keep your deployment running smoothly. These include the ability to add multiple administrators, run debugs, and send alert notifications based on exceeded thresholds, or in case the Cisco RPMS system is down.

For more detailed information on Cisco RPMS administration, including logging and cleaning up directories, refer to Chapter 2, "Cisco RPMS Administration" of the *Cisco Resource Pool Manager Server Configuration Guide*.

Oracle Database

Cisco RPMS maintains administrative information and persistent policy content in an Oracle database that can be shared with other applications and run on a dedicated server.

The Cisco RPMS system no longer requires Oracle replication when the database is on a separate server—Cisco RPMS provides its own polling and database triggers, so that data changes are automatically updated across an entire Cisco RPMS system.

Cisco RPMS caches all policy affecting information and does not rely on connection to the database for call processing. CDRs and reports are managed independently from the Oracle database.

User Administration

With Cisco RPMS, you can add multiple administrators to manage resources and dial services. Each administrator is identified by a username and password combination.

The following privileges are available for Cisco RPMS administrators. Each level includes the privileges of that level and all lower levels.

- Root:
  - Setting up an administrator
  - Administering and configuring an application
  - Configuring and viewing reports
- System administrator:
  - Administering and configuring an application
  - Configuring and viewing reports
- View only:
  - Viewing administration and configuration applications
  - Configuring and viewing reports
- Report and monitor:
  - Monitoring applications
  - Configuring and viewing reports

Debugging

One way for administrators to track issues is by using the debug feature. Debugging is a useful way to troubleshoot problems such as incoming calls not being accepted, or other calls that are being rejected by the system. With

debugging, you can use specific Cisco IOS commands depending on what you want to accomplish (resolve communication issues), or the symptoms for which you are debugging (troubleshooting, call analysis).

Warning **Using the debug feature generates a great deal of network traffic. As a result, your network performance can slow down dramatically. Take this into consideration when turning on the debug feature.**

You can use the Cisco IOS debugging commands for different reasons:

- Communication issues between the universal gateway and Cisco RPMS
- Customer profile issues

For more information on debugging, refer to the *Cisco Resource Pool Manager System Troubleshooting Guide* Chapter 4, "Universal Gateway IOS Debugging."

Alerts

Alerts keep administrators aware of any system errors or exceeded thresholds. If Cisco RPMS detects a system error that requires a shut down or when a configured threshold is exceeded, Cisco RPMS generates notification. The alerts are configured by the system administrator and are sent as e-mail notifications.

If alert logging is enabled, you must enter an e-mail server or source address.

- The mail server is the SMTP host configured to send out Cisco RPMS alert messages.
- The alert source address is the e-mail address of Cisco RPMS.

You must also specify an e-mail address (or multiple addresses, up to 20) to receive notification whenever a Cisco RPMS notification is generated.

Thresholds

Alert notifications are based partly on threshold settings; a threshold is simply a number or percent that tells Cisco RPMS at what point to send e-mail alerts. The system administrator configures the threshold, and when that limit is exceeded, the system sends out an e-mail informing the system administrator.

For example, if you set a session count threshold to notify you when 85 percent of a customer's allotted ports are in use, then as soon as the count reaches 85 percent, Cisco RPMS generates an e-mail notification that the threshold was exceeded.

For more information, see the "Overview: Thresholds" section on page 6-14.

Overview: Call Discrimination

With call discrimination, you can specify DNIS groups and call types to block and disconnect calls before they are assigned to Cisco UG resources, or before any other Cisco RPMS processing occurs. For example, use call discrimination to restrict a specific DNIS group to only modem calls by creating call discrimination settings for the DNIS group and the other call types (any, speech, digital, V.110, and V.120).

Note If a call type is not available in a RADIUS message, Cisco RPMS uses a call type of *any*.

Overview: Data Over Voice Bearer Services

Data over Voice Bearer Services (DOVBS) is a dial service that uses a customer profile to direct data calls with a speech call type to HDLC controllers.

This feature is pending RADIUS support in Cisco gateways and is not available for Cisco RPMS 2.0.1.

CHAPTER 3

Scaling Cisco RPMS Deployments

Topics in this chapter include:

Overview: How Does Cisco RPMS Scale?

Cisco RPMS 2.0.1 is designed to work in networks of varying sizes. You can tailor Cisco RPMS to work as an entry-level single server for a small deployment, or you can scale it out to manage a much more complex deployment of redundant policy server clusters in a large network or WAN. You can even implement a mixture of the two by starting small and growing the deployment as your needs change and scaling out as your network or customer base grows.

Scaling the Network

Cisco RPMS is designed for scalability as your network or needs grow. Distributing the deployment allows you to easily complete management tasks without bringing the entire system down, or stopping incoming calls.

Upgrading Software

As an example of scalability, let's describe what would happen in a network of Cisco AS5300s. One Cisco AS5300 can support 240 ports; so, for 120,000 ports, you need 500 Cisco AS5300s.

Theoretically, you *can* use one Cisco RPMS server, or an HA pair, to manage the entire network. But this might prove unwieldy for tasks such as upgrading the Cisco RPMS software; to load the new software, you must shut down the single server, interrupting all traffic routed to that server. As a result, no calls can connect during that time.

There is another workaround option. Instead of shutting down the existing server, you could add a new server to the network, then change the UG configurations to point to that new server. Eventually, after changing all 500 configurations, the network would point to the new service.

But neither of these options are optimal for running a network. A better way to plan for future growth is to deploy a distributed Cisco RPMS system. Although it is feasible for one RASER to manage this traffic load, a preferable method is to deploy multiple RASERs instead. By deploying four RASERs instead of just one, all the UGs are provisioned to point to those four RASERs.

The next time you want to upgrade the software, you can redirect the call set up operations from the UGs to where the RASERs direct AAA messaging. You can add a new Cisco RPMS server to the network, and simply change the configurations on the four RASERs so that all 120,00 ports point to new service.

You can also provision the RASER to automatically accept any preauthentication, taking preauth decisions out of service for all devices pointing to the RASER.

The UG can be provisioned to communicate with a list of AAA RADIUS servers. This provides fault tolerance in case one server is unreachable—the UG can redirect traffic to another RADIUS server in the list.

Cisco RPMS can leverage this feature for service upgrades in an operational network. Again, start with UG provisioned to communicate with four RASERs—more than enough to carry traffic load. To upgrade one of the RASERs, simply take the RASER out of service. The UG will dynamically redirect call control to the remaining RASERs. You can perform a software upgrade on the out-of-service RASER, and then when finished, put the RASER back in service to resume communication with the UG. This process repeats itself across each RASER until all RASERs have upgraded service (if necessary). Note that this upgrade is accomplished without shutting down operations.

Increased ISP Traffic

Distributing the network deployment serves another purpose as well. As a network provider needs increase, Cisco RPMS can address the increased demand. As the network grows, you can deploy more distributed components to share the increased processing load, and still manage the services across those components.

Different deployments have different call processing rates. For instance, a single server can manage 100,000 ports at 100 calls per second; by adding additional servers, you can increase the number of calls the system can handle.

How you deploy Cisco RPMS depends on your specific needs. This chapter describes different Cisco RPMS deployments, starting with a simple one, and growing increasingly larger and more complex.

Overview: How to Use Cisco RPMS

You are planning your Cisco RPMS deployment and tossing around ideas of how big of a deployment you need. You want it to be big enough to provide for your current needs, but also flexible enough to scale out as your needs or network grow. How do you provide for both?

Let's use a wholesaler named Wholesaler1 as an example of a Cisco RPMS customer. Wholesaler1 wants to provide access to its customers, Big ISP and LittleISP, so it allots 100,000 ports for each ISP. That means that both BigISP and LittleISP can call into the Wholesaler1 system and log in to the Internet—start sessions.

The 100,000 port count is more than enough for LittleISP, but BigISP is consistently busy, and needs lots of ports. BigISP regularly reaches its 100,000 port limit and may even try to exceed the 100,000 ports it has purchased from Wholesaler1. So, when BigISP's call number 100,001 dials in, Wholesaler1 must decide what to do. It might allow BigISP to access additional, previously configured overflow ports (usually at a premium price), or it can reject call 100,001 and all further BigISP calls until some existing BigISP calls end their sessions and free up ports.

To adequately plan for how many ports it needs, Wholesaler1 has to consider many things:

- How many ISP customers does Wholesaler1 actually have? In this case, will 200,000 ports be enough or too many for BigISP and LittleISP?
- Who uses those ports the most? Sporadic, infrequent users or heavy users? Is BigISP consistently surpassing its allotted 100,000 ports while LittleISP is consistently underusing them? Since ports are not dedicated, users choose from an available pool, which means Wholesaler1 oversells subscriptions. For example, it may have 200,000 BigISP users assigned to an available pool because it's unlikely all 200,000 will access them at the same time. This may not bother infrequent users like LittleISP, but heavy users like BigISP might agree to purchase more expensive, guaranteed access.
- How do the ISPs access those ports? For example, access by using ISDN versus modems (ISDN access might cost more) or access for speech versus data.

- How long are average sessions? Cisco RPMS processes a certain number of calls in the system at a time—the shorter the calls, the more users the system can process per hour (for example, 20 minute calls = three users per port per hour while 30 minute calls = two users).

Wholesaler1 must consider all these options when allotting ports. If customers like BigISP continue to grow and need access to more ports, then Wholesaler1's network will also continue growing to accommodate BigISP. Wholesaler1 might also experience network growth due to other factors, such as implementing new services or signing up new ISP customers.

Once Wholesaler1 decides how many ports to allot to the ISPs, it must also decide how to deploy Cisco RPMS to manage those ports.

Cisco RPMS comprises small, efficient components. Each self-contained piece handles a core functionality. This architecture gives a great deal of flexibility, since you can deploy the components in various combinations to meet performance and scale requirements.

The Cisco RPMS components include:

- Remote Access Service Router (RASER)
- Secure Web Server and SNMP Access
- Oracle Database
- High Availability (HA)

Large deployments might also include:

- Overflow server
- Service Level Agreement Server
- POP Manager
- Customer Profile Manager server

For more information on the components, see the "Defining Cisco Resource Policy Management System" section on page 1-1.

The following sections describe how you can run these components on a single server or distributed across multiple servers for flexibility and large service enforcement capacity.

Overview: Single or Dual Servers

A basic Cisco RPMS deployment can use a single server, or can introduce dual servers to provide redundancy and failure recovery. Following are descriptions and examples of both implementations.

Single Server

Figure 3-1 depicts a basic single Cisco RPMS host in a wholesale dial deployment. In this example, there is one RASER, raser1, with the IP address 171.12.40.1. UG always communicates directly with RASERs, so the term RASER is used when looking at Cisco RPMS from the UG perspective. For a single host deployment, the RASER also hosts all policy processing; the host is also referred to simply as the Cisco RPMS server.

Figure 3-1 also depicts a very basic call flow sequence diagram and the rudimentary Cisco IOS commands for provisioning a Cisco UG to enable preauthentication with Cisco RPMS.

Note In the figure, the UG is provisioned to direct all RADIUS messages to the RASER.

The Cisco RPMS RASER can proxy authentication and accounting messages directly to remote service sites, or route them to an existing AAA proxy (or list of proxies) in the wholesale network as depicted.

Figure 3-1 Single Cisco RPMS Host Wholesale Dial

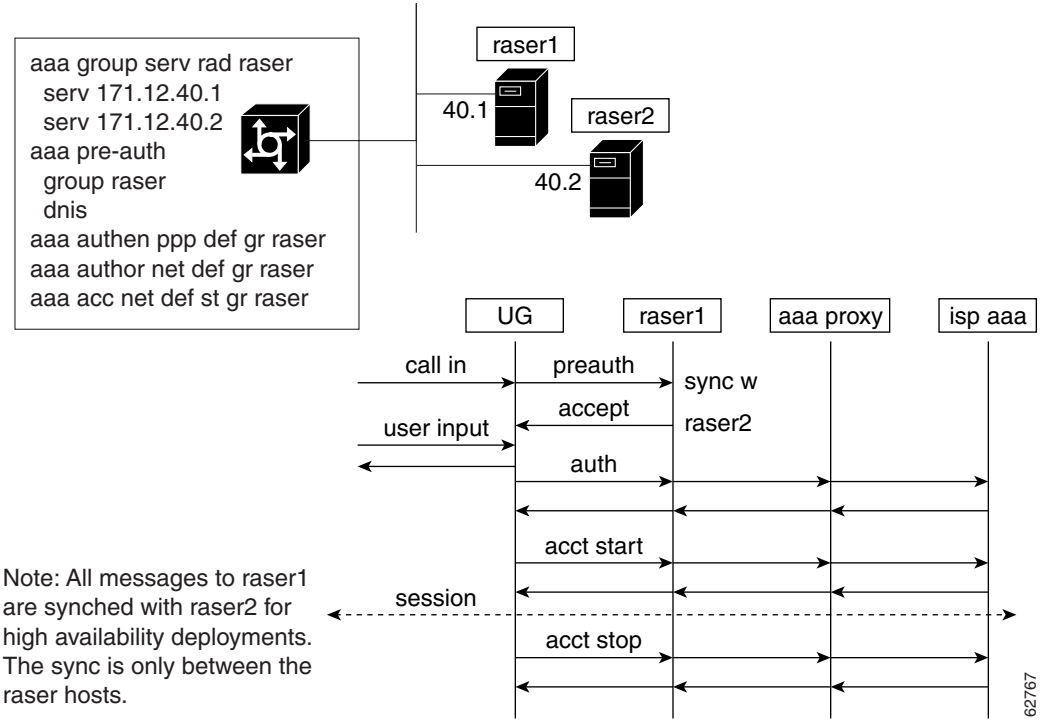

Redundant High Availability Dual Server

Figure 3-2 depicts redundant high availability (HA) dual Cisco RPMS hosts in a wholesale dial deployment. Deploying HA pairs ensures redundancy; if one server fails, the other server uses its stored cache to keep tracking calls. Each message received by an HA server is forwarded to its HA peer, ensuring that each server has an active, or hot, call state in the memory cache of managed calls.

Note The HA mechanism is used for providing redundancy to servers that must maintain in memory state.

Normally a RASER does not use HA, because pure RASER functionality is stateless. However, for dual server deployment, each server plays the role of a RASER as well as policy processing servers. The dual servers use HA to maintain the policy and active call state. The next sections describe scenarios in which stateful HA components can run on hosts independently from stateless RASERs.

Also depicted is a very basic call flow sequence diagram and the rudimentary Cisco IOS commands for provisioning a Cisco UG to enable preauthentication with the Cisco RPMS hosts.

In this example, there are two RASERs, raser1 with the IP address 171.12.40.1, and raser2 with the IP address 171.12.40.2. In the single host deployment, the UG directed all RADIUS messages to raser1. In this deployment the UG directs the messages to the RASER group. It doesn't matter which RASER the UG communicates with because raser1 synchronizes information with raser2, which caches the active call state in its memory.

The Cisco RPMS RASERs can proxy authentication and accounting messages directly to remote service sites, or route them to an existing AAA proxy (or list of proxies) in the wholesale network as depicted.

Figure 3-2 High Availability Dual Cisco RPMS Hosts Wholesale Dial

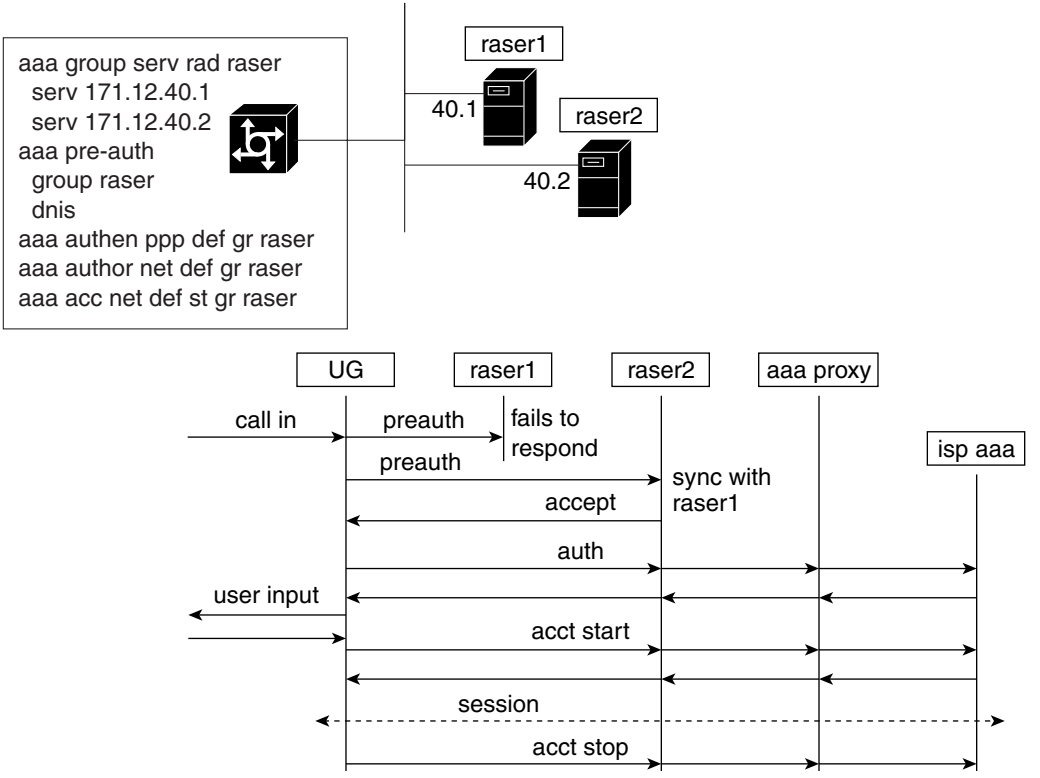

Figure 3-3 describes the failure recovery process when raser1 fails in an HA environment.

1. The UG directs the RADIUS messages to raser1, but raser1 fails to respond.
2. If raser1 does not respond, raser2 does by accepting the preauthentication.

   The two RASERs have remained in sync, so the raser2 uses its cached active call state, and the UG does not even notice an interruption.
3. raser2 accepts the preauthentication, and attempts to forward the information to raser1.

