Cisco SCMS Subscriber Manager User Guide

Version 3.0.5
OL-7199-05

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: DOC-719905=
Text Part Number: OL-7199-05

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

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

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

Modifying the equipment without Cisco's written authorization may result in the equipment no longer complying with FCC requirements for Class A or Class B digital devices. In that event, your right to use the equipment may be limited by FCC regulations, and you may be required to correct any interference to radio or television communications at your own expense.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

- Turn the television or radio antenna until the interference stops.
- Move the equipment to one side or the other of the television or radio.
- Move the equipment farther away from the television or radio.
- Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILTY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, *Packet*, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)

Printed in the USA on recycled paper containing 10% postconsumer waste.

*Cisco **SCMS Subscriber Manager** User Guide*

Copyright © 2002-2006 Cisco Systems, Inc.
All rights reserved.

CONTENTS

Preface ix

- Document Revision History ix
- Audience x
- Organization x
- Related Documentation xi
- Conventions xi
- Obtaining Documentation xii
  - World Wide Web xii
  - Documentation CD-ROM xii
  - Ordering Documentation xiii
  - Documentation Feedback xiii
- Obtaining Technical Assistance xiii
  - Cisco.com xiv
  - Technical Assistance Center xiv

General Overview 1-1

- The Cisco Service Control Concept 1-1
  - Service Control for Broadband Service Providers 1-2
- Cisco Service Control Capabilities 1-2
- The SCE Platform 1-3
- Service Configuration Management 1-4
- Management and Collection 1-4
  - Network Management 1-5
  - Subscriber Management 1-5
  - Data Collection 1-6

Introducing the Subscriber Manager 2-1

- Subscribers in the Cisco Service Control Solution 2-2
- Handling Subscribers 2-2

Flow of Subscriber Information 2-3
Number of Subscribers in the SM 2-4
SM Database 2-5
Subscriber ID 2-5
SM Fundamentals 2-6
The SM API 2-6
The SM Login Event Generators 2-6
Subscriber Introduction Mode: Push or Pull 2-6
Domains 2-8
Communication Failures 2-9
SM Cluster 2-10
SM Management 2-10
Subscriber Manager Fail-Over 2-11
Overview 2-12
Normal Operation 2-12
Fail-Over Topology 2-13
Fail-Over 2-15

Installation and Upgrading 3-1
Installation Overview 3-1
Contents of Distribution Files 3-2
Documentation 3-5
System Requirements 3-5
Installation Procedures 3-7
Typical Installation 3-8
Step 1: Extracting the Distribution Files 3-8
Step 2: Executing the install-sm.sh Script 3-9
Step 3: Setting the Password for the pcube User 3-9
Step 4: Configuring the System Memory Settings 3-10
Step 5: Rebooting 3-14
Step 6: Installing the SCA BB Package and LEG components 3-14
Step 7: Adding a User for PRPC Authentication 3-15
Verifying that the Installation was Successful 3-15
Configuring the Subscriber Manager 3-15

Additional Installation Procedures 3-16
Installing an SM Cluster 3-16
Installing the Subscriber Manager 3-17
Installing SM Cluster Agents 3-18
Troubleshooting the Installation 3-18
Installing an Application 3-19
System Changes Made by Installation Scripts 3-20
Logging Script Messages 3-20
pcube User and Group 3-20
Startup and Shutdown Scripts 3-20
Upgrading 3-21
Upgrade Procedure 3-21
upgrade-sm.sh Script 3-22
Data Duplication Procedure 3-24
Upgrading from Non-cluster Setup to Non-cluster Setup 3-24
Upgrading from Non-cluster Setup to Cluster Setup 3-24
Upgrading from a Cluster Setup version 3.x 3-26
Upgrading from a Cluster Setup version 2.x 3-26
Uninstalling 3-26
uninstall-sm.sh Script 3-27
Uninstalling VCS Agents 3-28

Configuration and Management 4-1
SM Management and Configuration Methods 4-1
Configuration File 4-1
Command-Line Utilities 4-2
Configuring a Subscriber Management Solution 4-3
Prerequisites 4-3
Step-by-Step Configuration Procedure 4-3
Using the CLU 4-20
Informative Output 4-21
Parsing CLU Operations and Options 4-21
Reloading the SM Configuration (p3sm) 4-23
Managing the SM (p3sm) 4-23

Managing Subscribers, Mappings, and Properties (p3subs) 4-25
Managing the Subscriber Database (p3subsdb) 4-28
Viewing and Connecting Network Elements (p3net) 4-30
Viewing Subscriber Domains (p3domains) 4-30
Managing the Cable Support Module (p3cable) 4-31
Installing an Application (p3inst) 4-32
Viewing Information of the PRPC Interface Server (p3rpc) 4-33
Managing a Cluster of Two SM Nodes (p3cluster) 4-34
Managing the User Log (p3log) 4-34
Viewing Statistics of the RADIUS Listener (p3radius) 4-35
Utilities 4-35

Configuration File Options A-1

Introduction A-1
Description of the Configuration File Options A-1
SM General Section A-2
SM High Availability Setup Section A-4
Subscriber State Persistency Section A-4
SM-LEG Failure Handling Section A-4
LEG-Domains Association Section A-5
Domain.XXX Section A-6
Default Domains Configuration Section A-7
Auto Logout Section A-8
Radius Listener Section A-9
Radius.NAS.XXX Section A-10
Radius.Property.Package Section A-11
Radius.Subscriber ID Section A-12
RPC.Server Section A-13
MPLS-VPN Section A-13
SCE.XXX Section A-14
FTP Section A-14
HTTP Tech-IF Section A-15
RDR Server Section A-16
Cable Adapter Section A-16

Data Repository Section A-17

Command-Line Utilities B-1

Introduction B-1

Description of the CLU Commands B-1

Informative Output B-2

Parsing CLU Operations and Options B-2

p3batch Utility B-3

p3cable Utility B-4

p3clu Utility B-5

p3cluster Utility B-5

p3db Utility B-6

p3domains Utility B-7

p3ftp Utility B-7

p3http Utility B-8

p3inst Utility B-8

p3log Utility B-10

p3net Utility B-11

p3radius Utility B-11

p3rpc Utility B-12

p3sm Utility B-13

p3subs Utility B-16

p3subsdb Utility B-19

CPE as Subscriber in Cable Environment C-1

Cable Support Module C-1

CM and CPE in the SM C-2

Static and Dynamic CMs C-3

Troubleshooting D-1

Using the Troubleshooting Chapter D-1

General Errors D-1

SM Not Running D-1

SM in Failure Mode D-2

General Setup Errors D-2

install-sm.sh Script–User is not Root D-2
install-sm.sh Script–User pcube Exists D-3
install-tt.sh Script D-3
install-dsn.sh Script D-3
TimesTen Database Setup Errors D-4
Introduction D-4
TimesTen DSN Configuration—Cannot Find Requested DSN D-5
TimesTen DSN Configuration—Data Source Name Not Found D-5
TimesTen Database Settings—Cannot Connect to Data Source D-6
TimesTen Configuration Error—Not Enough Memory D-7
TimesTen Configuration Error—Incorrect Memory Definitions D-7
TimesTen Configuration Error—Cannot Create Semaphores D-8
TimesTen Configuration Error—Cannot Read Data Store File D-9
TimesTen Configuration Error—Data Store Space Exhausted D-9
Network Management Command Line Utility (p3net) Errors D-10
First Connection—Operation Timed Out D-10
Status Error—Connection Down D-11
Status Error—Subscriber Management Down D-11
Subscriber Database Command Line Utility (p3subsdb) Errors D-12
CSV File Validation Error D-12
Cable Support Command Line Utility (p3cable) Errors D-13
CSV File Import Error D-13
Configuration Errors D-13
Network Management Errors D-14
Domain Errors D-15
PRPC Errors D-17
RADIUS Listener Errors D-17
Common Validation Errors D-17

Veritas Cluster Server Requirements and Configuration E-1
Veritas Cluster Server System Requirements E-1
Veritas Cluster Server Nodes on Remote Sites E-2
Replication Configuration Guidelines E-3
Replication Scheme Setup E-3

Replication Network Configuration E-3
Veritas Cluster Server Configuration Guidelines E-5
SM Cluster Resources Configuration E-5
Adding SM Cluster Resources E-8
Useful Operations E-15
Linking the Resources E-18
Saving and Closing Your Cluster Configuration E-18
Verifying That Service Group is Online E-18
SNMP Support E-19
Configuring the SnmpConsole Attribute E-22

Glossary of Terms GL-1

Index I-1

Preface

This document is a guide to the configuration and management of the Subscriber Manager (SM) application. It contains explanations of relevant concepts and terminology in addition to instructions on how to install, configure, and use the SM.

Note The guide assumes a basic familiarity with the telecommunications equipment and installation procedures, concepts of the Service Control Management Suite, the SCE 1000 and SCE 2000 platforms, and their related components.

This *Cisco SCMS Subscriber Manager User Guide* documents the application that processes subscriber records and maintains the subscriber database for use by the SCE platform. It contains instructions for installing the SM and the TimesTen database, editing and installing the configuration file, and managing the system using the Command-Line Utilities (CLU). It also specifies how to configure a Service Control deployment to support subscriber integration in its various forms.

Document Revision History

Cisco Service Center Release	Part Number	Publication Date
Release 3.0.5	OL-7199-05	November, 2006

Description of Changes

- Changes to the RPC.Server sections of the configuration file to include PRPC server security level. See *RPC.Server* (on page 4-12) and *RPC.Server Section* (on page A-13).
- Changes to p3rpc CLU. See *p3rpc Utility* (on page B-12).
- Addition of formatting rules for the subscriber ID. See *Subscriber ID* (on page 2-5).
- Addition of case insensitivity to the subscriber ID. See *SM General Section* (on page A-2).

Cisco Service Center Release	Part Number	Publication Date
Release 3.0.3	OL-7199-04	July, 2006

Description of Changes

- Changes to upgrade procedure. See *Upgrade Procedure* (on page 3-21).

- Updated installation software requirements to include Red Hat 4.

Cisco Service Center Release	Part Number	Publication Date
Release 3.0.3	OL-7199-03	May, 2006

Description of Changes

- Updated installation software requirements to include Solaris 10 and Red Hat 4.
- Updated installation procedures to reflect removal of support for TimesTen, DSNs, and Java.
- Added new section covering upgrade procedure from version 3.0 to 3.0.3. See *Upgrading from a Cluster Setup version 3.x* (on page 3-26).
- Changes to the procedure to uninstall VCS agents. See *Uninstalling VCS agents* (on page 3-28).
- Added new sections detailing the RDR Server configuration. See *RDR Server* (on page 4-13) *and RDR Server Sec*tion (on page A-16).
- Added new section detailing MPLS-VPN configuration. See *MPLS-VPN Section* (on page A-13).
- Added arguments for MPLS/VPN mappings. See *Managing Subscribers, Mappings, and Properties (p3subs)* (on page 4-25) and *p3subs Utility* (on page B-16)

Release 3.0	OL-7199-02	December, 2005

Audience

This guide is intended for the system administrator and the networking or computer technician who is responsible for the on-site installation and configuration of the Subscriber Manager (SM) application. It is also intended for the operator who will be responsible for daily operations of the SM.

Organization

This guide covers the following topics:

Chapter	Title	Description
Chapter 1	*General Overview* (on page 1-1)	Presents an overview of the Cisco Service Control Solution and the SCE platform
Chapter 2	*Introducing the Subscriber Manager* (on page 2-1)	Describes the Subscriber Manager solution, the handling of subscribers, the fundamentals and management of the SM application
Chapter 3	*Installation and Upgrading* (on page 3-1)	Describes the installation and basic configuration procedures for the SM and how to perform an upgrade of the SM

Chapter 4	*Configuration and Management* (on page 4-1)	Describes the SM configuration and management procedures and how to use the Command-Line Utilities (CLU)
Appendix A	*Configuration File Options* (on page A-1)	Provides a description of all parameters that can be configured by using the Subscriber Manager configuration file
Appendix B	*Command-Line Utilities* (on page B-1)	Describes the Command-Line Utilities (CLU) that the SM uses to manage and maintain the application
Appendix C	*CPE as Subscriber in Cable Environment* (on page C-1)	Deals with the special case when the Customer Premise Equipment (CPE) is considered as the subscriber in the Cisco Service Control solution for a cable environment, such as DSL and Wireless.
Appendix D	*Troubleshooting* (on page D-1)	Describes the error messages, probable causes, and solutions of the SM application
Appendix E	*Veritas Cluster Server Requirements and Configuration* (on page E-1)	Provides basic guidelines for the Veritas Cluster Server configuration in an SM cluster installation

Related Documentation

Use this ***SCMS Subscriber Manager*** User Guide in conjunction with the following Cisco documentation:

- *Cisco Service Control Application for Broadband (SCA BB) User Guide*
- *Cisco SCE 1000 or SCE 2000 platform User Guides*

Conventions

This document uses the following conventions:

Convention	**Description**
boldface font	Commands and keywords are in **boldface**.
italic font	Arguments for which you supply values are in *italics*.
[ ]	Elements in square brackets are optional.
{x | y | z}	Alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string, or the string will include the quotation marks.