   This information will be lost if raser1 is out of service. However, when raser1 comes back up, the HA mechanism forces a “sync” with raser2, so the two servers have identical, current information.

Figure 3-3 High Availability Cisco RPMS Failure Recovery

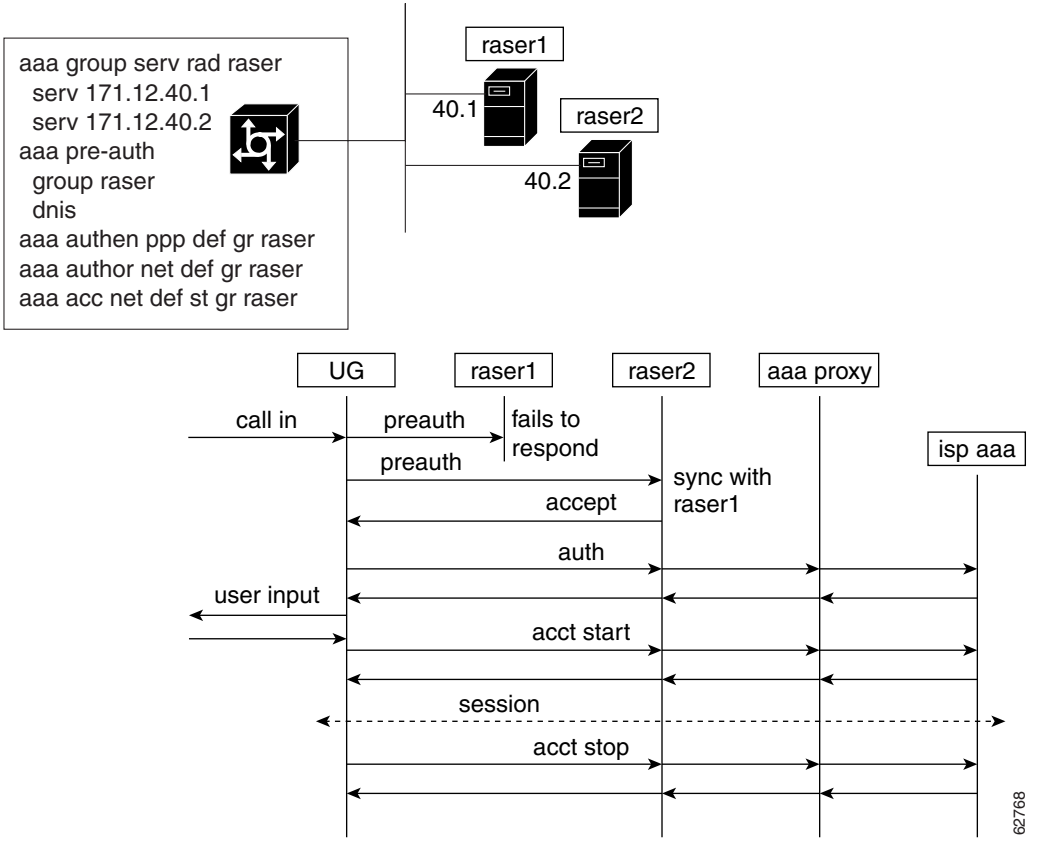

Overview: Partial Distribution

Cisco RPMS 2.0.1 introduces the option for partial distribution. You can install Cisco RPMS 2.0.1 so that it processes on one server, or distribute it to run across multiple server types:

- Cisco RPMS policy processing servers
- Cisco RASER servers (also referred to as Cisco RPMS proxies)

A Cisco RPMS policy processing server manages profile counts and enforces service level agreements. The policy processing server also contains the Web server. As such, customers connect to a Cisco RPMS policy processing server for Web access.

You can deploy the policy processing servers in high availability pairs to provide redundancy should one server fail.

Note With high availability, two Cisco RPMS servers act as a high availability pair or hot standby server pair. A Cisco RPMS server caches active call state in memory.

By providing an HA pair, this active state is mirrored across both servers. Should either server fail, the system continues to operate with the remaining server without losing state.

The Cisco RASER server manages RADIUS traffic between UGs, the Cisco RPMS policy processing servers, and a list of RADIUS AAA proxy servers. From the UGs, or other non-Cisco UGs in the network, the RASERs appear as RADIUS servers. These RASERs are completely stateless, so several can be deployed for RADIUS load sharing across the UG and for failover. The RASERs are stateless, so there is no reason to use HA across RASERs.

Stateless RASERs and Single Customer Policy Manager

Figure 3-4 depicts a Cisco RPMS partial distribution deployment and call flow. In Figure 3-4, the UG only sees the RASER group and not the Cisco RPMS Customer Profile Manager (CPM) policy processing server. The call flow only depicts calls to one RASER because this call flow is identical for either RASER for preauthentication instead of using the AAA proxy server.

Here is the call flow for Figure 3-4:

1. Every AAA message from the UG is directed to a RASER.
1. Each time a call setup comes in, the RASER tries to match the call to a specific customer. If the RASER finds a match, it sends a message to the CPM server.
2. The CPM server matches that call with its customer profile; if there is no customer profile or no match, it rejects the call, and the process ends.
3. If the CPM accepts the call, it passes an Accept message back to the Cisco RASER, which then relays the message to the UG.
4. The UG dedicates a resource (for example, a modem) to handle the accepted call, and the call connects.
5. Authentication, authorization, and accounting calls are all received directly from the UG by the RASER and forwarded to another AAA server. The accounting information is also sent to the CPM server when CDRs are enabled.

Figure 3-4 Cisco RPMS Partial Distribution Call Flow

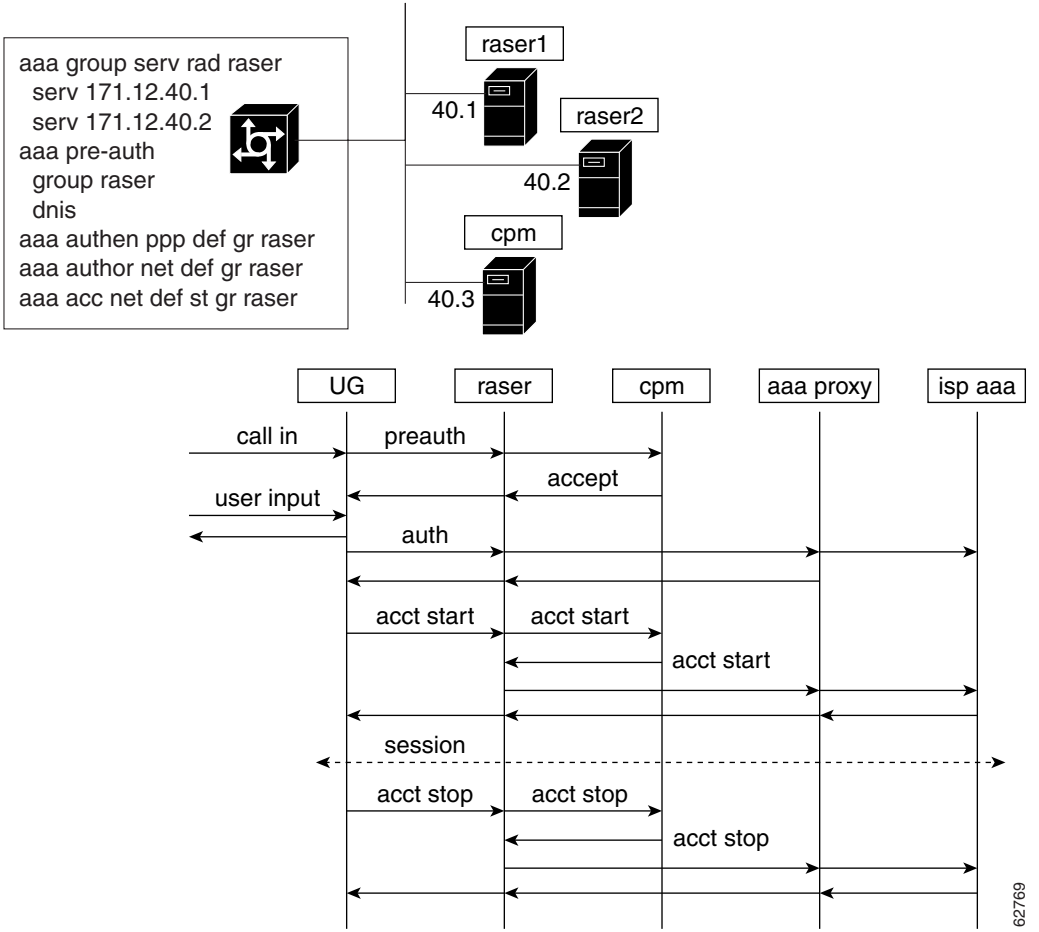

Stateless RASERs and Dual High Availability Customer Policy Manager

Figure 3-5 depicts a Cisco RPMS partial distribution deployment and call flow. In this example, there are two CPMs, cpm1 with the IP address 171.12.40.3 and cpm2 with the IP address 171.12.40.4. The UG only sees the RASER group, not the CPM. The call flow only depicts calls to one RASER because this call flow is identical for either RASER for preauthentication instead of using the AAA proxy server.

If the Call Succeeds

Here is the call flow for Figure 3-5:

1. Every AAA message from the UG is directed to a RASER.
1. Each time a call set up comes in, the RASER tries to match the call to a specific customer. If the RASER finds a match, it sends a message to the CPM server.
2. The CPM server matches that call with its customer profile; if there is no customer profile or no match, it rejects the call, and the process ends.
3. If the CPM accepts the call, it passes an Accept message back to the Cisco RASER, which then relays the message to the UG.
4. The UG dedicates a resource (for example, a modem) to handle the accepted call, and the call connects.

Authentication, authorization, and accounting calls are all received directly from UG by RASER and forwarded to another AAA server. The accounting information is also sent to the CPM server when CDRs are enabled.

Figure 3-5 Cisco RPMS Partial Distribution Call Flow for HA CPMs (Successful Call)

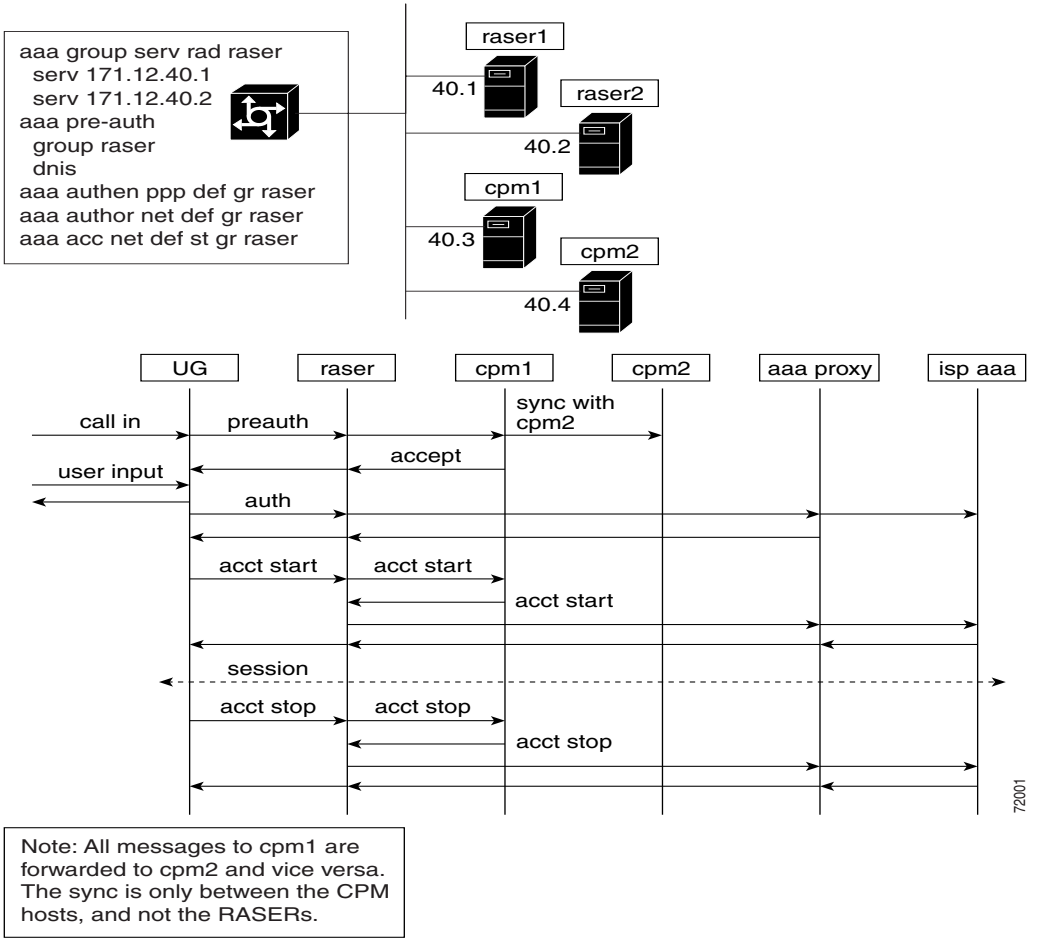

Note: All messages to cpm1 are forwarded to cpm2 and vice versa. The sync is only between the CPM hosts, and not the RASERs.

If the Call Fails

In this example, there are two CPMs, cpm1 with the IP address 171.12.40.3 and cpm2 with the IP address 171.12.40.4.

Figure 3-6 shows the recovery process if cpm1 fails in an HA environment:

1. The UG directs the RADIUS messages to raser1, which forwards the messages to cpm1.
2. cpm1 fails to respond.
3. If cpm1 does not respond, the RASER forwards the message to cpm2.
4. cpm2 responds, and accepts the preauthentication from the RASER.

   The two CPMs are synchronized, so cpm2 uses its cached active call state, and the UG does not even notice an interruption.
5. After cpm2 accepts the preauthentication, it attempts to forward the information to cpm1.

   This information is likely lost if cpm1 is out of service. However, when cpm1 comes back up, the HA mechanism forces a synchronization with cpm2, so the two servers have identical, current information.

Figure 3-6 Cisco RPMS Partial Distribution Call Flow for HA CPMs (Failed Call)

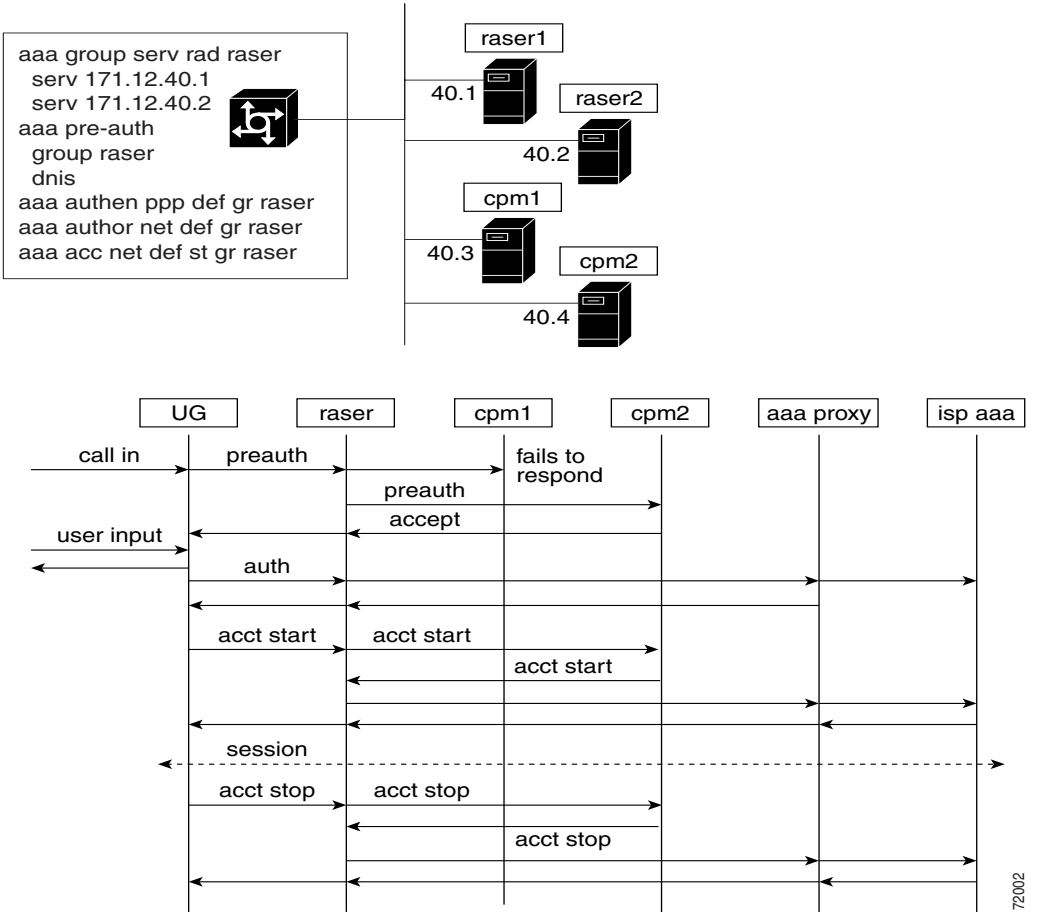

Overview: Full Distribution

In a wide-area network, multiple POPs can route SLA requests to dedicated customer SLA processors. An SLA processor is a process (or thread) that enforces the SLA associated with a particular customer.

In a fully distributed Cisco RPMS deployment, the following occurs:

- Customer-based processing threads run on the Customer Profile Management servers.
- Shared overflow pool management runs on POPM servers.
- RADIUS message routing between these policy processing agents and the UG is controlled by RASERs.

In large deployments, or ASAP deployments that require very high overall system call processing rates, multiple RASERs are deployed to scale and handle all UG incoming and outgoing RADIUS messages. CPM servers are deployed to scale and handle increased volumes of customer specific RADIUS content. POPM servers are deployed to scale and handle increased traffic throughput for specific trunk groups and shared overflow pool traffic.