Convention	Description
`screen` font	Terminal sessions and information that the system displays are in `screen` font.
`boldface screen` font	Information you must enter is in **`boldface screen`** font.
`italic screen` font	Arguments for which you supply values are in *`italic screen`* font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

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

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

Obtaining Documentation

The following sections provide sources for obtaining documentation from Cisco Systems.

World Wide Web

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

- *http://www.cisco.com*
- *http://www-china.cisco.com*
- *http://www-europe.cisco.com*

Documentation CD-ROM

Cisco documentation and additional literature are available in a CD-ROM package that ships 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 as 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/pcgi-bin/marketplace/welcome.pl

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

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can submit technical comments electronically. Click **Feedback** in the toolbar and select **Documentation**. After you complete the form, click **Submit** to send it to Cisco.

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:

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

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides *Cisco.com* (on page xiv) as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools. For Cisco.com registered users, additional troubleshooting tools are available from the TAC website.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information and resources at any time, from anywhere in the world. This highly integrated Internet application is a powerful, easy-to-use tool for doing business with Cisco.

Cisco.com provides a broad range of features and services to help customers and partners streamline business processes and improve productivity. Through Cisco.com, you can find information about Cisco and our networking solutions, services, and programs. In addition, you can resolve technical issues with online technical support, download and test software packages, and order Cisco learning materials and merchandise. Valuable online skill assessment, training, and certification programs are also available.

Customers and partners can self-register on Cisco.com to obtain additional personalized information and services. Registered users can order products, check on the status of an order, access technical support, and view benefits specific to their relationships with Cisco.

To access Cisco.com, go to http://www.cisco.com.

Technical Assistance Center

The Cisco Technical Assistance Center (TAC) website is available to all customers who need technical assistance with a Cisco product or technology that is under warranty or covered by a maintenance contract.

Contacting TAC by Using the Cisco TAC Website

If you have a priority level 3 (P3) or priority level 4 (P4) problem, contact TAC by going to the TAC website *http://www.cisco.com/tac*.

P3 and P4 level problems are defined as follows:

- P3—Your network is degraded. Network functionality is noticeably impaired, but most business operations continue.
- P4—You need information or assistance on Cisco product capabilities, product installation, or basic product configuration.

In each of the above cases, use the Cisco TAC website to quickly find answers to your questions.

To register for *Cisco.com* (on page xiv), go to *http://tools.cisco.com/RPF/register/register.do*.

If you cannot resolve your technical issue by using the TAC online resources, Cisco.com registered users can open a case online by using the TAC Case Open tool at *http://www.cisco.com/tac/caseopen*.

Contacting TAC by Telephone

If you have a priority level 1 (P1) or priority level 2 (P2) problem, contact TAC by telephone and immediately open a case. To obtain a directory of toll-free numbers for your country, go to *http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml*.

P1 and P2 level problems are defined as follows:

- P1—Your production network is down, causing a critical impact to business operations if service is not restored quickly. No workaround is available.
- P2—Your production network is severely degraded, affecting significant aspects of your business operations. No workaround is available.

CHAPTER 1

General Overview

This chapter provides a general overview of the Cisco Service Control solution. It introduces the Cisco Service Control concept and the Service Control capabilities. It also briefly describes the hardware capabilities of the Service Control Engine (SCE) platform and the Cisco specific applications that together compose the total Cisco Service Control solution.

This chapter contains the following sections:

- The Cisco Service Control Concept 1-1
- Cisco Service Control Capabilities 1-2
- The SCE Platform 1-3
- Service Configuration Management 1-4
- Management and Collection 1-4

The Cisco Service Control Concept

The Cisco Service Control solution is delivered through a combination of purpose-built hardware and specific software solutions that address various service control challenges faced by service providers. The SCE platform is designed to support classification, analysis, and control of Internet/IP traffic.

Service Control enables service providers to create profitable new revenue streams while capitalizing on their existing infrastructure. With the power of Service Control, service providers have the ability to analyze, charge for, and control IP network traffic at multigigabit wire line speeds. The Cisco Service Control solution also gives service providers the tools they need to identify and target high-margin content-based services and to enable their delivery.

As the downturn in the telecommunications industry has shown, IP service providers' business models need to be reworked to make them profitable. Having spent billions of dollars to build ever larger data links, providers have incurred massive debts and faced rising costs. At the same time, access and bandwidth have become commodities where prices continually fall and profits disappear. Service providers have realized that they must offer value-added services to derive more revenue from the traffic and services running on their networks. However, capturing real profits from IP services requires more than simply running those services over data links; it requires detailed monitoring and precise, real-time control and awareness of services as they are delivered. Cisco provides Service Control solutions that allow the service provider to bridge this gap.

Service Control for Broadband Service Providers

Service providers of any access technology (DSL, cable, mobile, and so on) targeting residential and business consumers must find new ways to get maximum leverage from their existing infrastructure, while differentiating their offerings with enhanced IP services.

The Cisco Service Control Application for Broadband adds a new layer of service intelligence and control to existing networks that can:

- Report and analyze network traffic at subscriber and aggregate level for capacity planning
- Provide customer-intuitive tiered application services and guarantee application SLAs
- Implement different service levels for different types of customers, content, or applications
- Identify network abusers who are violating the Acceptable Use Policy
- Identify and manage peer-to-peer, NNTP (news) traffic, and spam abusers
- Enforce the Acceptable Use Policy (AUP)
- Integrate Service Control solutions easily with existing network elements and BSS/OSS systems

Cisco Service Control Capabilities

The core of the Cisco Service Control solution is the purpose-built network hardware device: the Service Control Engine (SCE). The core capabilities of the SCE platform, which support a wide range of applications for delivering Service Control solutions, include:

- Subscriber and application awareness—Application-level drilling into IP traffic for real-time understanding and controlling of usage and content at the granularity of a specific subscriber.
  - Subscriber awareness—The ability to map between IP flows and a specific subscriber in order to maintain the state of each subscriber transmitting traffic through the SCE platform and to enforce the appropriate policy on this subscriber's traffic.

    Subscriber awareness is achieved either through dedicated integrations with subscriber management repositories, such as a DHCP or a Radius server, or via sniffing of Radius or DHCP traffic.
  - Application awareness—The ability to understand and analyze traffic up to the application protocol layer (Layer 7).

    For application protocols implemented using bundled flows (such as FTP, which is implemented using Control and Data flows), the SCE platform understands the bundling connection between the flows and treats them accordingly.
- Application-layer, stateful, real-time traffic control—The ability to perform advanced control functions, including granular BW metering and shaping, quota management, and redirection, using application-layer stateful real-time traffic transaction processing. This requires highly adaptive protocol and application-level intelligence.
- Programmability—The ability to quickly add new protocols and easily adapt to new services and applications in the ever-changing service provider environment. Programmability is achieved using the Cisco Service Modeling Language (SML).

Programmability allows new services to be deployed quickly and provides an easy upgrade path for network, application, or service growth.

- Robust and flexible back-office integration—The ability to integrate with existing third-party systems at the Service Provider, including provisioning systems, subscriber repositories, billing systems, and OSS systems. The SCE provides a set of open and well-documented APIs that allows a quick and robust integration process.
- Scalable high-performance service engines—The ability to perform all these operations at wire speed.

The SCE Platform

The SCE family of programmable network devices is capable of performing application-layer stateful-flow inspection of IP traffic, and controlling that traffic based on configurable rules. The SCE platform is a purpose-built network device that uses ASIC components and RISC processors to go beyond packet counting and delve deeper into the contents of network traffic. Providing programmable, stateful inspection of bidirectional traffic flows and mapping these flows with user ownership, the SCE platforms provide real-time classification of network usage. This information provides the basis of the SCE platform advanced traffic-control and bandwidth-shaping functionality. Where most bandwidth shaper functionality ends, the SCE platform provides more control and shaping options, including:

- Layer 7 stateful wire-speed packet inspection and classification
- Robust support for over 600 protocols and applications, including:
  - General—HTTP, HTTPS, FTP, TELNET, NNTP, SMTP, POP3, IMAP, WAP, and others
  - P2P file sharing—FastTrack-KazaA, Gnutella, BitTorrent, Winny, Hotline, eDonkey, DirectConnect, Piolet, and others
  - P2P VoIP—Skype, Skinny, DingoTel, and others
  - Streaming and Multimedia—RTSP, SIP, HTTP streaming, RTP/RTCP, and others
- Programmable system core for flexible reporting and bandwidth control
- Transparent network and BSS/OSS integration into existing networks
- Subscriber awareness that relates traffic and usage to specific customers

The following diagram illustrates a common deployment of an SCE platform in a network.

Figure 1-1: SCE Platform in the Network

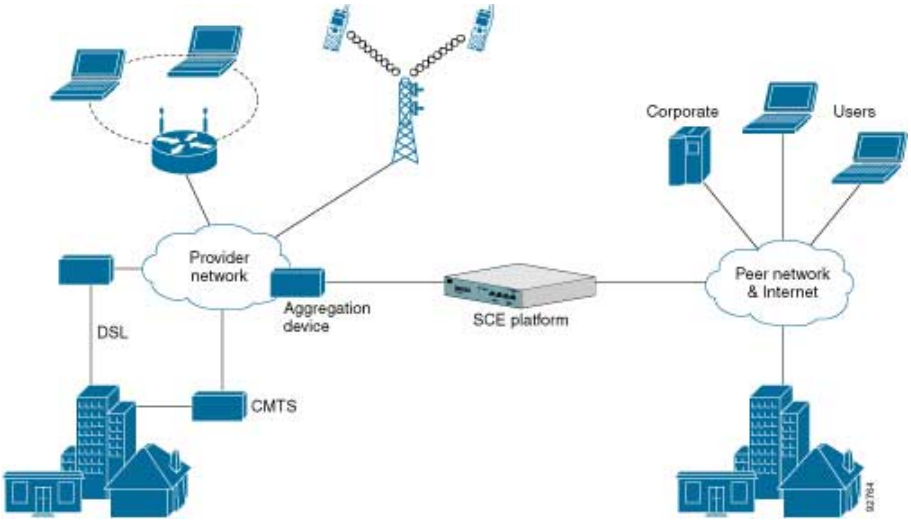

Service Configuration Management

Service configuration management is the ability to configure the general service definitions of a service control application. A service configuration file containing settings for traffic classification, accounting and reporting, and control is created and applied to an SCE platform. The SCA BB application provides tools to automate the distribution of these configuration files to SCE platforms. This simple, standards-based approach makes it easy to manage multiple devices in a large network.

Service Control provides an easy-to-use GUI to edit and create these files and a complete set of APIs to automate their creation.

Management and Collection

The Cisco Service Control solution includes a complete management infrastructure that provides the following management components to manage all aspects of the solution:

- Network management
- Subscriber management
- Service Control management

These management interfaces are designed to comply with common management standards and to integrate easily with existing OSS infrastructure.

Figure 1-2: Service Control Management Infrastructure

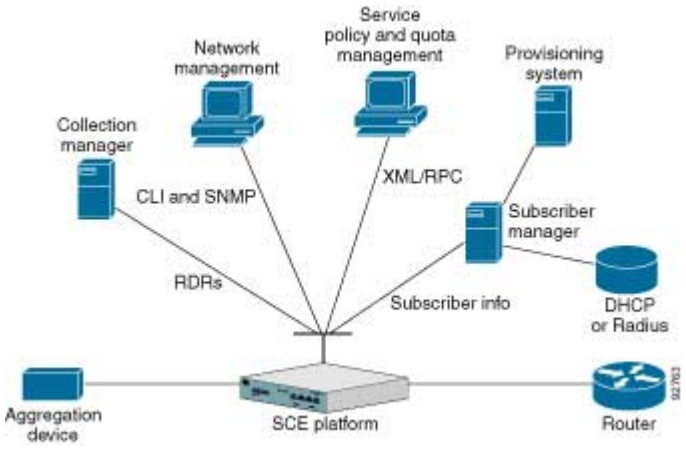

Network Management

Cisco provides complete network FCAPS (Fault, Configuration, Accounting, Performance, Security) Management.