A managed network can evolve from a partially to fully distributed Cisco RPMS system. The RASERs are a key service bridge between the hardware ports and the Cisco RPMS policy processing servers. In fact, if an initial deployment has several RASERs that are deployed where some are redundant, then policy processing can be incrementally added to the network as the network grows or service bandwidth increases, without necessarily reprovisioning the UG. Each RASER can manage traffic from hundreds of UGs. To extend or modify SLA management for the set of ports communicating with a RASER, you can configure the RASER to redirect RADIUS messages to an updated Cisco RPMS server. Because RASERs are stateless, a deployment with redundant RASERs can take a RASER out of service for upgrades with minimal impact to operations.

As the wholesale customer service expands through a larger part of the network, you can introduce more RASERs to capture traffic from the new ports and direct that traffic to existing policy processing servers. You also can migrate the SLA processors from one CPM host to run on a set of CPM servers—with each server dedicated to only a subset of customers.

Stateless RASERs and Separate Customer Policy Manager and POP Manager

The POPM and CPM servers each run various services through the network. The POPM, for example, is responsible for services, such as overflow pooling and active ports. The services running on the CPM include customer and VPDN session limits.

Figure 3-7 depicts a Cisco RPMS full distribution preauthentication call flow:

1. The UG sends a call to the RASER for preauthentication.
2. The RASER recognizes it as a call for a specific customer and sends it on to that customer's CPM server.
3. However, the CPM server has exceeded its threshold of calls; the CPM labels the incoming call as an overflow call and sends it back to the RASER.
4. The RASER forwards the overflow call to the POPM, which responds with an overflow acknowledgment.
5. The RASER then sends the call accept to the UG to dedicate a resource.

Figure 3-7 Cisco RPMS Full Distribution Pre-Auth Call Flow

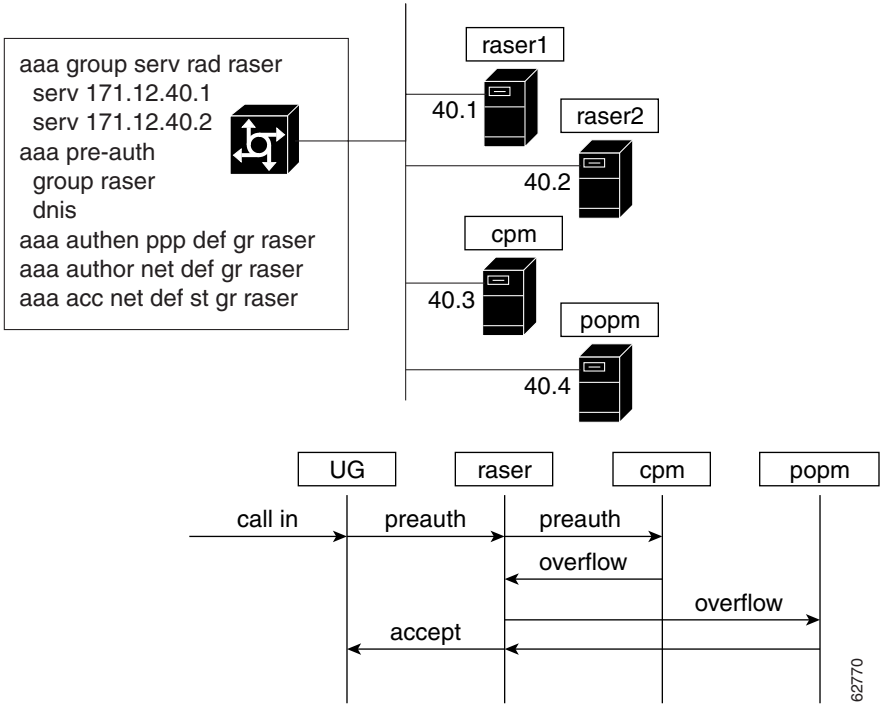

Stateless RASERs, Multiple Customer Policy Managers and Multiple POP Managers

Cisco RPMS allows you to use three major components to manage various customers on different hosts:

- RASERs
- CPMs
- POPMs

For example, in a wide-area network, you can deploy the components to manage one customer, BigISP, on one host, and other ISP customers on different hosts that are separate from the BigISP host. Even though the ISPs are on different hosts, both the BigISP and other ISPs in the WAN can access a shared pool of ports spanning multiple POPs or Virtual POPs (VPOPs).

In such a deployment, the POP or VPOP uses a POP manager to control the shared overflow pools for ports.

Note All calls come in from a specific POP. Because of that, the calls can only be considered for sharing ports that are affiliated with that POP.

Figure 3-8 depicts an abstract view of how to distribute the Cisco RPMS RASERs, CPMs, and POPMs with AAA servers to manage large networks spanning multiple POPs (or virtual POPs) while enforcing SLA limits across the multiple POPs.

Figure 3-8 Cisco RPMS Large Scale Deployment

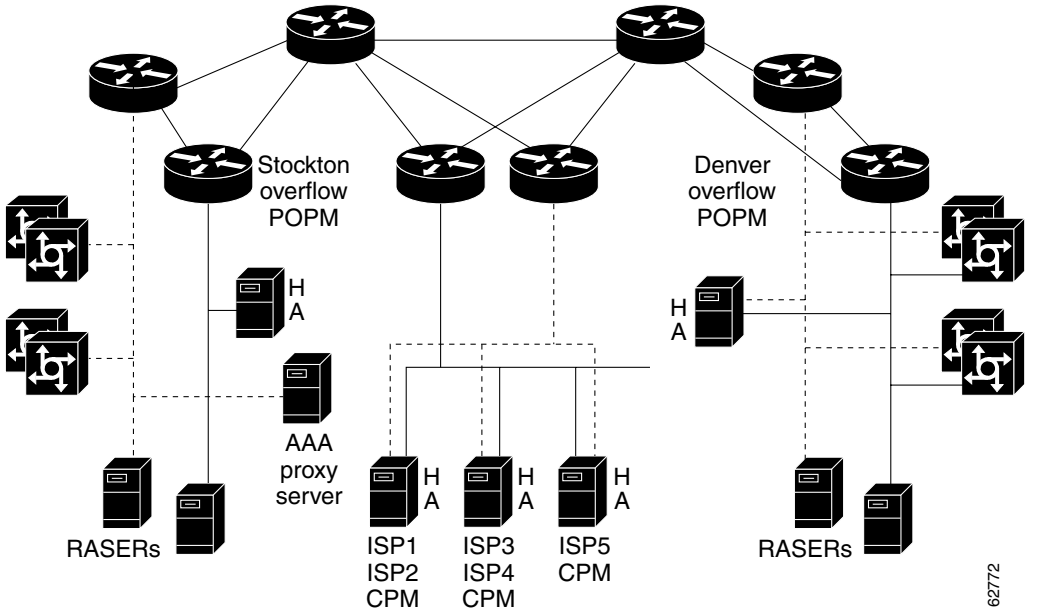

Fault Tolerance

Topics in this chapter include:

Overview: Fault Tolerance

Cisco RPMS allows you to build fault tolerance and resiliency into your dial service offerings. Fault tolerance and resiliency consist of the following features:

- Hot Standby
- Tolerance to Database Failures
- Cisco RPMS Autorestart
- Detection of Universal Gateway Failures
- Tolerance to AAA Server Failure

Hot Standby

Cisco RPMS provides the hot standby (or high availability) feature to protect against server failure. Any stateful Cisco RPMS component server, such as policy processors, can be deployed as a high availability (HA) pair. In this configuration, the servers continually synchronize with each other by exchanging messages.

Providing an HA pair means the active call counts and other network states are mirrored across both servers. If either server fails, the system continues to operate with the remaining server without losing state.

Both servers in a hot standby pair act as peers. There is no primary or backup server. So both servers are fully functional and independently capable of handling all the network traffic. Because of this, you can provision a network so that some devices (such as RASERs) communicate with one server, while others communicate with its peer. This type of provisioning can help with network load sharing.

Both servers in a hot standby pair replicate their state to the peer by exchanging messages over a reliable, TCP link, so that both servers receive the combined network traffic meant for the pair. The servers make their policy decisions based on the total network activity. For this reason, network load sharing does not mean less load or lower CPU utilization for the servers participating in hot standby configuration.

When installing a new server installation or after rebooting it, you can synchronize it with its peer at any time by using a CLI command. You can also configure HA servers to automatically synchronize with their peer at startup.

For more specific configuration information, refer to Chapter 5, “Configuring Cisco RPMS Fault Tolerance” in the *Cisco Resource Policy Management System Configuration Guide*.

Tolerance to Database Failures

Cisco RPMS needs database connectivity for configuration purposes only; database connectivity is not required for call processing. So, if the Oracle database fails, the Cisco RPMS still continues to accept UG requests.

When Cisco RPMS detects a database failure, it generates an e-mail. Additionally, an SNMP trap is generated when the connectivity to the database fails and when it is restored.

Cisco RPMS Autorestart

Cisco RPMS can detect server process failures and can automatically restart any Cisco RPMS processes that failed. The following components are monitored:

- Policy Processors—These processors enforce policies and also handle GUI requests for reports. When a policy processor fails, you must immediately restart it.
- RASER—This processor accepts RADIUS messages from the UG and forwards them to the appropriate destination (for example, the Cisco RPMS server or AAA proxy). Restarting this process is necessary for all call processing.
- Oracle RDBMS—Since the Oracle server is not distributed as a Cisco RPMS component, the Cisco RPMS system can detect Oracle RDBMS failure, but Cisco RPMS cannot automatically restart the Oracle database. The database administrator must manually restart Oracle. During this period, the DBServer detects an Oracle database failure and notifies the administrator.
- DBServer—If the DBServer shuts down when Oracle fails, the DBServer cannot restart until the database administrator brings up the Oracle database server. The autorestart process must verify that the Oracle database is operational before restarting the DBServer. The Oracle tool *tnsping* detects Oracle availability. If the DBServer fails, it automatically restarts provided the Oracle database is operational.
- FastTrack Web server—The FastTrack server consists of multiple http daemons that are monitored and restarted after a failure.
- Acme Web server (servlets)—The Acme Web server process is monitored and restarts after failure.

Detection of Universal Gateway Failures

Cisco RPMS implements a heartbeat checker mechanism that allows Cisco RPMS to test whether or not a UG is still active.

You can configure Cisco RPMS to use SNMP to automatically poll the UGs that are in the UG list. When polling, Cisco RPMS sends an SNMP Get request to each UG. If an SNMP agent is running on the UG, a Get request returns a message with either a time in hundredths of seconds that the encapsulated agent has been

running, or a "no such name" error message, which signifies that the agent and the UG are alive, regardless of the returned value. However, if a UG does not respond to the request, then Cisco RPMS resets all the corresponding active calls.

For information on configuring the heartbeat checker or for the heartbeat configuration definitions, refer to the "Overview: The Universal Gateway Heartbeat" section on page 2-31 in the *Cisco Resource Policy Management System Configuration Guide.*

Tolerance to AAA Server Failure

To enhance fault tolerance to AAA server failures, Cisco RPMS allows you to create a prioritized list of AAA servers. Cisco RPMS RASERs use this list to determine the destination of authorization and accounting messages received from the UG.

The RASERs forward messages to the AAA server with highest priority. If the RASERs detect that this AAA server has failed, they switch over to the server with the next highest priority. When the RASERs reach the end of this list, they continue from the top of the list again.

For more specific configuration information, refer to Chapter 5, "Configuring Cisco RPMS Fault Tolerance" in the *Cisco Resource Policy Management System Configuration Guide*.

CHAPTER 5

Cisco RPMS Deployment Scenarios

Topics in this chapter include:

Overview: Cisco RPMS Deployment Scenarios

You can use Cisco RPMS in various network deployment scenarios, by implementing different combinations of hardware or by supporting multiple services. You can tailor Cisco RPMS to work as a single server for a small deployment, or scale it to manage a much more complex deployment of redundant policy server clusters in a WAN. You can also implement a mixture of the two.

The network deployments types can include:

- Wholesale dial networks
  - Non-VPDN
  - VPDN
- Originating and terminating H.323 and SIP VOIP networks
- ASAP networks

Overview: Wholesale Dial Network Deployments

You can deploy Cisco RPMS in a wholesale dial network that comprises components such as:

- A UG
- Cisco RPMS
- AAA proxy server managed by the wholesaler
- Remote call control systems, such as an ISP's AAA server

Note The AAA proxy server is an optional component.

Wholesale dial allows wholesalers to lease equipment to several ISPs. The dial wholesaler provides the hardware equipment to the ISPs, including dial access equipment, such as the UG, to bridge the connection between the PSTN and the Internet. The ISPs manage their end customers and maintain their user names and passwords, which are enforced by the ISP's AAA server.

Implementing Cisco RPMS allows the wholesalers to enforce SLAs, so that the ISPs cannot adversely impact one another's access rights to the network resources.

You can deploy wholesale dial networks in several different ways, including:

- Distributed model
- Multiple POPs
- Single server deployment
- Dual server deployment
- Large-scale deployment

These various methods of deployment allow you to choose how to run Cisco RPMS processing, either over one server or across multiple server types. For more information on these deployment methods, see Chapter 3, "Scaling Cisco RPMS Deployments."

You can use Cisco RPMS to manage wholesale dial non-VPDN or VPDN services. Non-VPDN services use DNIS numbers (the called telephone numbers) or trunk groups to answer or reject calls based on the customer profile configuration. With VPDN services, Cisco RPMS determines whether to authenticate the call with a home gateway through a VPDN tunnel based on the type of VPDN services configured in Cisco RPMS.

Figure 5-1 shows an example of a Cisco RPMS wholesale dial deployment with the wholesaler's equipment on the left (the Cisco RPMS server, UG group, and optional AAA proxy server). The ISP's equipment is depicted on the right as the Customers A and B AAA servers.

In Figure 5-1, the deployment is a DNIS-based non-VPDN deployment with an optional AAA proxy server.

Figure 5-1 DNIS-Based Non-VPDN Dial Cisco RPMS Deployment with Optional AAA Proxy Server

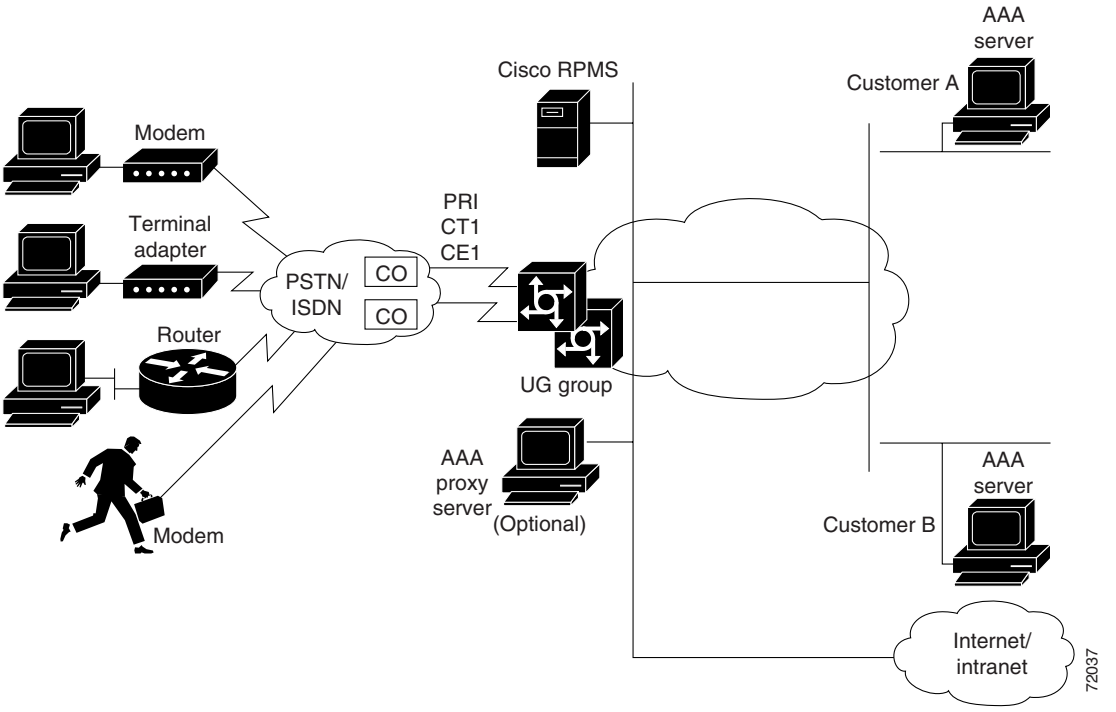

Figure 5-2 shows a basic call flow for a call request coming into a network deployment and the rudimentary Cisco IOS commands for provisioning a Cisco UG to enable preauthentication with the Cisco RPMS RASER.

Note The universal gateway is provisioned to direct all RADIUS messages to the RASER.

The Cisco RPMS RASER can proxy authentication and accounting messages directly to remote service sites, or route them to an existing AAA proxy (or list of proxies) in the wholesale network as shown.

Figure 5-2 Basic Call Flow Cisco RPMS Wholesale Dial Deployment

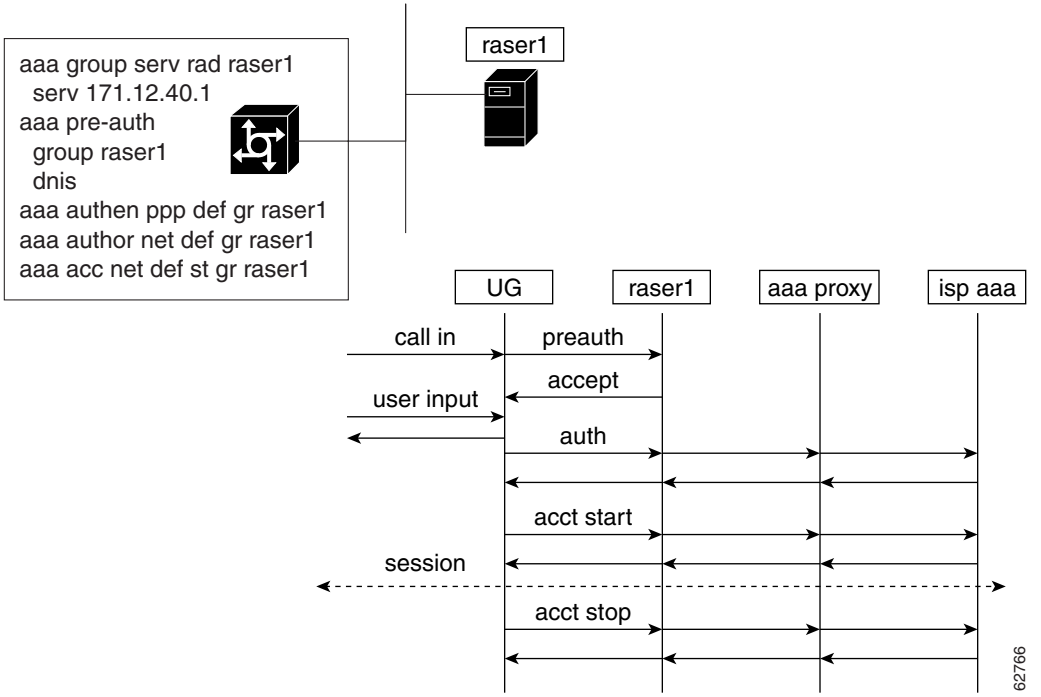

Single Server Deployment

A basic Cisco RPMS deployment can use a single server to host the RASER and an optional AAA proxy server. For small deployments, you can use one machine to host a single instance of each Cisco RPMS component, or a subset of the entire system, for easier management.

Figure 5-3 shows two single server deployments in a large WAN. In the figure, the CPM and the components are installed on the same machine as the RASER.

Note In Figure 5-3, there are two separate single server deployments: one on the left, and one on the right. In this scenario, both the deployments are independent of one another.

A single Cisco RPMS server can manage over 100,000 ports at over 100 calls per second and has only one entry in the UG configuration.

Note Although a single server can scale to manage large call throughput, there are many fault tolerance and service upgrade benefits to deploying multi-server deployments.

Figure 5-3 Cisco RPMS Wholesale Dial Single Server Deployments

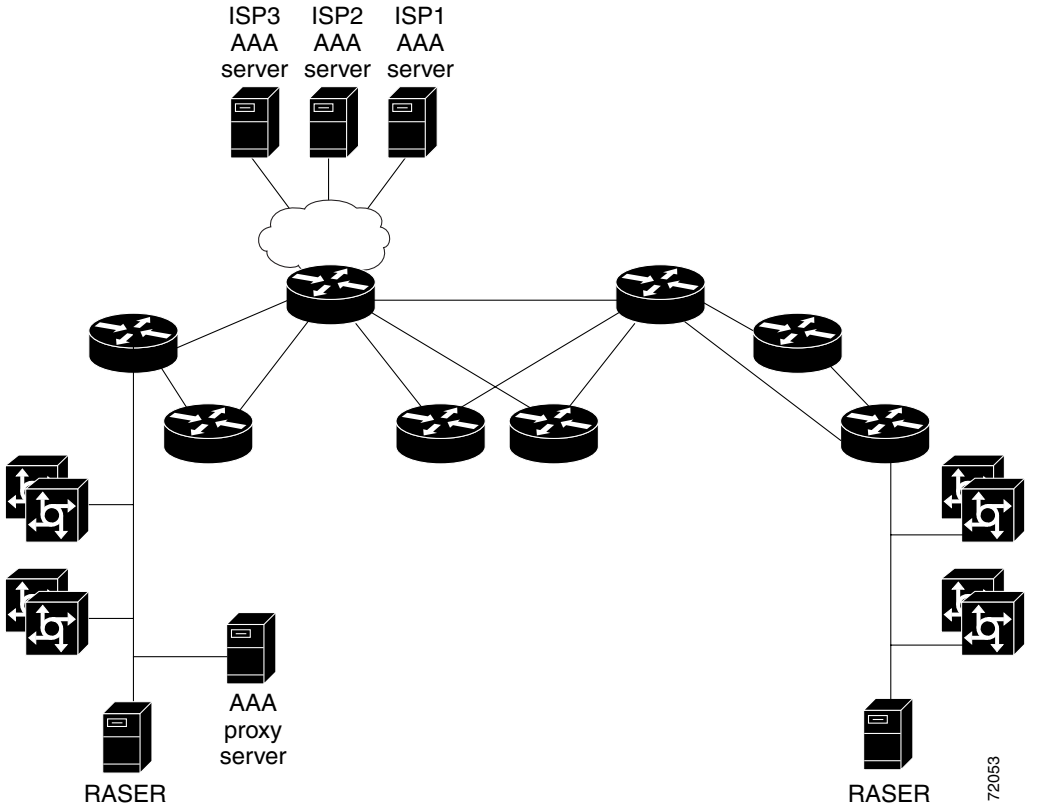

Figure 5-4 shows a a sample call flow for a single server deployment. It is a basic call flow for a call request coming into the network and the rudimentary Cisco IOS commands for provisioning the UG to enable preauthentication with the Cisco RPMS RASER.

Figure 5-4 Single Cisco RPMS Host Wholesale Dial

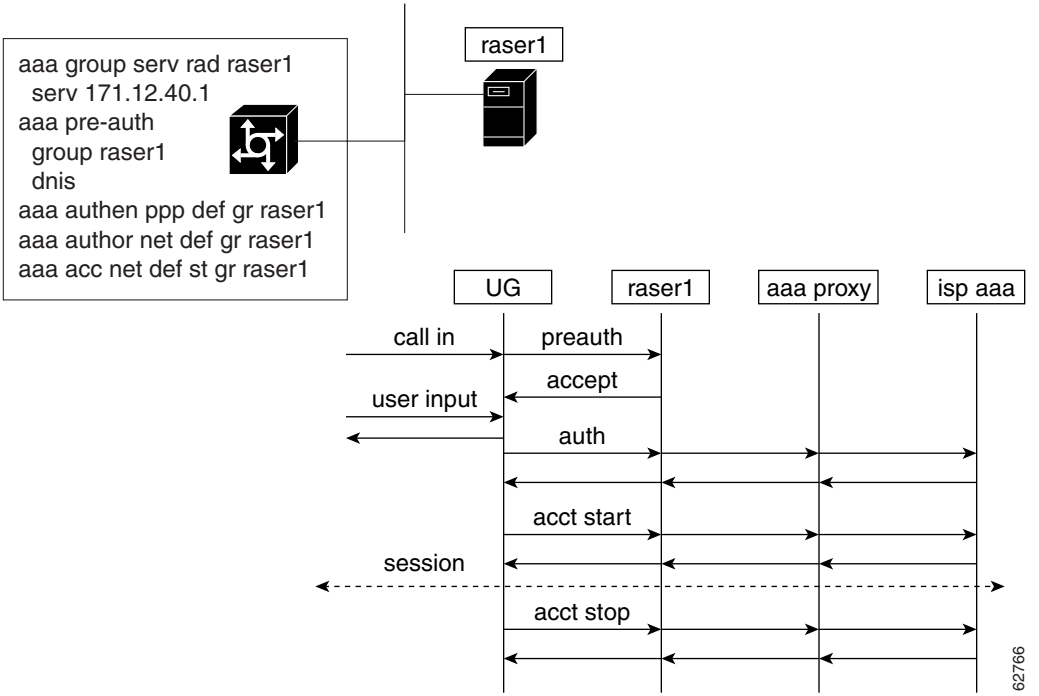

A call request comes into the deployment on the left side of the network in Figure 5-3.

Note The RASER at the bottom left of the figure is raser1.

Preauthentication Request

The steps for an accepted preauthentication request are:

1. The preauth request comes in through the UG on the left.
2. The UG routes the request to raser1.

3. raser1 (a single server) determines the customer profile that matches the request, such as ISP1, and verifies that the profile has not exceeded its limits. If the profile has not exceeded the limits, raser1 sends the call handling for that customer profile back to the UG.

Authentication Request

In this single server deployment, a AAA proxy server is also deployed.

The steps for an accepted authentication request are as follows:

1. The UG sends an authentication request to raser1.
2. raser1 forwards the request to the AAA proxy server.
3. The AAA proxy server determines which ISP the message is for and forwards it to that ISP's AAA server (such as ISP1).
4. The ISP1 AAA server accepts the authentication request and sends an accept message back to the AAA proxy server.
5. The AAA proxy sends the accept message back to raser1.
6. raser1 sends an accept message to the UG.
7. The UG starts a session.

Accounting

1. The UG sends an accounting request to raser1.
2. raser1 updates its call model and CDR records and forwards the request to the AAA proxy server.
3. The AAA proxy server determines which ISP the message is for, and forwards it to that ISP's AAA server (such as ISP1).
4. The ISP1 AAA server handles the request, and sends an acknowledgement message back to the AAA proxy server.
5. The AAA proxy sends the message back to raser1.
6. raser1 sends the message to the UG.

Dual Server Deployment

To ensure high availability, you can use a a dual Cisco RPMS server installation with a AAA proxy instead of a single server installation.

The dual (HA) server can also manage over 100,000 ports over 100 calls per second, but the UGs must have two entries in the UG configuration. Not all of the UG configurations need the two Cisco RPMS servers entered in same order. For instance, half could point to RPMS1 as “primary” and point to RPMS2 as “secondary.” The other half could be swapped.

Figure 5-5 shows two dual server HA deployments in a large WAN. This figure also has a AAA proxy server for each deployment, which like the other components, is installed on a machine separate from the RASER.

Note In Figure 5-5, there are two separate dual server deployments: one on the left, and one on the right. In this scenario, both the deployments are independent of one another.

Figure 5-5 Cisco RPMS Wholesale Dial Dual Server Deployments

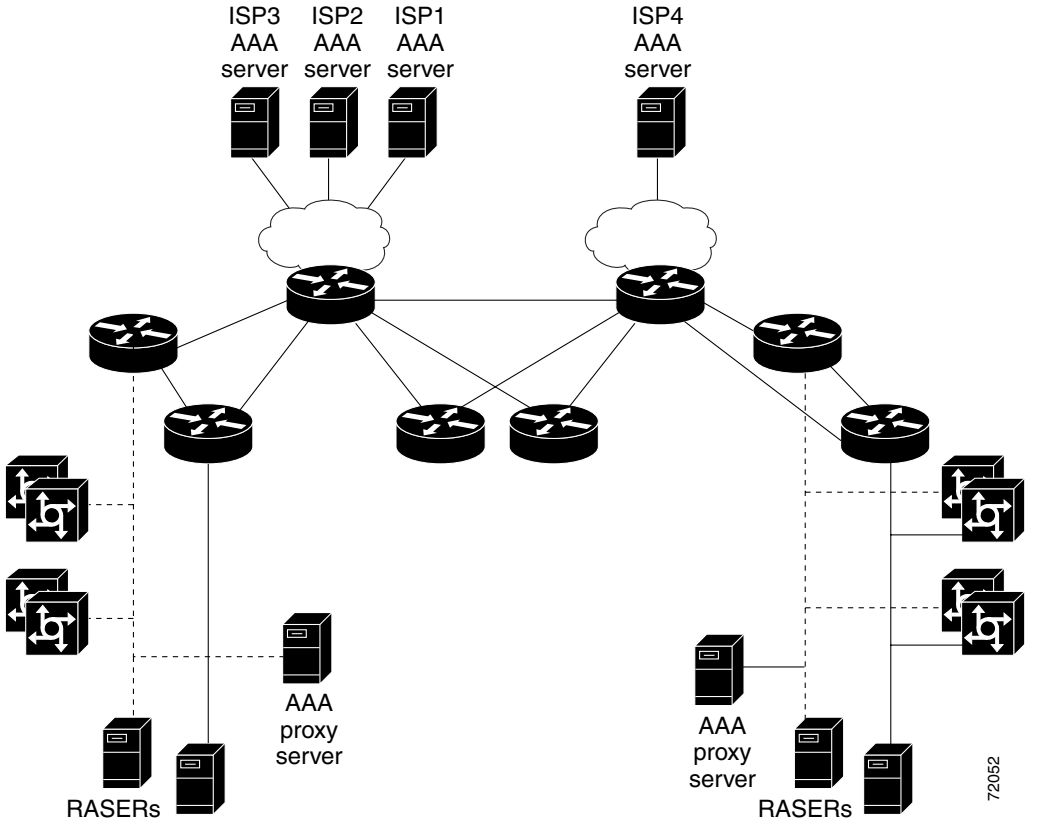

Figure 5-6 shows a a sample call flow for a dual server HA deployment.

Figure 5-6 High Availability Dual Cisco RPMS Hosts Wholesale Dial

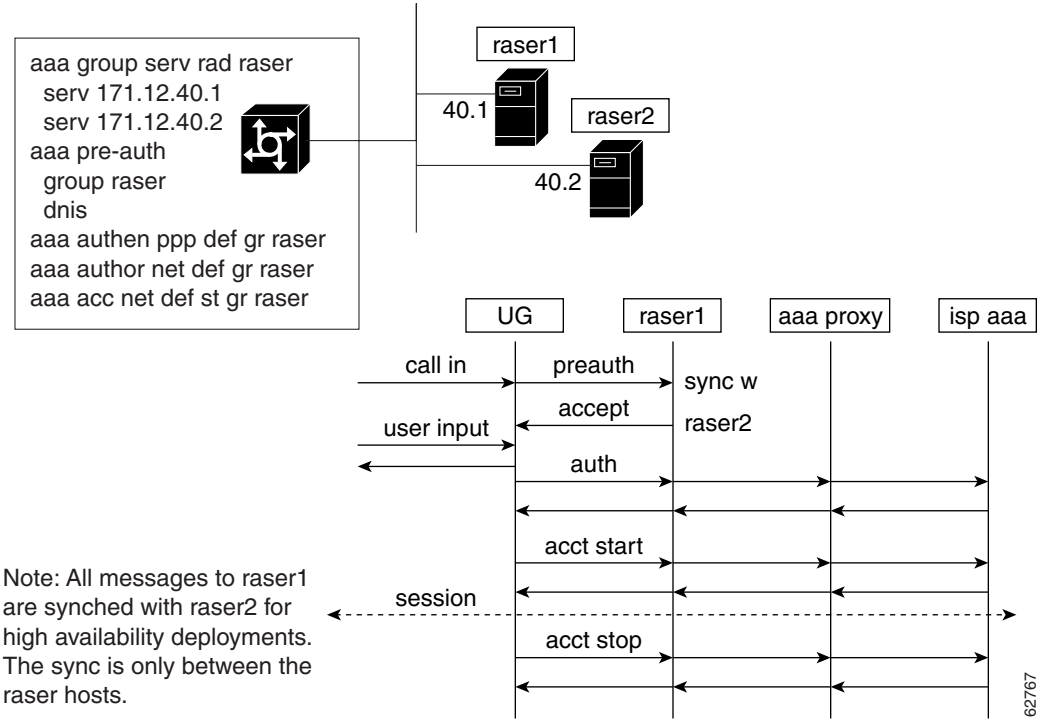

Note raser1 and raser2 are part of the RASER group on the left.

Preauthentication Request

For example, the steps for an accepted preauth request for a call request coming into the deployment on the left of Figure 5-5 are:

1. The preauth request comes in through the UG on the left.
2. The UG routes the request to raser1.
3. raser1 synchronizes the message with raser2, so that both servers share the same information.

4. raser1 determines which customer profile matches the request, such as ISP1, and sends the call handling for that profile back to the UG.

Authentication Request

In this dual server deployment, the AAA proxy server is also deployed.

The steps for an accepted authentication request for this call flow are as follows:

1. The UG sends an authentication request to raser1.
2. raser1 forwards the request to the AAA proxy server.
3. The AAA proxy server determines which ISP the message is for, and forwards it to that ISP's AAA server (such as ISP1).
4. The ISP1 AAA server accepts the authentication request and sends an accept message back to the AAA proxy server.
5. The AAA proxy sends the accept message back to raser1.
6. raser1 sends an accept message to the UG.
7. The UG starts a session.

Accounting

1. The UG sends an accounting request to raser1.
2. raser1 synchronizes the message with raser2, so that both servers share the same information.
3. raser1 updates its call model and CDR records and forwards the request to the AAA proxy server.
4. The AAA proxy server determines which ISP the message is for, and forwards it to that ISP's AAA server (such as ISP1).
5. The ISP1 AAA server handles the request and sends an acknowledgement message back to the AAA proxy server.
6. The AAA proxy sends the message back to raser1.
7. raser1 sends the message to the UG.

Multiple Points of Presence

You can deploy Cisco RPMS in a wholesale dial network comprised of many components including universal gateways, RASERs, AAA proxy servers managed by the wholesaler, and remote call control systems, such as an ISP's AAA server. Figure 5-7 shows these components in a large wholesale dial WAN with two POPs: Stockton and Denver.

Note In Figure 5-7, all the equipment belongs to the wholesaler.

Figure 5-7 Cisco RPMS Wholesale Dial Multiple POP Deployment

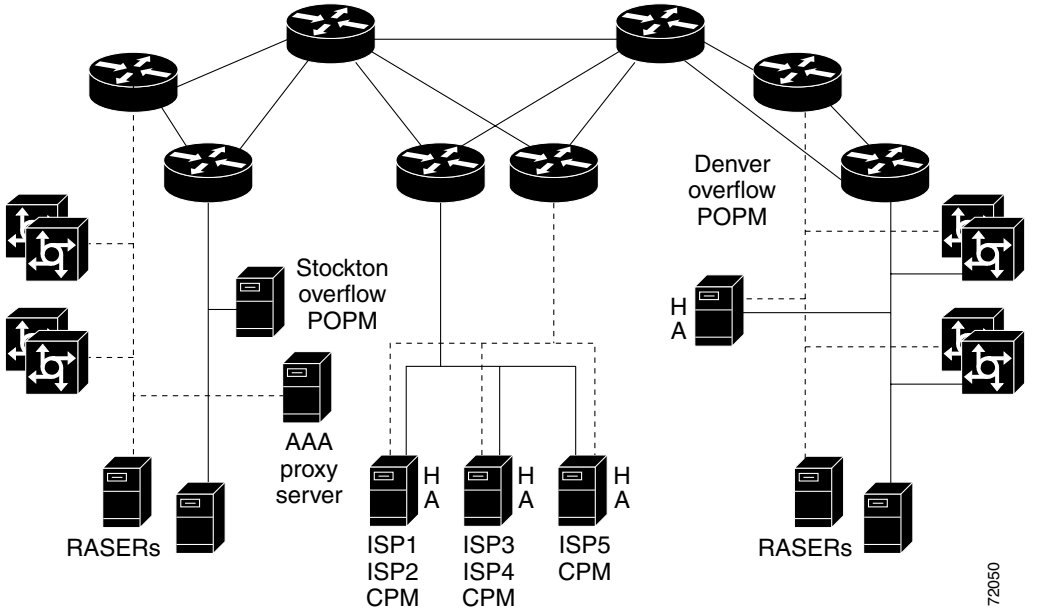

Large-Scale Deployments

Figure 5-8 shows a large-scale overview of a Cisco RPMS deployment network. Like the previous figures, this image shows both the wholesaler and ISP's equipment. The ISP equipment includes ISP1, ISP2, ISP3, and ISP4 AAA servers connecting to the wholesaler's equipment through the cloud.