Two interfaces are provided for network management:

- Command-Line Interface (CLI)—Accessible through the Console port or through a Telnet connection, the CLI is used for configuration and security functions.
- SNMP—Provides fault management (via SNMP traps) and performance monitoring functionality.

Subscriber Management

Where the SCA BB enforces different policies on different subscribers and to track usage on an individual subscriber basis, the Cisco Service Control Management Suite (SCMS) Subscriber Manager (SM) may be used as middleware software for bridging between the OSS and the SCE platforms. Subscriber information is stored in the SM database and can be distributed between multiple platforms according to actual subscriber placement.

The SM provides subscriber awareness by mapping network IDs to subscriber IDs. It can obtain subscriber information using dedicated integration modules that integrate with AAA devices, such as Radius or DHCP servers.

Subscriber information may be obtained in one of two ways:

- Push Mode—The SM pushes subscriber information to the SCE platform automatically upon logon of a subscriber.
- Pull Mode—The SM sends subscriber information to the SCE platform in response to a query from the SCE platform.

Data Collection

The Cisco Service Control solution generates usage data and statistics from the SCE platform and forwards them as Raw Data Records (RDRs), using a simple TCP-based protocol (RDR-Protocol). The Cisco Service Control Management Suite (SCMS) Collection Manager (CM) software implements the collection system, listening in on RDRs from one or more SCE platforms and processing them on the local machine. The data is then stored for analysis and reporting functions, and for the collection and presentation of data to additional OSS systems such as billing.

Introducing the Subscriber Manager

This chapter contains the following sections:

- Subscribers in the Cisco Service Control Solution 2-2
- Handling Subscribers 2-2
- SM Fundamentals 2-6
- SM Management 2-10
- Subscriber Manager Fail-Over 2-11

The Subscriber Manager (SM) is a middleware software component that supplies subscriber information for multiple SCE platforms in deployments where dynamic subscriber awareness is required. It does this in one of two ways:

- By pre-storing the subscriber information
- By serving as a stateful bridge between an AAA system or a provisioning system and the SCE platforms

The SCE platforms use subscriber information to provide subscriber-aware functionality, per-subscriber reporting, and policy enforcement.

Some Cisco Service Control solutions can also operate without subscriber awareness:

- Subscriber-less—Control- and link-level analysis functions are provided at a global device resolution.
- Anonymous subscriber—The system dynamically creates “anonymous” subscribers per IP address. User-defined IP address ranges may then be used to differentiate between anonymous subscribers policies.
- Static subscriber awareness—Subscriber awareness is required, but allocation of network IDs (mainly IP addresses) to subscribers is static.

In these three modes, the SCE platform handles all subscriber-related functionality and an SM module is not required.

Note Starting with SM version 2.2, you can configure the SM to operate either with or without a cluster of two SM nodes. The added functionality when operating in a *cluster topology* provides powerful new features such as *fail-over* and *high availability*. (For definitions of italicized terms, see the *Glossary of Terms*.) The information in most of this chapter is applicable whether using a cluster or not. However, for clarity, information that is applicable *only* when using a cluster is presented separately at the end of this chapter, in the section *Subscriber Manager Fail-Over* (on page 2-11).

Subscribers in the Cisco Service Control Solution

A *subscriber* is defined as a managed entity on the subscriber side of the SCE platform, to which accounting and policy are applied individually. The subscriber side of the SCE platform is the side of the SCE platform that points to the access or downstream part of the topology, as opposed to the network side, which points to the core of the network.

Handling Subscribers

The SM addresses the following issues in allowing dynamic subscriber awareness:

- Mapping—The SCE platform encounters flows with network IDs (IP addresses) that change dynamically, and it requires dynamic mapping between those network IDs and the subscriber IDs. The SM database contains the network IDs that map to the subscriber IDs. This is the main functionality of the SM.
- Policy—The SM serves as a repository of policy information per subscriber. The policy information may be preconfigured to the SM, or dynamically provisioned when the mapping information is provided.
- Capacity—The SCE platform or platforms may need to handle (over time) more subscribers than they can concurrently hold. In this case, the SM serves as an external repository for subscriber information, while the SCE platform is introduced only with the online or active subscribers.
- Location—The SM supports the functionality of sending subscriber information only to the relevant SCE platforms, in case such functionality is required. This is implemented using the domains mechanism or the Pull mode (see *Pull Mode* (on page 2-7)).

The SM database (see *SM Database* (on page 2-5)) can function in one of two ways:

- As the only source for subscriber information when the SM works in standalone mode
- As a subscriber information cache when the SM serves as a bridge between a group of SCE devices and the customer Authentication, Authorization, and Accounting (AAA) and Operational Support Systems (OSS).

Flow of Subscriber Information

The following figure shows the flow of subscriber information through the SM.

Figure 2-1: SM General Architecture

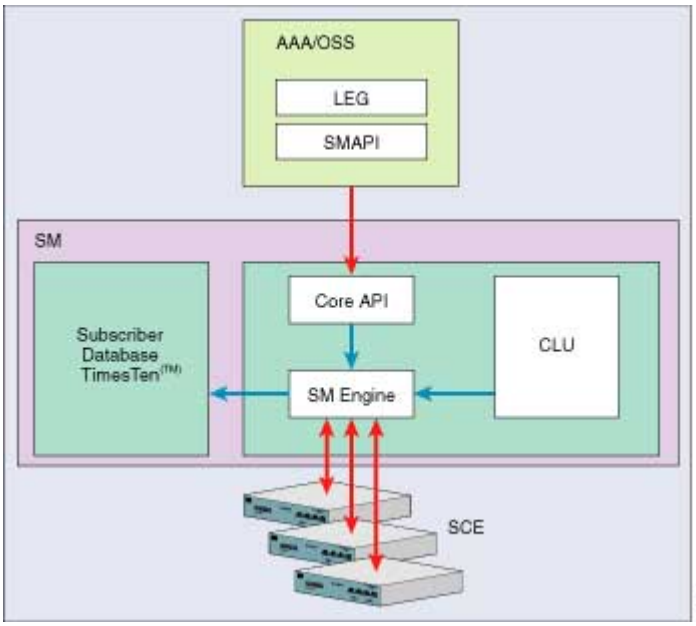

The flow takes place as follows:

- Subscriber information enters the SM in one of two ways:
  - Automatically upon the subscriber going online—A Login Event Generator (LEG) software module that integrates with the customer AAA system (such as DHCP Server, RADIUS, or Network Access System (NAS)) identifies a subscriber login event, and sends it to the SM by using the SM API.
  - Manual setup—Subscriber information is imported into the SM from a file or by using the Command-Line Utilities (CLU).

  Automatic and manual modes can be combined. For example, all subscribers may be loaded to the SM via manual setup, and a subset of the subscriber record (domain, network ID, and so on) changed automatically through the SM API.
- In automatic mode, the SCMS SM Java or C/C++ APIs are used for delegating subscriber information to the SM (see the *Cisco SCMS SM Java API Programming Guide* or the *Cisco SCMS SM C/C++ API Programming Guide*).
- The SM Engine:
  - Stores subscribers in the subscriber database
  - Introduces subscriber information to SCE Platforms

  The information may be passed automatically to the SCE platform, or it may reside in the SM database until requested by the SCE platform.

The SM may be configured with more than one SCE platform. These SCE platforms may be grouped into domains. Each domain represents a group of SCE platforms that serve the same group of subscribers.

Number of Subscribers in the SM

The subscribers of the service provider may be divided into the following logical types (at any given moment):

- Offline subscriber—A subscriber that currently does not have any IP address and as such does not generate any IP traffic. Such subscribers are not stored in the SCE platform.
- Online subscriber—A subscriber that is currently online.

  At any particular time, a certain number of online subscribers will be idle, that is, connected to the service provider but not generating any IP traffic.
- Active subscriber—An online subscriber that is actually generating IP traffic (such as by browsing the Internet or downloading a file).

In addition, the total number of subscribers is all the subscribers whose IP traffic might be traversing through the SCE platforms in a specific deployment.

There are four general scenarios for a network system using the SCE platforms:

- Total number of subscribers can be statically stored in a single SCE platform.

  This is the simplest, most reliable scenario. It may not require the use of the SM.
- Total number of subscribers exceeds the capacity of the SCE platform, but the number of online subscribers predicted at any time can be statically stored in the SCE platform.

  It is recommended to use the SM in Push mode. See *Push Mode* (on page 2-7).
- Number of online subscribers exceeds the capacity of the SCE platform, but the number of active subscribers predicted at any one time can be statically stored in the SCE platform.

  The SM must be used in Pull mode. See *Pull Mode* (on page 2-7).
- Number of active subscribers predicted at any one time exceeds the capacity of the SCE platform.

  Multiple SCE devices must be installed to divide the subscribers among the SCE platforms. If the system is divided into domains (see *Domains* (on page 2-8)), so that the SM knows in advance to which SCE platform a particular subscriber should be sent, Push mode may be used. Otherwise, Pull mode is required.

For specific scenarios using the SM with multiple servers and/or SCE platforms, see *System Configuration Examples* (on page 4-15).

SM Database

The SM uses a commercial relational database from TimesTen, optimized for high performance and with a background persistency scheme. The In-Memory Database efficiently stores and retrieves subscriber records.

A subscriber record stored in the SM Database (SM-DB) consists of the following components:

- Subscriber name (key)—A string identifying the subscriber in the SM. Maximum length: 40 characters. This can be case-sensitive or case-insensitive depending on the configuration file. By default, the database is case-sensitive. If the database is case-insensitive, the SM will convert the name to lower case when updating or querying the database.
- Domain (secondary key)—A string that specifies which group of SCE devices handles this subscriber.
- Subscriber network IDs (mappings)—A list of network identifiers, such as IP addresses or VLANs. The SCE uses these identifiers to associate network traffic with subscriber records.
- Subscriber policy—A list of properties that instruct the SCE what to do with the network traffic of this subscriber. The content of this list is application specific.
- Subscriber state (for example, quota used)—A field that encodes the subscriber state, recorded by the last SCE, to handle the network traffic of this subscriber.

You can access the subscribers using one of two indexes:

- Subscriber name
- Subscriber name + domain

Note that in cluster redundancy topology, the active machine database replicates the subscriber data to the standby machine database. For additional information, see the *Subscriber Manager (SM) Fail-Over* (on page 2-11) section.

Subscriber ID

The Subscriber ID is a string representing a subscriber that is a unique identifier for each subscriber from the customer perspective. For example, it may represent a subscriber name or a CM MAC address. This section lists the formatting rules of a subscriber ID.

It may contain up to 64 characters. The following characters may be used:

Alphanumerics	$ (dollar sign)	. (period or dot)	_ (underscore)
- (minus sign or hyphen)	% (percent sign)	/ (slash)	~ (tilde)
! (exclamation mark)	& (ampersand)	: (colon)	' (apostrophe)
# (number sign)	() (parentheses)	@ (at sign)	

For example:

```
String subID1="john";
String subID2="john@yahoo.com";
String subID3="00-B0-D0-86-BB-F7";
```

SM Fundamentals

The SM API

Use the SM API for:

- Altering the fields of an already existing subscriber record
- Setting up new subscribers in the SM
- Performing queries

The SM API is provided in C, C++, and Java. It serves as the bottom-most layer of every LEG.

SM API programmer references are provided in the *Cisco SCMS SM C/C++ API* and the *Cisco SCMS SM Java API* Programmer Guides.

The SM Login Event Generators

The SM Login Event Generators (LEGs) are software components that use the SM API to generate subscriber-record update messages (such as login/logout) and send them to the SM. LEGs are usually installed with AAA/OSS platforms, or with provisioning systems. They translate events generated by these systems to Cisco Service Control subscriber update events.

The unique functionality of each LEG depends on the specific software package with which it interacts. For example, RADIUS LEGs, DHCP LEGs, or some provisioning third party system LEGs may be implemented. LEGs can set up subscribers or alter any of the fields of an existent subscriber record.

You can connect multiple LEGs to a single SM. Conversely, a single LEG can generate events for multiple domains.

Subscriber Introduction Mode: Push or Pull

As illustrated in the figure in *Flow of Subscriber Information* (on page 2-3), the SM introduces subscriber data to the SCE platforms. This operation functions in one of two modes:

- Push—This is the simpler and recommended mode.
- Pull—Use this mode only in special cases, as explained below.

Push or Pull mode is configured for the entire SM system.

For information detailing the configuration of the subscriber integration modes, see the *SM General Section* (on page A-2).

Push Mode

In Push mode, immediately after adding or changing a subscriber record, the SM distributes, or pushes, this information to the relevant SCE platforms, as determined by the subscriber domain. When the subscriber starts producing traffic through the SCE platform, it is ready with the required subscriber information.

In some scenarios, factors such as capacity limitations make it impossible to use Push mode.

Note Use Push mode only if all online subscribers associated with a domain can be loaded simultaneously into all the SCE platforms in the domain.

Pull Mode

In Pull mode, the SCE platforms are not notified in advance of subscriber information. When an SCE platform cannot associate the IP traffic with a subscriber, it will request, or pull, the information from the SM.

The advantage of Pull mode is that there is no need to know in advance which SCE platform serves which subscriber.

The disadvantages of Pull mode are:

- Increased communication in the SM-SCE link
- Increased load on the SM, as it processes incoming requests from both the SCE device and the LEG.

Note By default, the SCE does not request subscriber information from the SM. You must configure anonymous groups in the SCE for the set of IP ranges that should be requested from the SM. See the SCE User Guide for more details on anonymous subscriber groups.

Note Pull mode must be used when all online subscribers associated with a domain exceed the capacity of the SCE platforms in the domain (but the number of active subscribers can still be loaded into the SCE platforms in the domain).

The following table summarizes the differences between the Push mode and Pull mode:

Table 2-1 Differences Between Push Mode and Pull Mode

Aspect of Use	Push Mode	Pull Mode
When to use	For simple provisioning of subscriber information to the SCE platform	For real-time, on-demand subscriber information retrieval Used in large scale deployments: • When there is no way of knowing from the IP assignment process which SCE platform will be serving a particular subscriber • When the required number of logged-in subscribers is greater than the number of concurrently active subscribers that the SCE platform can handle
Functional flow at access time	• Subscriber network login or access • From subscriber information to LEG to SM • From SM to the relevant SCE platforms	• Subscriber network login or access • From subscriber information to LEG to SM (hold in the SM database) • When the subscriber starts producing traffic that traverses the SCE platform • SCE platform asks for the subscriber information • From SM (SM database) to SCE platform
Subscriber information at the SCE platform	SCE platform always has current subscriber information: • Immediate policy enforcement • Real-time system architecture	SCE gets subscriber information on demand

Domains

The SM provides the option of partitioning SCE platforms and subscribers into subscriber domains.

The motivation for the domains concept is for enabling a single SM to handle several separate network sections, and for better control of subscriber introduction to the SCEs.

A subscriber domain is a group of SCE platforms that share a group of subscribers. The subscriber traffic can pass through any SCE platform in the domain. A subscriber can belong to only a single domain. Usually a single SCE platform serves a subscriber at any given time.

Domains are managed differently in the Push and Pull modes:

- In Push mode, all the subscribers in a subscriber domain are sent to all SCEs in the domain. The main reason for the number of SCE platforms in a single domain is redundancy.
- In Pull mode, the pull requests are handled only for subscribers in the domain of the pulling SCE platform. In Pull mode, usually a single domain covers all the subscribers.

The system is configured with one default subscriber domain called “subscribers”. When adding an SCE platform to the SM, it is automatically added to this default domain, unless otherwise specified. Subscribers are also associated with this default subscriber domain, unless otherwise specified. To associate a subscriber with a different domain, first define this domain in the configuration file, and then explicitly specify it when adding the subscriber to the SM. To associate an SCE platform with a non-default subscriber domain, edit and reload the configuration file. For more information, see *Chapter 4, Configuration and Management* (on page 4-1).

Communication Failures

A communication failure may occur either on the LEG-SM communication link or on the SM-SCE communication link. A communication failure may occur due to a network failure or because the SCE, SM, or LEG has failed. High availability and recovery from an SM failure are discussed in *SM Cluster* (on page 2-10).

When configuring the system, you should consider three issues related to communication failures:

- Communication failure detection—A timeout after which a communication failure is announced
- Communication failure handling—The action to be taken when communication on the link fails
- Communication failure recovery—The action to be taken when communication on the link resumes

Failure Detection Mechanism

Either one of two mechanisms detects a communication failure:

- Monitoring the TCP socket connection state. All peers do the monitoring.
- Using a keep-alive mechanism at the PRPC protocol level

Failure Handling Mechanism

There are two configuration options for handling communication failures:

- Ignore communication failures
- Erase the subscriber mappings in its database and start handling flows without subscriber awareness

Erasing the mappings in the database is useful when you want to avoid incorrect mappings of subscribers to IP addresses. This configuration is implemented by requesting to clear all mappings upon failure.

Failure Recovery Mechanism

The SM recovers from communication failures by resynchronizing the SCE platform with the SM database.

SM Cluster

The SM supports high availability using the Veritas Cluster Server (VCS) technology. In a high availability topology, the SM software runs on two machines, designated as the active machine and the standby machine. Subscriber data is continuously replicated from the active to the standby machine, ensuring there is minimal data loss in case of active SM failure. When the active machine fails, the standby machine discovers the failure and becomes active. For additional information, see the *Subscriber Manager Fail-Over* (on page 2-11) section.

SM Management

SM management includes configuration, fault management, logging management, and performance management.

Configure the SM using the following:

- Configuration file (*p3sm.cfg*)—For setting all configuration parameters of the Subscriber Manager.

Note Changes that you make in the configuration file take effect only when you load the configuration file using the Command-Line Utilities (CLU) or when you restart the SM.

  For a detailed description of this file, see *Appendix A, Configuration File Options* (on page A-1).

- Command-Line Utilities (CLU)—For ongoing subscriber management and monitoring of the SM. CLU commands are shell tools that you can use to manage subscribers, install or update applications, retrieve the user log, and load the configuration file when updated.

  For a complete description of the Command Line Utilities, see *Appendix B, Command-Line Utilities* (on page B-1).

  The CLU can be invoked locally, through a Telnet (or SSH) session to the SM hosting platform.

Use the SM user log files for logging, fault, and performance management. The log file contains information regarding system events, failures, and periodic system performance reports.

Subscriber Manager Fail-Over

You can configure the SM to operate with or without a cluster. The added functionality when operating in a *cluster topology* provides powerful new features such as *fail-over* and *high availability*. (For definitions of italicized terms, see the *Glossary of Terms*.)

This section describes topics that are related to using the Subscriber Manager together with clusters and redundancy.

As the Subscriber Manager plays a critical role in the Cisco ***SCA BB*** solution that is deployed in tier-one service provider environments, it also supports, starting with SM version 2.2, a fail-over operational mode. This feature minimizes system downtime that is caused by SM failure (as further discussed in *Overview* (on page 2-12)).

This section introduces various concepts that are related to using a cluster of two SM nodes in a fail-over operational mode.

Note For the purposes of this section, it is assumed that the reader is familiar with the Veritas Cluster technology.

Overview

The fail-over scheme that is implemented in the SM is based on the Veritas cluster technology. The cluster includes two machines, each of them running SM TimesTen and Veritas software. The Veritas Cluster Server (VCS) software consolidates the SMs and exposes a single entity by providing a single virtual IP address for the entire cluster.

The cluster software distinguishes an active and a standby machine: the active machine "owns" the virtual IP address and all network connections, while the standby machine is passive until a fail-over occurs. At fail-over, the IP address is passed from the failing server to the backup server, which becomes activated and re-establishes all network connections.

When a fail-over occurs, the LEGs lose their connection with the failed SM, and reconnect to the activated (backup) SM and retransmit their uncommitted messages. The activated SM connects to the SCE platforms and performs an SCE resynchronization.

The TimesTen database replication agent constantly replicates the SM database from the active node to the standby node. This enables a fast fail-over from one SM to another, since the subscriber data in the activated machine is always valid. The two SM nodes do not communicate except for passing the subscriber data.

The VCS uses software components called "cluster agents" to monitor and control the state of resources such as Network Interface Cards (NICs), disks, IP addresses, and processes. Cisco supplies cluster agents to monitor the SM and the TimesTen database daemon and replication agent.

As part of the cluster operation, the TimesTen database daemon and replication agents are up and running regardless of the fail-over state. The SM Veritas agent monitors the daemon and the replication agent process. In case one of them fails, a fail-over takes place.

Note The SM software configuration on both the active and the standby machines must be *identical*. Apply the same configuration file and the same application PQI module to both machines.

The following sections describe these concepts in further detail.

Normal Operation

The two SM nodes operate in hot-standby mode, where at any given time one node (the active node) receives and processes all the SM events, while the other node (the standby node) waits and is ready to go into operation on fail-over. For enabling seamless fail-over and for minimizing the fail-over time, the two SM nodes operate without an external storage device.

During the normal operation of the cluster, the active node (selected by the cluster):

- Performs all SM functionality of a non-cluster environment
- Provides "health" information for the cluster agent
- Periodically replicates its subscriber database to the standby node

On the standby node, both the SM and the TimesTen software are running:

- The SM is fully configured. (It is applied with the same configuration file and PQI application module as the active node, but does not interfere with the active node's work.)
- The SM connects to the TimesTen database, but does not connect to the LEG and the SCE devices.
- The TimesTen software is operating as a replication client for the subscriber database, receiving and applying updates from the active node's TimesTen software.

Fail-Over Topology

The following figure depicts an SM cluster configuration in a topology with a redundant AAA server and two SCE 2000 platforms that are cascaded for redundancy.

Figure 2-2: SM Cluster Configuration for Fail-Over Topology

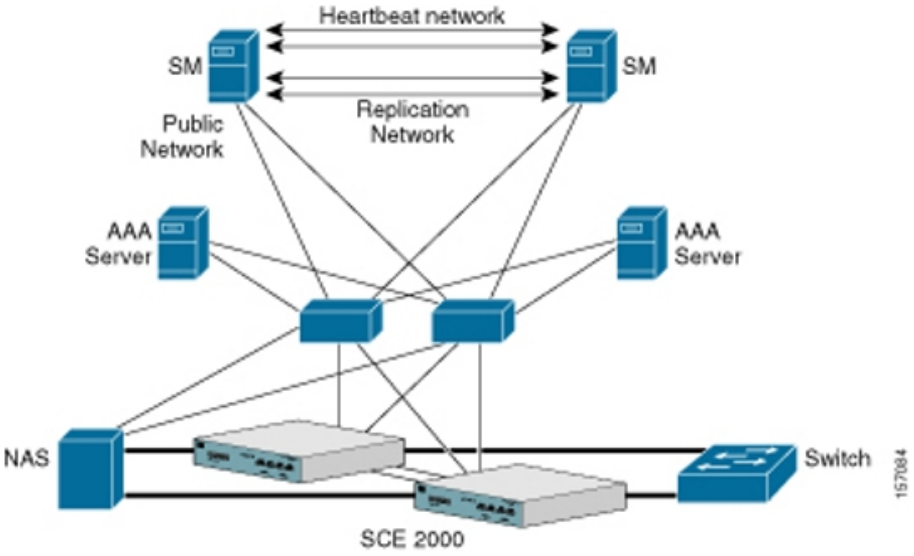

As already mentioned, an SM fail-over topology includes two SM nodes connected in a cluster scheme.

Two dedicated (private) redundant networks interconnect the two nodes:

- Heartbeat network—Used by the Veritas Cluster Server to perform cluster monitoring and control.
- Replication network—Used by the replication process to pass the subscriber records.

The two nodes should be located in the same site, where the heartbeat network is implemented using back-to-back connectivity between the two nodes or via redundant switches. Each node in the cluster has redundant network paths (NICs) connecting it to all of the external entities with which the SM communicates (AAA, LEG, SCE).

Each node in the cluster has a minimum of six Ethernet NICs, where

- Two NICs are used for the (private) heartbeat network
- Two NICs are used for the (private) replication network
- Two NICs are used for the public network (connectivity to SCEs and LEGs, and management of the SM).

The cluster has a virtual IP (VIP) address used for communication with the external entities. Each node in the cluster has also an IP address for administration of the node/cluster, as well as an IP address for replication use.

Upon failure of the primary NIC of the public network, there is a fail-over to the secondary NIC on the same node, keeping the same IP addresses (VIP1), with no fail-over of the cluster. Upon failure of the primary NIC of the replication or heartbeat networks, there is fail-over to the secondary NIC on the same node, keeping the same IP addresses (VIP2 and VIP3), with no fail-over of the cluster.

The following diagram illustrates the usage of the regular and virtual IP addresses used in cluster configuration:

- Administration of the nodes uses IP1/IP2 and IP3/IP4 respectively.
- The cluster IP address for external clients over the public network uses VIP1.