Large-scale deployments have many advantages. In such a deployment, Cisco RPMS manages the traffic across the different regions by using the AAA proxy servers and RASERs. The Cisco RPMS components allow the wholesaler to route the traffic over all the regions, or to any other POPs that are deployed in the WAN.

The introduction of global CPM servers and the local POPM servers —one in Stockton and one in Denver—in Figure 5-8 shows how Cisco RPMS can scale, so that global policies are applied for ISP usage spanning the POPs. In a deployment such as Figure 5-8, Cisco RPMS can manage traffic from an ISP such as ISP4 originating from either the Stockton or the Denver UGs.

A large-scale deployment also has another advantage: it allows multiple methods for users to connect. For example, users can access the network by using VPDN, non-VPDN, ISDN, or modem services.

Note You can also see the point where the wholesaler's equipment meets the ISP's equipment in Figure 5-1; in that image, the wholesaler's equipment is labeled as the UG group.

Figure 5-8 Large Scale Cisco RPMS Wholesale Dial Deployment

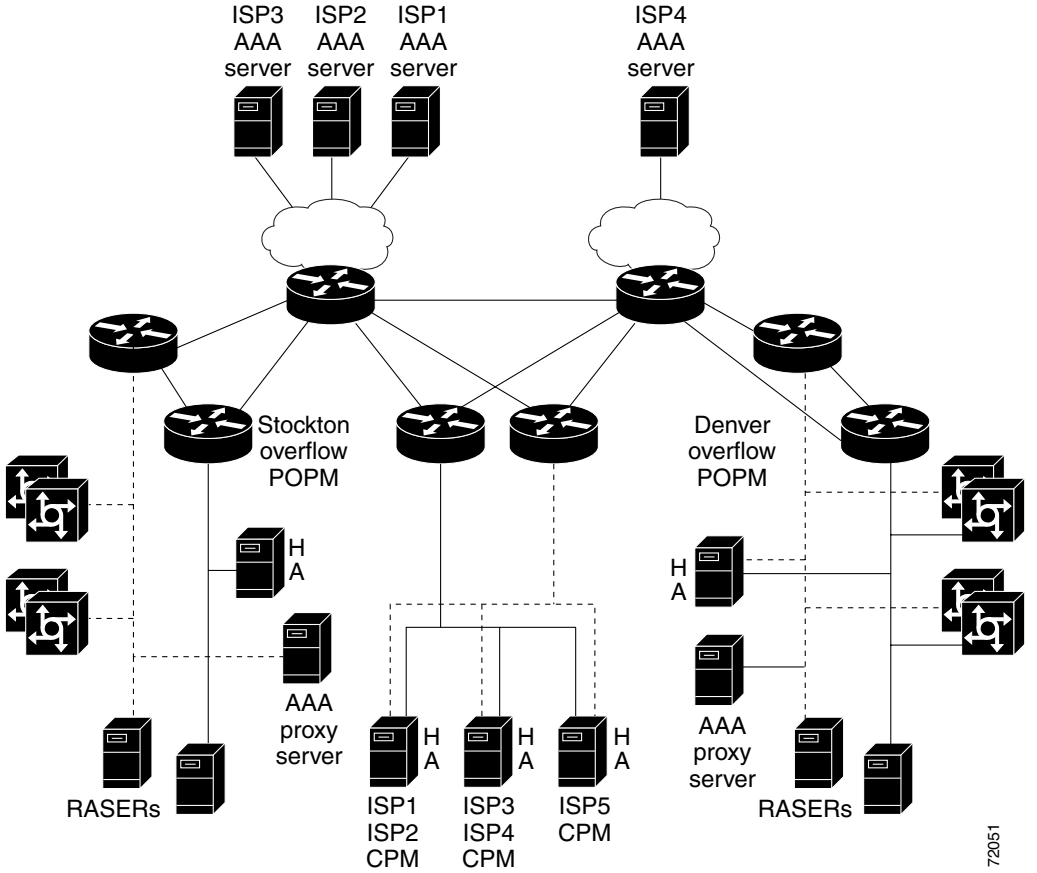

Figure 5-9 shows a a sample preauthentication call flow for a large scale deployment.

Figure 5-9 Cisco RPMS Large Scale Deployment Pre-Auth Call Flow

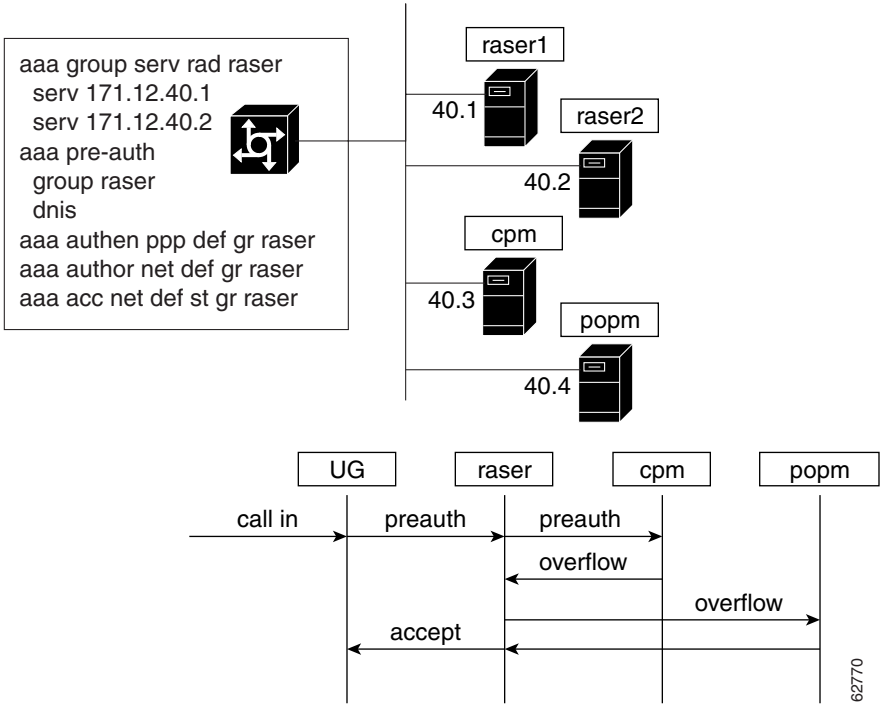

Preauthentication Request

The steps for an accepted shared overflow preauth request for this call flow are:

1. The preauth request comes in through the Stockton UG to Stockton RASERs or through Denver UG to Denver RASERs.
2. The UG routes the request to raser1.
3. raser1 determines to which ISP CPM server it should forward the request, such as ISP2, and then forwards it to ISP2 CPM. Note that both Stockton and Denver RASERs could route to a single (or dual) ISP2 CPM.
4. In this flow, ISP2 CPM tentatively accepts overflow, labels the call as an overflow call, and sends it back to raser1.

5. raser1 sends the overflow call request to the POPM that manages trunks for this call. Stockton RASERs would route to one of the Stockton POPM servers and Denver RASERs would route to a Denver POPM server.
6. The POPM accepts the overflow preauth request, and sends the accept request back to UG.
7. The UG starts a session.

For more dial call detail, such as reject call flows, see the *Cisco Resource Policy Management System Configuration Guide*.

VPDN

Cisco RPMS deployments can also manage VPDN traffic. By using VPDN and placing the AAA servers outside the Cisco RPMS network architecture, telephone companies and other dial service providers can remain independent.

If the universal gateway is configured for VPDN service, it sends a request to Cisco RPMS for VPDN information after answering a call. Cisco RPMS determines whether to authenticate the call with a home gateway through a VPDN tunnel based on the type of VPDN services configured in Cisco RPMS.

Cisco RPMS supports the following types of VPDN services:

- Dialed number information service (DNIS) VPDN dial service
- Domain name VPDN dial service
- VPDN request forwarding to an external AAA server

In a VPDN dial network, the RASER determines whether the call is DNIS-based (the username string contains the value “dnis”) or domain-based (the user ID contains a domain name).

DNIS-based VPDN uses DNIS numbers or trunk groups to answer or reject calls based on the customer profile configuration. Using the DNIS number, VPDN routes the session to the appropriate home gateway. Domain-name VPDN utilizes both a username and domain name to route the session to the appropriate home gateway.

Wholesale dial services rely on VPDN data specified in the AAA server or the Cisco RPMS VPDN group. To manage the tunnel availability, Cisco RPMS uses:

- VPDN session limit
- VPDN session overflow limit

For increased VPDN tunnel management, Cisco RPMS allows you to set an IP endpoint session count limit for each IP endpoint. IP endpoints are configured for VPDN groups.

You can easily integrate Cisco RPMS into existing VPDN dial service offerings. To convert existing VPDN users, configure Cisco RPMS for the required user configurations by using DNIS groups, trunk groups, customer profiles or other Cisco RPMS configuration settings.

Note Cisco RPMS will not use resources unless they are configured. Calls with DNIS numbers not configured for a DNIS group and not accepted by the default DNIS group and call type configuration are rejected.

For more detailed configuration and VPDN information, refer to *Cisco Access VPN Solutions Using Tunneling Technology.*

VPDN Groups

VPDN groups contain the following information required to establish a tunnel with a home gateway:

- The maximum number of sessions and overflow sessions
- The IP addresses of the home gateways
- The DNISs that belong to the VPDN group
- Threshold settings

VPDN With Optional AAA Proxy Servers

By using VPDN and placing the AAA proxy servers outside the Cisco RPMS network architecture, telephone companies and other dial service providers can remain independent.

Note Though AAA servers are optional, Cisco recommends using them in deployments.

Figure 5-10 VPDN Deployment with Optional AAA Proxy Server

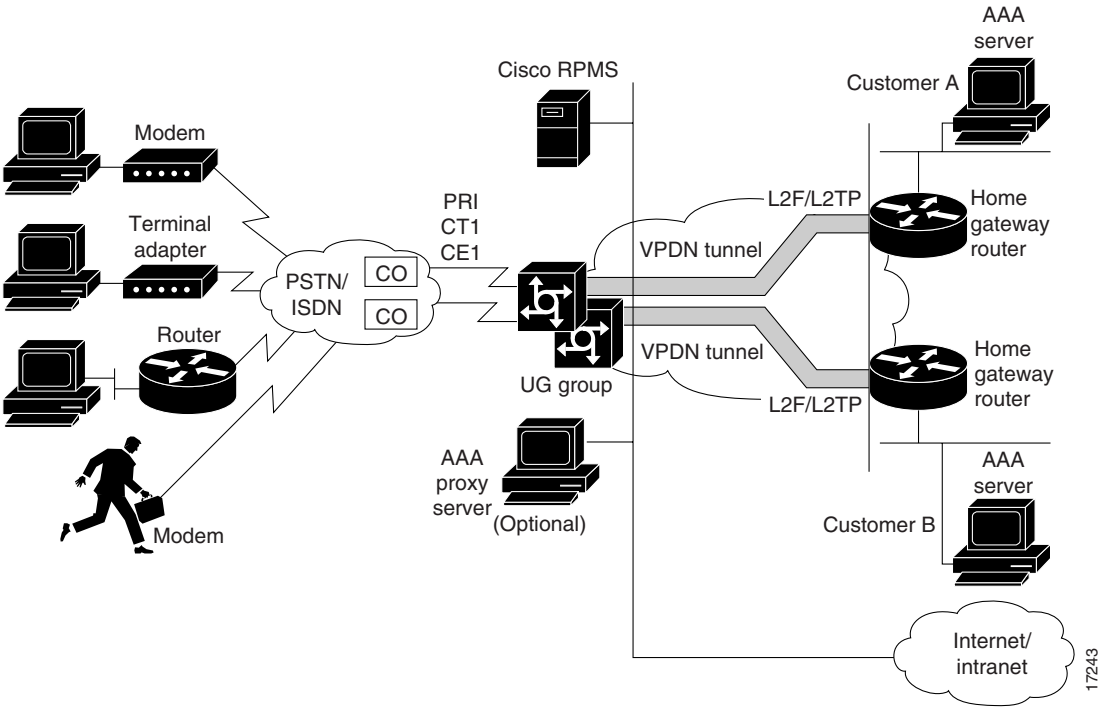

Overview: Voice Over IP Deployments

With the introduction of Cisco RPMS ASAP, you can also associate VOIP calls to policies by using Cisco H.323 Gatekeeper Clear Tokens, originating IP network masks, and various Cisco H.323 and SIP VSAs.

In a VOIP deployment, Cisco RPMS ASAP interprets and enforces universal port policy management and preauthentication for SLAs by using the RADIUS protocol.

Cisco RPMS ASAP determines which specific customer SLA policy to apply to a call based on the DNIS. Because you cannot associate dial and voice calls with the same DNIS, the Cisco RPMS customer profiles are accessed using independent dial or voice DNIS groups.

Voice calls accepted by Cisco RPMS ASAP continue a setup sequence similar to dial calls. The response from Cisco RPMS ASAP is returned to the calling entity, such as an ISDN or SIP call signaling interface, which then proceeds with the regular call flow. Calls rejected by Cisco RPMS ASAP apply the error codes or rejection reasons specified by the signaling entity.

Incoming H.323 or SIP VOIP Networks

Figure 5-11 shows an abstraction for a simple wholesale H.323 VOIP network. The components used for call control in this network include the universal gateway, Cisco RPMS, gatekeepers and a voice billing server managed by the wholesaler.

Note The voice billing server is optional.

Figure 5-11 Call Control for Incoming H.323 VOIP

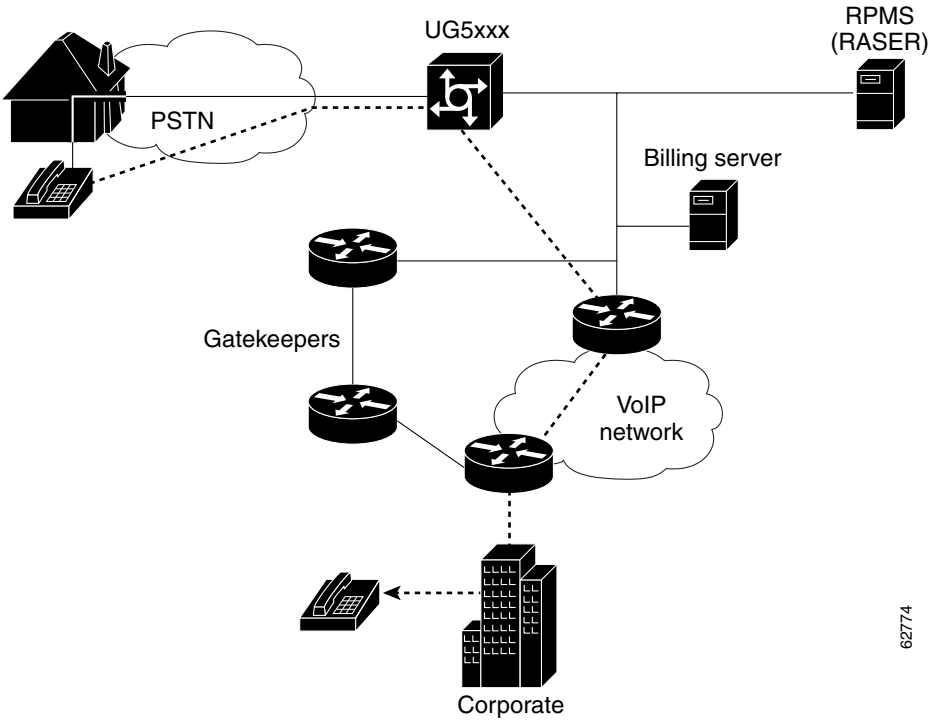

In this scenario, call control is applied to the call entering the VOIP wholesale network from the PSTN, or to the call arriving from a VOIP IP network. The call control for incoming PSTN voice call is about the same as for incoming dial. The call control sequence for VOIP includes interaction with the gatekeeper for H.323 VOIP or SIP proxy server for SIP.

Figure 5-12 shows a very basic call flow sequence diagram and rudimentary Cisco IOS commands for provisioning a Cisco UG in a VOIP network to enable preauthentication with Cisco RPMS. Note that the universal gateway directs all RADIUS messages to the RASER.

Figure 5-12 Pre-Auth Incoming H.323 VOIP Call Flow

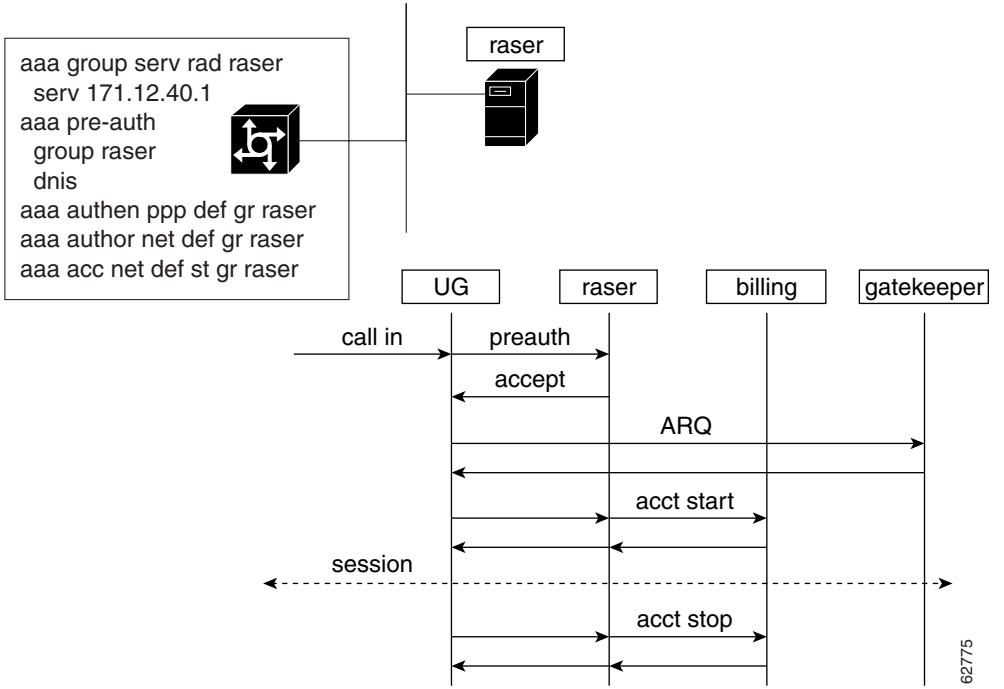

H.323-Based Voice Termination from an Internet Telephony Service Provider or from a Telephony - Application Service Providers

Figure 5-13 shows a VOIP call received from an ITSP or Telephony - Application Service Provider (T-ASP). The call travels through the wholesaler's IP network, and out the PSTN cloud.

In this network, Cisco RPMS ASAP manages the port policies, as well as preauthentication and authentication requests. The components used for call control in this network include the universal gateway, Cisco RPMS, gatekeepers, AAA proxy server.

Information about the call is sent in a preauthentication request from the UG to Cisco RPMS ASAP. Cisco RPMS locates the appropriate SLA and accepts or rejects the request based on the SLA limits.

If AAA is configured, authentication and authorization take place. Call information is sent from the UG to the Cisco RPMS and the ISP's AAA server, including the resource selected to handle the call and billing information.

Figure 5-13 H.323-Based Voice Termination from ITSP or T-ASP

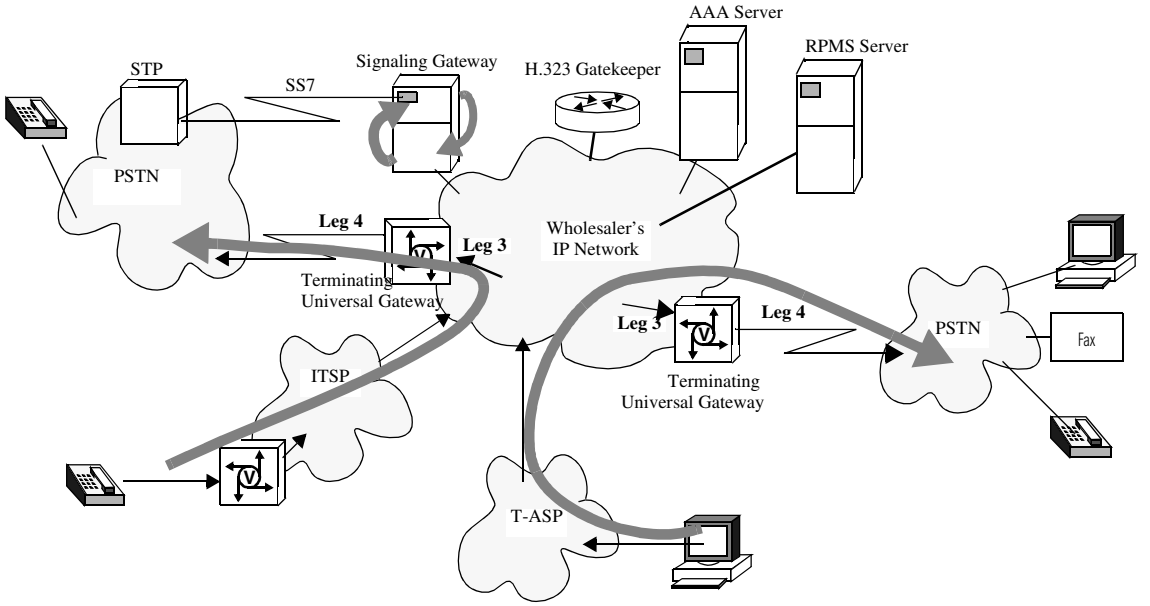

SIP-Based Voice Termination from an Internet Telephony Service Providers or from a Telephony - Application Service Providers

In this scenario, a voice call from a SIP telephone or SIP terminal is sent from an ITSP to an ISP. The Cisco SIP Proxy Server (CSPS) chooses a UG to forward the SIP INVITE to, based on its own routing mechanism.

Again, Cisco RPMS is used in this VOIP deployment for checking SLA limits, for managing the port policies, and for preauthentication and authentication purposes.

Figure 5-14 SIP-Based Voice Termination from ITSP

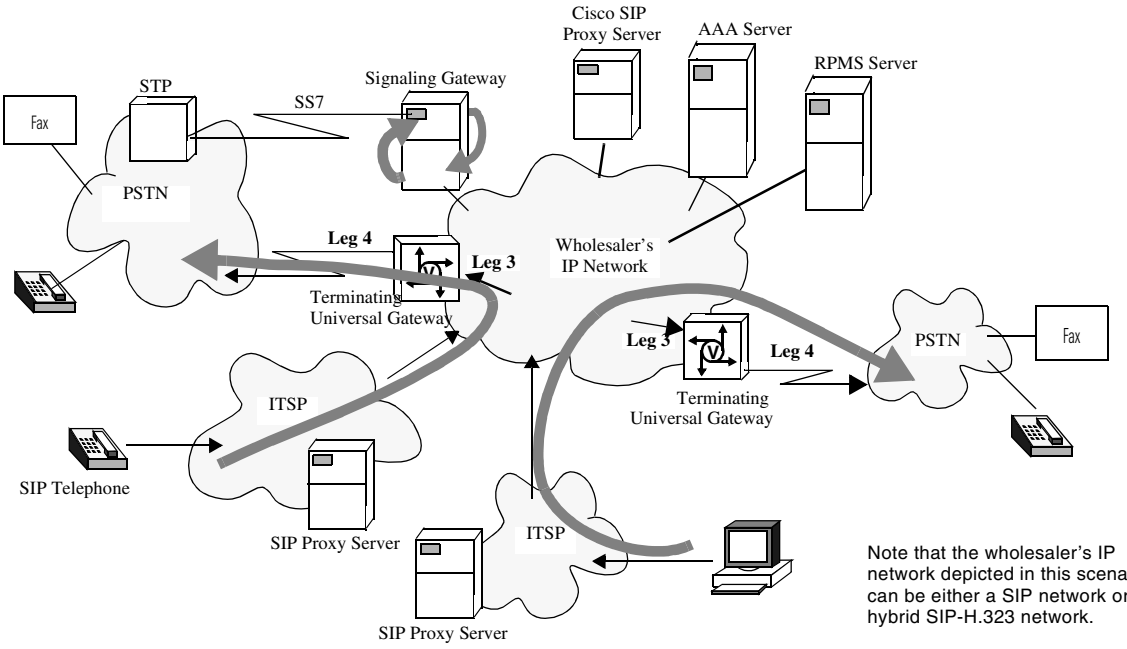

H.323-Based Prepaid Voice

In this scenario, the end customer dials into a prepaid voice service from the PSTN. Cisco RPMS checks the SLA limits, manages the port policies, and accepts or rejects preauthentication and authentication requests.

Cisco RPMS is also used for billing identification. In this scenario, it queries the caller for a billing identification PIN, which is sent to the prepaid billing server by Cisco RPMS.

Figure 5-15 H.323-Based Prepaid Voice

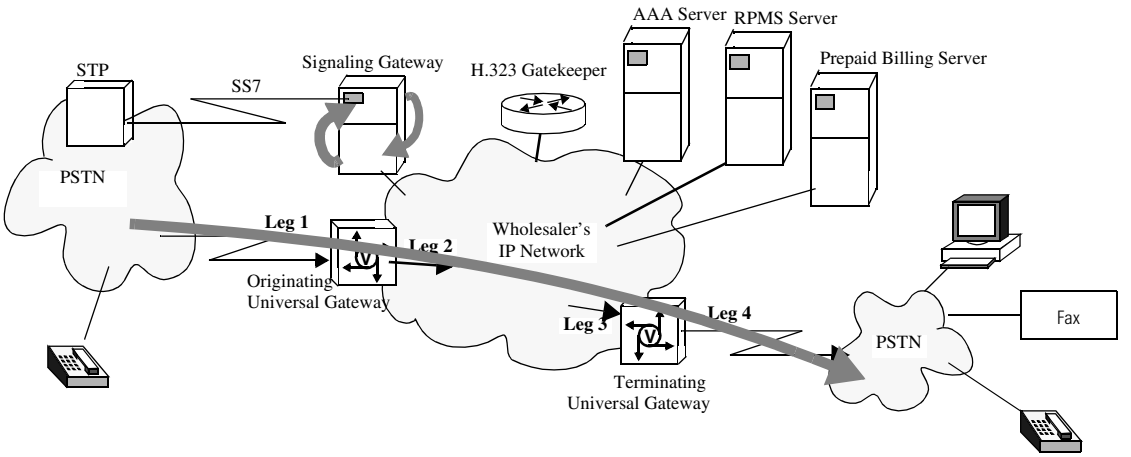

Overview: Any Service, Any Port Deployments

The Cisco ASAP solution is a network architecture that combines both voice and dial services on a single network. Cisco RPMS is designed to be integrated into the Cisco ASAP solution, to manage both voice and dial access to universal ports from the Cisco UGs.

Cisco RPMS is an integral component of the ASAP solution with its capability to manage voice and dial access to the universal ports on the Cisco UGs on a call-by-call basis. Cisco RPMS will support the Cisco ASAP network solution in mid-2002, and requires an upgrade for both Cisco RPMS and the Cisco IOS release.

Figure 5-16 provides a general overview of how Cisco RPMS is deployed in a Cisco ASAP environment. The Cisco RPMS corresponds with the images in the "Overview: Voice Over IP Deployments" section on page 5-21, depending on the type of deployment used, either dial (VPDN) or voice (H.323 or SIP).

Voice call provisioning and solutions are described more thoroughly in the Cisco RPMS Voice and the Cisco ASAP Solution Guides located at:

http://www.cisco.com/univercd/cc/td/doc/product/access/solution/asap/

Figure 5-16 Cisco ASAP Solution Architecture

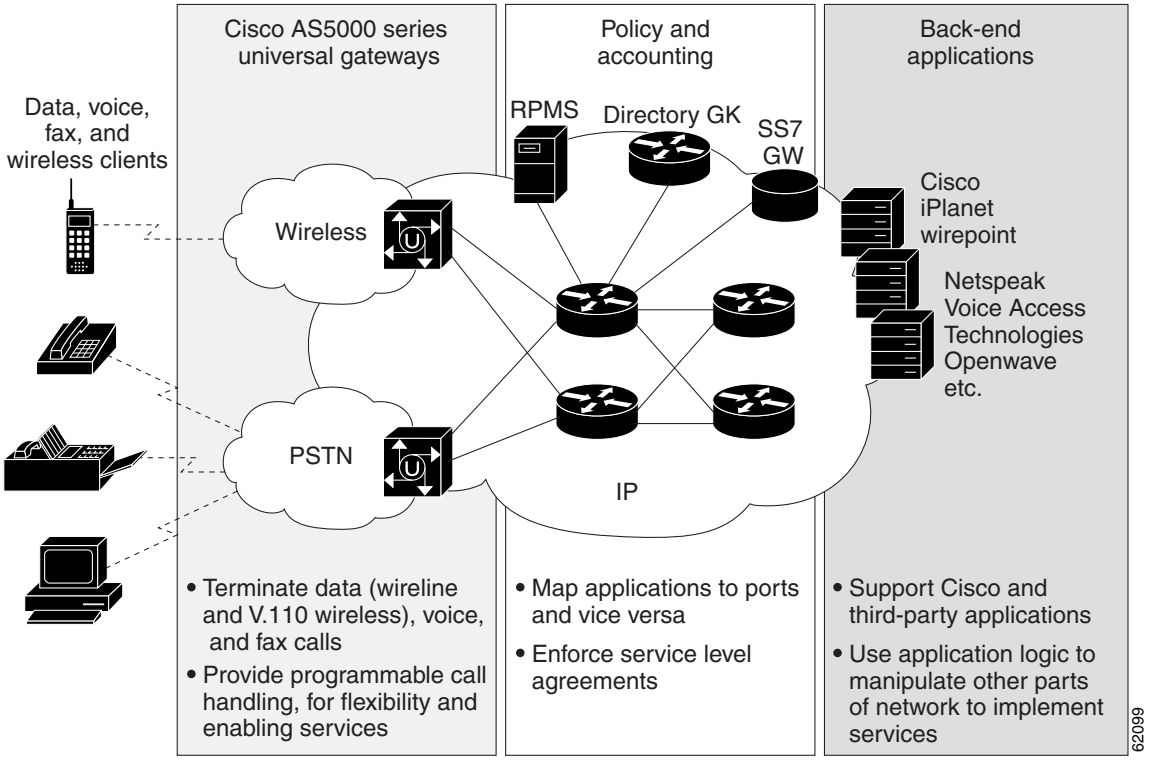

CHAPTER 6

Cisco RPMS Building Blocks

Topics in this chapter include:

Overview: Cisco Resource Policy Management System Building Blocks

Cisco RPMS works by implementing service level agreements (SLAs), or contracts, between wholesale and retail service providers. The SLAs are set up by creating customer profiles, call discriminations, VPDN services and shared overflow pools from the following building blocks:

- DNIS Groups
- Call Types
- Trunk Groups

After you have created the building blocks, you can use them to configure an SLA.

DNIS Groups

A DNIS group is a configured list of DNIS numbers corresponding to the numbers dialed by particular customers, service offerings, or both. Cisco RPMS checks the DNIS number of inbound calls against the configured DNIS groups or the default DNIS group. If a match is found, the configured information in the customer profile to which the DNIS group is assigned is used. If a match is not found, the default DNIS group and default customer profile are used. If a default customer profile is not configured, the call is rejected.

For an 800, 888, or similar number, the DNIS number the telco switch delivers to the UG is different from the user-dialed number. Cisco RPMS provides an optional reference number in each DNIS group for numbers requiring public network database lookup and results in a PSTN-routable DNIS.

For information on configuring a DNIS group, refer to the "Task 1: Creating a DNIS Group" section on page 3-4 in the *Cisco Resource Policy Management System Configuration Guide*.

Trunk Groups

Trunk groups are assigned to customer profiles with a call type for resource allocation and management.

Cisco RPMS maps incoming call requests to a specific customer profile by DNIS and call type (any, digital, speech, v.110, and v.120), and originating trunk.

Trunks and trunk groups allow Cisco RPMS users to manage multiple calls from different areas, by using more than one customer profile. For example, one POP can handle calls from multiple geographical regions and a VPDN customer can enforce counts for those different areas.

Cisco RPMS uniquely identifies a trunk by the following parameters: originating UG, and the port and slot number the trunk terminates on. The information tells you from which area the call came, and this forms a trunk. A trunk group comprises several trunks. Without the UG slot and port information, the DNIS and call type maps to just one customer profile, and you cannot distinguish from where the call came.

For information on configuring trunk groups, refer to "Task 1: Creating a Trunk Group" section on page 3-9 in the *Cisco Resource Policy Management System Configuration Guide*.

Call Types

Customer profiles use call types to identify which default customer profile to use for an incoming call. For DNIS groups, when multiple default customer profiles are used, the call type of the DNIS group identifies which default customer profile to use for an incoming call.

When calls originate from ISDN and CAS (CT1, CT3, and CE1), their call types assign calls to the appropriate DNIS group. Call types for ISDN are based on Q.931 bearer capability. Call types for CAS are assigned on static channel configuration or dynamically based on DNIS.

Supported call types are:

- any
- speech (for modem calls)
- digital (for ISDN data calls)

- V.110
- V.120

A call type may not be available in RADIUS requests originating from non-Cisco UGs. Because of this, Cisco RPMS uses a call type value equal to *any,* which is similar to a wildcard, but has less precedence than any other call type.

Cisco RPMS first tries to map a call to a customer profile associated with a specific call type. If no such customer is found, or a call type value is not available in the Access Accept message, the call is mapped to the customer profile associated with the *any* value. If no profile is associated with the *any* value, the call is rejected.

For information on call types, refer to the "Overview: Call Type" section on page 3-13 in the *Cisco Resource Policy Management System Configuration Guide*.

Overview: Service Level Agreements

Once the core building blocks are created, you can use the following components to build SLAs:

- Customer Profiles
- Call Discrimination
- VPDN
- Overflow Pools

For information on configuring SLAs, refer to the "Overview: Building Service Level Agreements" section on page 4-3 in the *Cisco Resource Policy Management System Configuration Guide*.

Customer Profiles

A standard customer profile is a set of parameters created for a specific service provider customer. The parameters are configured by the Cisco RPMS administrator and are based on the DNIS, call type and trunk.

You can assign configured DNIS groups, trunk groups, and IP groups to customer profiles. These customer profiles are selected by matching the incoming call characteristics (DNIS, call type, trunk) with the combinations of the entries within these associated groups.

More specifically, a customer is associated with a particular DNIS group, call type, and trunk group tuple. For example, you can associate a customer with:

- DNIS group *Dnis Group 1* (contains dnis 5551212)
- Call type *speech*
- Trunk group *Trunk Group 1* (contains trunk from UG 1.2.3.4, slot 0, port 1)

Subsequently, all call requests with the DNIS 5551212, call type speech, and originating from UG 1.2.3.4, slot 0, port 1 are mapped to this customer.

After the call is answered, customer profiles can also be associated with VPDN groups, so the configured VPDN sessions, home gateway information, and other data necessary to set up or reject a VPDN session are applied to the answered calls. The VPDN group data includes:

- Associated domain name or DNIS and trunk group
- Home gateway information

Customer profiles permissions might include the level of service users have, or whether the system even responds to users in certain situations. For instance, high level users can connect using higher modem speeds like 56K, while low level users connect at low speeds, such as 28K. If a low level user tries to access at 56K, the system rejects the call.