Figure 2-3: Regular and Virtual IPs in Cluster Configuration

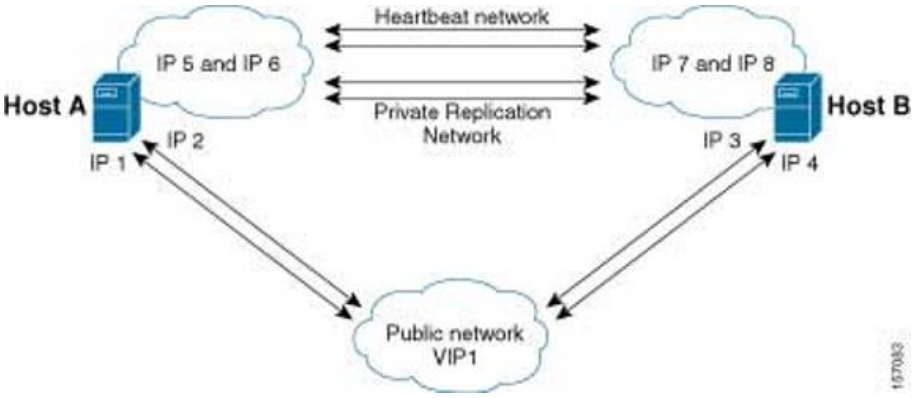

For further information about replication IP configuration, see *Appendix E, Veritas Cluster Server Requirements and Configuration* (on page E-1).

Fail-Over

Fail-Over Operation

During normal operation, the Veritas Cluster Server mechanism automatically selects one of the SM servers to be active and the other to be standby.

The active SM server performs all the normal SM functionality. The two servers maintain the heartbeat mechanism between them, and the active server continuously replicates the subscriber database to the standby server's database.

The standby SM server acts as a hot-standby machine, so it is completely ready for taking over (becoming activated) in a minimal fail-over time.

The following types of failures trigger the fail-over mechanism:

- SM application failure, including failure of the TimesTen database.
- Failure of the TimesTen daemon of the TimesTen replication process.
- SUN server failure, due to failure of one of the resources of the server; for example, failure of both of the public network NICs.
- Manual activation of fail-over.

Note Communication failure does *not* cause a fail-over if there is a redundant NIC. Therefore, because each SUN machine has two NICs for connecting to external devices, a failure of one of the NICs merely causes switching to the redundant NIC, without activating the fail-over mechanism.

After detecting a failure, the standby SM becomes activated, and the following occurs:

- The activated SM takes over the IP resources of the virtual IP mechanism.
- The LEGs reconnect to the activated SM.
- The activated SM creates IP connections with the SCEs and resynchronizes with them.
- The activated SM starts processing information that is sent from the different LEGs and forwards it to the SCEs.

Recovery

Different types of failures require different triggering for the recovery procedure. Some failures may recover automatically such as intra-node ports link-failure, which recovers automatically when the link revives, while others may need manual intervention.

Recovery may take place when an SM that experienced a failure is self-recovered or after it was replaced (if needed). The purpose of the recovery procedure is to take the cluster back to a fully functional mode. When the recovery procedure ends, the behavior is the same as it was after installation.

The failed SM server is recovered manually or automatically, according to the type of failure that occurred. The recovery procedures, and when they are used, are described in the following sections.

Machine Reboot

Recovering from a machine reboot is a fully automatic recovery process, where the failed SM server reboots, and after establishing a connection with the other server and synchronizing the databases, the cluster of the two SM servers is ready again for fail-over operation.

The following steps are automatic steps:

Step 1 The reboot process is run on the node.

Step 2 VCS makes the node standby.

Step 3 The node boots.

Step 4 VCS establishes intra-node communication and the new node joins the cluster.

Step 5 The TimesTen database replication process is started from the point before the reboot.

The SM in the recovered server is ready after the database recovery process is running and the SM moves from Init state to Standby state.

Replacing the Server

Replacing the server is necessary when the machine has an unrecoverable physical failure. A new machine that is installed with fresh SM, TimesTen, and VCS installations replaces the server.

Replacing the server is a manual recovery, where the failed SM server is physically replaced. After connecting the new SM server to the network, configuring it and synchronizing the two databases, the cluster of the two SM servers is ready again for fail-over operation.

To manually replace the server:

Step 1 Connect a new server to the inter-node ports and intra-node ports (but leave the network ports disconnected).

Step 2 Basic network and cluster configurations—perform manually (the first time).

Step 3 Copy the configuration files from the active node.

Use the following CLU command if you only need to copy the *p3sm.cfg* file:

p3sm –-load-config –-remote=<*NEW-SM_IP*>.

Step 4 Perform the TimesTen database duplication operation. See *Database Duplication Recovery* (on page 2-17).

Step 5 Start the VCS operation on the recovered node.

Step 6 Connect the network ports.

The SM in the recovered server is ready after the database recovery process is completed and the SM moves from Init state to Standby state.

Database Duplication Recovery

Database duplication recovery is a manual recovery, which is needed when the standby node database loses synchronization with the active node database. Loss of synchronization can occur when one of the SM machines is replaced or when the replication process on the active node fails to replicate all of the data inserted to its database (replication NICs were disconnected).

To perform the database duplication recovery (in standby node):

Step 1 Stop the cluster server (VCS) monitoring of the resources. Use the VCS CLU **`hastop -local`** command to stop the VCS.

Step 2 Stop the SM, so it will not be affected by clearing the database. Use the CLU command **`p3sm --stop`**.

Step 3 Stop the replication agent. Use the CLU command **`p3db --rep-stop`**.

Step 4 Destroy the database. Use the CLU command **`p3db --destroy-rep-db`**.

Step 5 Duplicate the remote database to the local machine. Use the CLU command **`p3db --duplicate`**.

Step 6 Start the cluster server monitoring of the resources (use the VCS CLU **`hastart`** command), which will automatically start the replication process and the SM.

Database Duplication Recovery Management

The two SM servers are configured using Command-Line Utilities and a configuration file (see *Chapter 4, Configuration and Management* (on page 4-1) and *Configuring a Subscriber Management Solution* (on page 4-3)). The actual configuration is performed for the active SM and then manually replicated for the standby SM.

To perform configuration duplication:

Step 1 Establish an FTP connection between the active and standby machines.

Step 2 Copy the configuration file from *`~pcube/sm/server/root/config/p3sm.cfg`* on the active node to the standby node, and apply the SM configuration file by using the CLU command **`p3sm --load-config`**.

Alternatively, you can replicate the SM configuration file to the standby machine by running the CLU command **`p3sm --load-config --remote=<`***`NEW-SM_IP`***`>`** on the active machine.

Step 3 Copy the application PQI file you installed on the active node to the standby node.

Step 4 Install the PQI file. Use the CLU command **`p3inst --install -f <PQI file path>`**.

Step 5 If you have made changes in the database-related configuration files, copy the files to */etc/system* (for **Solaris**) or to */etc/sysctl.conf* (for **Linux**), and */var/TimesTen/sys.odbc.ini* from the active node to the standby node.

Note If you perform Step 5 (copying files), a reboot of the standby node is required.

Note If the database is located in different directories in the two nodes, then the files *sys.odbc.ini* in both nodes are not identical and the actual parameter changed in the file must be copied.

Configure and administer the Veritas Cluster Server using Veritas tools.

Notifications are enabled through SNMP traps that the Veritas Cluster Server provides. The Veritas Cluster Server supports SNMP traps such as:

- Fatal failure detected (local or remote)
- Secondary node starts fail-over procedure
- Secondary node is operational (end of fail-over)

CHAPTER 3

Installation and Upgrading

This chapter describes how to install the Cisco Service Control Management Suite Subscriber Manager (SCMS SM), in addition to how to upgrade and uninstall. This chapter also discusses topics related to installation, upgrading, and uninstalling.

This chapter contains the following sections:

- Installation Overview 3-1
- Contents of Distribution Files 3-2
- Documentation 3-5
- System Requirements 3-5
- Installation Procedures 3-7
- Typical Installation 3-8
- Verifying that the Installation was Successful 3-15
- Configuring the Subscriber Manager 3-15
- Additional Installation Procedures 3-16
- Troubleshooting the Installation 3-18
- Installing an Application 3-19
- System Changes Made by Installation Scripts 3-20
- Upgrading 3-21
- Uninstalling 3-26

Installation Overview

Installing the SM is an automated process. It consists of executing an installation script residing on the root of the SM distribution files supplied by Cisco.

Note For Solaris: The procedure also requires modifying the `/etc/system` file. Do this manually or use some other automated utility.

Note For Linux: The procedure also requires modifying the `/etc/sysctl.conf` file. Do this manually or use some other automated utility.

The installation procedure installs the following components:

- SM and Command-Line Utilities (CLU)
- TimesTen database and DSN
- Java Runtime Environment (JRE)
- SM Veritas Cluster Agents

The installation procedure also includes:

- Setting up a pcube user and group
- Adding startup and shutdown scripts
- System configuration for TimesTen (performed manually or using a script)
- Replication scheme setting (performed by running a CLU). (Relevant only for cluster setups).

After completing installation and configuration, you can use the SM to introduce subscribers to the system.

Contents of Distribution Files

The SCMS SM components are supplied in three distribution files:

- SM for Solaris
- SM for Linux
- Login Event Generators (LEGs)

Each distribution file is supplied as a tar file, which is compressed by gzip and has an extension of `.tar.gz`. The following table lists the contents of the SM installation distribution files for Solaris and Linux.

Table 3-1 Contents of SM Distribution Files

Path	Name	Description
DIST_ROOT		Cross-platform files
	hooks.sh	User-defined function for upgrade
	install	Typical installation procedure description
	install-dsn.sh	TimesTen DSN configuration script
	installjava.sh	JRE installation script
	install-sm.sh	SM installation script
	install-tt.sh	TimesTen installation script

Path	Name	Description
	install-vcs-agents.sh	VCS agents installation script
	linux-def.sh	Linux-specific definitions (only in the Linux distribution file)
	solaris-def.sh	Solaris-specific definitions (only in the Solaris distribution file)
	MANIFEST	CD information
	p3sm.sh	Startup and shutdown script
	Prerequisites	System minimal requirements list
	sm-common.sh	General installation script
	sm-dist.tar.gz	SM distribution
	tt-sysconf.sh	TimesTen system configuration script
	uninstall-sm.sh	SM uninstall script
	upgrade-sm.sh	SM upgrade script
	vcs-agents-dist.tar.gz	VCS agents distribution
DIST_ROOT /Java/		Java Runtime Environment files
	j2re1.4.2_05-linux.tar.gz	JRE for Linux (only in the Linux distribution file)
	j2re1.4.2_05-solaris.tar.gz	JRE for Solaris (only in the Solaris distribution file)
	LICENSCE	JRE license
DIST_ROOT /TimesTen/		TimesTen files
	pqb_resp_uninst.txt	Response file for TimesTen uninstall
	pqb-odbc-ini.txt	Open Database Connectivity (ODBC) definitions
	pqb-response50.txt	Response file for TimesTen installation
	pqb-sys-odbc-ini.txt	Open Database Connectivity (ODBC) definitions
	TT5125LinuxRH32.tar.Z	TimesTen for Linux (only in the Linux distribution file)
	TT5125Sparc64.tar.Z	TimesTen for Solaris 64-bit (only in the Solaris distribution file)

The following table lists the contents of the LEG distribution file:

Table 3-2 Contents of the LEG Distribution File

Path	Name	Description
DIST_ROOT		Cross-platform files

Path	Name	Description
	MANIFEST	Distribution information
DIST_ROOT/bgp_leg		Border Gateway Protocol (BGP) LEG files
	bgp_leg.tar.gz	BGP LEG distribution
	Install	LEG installation procedure description
	Install-bgp-leg.sh	BGP LEG installation script
	linux-def.sh	Linux-specific definitions
	sm-common.sh	General installation script
	solaris-def.sh	Solaris-specific definitions
DIST_ROOT/cnr_leg		Cisco Network Register (CNR) LEG files
	cnr-leg-dist.tar.gz	CNR LEG distribution
	Install	LEG installation procedure definitions
DIST_ROOT/Lease_Query_Leg		Lease Query LEG files
	dhcp_forwarder.tar.gz	DHCP Forwarder distribution
	Install	LEG installation procedure description
	install-forwarder.sh	DHCP Forwarder installation script
	linux-def.sh	Linux-specific definitions
	sm-common.sh	General utility script
	solaris-def.sh	Solaris-specific definitions
DIST_ROOT/Lease_Query_LEG/sce		Lease Query LEG SCE files
	leaseq.pqi	DHCP Lease Query LEG distribution
	dhcp_pkg.cfg	Default configuration file for package association
DIST_ROOT/Lease_Query_LEG/sm		Lease Query LEG SM files
	leaseq.pqi	DHCP Lease Query LEG distribution
DIST_ROOT/rdr_dhcp_leg		SCE-Sniffer DHCP LEG files
	Install	LEG installation procedure description
	rdrdhcp.pqi	SCE-Sniffer DHCP LEG distribution
DIST_ROOT/rdr_radius_leg		SCE-Sniffer RADIUS LEG files
	Install	LEG installation procedure distribution
	rdradius.pqi	SCE-Sniffer RADIUS LEG distribution
DIST_ROOT/sce_api		SCE Subscriber API files

Path	Name	Description
	readme	API setup procedure description
	Sce-java-api-dist.tar.gz	API distribution
DIST_ROOT/sm_api		SM API files
	readme	API setup procedure description
	sm-c-api-dist.tar.gz	C API distribution
	sm-java-api-dist.tar.gz	Java API distribution

Documentation

The SM installation distribution file contains the following documents:

- Manifest—Contains the version and build numbers for all components from which the distribution files were built
- Install—The SCMS SM typical installation procedures
- Prerequisites—Minimal system requirements for installation of the SM

System Requirements

You can install the SM on the following platforms:

- Solaris—SUN SPARC machine running Solaris. See Table 3-3, Minimal System Hardware Requirements, and Table 3-4, Solaris Minimal System Software Requirements.
- Linux—Machine with Intel-based processor running Linux. See Table 3-3, Minimal System Hardware Requirements, and Table 3-5, Red Hat Minimal System Software Requirements.