Customer profiles contain the following components:

- Customer profile name and description—Name and description of the customer.
- Session Count Limit—Maximum number of standard guaranteed sessions
- Oversubscription Limit—Maximum number of guaranteed sessions above the session count limit
- Shared Overflow Limit—Maximum number of shared overflow sessions above the total guaranteed limit (session + oversubscription limit)
- Threshold settings for the percentage of sessions, overflow sessions, and rejected sessions. E-mail alerts are generated when the threshold limits have been exceeded.

- Assigned DNIS groups containing a list of DNIS numbers that belong to a customer or service type.
- Assigned trunk groups containing a list of UG IP addresses and slot or port numbers that belong to a customer or service type.
- Assigned VPDN groups containing information to establish a tunnel with a home gateway.

The customer profiles are broken into two parts: incoming and outgoing call information. The incoming side of the customer profile determines if the call is answered by using parameters such as the dialed number (DNIS), call type and trunk from the assigned DNIS group, call type, trunk group, limits and thresholds. As sessions start and end, session counters increase and decrease respectively.

The outgoing side of the customer profile directs the answered call to the appropriate destination:

- To a local authentication, authorization, and accounting (AAA) server of non-VPDN dial applications, and Internet or intranet access.
- To a tunnel established to a wholesale (VPDN) dial customer's home gateway by using L2TP or L2F VPDN technology.

For information on configuring customer profiles, refer to "Task 1: Creating a Customer Profile" section on page 4-4 in the *Cisco Resource Policy Management System Configuration Guide*.

Default Customer Profiles

Default customer profiles are identical to standard customer profiles, except that they do not have any associated DNIS groups or trunk groups. This is accomplished by using the reserved keyword **default** for the DNIS group or trunk group.

Default customer profiles provide session counting and resource assignment to incoming calls that do not match the configured DNIS groups. Non-VPDN dial services and domain name-based VPDN services use default customer profiles.

A customer profile with the default DNIS, call type, or specific trunk groups has calls mapped to itself based on the trunk it came in on and the call type, but independent of the DNIS. A customer profile with DNIS, call type, or default trunk group has calls mapped to itself based on the DNIS and call type, but independent of the trunk.

You can configure multiple default customer profiles. For DNIS groups, when multiple default customer profiles are used, the call type (any, speech, digital, V.110, V.120) of the default DNIS group identifies which default customer profile to use for an incoming call. At most, four default profiles (one for each call type) can be configured on each Cisco RPMS server. For trunk groups, when multiple default customer profiles are used, UG IP or slot and port are used.

If default customer profiles are not defined, then calls that do not match a DNIS group in a customer profile are rejected.

There are different reasons why you might use a default customer profile instead of a standard profile. When configuring a standard customer profile, you must list all the DNIS numbers in one or more DNIS groups, and then associate them with the customer profile. The default customer profile is useful for mapping all other non-associated DNIS numbers to a customer profile.

The second reason you might use a default profile is that while using domain-based VPDN, the VPDN connection is decided by the domain name and not by the DNIS number. In this case, you can define a default customer profile and add the allowed domain names to its VPDN profile.

Call Discrimination

With call discrimination you can prevent specific call types from accessing resources on the UG. For example, use call discrimination to restrict a specific DNIS group to only modem calls by creating call discrimination settings for the DNIS group and the other call types (any, speech, digital, V.110, and V.120).

Note If a call type is not available in a RADIUS message, Cisco RPMS uses a call type of *any*.

To create a call discrimination, you need a list of DNIS groups and call types not allowed to access the DNIS groups.

For information on configuring call discrimination, refer to "Overview: Configuring Call Discriminations" section on page 4-9 in the *Cisco Resource Policy Management System Configuration Guide*.

VPDN

Cisco RPMS 2.0.1 supports VPDN tunneling through RADIUS for various UGs and NASs. If the UG is configured for VPDN service, it sends a request to Cisco RPMS for VPDN information after answering a call. Cisco RPMS determines whether to authenticate the call with a home gateway through a VPDN tunnel based on the type of VPDN services configured in Cisco RPMS. Cisco RPMS supports the following types of VPDN services:

- Dialed number information service (DNIS) VPDN dial service
- Domain name VPDN dial service
- VPDN request forwarding to external AAA server

To identify VPDN calls, the UG sends two request types as a RADIUS Access Request message. The first type, user authorization requests, are processed by a AAA server, while VPDN session request messages are processed by a Cisco RPMS server.

You must use a RADIUS proxy in a VPDN solution to distinguish between these requests. Directing VPDN session request information to Cisco RPMS allows you to manage tunnel related counters, and directing regular user authorization messages to the AAA server authorizes a user session.

For information on configuring VPDN, refer to "Overview: Configuring VPDN Services" section on page 4-11 in the *Cisco Resource Policy Management System Configuration Guide*.

VPDN Groups

VPDN groups contain the information required to establish a tunnel with a home gateway, the maximum number of sessions and shared overflow sessions, the IP addresses of the home gateways, the DNISs that belong to the VPDN group, and threshold settings.

VPDN groups contain DNIS-number or domain-name assignments, so that answered calls can be assigned the configured VPDN home gateway information and other data necessary to set up or reject the VPDN session.

For DNIS-based VPDN dial service, VPDN groups are assigned to customer profiles based on the incoming DNIS number and the configured DNIS groups for the initial resource allocation and customer session limits.

For domain-based VPDN dial service, VPDN groups are assigned to the default customer profile with the matching call type and default DNIS group assignment. The default customer profile contains the initial resource allocation and customer session limits.

You can specify multiple home gateway IP endpoints for a VPDN group. If you specify three or more IP endpoints, load balancing is used to ensure that traffic is distributed across the IP endpoints. This load balancing method obtains the first three available IP endpoints listed in the Valid IP Endpoints page and selects the first available IP endpoint as the one to use. The following request obtains the next three IP endpoints from the list, and uses the first available. Cisco RPMS also rotates available IP endpoints within the groups of three, using the second and third available IP endpoints before using the first IP endpoint a second time. For example:

Available IP endpoints:

IP Endpoint1, IP Endpoint2, IP Endpoint3, IP Endpoint4, IP Endpoint5, IP Endpoint6

First Home Gateway Request:

IP Endpoint1, IP Endpoint2, IP Endpoint3

Selected: IP Endpoint1

Second Home Gateway Request:

IP Endpoint4, IP Endpoint5, IP Endpoint6

Selected: IP Endpoint4

Third Home Gateway Request:

IP Endpoint2, IP Endpoint3, IP Endpoint1

Selected: IP Endpoint2

Note For a multichannel ISDN client to have proper routing on its second B channel in a multiple home gateway VPDN environment, configure Stacked Group Bidding Protocol (SGBP) on the home gateways. Enable VPDN multihop to establish a single logical termination point for the client.

Trunk-Based VPDN Groups

In Cisco RPMS Release 2.0, VPDN groups could be distinguished only by their DNIS number or domain name. Cisco RPMS Release 2.0.1 introduces trunk-based VPDN groups, wherein you can assign a VPDN group specific to a customer profile, or specific to the trunk group affiliated with that profile.

Trunk-based VPDN groups help you manage networks by associating a domain name or DNIS number with a configured trunk group. For example, let's say you had two trunk groups, San Francisco and San Jose. The VPDN group could manage VPDN calls from the domain "cisco.com" originating from the San Francisco trunk group, or VPDN calls originating from the San Jose trunk group.

You also can use the wildcard character, an asterisk (*), for the DNIS number or domain name. However, the wildcard must replace the entire domain name or DNIS number; partial wildcarding is not allowed.

To associate the VPDN group and trunk group, Cisco RPMS searches for a matching VPDN group by using a specific domain name and trunk group. If it cannot, Cisco RPMS further attempts matches using the following parameters:

- Domain name and default trunk group
- Any domain name (using the wildcard) and a specific trunk group
- Any domain name and default trunk group

For more information, refer to the"Task 3: Associating a DNIS" and"Task 2: Associating a Domain Name" sections in Chapter 4, "Building Service Level Agreements" of the *Cisco Resource Policy Management System Configuration Guide*.

Subscription and Oversubscription Limits

Cisco RPMS provides hard subscription and soft oversubscription limits for customer profiles. As such, a wholesaler can negotiate and enforce a variety of SLAs throughout the network, ranging from hard or guaranteed limits, to soft or shared first come, first serve access to pools of resources or combinations of both.

A primary SLA is usually a guaranteed subscription agreement between the wholesaler and an ISP (or other customer). An oversubscription, or soft limit, can be best effort, charging the ISP a premium for any calls accepted beyond the guaranteed limit. The oversubscription also can be shared: multiple customers competing on a first come, first serve basis for a virtual pool of resources. Any customer who exceeds a guaranteed subscription limit and overflows to this pool competes for access to the remaining shared resources in a pool. As a result, any customer session accepted in an overflow pool counts against both that individual customer's oversubscription limit, and the shared session limit associated with the shared overflow pool. With guaranteed customer limits, only sessions for that customer count against the guaranteed session limit for that customer alone.

Overflow Pools

Cisco RPMS is used to pool resources across a wholesale network and enforce limits so that no one customer can oversubscribe their resource usage and adversely impact another guaranteed service agreement. In the simplest deployment, this means that a wholesaler can simply provision a pool of resources equivalent to the sum of all guaranteed access limits. Any customer can access any resource from the pool, and Cisco RPMS ensures that no customer exceeds the guaranteed limit. This ability to pool resources when using Cisco RPMS alleviates the need to set aside dedicated equipment on a per customer basis.

Additionally, Cisco RPMS can provide shared overflow pools. The wholesaler can actually have more resources within the network than the sum of guaranteed limits. Additional resources are just one reason the wholesaler may contract shared oversubscription. Customers who exceed their guaranteed session limits can access the additional resources on a first come, first serve basis.

With Cisco RPMS, you can control overflow access through shared overflow pools. To create an overflow pool, you must give it a name, and then associate one or more trunk groups to it. The basic idea behind Cisco RPMS is to allow a pool

of resources to be shared across customers; overflow pools extend that idea for oversubscription calls. Therefore, shared overflow pools allow wholesalers to maintain and enforce SLA limits that span multiple customers.

Note A trunk group cannot be associated to more than one overflow pool.

An overflow pool has the following limits to enforce SLAs:

- The total number of overflow calls that are allowed for the overflow pool.
- The percentage of overflow calls that are allowed for different service types. This allows you to split the overflow pool between dial and voice calls. For example, you can set both options to 100 percent so that no preference is given to one service type over another. You can also vary the service mix, so that dial service has access to any 60 percent, and voice has access to any 40 percent. After 60 percent of the resources have active dial calls, new dial calls are rejected to assure that 40 percent are reserved for voice overflow. The service mix is configurable.

During call processing, a customer profile is selected based on the DNIS, trunk group, IP group and the call type. A service type of dial or voice is associated with the call based on the DNIS group to which the call was mapped. The default DNIS is dial.

Note Any call matching IP Group is assumed to be voice.

The service type provides an additional method to limit access to shared overflow pools based on service type. An overflow call for a customer is accepted (provided it is within the customer's overflow limits) only if the following two criteria are met:

- Total call count of the overflow pool associated with the trunk group is within the overflow pool size.
- Total call count for the service type is within the service type limit for that overflow pool.

For information on configuring overflow pools, refer to the "Overview: Configuring Overflow Pools" section on page 4-14 of the *Cisco Resource Policy Management System Configuration Guide*.

Overview: Thresholds

Thresholds are used for notification; when a threshold is exceeded, Cisco RPMS sends a notification. This feature can be used to determine whether a particular customer is hitting the threshold too often and needs to buy more ports from the wholesaler to avoid paying costly premiums for the oversubscription ports.

Thresholds are configured in the customer profile. When port allotments are about to be or have been exceeded, e-mail notification is provided to the wholesaler.

Example of Threshold Settings

Table 6-1 shows the configurable threshold fields in a customer profile and sample settings for those thresholds.

Table 6-1 Sample Threshold Settings

Field Name	Sample Settings
Customer Name:	BigISP
Description:	cust on Wholesaler1
Session Count Limit:	100,000
Session Count Threshold:	85%
Oversubscription Limit	10,000
Shared Overflow Limit:	10,000
Shared Overflow Threshold:	85%
Shared Overflow Call Rejection Threshold:	10%
Call Rejection Threshold:	85%

Let's walk through this configuration, for the customer *BigISP*. The BigISP profile has a session count limit of 100,000 which means that BigISP has 100,000 ports available to use. The session count threshold is set to 85 percent, so when 85,000 ports are used, Cisco RPMS notifies Wholesaler1 by e-mail or page. This is useful for determining whether or not BigISP has purchased enough ports; if BigISP consistently hits the threshold, it may need to purchase more ports.

BigISP has an oversubscription limit of 10,000 so that an additional 10,000 oversubscription ports can be used, but at a premium rate.

BigISP also has a shared overflow limit of 10,000. This means that if 110,000 ports are used, Wholesaler1 can still take additional BigISP calls (10,000 in this case), but BigISP will pay a premium price for those calls. Again, the shared overflow threshold is set to 85 percent, so when 8,500 of the overflow ports are used, another alert goes out to the administrator.

The session count limit (100,000) plus the oversubscription limit (10,000) plus the shared overflow limit (10,000) equals the maximum number of sessions allowed (120,000). Any call after number 120,000 is automatically rejected until current users end their sessions and free up ports. . Figure 6-1 shows these different thresholds.

Figure 6-1 Threshold Settings for a Session Count of 100 with a Shared Overflow of 10

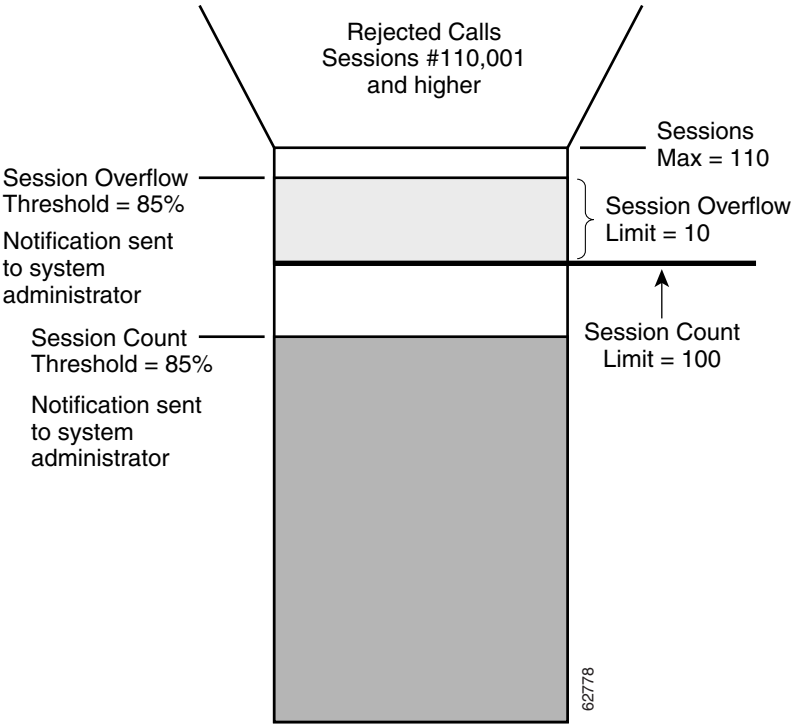

For information on configuring thresholds, refer to the "Overview: Thresholds" section on page 6-14 in the *Cisco Resource Policy Management System Configuration Guide*.

Customer Levels

With call thresholds, customer profiles have may have subgroups associated with them. These subgroups include:

- VPDN groups
- DNIS groups
- Trunk groups
- Call types

You can configure Cisco RPMS to generate e-mail notification when the following thresholds settings have been exceeded:

- From customer profiles:
  - Session count threshold—Percentage of the session count limit.
  - Shared overflow threshold—Percentage of the shared overflow limit.
  - Shared overflow call rejection threshold—Number of shared overflow call rejections as a percentage of the number of attempted shared overflow calls since the last system reset.
  - Call rejection threshold—Number of calls rejected per customer profile as a percentage of the number of attempted calls since the last system reset.
- From VPDN groups:
  - Session count threshold—Percentage of the session count limit.
- From VPDN IP endpoints:
  - Session count threshold—Percentage of the IP endpoint session count limit.

After exceeding a threshold, no additional e-mail is sent until the threshold is reactivated. The Alert Low Water Percentage parameter specifies the percentage low-point value to which that threshold must return to reactivate the threshold setting. This prevents e-mail storms when a threshold is continually reached without much variation.

For information on customer levels, refer to the "Customer Levels" section on page 6-16 in the *Cisco Resource Policy Management System Configuration Guide*.

Overview: Reports

Cisco RPMS records report data for network dial service analysis and troubleshooting. Report data is a snapshot of the current activity and appears on screen.

You can also filter reports to remove unnecessary rows of information and allow for small manageable reports.

You can view reports in the Cisco RPMS GUI or save the reports to file. You can create the following types of reports:

- Active Call—Monitor current call connection data, including connection time, connection speeds, UG, UG port, DNIS number, and VPDN data.
- Customer Profile Report—See customer profile status by checking the total number of active calls, oversubscription calls, shared overflow calls, call limits, and number of calls rejected. You can check current call activity through the call count, tunnel count, and VPDN sessions. You can verify UG port availability and resource configuration by referencing the number of calls rejected because of resources not available, rejected VPDN sessions, and calls rejected due to size limit.
- DNIS Report—You can check the number of calls using the DNIS. You can see DNIS numbers and DNIS group assignments.
- DNIS Group Report—You can verify DNIS group configuration and monitor DNIS group usage.
- Domain Name Report—You can verify domain assignments to customer profiles and monitor VPDN tunnel counts and session rejections, monitoring current domain usage and VPDN active sessions.
- VPDN Group Report—You can monitor the current VPDN sessions by VPDN group. You can verify VPDN group configurations and assignments to customer profiles and troubleshoot session and call rejections due to VPDN session limits and resource availability.
- IP Endpoints Report—You can verify the number of sessions between the UG and the home gateway on this set of IP endpoints.
- Tunnel Report—You can see the status of all active calls on the VPDN network and monitor IP tunnel sessions by DNIS group and domain name. In addition, verify VPDN DNIS and domain configurations and session usage.