The machine should conform to the system requirements listed in the following tables.

Note The specifications listed in Table 3-3 are minimal. They should be verified in order to guarantee specific performance and capacity requirements.

Table 3-3 Minimal System Hardware Requirements

Item	Requirement
CPU	• SUN SPARC, 64-bit, minimum 500 MHz (for Solaris) • INTEL processor, 32-bit, minimum 1 GHz (for Linux Red Hat)
RAM	Minimum 1 GB; see the RAM and Memory Configuration Parameters Versus Number of Subscribers table in *Configuring the System Memory Settings* (on page 3-10).

Item	Requirement
Free Disk Space	Minimum 3 GB total, of which: • Minimum 1 GB free on partition where VARDIR (SM database repository) is installed • Minimum 0.5 GB free on partition where PCUBEDIR (SM files) is installed • Minimum 200 MB free on partition where */tmp* is mounted
Network Interface	Depends on whether or not the configuration includes a cluster: • No cluster—One (1) 100BASE-T Ethernet • With cluster—Six (6) 100BASE-T Ethernet
CD-ROM drive	Recommended

Note For the hardware and software system requirements for the Veritas Cluster Server, see *Veritas Cluster Server System Requirements* (on page E-1).

Table 3-4 Solaris System Software Requirements

Item	Requirement
OS	Solaris 5.8 64-bit build 04/01 or later; currently, only 64-bit versions of Solaris 5.8 and 5.9 are supported. Solaris Core Installation
System Packages	Mandatory: • SUNWbash—GNU Bourne-Again shell (bash) • SUNWgzip—GNU Zip (gzip) compression utility • SUNWzip—Info-Zip (zip) compression utility • SUNWlibC—Sun Workshop Compilers Bundled libC • SUNWlibCx—Sun WorkShop Bundled 64-bit libC • sudo (superuser do) package Optional: • SUNWadmap—system administration applications • SUNWadmc—system administration core libraries

Note It is strongly recommended to apply the latest patches from SUN. You can download the latest patches from the SUN patches website.

Table 3-5 Red Hat System Software Requirements

Item	Requirement
OS	Red Hat Enterprise Linux AS/ES 3.0/4.0; currently, only 32-bit versions are supported. Red Hat Core Installation
System Packages	Mandatory: • GNU Bourne-Again shell (bash-2.05b-29.i386.rpm) • GNU Data Compression Program (gzip-1.3.3-9.i386.rpm) • File compression and packaging utility (zip-2.3-16.i386.rpm) • Standard C++ libraries for Red Hat Linux 6.2 backward compatibility (compat-gcc-7.3-2.96.122.i386.rpm) • sudo (superuser do) package For integrating with the C API: • GNU cc and gcc C compilers (gcc-3.2.3-20.i386.rpm) • C++ support for the GNU gcc compiler (gcc-3.2.3-20.i386.rpm)

Note It is strongly recommended to apply the latest patches from Red Hat.

Installation Procedures

All installations can be performed by executing one of the installation scripts located on the *root* of the SM distribution file.

You can choose to install the SM, TimesTen, and Java separately. In most cases, the SM installation script is the only script needed for completing the installation.

Each installation script displays messages describing the significant steps that are being performed. These messages are also sent to the system log for future reference. See *Logging Script Messages* (on page 3-20) for more information about the system log messages.

If you try to install the SM on a machine on which the SM is currently running, or to a directory in which the SM is already installed (even if not running), the operation will fail and you will be requested to upgrade the SM. See *Upgrading* (on page 3-21).

Typical Installation

This section assumes that you want to install the SM and TimesTen components, in addition to the Java Runtime Environment (JRE), and that you want to perform the necessary system configurations in order for these components to work.

There are seven steps for the installation, which are further described in the following sections:

Step 1 Extracting the distribution files.

Step 2 Executing the *install-sm.sh* script. The root user must invoke this script.

Step 3 Setting a password for the user **pcube**.

Step 4 Configuring the system memory settings.

Step 5 Rebooting the computer.

Step 6 Installing the SCA BB package and Login Event Generators

Step 7 Add a user for PRPC authentication

Note In a high availability setup (see *SM Cluster* (on page 2-10)), you must install the SM Cluster VCS agents. See *Installing SM Cluster Agents* (on page 3-18).

Step 1: Extracting the Distribution Files

Before you can install the SM, you must first load and extract the distribution files on the installed machine or in a directory that is mounted to the installed machine.

To extract the distribution files:

Step 1 Download the distribution files from the Cisco web site.

Step 2 Use FTP to load the distribution files to the SM.

Step 3 Unzip the files using the following command:

```
gunzip SM_dist_3.0_B<build number>.tar.gz
```

Step 4 Un-tar the tar the file using the following command:

```
tar –xvf SM_dist_3.0_B<build number>.tar
```

Step 2: Executing the install-sm.sh Script

Note Before starting the installation, make sure that disk space requirements listed in the *System Requirements* (on page 3-5) are satisfied.

To execute the install-sm.sh script:

From your workstation shell prompt, move to the directory to where the distribution file was extracted and run the following command:

`# install-sm.sh` [-d *install-directory*]

In this script, *install-directory* is the location in which you want to install the SM.

If you use the default install location, `/opt/pcube`, you need not specify the **-d** flag.

For additional information about the script operation, see *Installing the Subscriber Manager* (on page 3-17).

Step 3: Setting the Password for the pcube User

After the installation script has completed successfully, set the password for the pcube user by running the following command:

`# passwd pcube`

Note It is important to remember the password you have selected.

Step 4: Configuring the System Memory Settings

System Memory Settings

Without Quota Manager

Set the system memory configuration requirements according to the maximum number of subscribers. The following table lists the recommended memory configuration values based on the number of supported subscribers. The settings apply when the Quota Manager is disabled

Table 3-6 RAM and Memory Configuration Parameters Versus Number of Subscribers

Maximum Number of Subscribers	Required RAM	SM Process Memory Setting	(TimesTen Memory Settings) Shared Memory	(TimesTen Memory Settings) PermSize	(TimesTen Memory Settings) TempSize
100,000	1 GB	256 MB	512 MB	200 MB	100 MB
500,000	2 GB	512 MB	1024 MB	512 MB	256 MB
1,000,000	3 GB	768 MB	1280 MB	768 MB	256 MB
2,000,000	4 GB	1280 MB	2048 MB	1536 MB	256 MB
3,000,000	5 GB	1792 MB	2560 MB	2048 MB	256 MB
4,000,000	6 GB	2048 MB	3328 MB	2816 MB	256 MB

Note The required RAM in the table is calculated for 40 SCE connections per SM. For each additional SCE you should add an additional 25 MB for the required RAM and SM process memory setting.

Description of the table columns:

- Maximum Number of Subscribers—The maximum number of subscribers that the SM has to support. For additional information about the maximum number of subscribers configuration, see *Configuring the Maximum Number of Subscribers* (on page 3-11).
- Required RAM—The RAM requirement for the machine running the SM.
- SM Process Memory Setting—The required memory configuration for the SM process itself. For additional information about the SM process memory configuration, see *Configuring the SM Process Memory Settings* (on page 3-14).
- The configuration required for TimesTen to run correctly. For additional information, see *Configuring the System for TimesTen* (on page 3-12).

If the previous table does not list the maximum number of subscribers that you require, use the settings specified for the next higher value of Maximum Number of Subscribers. For example, for 1,200,000 subscribers, use the values specified for 2,000,000 subscribers (4 GB of RAM, and so on).

With Quota Manager

The following table lists the recommended memory configuration values based on the number of supported subscribers. The settings apply when the Quota Manager is enabled

Table 3-7 RAM and Memory Configuration Parameters Versus Number of Subscribers

Maximum Number of Subscribers	Required RAM	SM Process Memory Setting	(TimesTen Memory Settings) Shared Memory	(TimesTen Memory Settings) PermSize	(TimesTen Memory Settings) TempSize
500,000	3 GB	512 MB	1280 MB	768 MB	256 MB
1,000,000	3 GB	768 MB	1792 MB	1280 MB	256 MB
2,000,000	5 GB	1280 MB	3072 MB	2560 MB	256 MB
3,000,000	7 GB	1792 MB	4096 MB	3584 MB	256 MB
4,000,000	8 GB	2048 MB	5376 MB	4864 MB	256 MB

Note The required RAM in the table is calculated for 20 SCE connections per SM. For each additional SCE you should add an additional 50 MB for the required RAM and SM process memory setting.

Configuring the Maximum Number of Subscribers

There is a limit to the maximum number of subscribers that can be stored in the Subscriber Manager database. The limit is 4,000,000 subscribers for Solaris and 2,000,000 subscribers for Linux.

The Subscriber Manager default configuration supports a maximum of 800,000 subscribers.

To increase this number:

Step 1 Add the following line to the `[Data Repository]` section of the *`p3sm.cfg`* configuration file:

`max_number_of_subscribers=`*number*

Step 2 Restart the SM.

Note In cluster setups, perform this on both SM machines.

Configuring the System for TimesTen

TimesTen requires that certain changes be made in the system kernel configuration file (*/etc/system* in Solaris and */etc/sysctl.conf* in Linux). These changes increase the shared memory and semaphore resources on the Solaris machine from their defaults. For additional information regarding these changes, refer to the TimesTen documentation.

Note It is recommended that you review the */etc/system* or the */etc/sysctl.conf* file before running the *tt-sysconf.sh* script, because the script overwrites the current file settings with the values listed in the "Making the changes manually" procedure. If you want to keep some or all of the current file settings, edit the system configuration file and perform the changes manually.

Configuring the System Kernel Configuration File

TimesTen requires that certain changes be made in the operating system kernel configuration file:

- For Solaris, modify file */etc/system*.
- For Linux, modify file */etc/sysctl.conf*.

These changes increase the shared memory and semaphore resources on the machine from their defaults.

Making the changes automatically:

Note It is recommended that you review the system configuration file before running the *tt-sysconf.sh* script, because the script overwrites the current file settings with the values listed in the "Making the changes manually" procedure. If you want to keep some or all of the current file settings, edit the configuration file by performing the changes manually.

To make the required changes automatically, run the *tt-sysconf.sh* script file. The root user must invoke this script file, without arguments, as follows:

```
# tt-sysconf.sh
```

Making the changes manually:

Editing the configuration file manually is required when you require support for more than 100,000 subscribers in the SM. Your system's sizing requirements only affect the shared memory size. To determine the correct configuration values for your system, see the table in *System Memory Settings* (on page 3-10).

- For Solaris, make the required changes manually by adding the following lines to the */etc/system* file and configuring the shared memory size:

```
*---- Begin settings for TimesTen
set semsys:seminfo_semmni = 20
set semsys:seminfo_semmsl = 100
set semsys:seminfo_semmns = 2000
set semsys:seminfo_semmnu = 2000
set shmsys:shminfo_shmmax = 0x20000000
*---- End of settings for TimesTen
```

- For Linux, make the required changes manually by adding the following lines to the */etc/sysctl.conf* file and configuring the shared memory size:

```
*---- Begin settings for TimesTen
kernel.shmmax = 536870912
kernel.sem = “SEMMSL_250 SEMMNS_32000 SEMOPM_100 SEMMNI_100
*---- End of settings for TimesTen
```

Configuring /var/TimesTen/sys.odbc.ini

Some installations might require changing TimesTen parameters so that the database will run as desired. However, do not make any changes if the default values suit your requirements.

Setting the multi-processor optimization:

If your system is a multi-processor machine, the value of the ***SMPOptLevel*** parameter of the Pcube_SM_Repository in the *sys.odbc.ini* file should be set to 1. Otherwise, it should be set to 0 or not set at all. The installation script automatically sets this parameter according to the number of available processors.

Setting the database size:

If your system needs to support more than 100,000 subscribers, set the values of parameters ***PermSize*** and ***TempSize*** of the Pcube_SM_Repository in the *sys.odbc.ini* file. See *System Memory Settings* (on page 3-10).

For example:

```
PermSize=500
TempSize=150
```

Note Solaris—Remember to set the value of parameter ***shmsys:shminfo_shmmax*** in the */etc/system* file to be larger than the sum of ***PermSize*** and ***TempSize***.

Note Red Hat—Remember to set the value of parameter *kernel.shmmax* in the */etc/sysctl.conf* file to be larger than the sum of *PermSize* and *TempSize*.

Configuring the SM Process Memory Settings

By default, the SM process uses 256 MB of RAM memory. However, in certain application component configurations, the SM process needs to allocate additional memory to work correctly. Setting an environment variable called *PCUBE_SM_MEM_SIZE* with the desired memory size (in megabytes) instructs the SM start-up scripts to allocate the defined memory size for the SM process.

You can set the memory size value for this environment variable for the user **pcube**, or you can configure the desired process memory size in the *sm.sh* file located in the root directory of the user **pcube** (*~pcube/sm.sh*).

The following example, which shows a line in the *sm.sh* file, defines a memory size of 512 MB for the SM process:

```
PCUBE_SM_MEM_SIZE=512
```

Note To prevent performance degradation because of memory swapping, make sure that the machine has enough RAM for the SM process, the SM database, and all of the other applications running on this machine.

Note To determine the correct memory values for your installation, see *System Memory Settings Versus* (on page 3-10).

Step 5: Rebooting

Reboot the computer to complete the installation.

Step 6: Installing the SCA BB Package and LEG components

Depending on the integration type, you might need to install the SCA BB package on the SM or install Login Event Generator (LEG) modules.

To perform the installation, use the `p3inst` command-line utility. For example:

```
> p3inst --install –file=eng30.pqi
```

For additional information, see *Installing an Application* (on page 3-19).

Step 7: Adding a User for PRPC Authentication

It is necessary to add a user for PRPC authentication because the SCA BB application requires a username and password when connecting to the SM.

To add a user for PRPC authentication, use the `p3rpc` command-line utility. For example:

```
> p3rpc --set-user --username=<username> --password=<password>
```

Verifying that the Installation was Successful

To verify that the installation was successful, run a CLU utility, such as the `p3sm` command that displays general information about the SM.

To verify that the SM installation was successful:

From your workstation shell prompt, change to the *~pcube/sm/server/bin* directory, and type:

```
> p3sm --sm-status
```

The above command displays the current status of the SM.

Note Wait a few minutes after the installation before running this command to allow the SM to become operational.

The output of this command should indicate that the SM is running.

In case of errors during installation, the command will output a description of these errors.

Configuring the Subscriber Manager

After installing the SM, you can configure the SM to your specific needs. In particular, you should address the following parameters at this point:

- topology—Cluster or standalone
- introduction_mode—Pull or push
- support_ip_ranges—Whether IP-ranges should be used in the installed setup

To configure the SM, edit its configuration file, *p3sm.cfg*, using any standard text editor. The configuration file is described in detail in *Chapter4, Configuration and Management* (on page 4-1), and in *Appendix A, Configuration File Options* (on page A-1). After you finish editing the *p3sm.cfg* configuration file, use the **p3sm** utility for updating the SM with the new settings:

To load the SM with a new configuration file (*p3sm.cfg*):

From your workstation shell prompt, type:

```
> p3sm --load-config
```

The configuration file is loaded and the SM configuration is updated accordingly.

Additional Installation Procedures

The following procedures complement the ones described in Typical Installation:

- *Installing an SM Cluster* (on page 3-16)—Should be used if installing two SM nodes for the first time.
- *Installing the Subscriber Manager* (on page 3-17)—Should be used if installing the SM database to a location other that the default, or if the user **pcube** is already defined in the system.
- *Installing SM Cluster Agents* (on page 3-18)—Must be installed when in a High Availability setup where Veritas Cluster Server (VCS) is used.

Installing an SM Cluster

The installation of an SM cluster is very similar to installing the SM on two machines.

To install an SM Cluster:

Step 1 Before installing the SM cluster, you must first install the Veritas Cluster Server software on both nodes.

Step 2 Install the SM on both machines, as described in *Installing the Subscriber Manager* (on page 3-17).

Step 3 Configure the SM topology parameter to the cluster, as described in *Configuring the Subscriber Manager* (on page 3-15).

Step 4 Configure the replication scheme for the data-store replication to the redundant machine.

Run the CLU: **p3db --set-rep-scheme**.

Step 5 Install the SM VCS agents, as described in *Installing SM Cluster Agents* (on page 3-18).

Step 6 Configure the VCS, as described in *Appendix E, Veritas Cluster Server Requirements and Configuration* (on page E-1).

Installing the Subscriber Manager

Note This installation is customizable.

To execute the *install-sm.sh* script:

From your workstation shell prompt, enter the following command:

 # install-sm.sh [*command options*]

The following table lists the command options.

Table 3-8 Options for install-sm.sh

Options	Description
-d	Specifies the install directory for ~*pcube*. This directory must *not* be an existing directory. This directory must be specified as a full pathname that begins with "/". Default: */opt/pcube*
-o	Specifies the existing home of user pcube as the install directory. Note that the options **–d** and **–o** are mutually exclusive.
-v	Specifies the directory for data storage. This directory must *not* be an existing directory. This directory must be on a partition with at least 1 GB of free space. This directory must be specified as a full pathname that begins with "/". Default: *InstallDirectory/var*
-h	Prints a help message and exits.

The script performs the following steps:

- Checks for validity of arguments and sufficient disk space.
- Adds (or verifies the existence of) a user **pcube** and a group **pcube**.
- Populates the **pcube** home directory with SM and CLU directory structure.
- Invokes the JRE installation script with **pcube** home as the target directory. The JRE installation does not affect any existing Java installations.
- Invokes the TimesTen installation script with **pcube** home as the target directory.

- Creates the SM DSN for TimesTen with **pcube** home as the target directory. It is possible to install the SM DSN for TimesTen in a specified directory by using the -v option.
- Creates startup and shutdown scripts in */etc*.
- Creates the shell preamble *~pcube/sm.sh*, which contains environment variables that depend on the actual folder in which the SM was installed.

EXAMPLE:

The following example installs the SM and CLU to a directory named */usr/local/pcube*, using the default data storage directory definition.

```
# install-sm.sh –d /usr/local/pcube
```

Installing SM Cluster Agents

The installation distribution file contains a set of customized Veritas Cluster Agents for supporting monitoring and controlling of SM-related resources in cluster topology. You must install the cluster agents under the VCS *bin* directory.

To install VCS agents:

From your workstation shell prompt, type:

 # install-vcs-agents.sh [*command-options*]

The following table lists the command options.

Table 3-9 Options for install-vcs-agents.sh

Options	Description
-d	Specifies the installation directory for the agents, which must be the *bin* directory of the VCS. This directory must be an existing directory. This directory must be specified as a full pathname that begins with '/'. Default: */opt/VRTSvcs/bin*
-h	Prints a help message and exits.

The script performs the following steps:

- Checks that the installation directory exists.
- Extracts the agent distribution file to the specified directory.
- Copies the VCS default-script-agent-executable from the installation directory to all agent directories.

Troubleshooting the Installation

For troubleshooting the installation, see *Appendix D, Troubleshooting* (on page D-1).

Installing an Application

An application can be installed on the SM in order to customize the components. You can also upgrade an existing application to a new version, or return to a previous version (rollback) of an application. Use the `p3inst` utility to install or uninstall an application (PQI file).

Note You must run the p3inst utility as user **pcube**. The script is located in the *~pcube/sm/server/bin* directory.

For additional details of how to install a specific application such as SCA BB, refer to the application installation guide.

To install or uninstall an application:

From your workstation shell prompt, enter the following command with the appropriate parameters:

> p3inst *operation filename* [*installation/upgrade parameters*]

The following table lists the **p3inst** operations.

Table 3-10 p3inst Operations

Option	Description
--install	Installs the specified application PQI file to the SM. It may be necessary to specify arguments for the installation procedure in the command line.
--uninstall	Uninstalls the specified application from the SM.
--upgrade	Upgrades the specified application. It may be necessary to specify arguments for the upgrade procedure in the command line.
--rollback	Returns the specified application to the previous version.
--describe	Displays the contents of the specified application file.
--show-last	Lists the last installed PQI file

EXAMPLE 1:

The following example shows how to install the specified installation file to the device.

```
> p3inst --install --file=eng216.pqi
```

EXAMPLE 2:

The following example shows how to uninstall the specified installation file from the device.

```
> p3inst --uninstall –-file=oldInstallation.pqi
```

System Changes Made by Installation Scripts

This section describes the system changes applied automatically by the SM installation. The SM installation adds a dedicated user and group, and startup and shutdown scripts.

Logging Script Messages

Script messages are logged into the system log in the following manner:

- For Solaris—The installation scripts log all their messages into the system log, which is usually the file located at */var/adm/messages*. The messages are logged to the *user.info* syslog category.
- For Linux—The installation scripts log all their messages into the system log, which is usually the file located at */var/log/messages*. The messages are logged to the *user.info* syslog category.

pcube User and Group

During installation, a user named **pcube** is created (unless it already exists) with its own group. This user owns all installed SM and CLU files. The user home directory is the installation directory selected during installation. For security purposes, the user is initially created with a locked password. You must assign a new password.

Startup and Shutdown Scripts

The SM is started on boot to run level 2, and is stopped when leaving this run level (for example, when the machine is shut down).

The installer installs the following files for startup and shutdown:

- For Solaris:

```
-rwxr--r--   1 root   other   /etc/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc0.d/K44p3sm -> /etc/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc1.d/K44p3sm -> /etc/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc2.d/S92p3sm -> /etc/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rcS.d/K44p3sm -> /etc/init.d/p3sm
```

- For Linux:

```
-rwxr--r--   1 root   other   /etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc.d/rc0.d/K44p3sm ->
/etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc.d/rc1.d/K44p3sm ->
/etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc.d/rc2.d/S92p3sm ->
/etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc.d/rc3.d/S92p3sm ->
/etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other   /etc/rc.d/rc5.d/S92p3sm ->
/etc/rc.d/init.d/p3sm
lrwxrwxrwx   1 root   other  /etc/rc.d/rc6.d/K44p3sm ->
/etc/rc.d/init.d/p3sm
```

The TimesTen installer creates similar startup and shutdown scripts.

Upgrading

Subscriber Manager version 3.0.3 supports several types of upgrade procedures, according to the SM version that was previously installed and the requirement (or lack of requirement) for fail-over in the new installation.

In some topologies, such as cluster, the upgrade procedure does not consist of running a script, because it is required to have minimum down time.

The following sections describe the various procedures needed for upgrading in different topologies and for moving from one topology to another.

Upgrade Procedure

To upgrade the Subscriber Manager:

Step 1 Extract the distribution files as described in *Step 1: Extracting the Distribution Files* (on page 3-8).

Step 2 Login as the *root* user.

Step 3 If upgrading from version 2.x and using a cluster setup, uninstall VCS agents as described in *Uninstalling VCS agents* (on page 3-28). The resource names to use are PcubeSm, OnOnlyProcess, and TimesTenRep.

Step 4 If using a cluster setup, stop the VCS monitoring of the SM, by running the following VCS CLU command from */opt/VRTSvcs/bin*:

```
#./hastop –local
```

Step 5 If upgrading from version 2.x, disable the state exchange between the SM and the SCE platform by editing the SM configuration file (*p3sm.cfg*) and set **save_subscriber_state=false**, then load the configuration file using the following command:

Note You must use this CLU as user **pcube**.

```
> p3sm --load-config
```

Step 6 Run the *upgrade-sm.sh* script from the distribution root directory.

Step 7 If upgrading from version 2.x, drop the old replication scheme and restart the SM.

Note You must use these CLUs as user **pcube**.

```
> p3db --drop-rep-scheme
> p3sm --restart
```

Step 8 Perform the specific upgrade instructions of your application or LEGs. For additional information, see *Installing an Application* (on page 3-19).