- Recent Call Report—You can track past call disconnections and rejections for troubleshooting and resource usage and monitor past call rates and ensure effective resource utilization for dial services.
- Overflow Pool Report—You can see details such as the number of overflow calls that are accepted and rejected against different service types for each defined overflow pool.

For more information on Cisco RPMS reports, refer to Chapter 6, "Reporting and Accounting" in the *Cisco Resource Policy Management System Configuration Guide.*

Overview: Sending Vendor Specific Attributes with Access Accept Messages

Cisco RPMS allows you to include RADIUS vendor specific attributes (VSAs) with the Access Accept messages sent to the UG. With RADIUS VSAs, you can:

- Add VSA information about various UG vendors
- Associate selected VSAs with customer profiles for various call types
- Specify values to be sent in Access Accept messages for the selected VSAs

You must first add the vendors by configuring them at the Adminstration>RADIUS screen with the following information:

- A name
- RADIUS ID
- Description
- The VSA details about that vendor

The UG Type entry specifies the RADIUS translator handling the RADIUS requests and responses for the UG, and is required for handling differences between UGs from various vendors. The Cisco translators are available when you first install Cisco RPMS, and you can add new translators through the command-line interface.

After adding vendors, you must associate the VSAs to customer profiles, based on the call type. You must also give the value for the VSAs to send to the UGs.

Building Vendor Specific Attributes for Modem Management

Cisco RPMS includes a feature to send VSAs with an Access Accept message. Modem configurations to the Cisco UG in RADIUS are handled by using the Cisco VSA Attr[26]. The VSA contains a string of the following syntax:

```
preauth:modem-service= modem min-speed <x> max-speed <y> modulation
<z> error-correction <a> compression <b>"
```

Table 6-2 shows possible values for various options:

Table 6-2 Modem Configuration VSA Options

Command/Definition	Argument/Options
min-speed:Minimum Speed	<300-56000>, any
max-speed:Maximum Speed	<300-56000>, any
modulation:Modulation	K56Flex, v22bis,v32bis,v34,v90,any
error-correction	lapm, mnp4
compression	mnp5, v42bis

An example of a VSA:

```
preauth:modem-service=modem min-speed 300 max-speed 28000 modulation
v32bits
```

For information on configuring the modem to send VSAs with an Access Accept message, refer to the "Building Vendor Specific Attributes for Modem Management" section on page 2-13 in the *Cisco Resource Policy Management System Configuration Guide*.

Building Vendor Specific Attributes to Control Authentication Type

You can use Cisco RPMS to send VSAs in Access Accept messages to determine whether or not subsequent authentication is performed on calls accepted within a particular customer profile.

If subsequent authentication is configured, you can use an additional VSA to determine the type of authentication to use. These two VSAs contain strings of the following syntax:

```
preauth:auth-required=1
preauth:auth-type=pap
```

Table 6-3 displays possible values for various options:

Table 6-3 Subsequent Authentication VSA Options

Command/Definition	Argument/Options
auth-required	1 for authentication, 0 for no authentication
auth-type	pap, chap, or ms-chap

For information on configuring the modem to determine whether or not subsequent authentication is performed on calls accepted within a particular customer profile, refer to the "Building Vendor Specific Attributes to Control Authentication Type" section on page 2-14 in the *Cisco Resource Policy Management System Configuration Guide*.

Overview: Universal Gateway IP Address Groups

Cisco RPMS ASAP introduces universal gateway IP address groups. You will be able to use the groups to track voice calls originating from an IP address.

Note The IP address group features are the same for both ASAP and for wholesale dial, but ASAP will provide additional features.

Universal gateway IP address groups contain a list of IP addresses that define all points of ingress from an ITSP into a wholesale SP's network. After defining the IP address group, you can associate it with a customer profile.

Customer profiles can contain IP address groups in addition to DNIS groups. An IP address group is similar to a DNIS group and identifies a wholesale customer. However, since the IP address groups are always used for voice calls, you do not have to specify a call type while associating the IP address group with a customer profile.

The IP address used for finding a customer is taken from a VSA included with the RADIUS Access-Request message, which must follow the preauthentication syntax of username=dnis for Cisco RPMS to distinguish it as a preauthentication request.

When Cisco RPMS receives a RADIUS preauthentication message, it searches for a VSA to use for finding an IP address group. If a VSA is found, it is used to find a customer profile with the IP address in one of its associated IP groups. If the VSA is not found, or the VSA does not map to any customer profile, Cisco RPMS continues with the DNIS-based customer profile lookup.

Links for Configuring Cisco RPMS Deployment Scenarios

Topics in this chapter include:

Configuring Cisco RPMS Deployment Scenarios

This chapter provides links to the tasks required to provision the deployment scenarios defined in Chapter 5, "Cisco RPMS Deployment Scenarios." These scenarios include:

- Non-VPDN dial

  For more information on configuring Cisco RPMS in a VPDN environment, refer to the "Overview: DNIS-Based Non-VPDN Dial Service" section of the *Provisioning Cisco RPMS in a Wholesale Dial Network* guide.

- VPDN dial

  For more information on configuring Cisco RPMS in a VPDN environment, refer to "Overview: Configuring VPDN" section of the *Provisioning Cisco RPMS in a Wholesale Dial Network* guide.

- Incoming VOIP

  For more information on incoming VOIP, refer to the "Overview: Voice Over IP Deployments" section on page 5-21. Tasks for configuring Cisco RPMS in a VOIP environment will be available in the next release of Cisco RPMS.

- Terminating VOIP

  For more information on terminating VOIP, refer to the "Overview: Voice Over IP Deployments" section on page 5-21. Tasks for configuring Cisco RPMS in a VOIP environment will be available in the next release of Cisco RPMS.

- Call discrimination

  For more information on call discrimination, refer to refer to "Overview: Configuring Call Discriminations" section on page 4-9 in the *Cisco Resource Policy Management System Configuration Guide*.

The set of tasks for each scenario has been divided between Cisco RPMS administration and configuration.

Cisco RPMS Tasks in a Non-VPDN Wholesale Dial Environment

When using non-VPDN, you must complete the following tasks:

- Cisco RPMS Administration Tasks for a Non-VPDN Wholesale Dial Environment
- Cisco RPMS Configuration Tasks for a Non-VPDN Wholesale Dial Environment

Cisco RPMS Administration Tasks for a Non-VPDN Wholesale Dial Environment

Cisco RPMS administration includes configuring the universal gateway and AAA server lists.

Universal Gateway Lists

You must enter all the UGs communicating with the Cisco RPMS RASERs in the universal gateway list.

For more information, refer to the "Overview: Communicating with Universal Gateways" section on page 2-16 of the *Cisco Resource Policy Management System Configuration Guide*.

AAA Lists

In the AAA list, enter all AAA servers and AAA proxies to which the RASERS will forward user authentication and accounting requests.

For more information, refer to the "Overview: AAA Servers" section on page 2-19 of the *Cisco Resource Policy Management System Configuration Guide*.

Cisco RPMS Configuration Tasks for a Non-VPDN Wholesale Dial Environment

When configuring Cisco RPMS for a non-VPDN scenario, you must configure the following components:

- DNIS and DNIS Groups

  For more information, refer to the "Overview: Configuring DNIS Groups" section on page 3-4 of the *Cisco Resource Policy Management System Configuration Guide*.

- Trunks and Trunk Groups

  For more information, refer to the "Overview: Configuring Trunk Groups" section on page 3-8 of the *Cisco Resource Policy Management System Configuration Guide*.

- Overflow Pools

  For more information, refer to the "Overview: Configuring Overflow Pools" section on page 4-14 of the *Cisco Resource Policy Management System Configuration Guide*.

- Customer Profiles

  For more information, refer to the "Overview: Configuring Customer Profiles" section on page 4-3 of the *Cisco Resource Policy Management System Configuration Guide*.

Cisco RPMS Tasks in a VPDN Wholesale Dial Environment

When providing VPDN, follow the same administrative and configuration tasks as for non-VPDN, as well as the following additional tasks.

Cisco RPMS Administration Tasks for a VPDN Wholesale Dial Environment

For Cisco RPMS administration using VPDN, complete the Universal Gateway Lists and AAA Lists tasks from the Cisco RPMS Administration Tasks for a Non-VPDN Wholesale Dial Environment section.

Cisco RPMS Configuration Tasks for a VPDN Wholesale Dial Environment

For Cisco RPMS configuration using VPDN, complete the DNIS and DNIS Groups, Trunks and Trunk Groups, Overflow Pools and Customer Profile tasks from the previous sections. Additionally, you must configure VPDN Groups.

VPDN Groups

To create any new VPDN groups for DNIS-based VPDN, enter the appropriate DNIS numbers in the domain or DNIS data portion of the VPDN group.

For more information, refer to the "Overview: Configuring VPDN Services" section on page 4-11 of the *Cisco Resource Policy Management System Configuration Guide*.

For domain-based VPDN, enter the appropriate domain names in the domain or data portion of the VPDN group.

For more information, refer to the "Overview: Domain Name VPDN Dial Service" section on page 4-11 of the *Cisco Resource Policy Management System Configuration Guide*.

To load balance the VPDN session between multiple home gateways, enter the home gateway IP addresses in the IP Endpoint section of the VPDN group.

For VPDN request forwarding, do not enter the respective domains or DNIS in any VPDN group.

For more information, refer to the "Overview: VPDN Request Forwarding to an External AAA Server" section on page 4-14 of the *Cisco Resource Policy Management System Configuration Guide*.

Associating a VPDN Group with a Customer Profile

Once the VPDN group is defined, you can associate the VPDN group with the respective customer profile. If the service does not require enforcing limits at preauthentication (for example, no limits are required for the customer profile), but only for VPDN sessions, associate the VPDN group with the default customer profile, or set the reserved and overflow session limit for the customer to unlimited.

For more information, refer to the "Task 3: Adding a VPDN Group to a Customer Profile" section on page 4-13 of the *Cisco Resource Policy Management System Configuration Guide*.

Trunk-Based VPDN Groups

You also can use trunk-based VPDN groups to manage the network. To do so, you must associate the domain name or DNIS number with a configured trunk group.

For more information, refer to the of the *Cisco Resource Policy Management System Configuration Guide*.

Incoming H.323 or SIP VOIP

For incoming voice calls, the administration and configuration of Cisco RPMS is the same as for non-VPDN wholesale dial.

Terminating H.323 or SIP VOIP

For terminating voice calls, the Cisco RPMS configuration involves:

- Cisco RPMS Administration for Terminating H.323 or SIP VOIP
- Cisco RPMS Configuration for Terminating H.323 or SIP VOIP

Cisco RPMS Administration for Terminating H.323 or SIP VOIP

The Cisco RPMS administration tasks for a terminating H.323 or SIP VOIP deployment require configuring the client list. If you have deployed an optional billing server, you must also configure the billing server lists.

Client List

You must enter all H323 gatekeepers or SIP proxies communicating with the Cisco RPMS RASER in the Cisco RPMS client list.

Billing Server Lists

If deploying a billing server, enter all billing servers to which the RASERS will forward user authentication and accounting requests in the AAA servers list or billing server list, respectively.

For more information, refer to the "Overview: AAA Servers" section on page 2-19 of the *Cisco Resource Policy Management System Configuration Guide*.

Cisco RPMS Configuration for Terminating H.323 or SIP VOIP

DNIS and DNIS Groups

Enter any new DNIS numbers or DNIS groups. Select the service type voice when configuring the DNIS group. If the scenario does not require using a specific DNIS (for example, service level agreements are enforced by using trunks only), then no DNIS configuration is required.

For more information, refer to the "Overview: Configuring DNIS Groups" section on page 3-4 of the *Cisco Resource Policy Management System Configuration Guide*.

Note For SIP proxies, configuring DNIS is not required.

Voice Over IP Address Groups

Voice Over IP address groups are used to track voice calls originating from IP. The IP address groups contain a list of IP addresses defining all points of ingress from an ITSP into a wholesale service provider's network. To use the IP address group, you would enter any new IP addresses and groups.

Note IP address groups will be introduced in the Cisco RPMS ASAP release.

For more information, refer to the "Voice Over IP Address Groups" section on page 2-16.

Customer Profiles

Once the DNIS numbers and IP groups have been defined, you can define the customer profile. Associate the IP group and DNIS group tuple with the customer. If using SIP proxies, always associate the IP group with the default DNIS group.

For more information, refer to the "Overview: Configuring Customer Profiles" section on page 4-3 of the *Cisco Resource Policy Management System Configuration Guide*.

Call Discrimination

Call discrimination prevents calls with a specified DNIS and call type from accessing resources on the UGs. To configure discrimination, simply add entries to the call discrimination table in Cisco RPMS.

For more information, refer to the "Overview: Configuring Call Discriminations" section on page 4-9 of the *Cisco Resource Policy Management System Configuration Guide*.

Helpful Links

This appendix provides helpful links to specific Cisco universal gateways and Cisco IOS releases used in a Cisco RPMS deployment. It also provides a link to the Cisco ASAP information.

Topics in this chapter include:

Overview: Universal Gateway Links

For information about Cisco universal gateways, use the following URLs.

Cisco AS5300

For information about the Cisco AS5300, go to:

http://www.cisco.com/pcgi-bin/Support/PSP/psp_view.pl?p=Hardware:AS5300

Cisco AS5350

For information about the Cisco AS5350, go to:

http://www.cisco.com/pcgi-bin/Support/PSP/psp_view.pl?p=Hardware:AS5350

Cisco AS5400

For information about the Cisco AS5400, go to:

http://www.cisco.com/pcgi-bin/Support/PSP/psp_view.pl?p=Hardware:AS5400

Cisco AS5800

For information about the Cisco AS5800, go to:

http://www.cisco.com/pcgi-bin/Support/PSP/psp_view.pl?p=Hardware:AS5800

Overview: Cisco IOS Release Links

For information about Cisco IOS Release 12.2, go to the following URL:

http://www.cisco.com/warp/public/732/releases/release122.shtml

Overview: Cisco ASAP Network Solution Links

For information about the Cisco ASAP Solution, go to:

http://www.cisco.com/warp/public/779/servpro/solutions/remote/asap.html

To read the Cisco ASAP 1.0 documentation, go to:

http://www.cisco.com/univercd/cc/td/doc/product/access/solution/asap/

GLOSSARY

A

AAA Authentication, authorization, and accounting. AAA is a suite of network security services which provides the primary framework through which access control is set up on a Cisco router or access server.

ACS Access control server.

ADSL Asymmetrical digital subscriber line.

APM AccessPath Manager; a Cisco dial stack management system, now called Cisco Access Manager (CAM).

AR Cisco Access Registrar.

ASAP Any Service, Any Port.

AV Attribute value.

B

BC Bearer capability.

C

CAM Cisco Access Manager, a Cisco dial stack management system; formerly APM.

CAS Channel associated signaling.

CDR Call detail record.

CHAP Challenge Handshake Authentication Protocol.

CLI Command-line interface.

CLID Calling line identification.

CO Central office; telephone company phone switch.

CoS Class of service.

CPM Customer Profile Manager.

CSD Circuit switch data (for example, ISDN data call).

CSV Comma-separated value; a nonproprietary comma-delimited file used to transfer data into spreadsheet programs.

CT1 Channelized T1 trunk service.

D

DNIS Dialed number information service (the called telephone number).

DOVBS Data over Voice Bearer Service. Not available in Cisco RPMS Release 2.0.1.

DS0 A single 64K/56K channel in a PRI or CT1.

G

GRS Global Roaming Server.

GKTMP Cisco Gatekeeper Transaction Message Protocol.

GUI Graphical User Interface.

H

H.323 An International Telecommunication Union (ITU-T) standard that describes packet-based video, audio, and data conferencing.

HA High availability.

HDLC High-Level Data Link Control.

HTML Hypertext Markup Language.

HTTP Hypertext Transfer Protocol; the protocol for moving hypertext files across the Internet.

I

ISDN Integrated Services Digital Network; a digital circuit-switched telephony infrastructure.

ISP Internet service provider.

L

LNS L2TP Network Server.

LRQ Location Request.

M

MIB Management Information Base; an SNMP-compliant description of a network device.

N

NAS Network access server. Refers to non-Cisco equipment.

NEBS Network Equipment Building Systems; a standard for switches and other equipment.

NOC Network operations center.

P

PAP Password Authentication Protocol.

POP Point of Presence.

POPM Point of Presence Manager.

POTS Plain old telephone service; the acronym is used in the term “POTS-routable,” which refers to the ability of the telephone system to deliver a call to a particular switch and return a particular DNIS.

PPM Port Policy Management.

PRI Primary Rate Interface, an ISDN service.

PSTN Public Switched Telephone Network.

Q

QoS Quality of service.

R

RADIUS Remote Authentication Dial-In User Service; the protocol used for communication between the UG and the RPMS.

RASER Remote Access Service Router .

RFC Request For Comments; a document published by the InterNIC describing a proposed Internet standard.

RLM Redundant Link Manager.

RPM Resource Pool Manager; the standalone resource pool management application that resided on the UG in Cisco RPMS 1.x.

RPMS Resource Policy Management System.

RSP Regional service provider.

S

SGBP Stack Group Bidding Protocol.

SIP Session Initiation Protocol.

SLA Service level agreement.

SNMP Simple Network Management Protocol.

SQL Structured Query Language.

SSL Secure Socket Layer; a security protocol that provides privacy over HTTP.

T

TA (ISDN) terminal adapter.

TCP/IP Transmission Control Protocol/Internet Protocol.

TMN Telecommunications Management Network; a standard for management of data networks.

U

UCP User Control Point (formerly Network Control Point [NCP]).

UG Universal Gateway.

V

VPDN Virtual private dial-up network.

VOIP Voice over IP.

VPOP Virtual Point of Presence.

VSA Vendor specific attribute.

A

B

C

D

E

F

G

H

I

L

M

N

O

P

R

S

T

U

V

W