Step 9 If upgrading from version 2.x and using a cluster setup, configure the replication scheme for the data-store replication to the redundant machine.

Note You must use this CLU as user **pcube**.

```
> p3db --set-rep-scheme
```

Step 10 If using a cluster setup, install the SM VCS agents, as described in *Installing SM Cluster Agents* (on page 3-18).

Step 11 If upgrading from version 2.x and using a cluster setup, configure the VCS, as described in *Veritas Cluster Server Requirements and Configuration* (on page E-1)

Step 12 If using a cluster setup, restart the VCS monitoring by running the following VCS CLU command from */opt/VRTSvcs/bin*:

```
#./hastart
```

Step 13 (Optional) If upgrading from version 2.x, remove any obsolete subscriber state information, by running the SM CLU as **pcube** user:

```
> p3subsdb --clear-all-states
```

Step 14 (Optional) If upgrading from version 2.x, remove any obsolete subscriber properties, by running the SM CLU as **pcube** user:

```
> p3subsdb --export –o <csv-file>
> p3subsdb --clear-all
> p3subsdb --import –f <csv-file>
```

or

```
> p3subsdb --remove-property --property=prop
> p3sm --resync –n SCE_NAME
```

upgrade-sm.sh Script

Subscriber Manager version 3.0.3 provides an upgrade script that implements an upgrade from previous versions. The upgrade procedure script preserves the subscriber database and the entire SM configuration, including network elements, domains, and application-specific components.

Note For Solaris: Previous versions of the SM on Solaris used a 32-bit or 64-bit Java Virtual Machine (JVM) and database. By default, SM version 3.0.3 is installed with a 64-bit JVM and database. There is no choice as to whether to upgrade to 64-bit.

Note For Linux: Upgrades on Linux systems are only from SM 2.5.x and 3.x releases. The Linux platform is used only with a 32-bit JVM and database.

For additional information, see the following sections:

- *Upgrading from Non-cluster Setup to Non-cluster Setup* (on page 3-24)
- *Upgrading from Non-cluster Setup to Cluster Setup* (on page 3-24)
- *Upgrading from a Cluster Setup version 2.x* (on page 3-26)

To execute the *upgrade-sm.sh* script:

From your workstation shell prompt, enter the following command:

 # upgrade-sm.sh [*command-options*]

The following table lists the command options.

Table 3-11 Options for upgrade-sm.sh

Options	Description
-d	Destroys the database during the upgrade of TimesTen.
-h	Shows this message.

The script performs the following steps:

- Detects existing SM version.
- Detects new SM version.
- Verifies that Java is installed on the machine.
- Verifies that the user **pcube** exists.
- Verifies that an SM of version 2.2 and later is present on the system.
- Stops the current SM (if running).
- Backs up existing contents of the subscriber database to an external file.
- Removes the TimesTen database.
- Backs up SM configuration files.
- Installs the updated versions of SM and TimesTen.
- Invokes a separate program for upgrading the SM and database configuration files.
- Restores the backed up contents of the subscriber database (unless specified otherwise).
- Starts the upgraded SM.

Note To complete the upgrade process of the SM, you are required to follow the upgrade process instructions of your application and LEGs as described in the specific user guides. In general, you must run the `p3inst` CLU to upgrade or re-install your application or LEG PQI files.

EXAMPLE:

The following example upgrades the SM, keeps the current database, and does not pause the upgrade for PQI installation.

```
# upgrade-sm.sh
```

The following example upgrades the SM and removes the current database.

```
# upgrade-sm.sh –d
```

Note An SM reboot is not required after the upgrade procedure.

Data Duplication Procedure

The data duplication procedure enables the user to duplicate or copy the entire database from one machine to the other, and then keep the databases synchronized by running the replication agent at the end. Some of the upgrade procedures described in the following sections use this procedure.

For details of the procedure, see *Database Duplication Recovery* (on page 2-17).

Upgrading from Non-cluster Setup to Non-cluster Setup

To upgrade from a non-cluster setup to a non-cluster setup, follow the instructions described in *Upgrade Procedure* (on page 3-21).

Upgrading from Non-cluster Setup to Cluster Setup

This section describes the basic procedure for upgrading from a non-cluster setup to a cluster setup, from version 2.2 and up.

Note This procedure attempts to minimize the SM downtime as much as possible. Therefore, if subscriber service is not an issue, use instead the procedure for installing a new machine and upgrading a new machine.

In the following procedure, suppose that SM-A is the original SM machine running SM version 2.2 and later, and SM-B is the new SM machine being added for redundancy.

To upgrade from a non-cluster setup to a cluster setup:

Step 1 Install the VCS on both machines.

Step 2 Install SM-B:

a) Install SM on the new machine by using the install-sm.sh script.

b) Install the application and LEG PQIs on the SM using the CLU `p3inst` command.

Step 3 Upgrade SM-A:

Follow the instructions described in *Upgrade Procedure* (on page 3-21).

Note From this step until the upgrade procedure is completed, there is no SM to handle subscribers.

Step 4 Replicate the SM configuration from SM-A to SM-B.

Copy the *p3sm.cfg* configuration file manually from SM-A to SM-B. To load the configuration file, see *Reloading the SM Configuration (p3sm)* (on page 4-23).

Step 5 Duplicate the subscriber database, as described in *Data Duplication Procedure* (on page 3-24).

Configure the replication scheme for the data store replication to the redundant machine.

Note This CLU must run on both machines, and as user **pcube**.

```
> p3db --set-rep-scheme
```

Step 6 Create a cluster:

a) Configure SM-A and SM-B to support a cluster.

b) Make SM-B standby. Use the CLU command `p3cluster --standby`.

c) Ensure that SM-A is active. Use the CLU command `p3cluster --active`.

d) Configure the VCS.

e) Run the VCS on the setup.

Step 7 Configure the LEG applications to send logins to the cluster virtual IP.

Upgrading from a Cluster Setup version 3.x

This section describes the basic procedure for upgrading from a cluster setup to a cluster setup, from SM version 3.0 and up.

Note This procedure does not have a service down time.

In the following procedure, SM-A is the active SM machine and SM-B is the standby SM machine.

To upgrade from a cluster setup to a cluster setup without service down time:

Step 1 Upgrade SM-B by following the instructions described in *Upgrade Procedure* (on page 3-21).

Step 2 Manually trigger a failover using the Veritas cluster manager and wait until SM-B becomes active and SM-A becomes standby.

Step 3 Upgrade SM-A by following the instructions described in *Upgrade Procedure* (on page 3-21).

Upgrading from a Cluster Setup version 2.x

This section describes the basic procedure for upgrading from a cluster setup to a cluster setup, from SM versions 2.x.

Note This procedure has a service down time.

To upgrade from a cluster setup, follow the instructions described in *Upgrade Procedure* (on page 3-21). Each step should be applied on both machines.

Uninstalling

To uninstall the Subscriber Manager:

Step 1 If using a cluster setup, stop the VCS monitoring of the SM, by running the following VCS CLU command from */opt/VRTSvcs/bin*:

```
#./hastop -local
```

Step 2 Run the uninstall script from the SM distribution root directory:

```
-#./uninstall-sm.sh
```

(For additional information, see *uninstall-sm.sh Script* (on page 3-27)).

Step 3 (Optional) If using a cluster setup, remove the Veritas Cluster agents as described in *Uninstalling VCS agents* (on page 3-28). Remove the following resource names: OnOnlyProcess, SubscriberManager, and TimesTenRep.

Step 4 (Optional) Remove the **pcube** user, by running the following command:

```
# userdel -r pcube
```

Note If you chose to keep TimesTen installed, do *not* remove the **pcube** user.

uninstall-sm.sh Script

To execute the *uninstall-sm.sh* script:

From your workstation shell prompt, enter the following command:

 # uninstall-sm.sh [*command-options*]

The following table lists the command options:

Table 3-12 Options for uninstall-sm.sh Script

Options	Description
-n	Do not remove SM database.
-h	Shows the help message

The script performs the following steps:

- Stops the SM.
- Stops the replication agent (in cluster setups) if the –n flag is not used
- Destroys the data-stores if the –n flag is not used
- Uninstalls the TimesTen database.
- Removes the SM directories and boot files.
- Removes the Java that was installed as part of the SM installation.

Uninstalling VCS Agents

To uninstall the VCS agents:

Step 1 Remove the VCS agents by using the Veritas Cluster Manager, or by using the following CLU: (The resource names in your system might have different names, use **`hares –list`** to see the existing resource names).

```
hares –delete TimesTenDaemon
hares –delete SM
hares –delete ReplicationAgent
```

Step 2 Remove the VCS resource types by using the following CLU: (The type names in your system might have different names, use **`hatype –list`** to see the existing type names).

```
hatype –delete OnOnlyProcess
hatype –delete SubscriberManager
hatype –delete TimesTenRep
```

Step 3 Delete the VCS agent from the disk by entering the following command:

```
rm –rf /opt/VRTSvcs/bin/OnOnlyProcess
rm –rf /opt/VRTSvcs/bin/SubscriberManager
rm –rf /opt/VRTSvcs/bin/TimesTenRep
```

Note Repeat the procedure above for each additional Veritas Cluster agent that you wish to remove.

Configuration and Management

This chapter describes how to configure and manage the SM.

This chapter contains the following sections:

- SM Management and Configuration Methods 4-1
- Configuring a Subscriber Management Solution 4-3
- Using the CLU 4-20

SM Management and Configuration Methods

Configure and manage the Subscriber Manager using:

- *Configuration File* (on page 4-1)—Used for configuring the parameters that define the system behavior.
- *Command-Line Utilities* (on page B-1)—Used mainly for viewing and for subscriber management.

The configuration file and Command-Line Utilities (CLU) give you complete control over the SM; including subscriber management, database management, and SCE network configuration and management.

Configuration File

The SM uses a configuration file, `p3sm.cfg`, located in `~pcube/sm/server/root/config/p3sm.cfg`. For a detailed description of the configuration file parameters, see *Appendix A, Configuration File Options* (on page A-1).

The configuration file, together with the *Command-Line Utilities* (on page B-1), is used for configuring all the parameters that define the behavior of the SM application.

The configuration file contains the following types of parameters:

- General, system-wide parameters, such as subscriber state saving, persistency, subscriber introduction mode (Pull mode or Push mode), and topology
- Parameters for handling SM-LEG connection failure
- Parameters for handling SM-SCE connection failure
- Parameters for SCE platform configuration

- Parameters for domain configuration
  - Associating domains and SCE platforms
  - Specifying domain aliases
  - Specifying domain properties
- Auto-logout parameters, for controlling automatic logout of subscribers after timeout
- Parameters for RADIUS Listener configuration
  - Specifying NAS configuration
  - Specifying properties configuration
- Parameters for FTP, HTTP, and PRPC server configuration
- Parameters for Cable Adapter configuration
- Parameters for configuring SM operation with the TimesTen database

Usually, the parameters in the configuration file are specified once when setting up the system, and are valid throughout the system lifetime. To modify the configuration file parameters, edit the file using any text editor and reload it using the CLU (see *Reloading the SM Configuration (p3sm)* (on page 4-23)). The configuration file can be loaded on starting or restarting the SM and by explicitly running the CLU command.

The configuration file is designed so that the same configuration file can be used in multiple SM applications of a high availability setup. This enables the user to replicate the configuration by simply copying the file from one machine to another.

Command-Line Utilities

The SM provides a set of Command-Line Utilities (CLU), which you use, together with the configuration file, to configure the parameters that might change during the operation of the SM.

The CLU enables the user to configure the SM using shells installed on the machine. CLU commands are executable only when the user is logged in to the machine using the **pcube** user account, which is always installed (see *Chapter 3, Installation and Upgrading* (on page 3-1)). The CLU is used mainly for viewing and for subscriber management.

In high availability setups, you cannot use the CLU to perform subscriber management operations on the standby SM. Moreover, the standby SM refreshes the database before performing subscriber display operations, so the operation takes longer (than for the active SM). Therefore, it is recommended to perform all subscriber operations on the active SM.

This chapter explains how to perform various tasks using the appropriate CLU, but it does not describe the CLU in detail. For a complete, detailed description of the CLU, see *Appendix B, Command-Line Utilities* (on page B-1).

Configuring a Subscriber Management Solution

This section explains the procedure for configuring a Cisco Service Control deployment consisting of several SCE platforms and Subscriber Manager (SM) systems in order to make it ready for subscriber integration.

This section uses the terminology and tools explained in previous chapters and, when needed, terms and configuration tools explained in the SCE 1000 and SCE 2000 User Guides.

Prerequisites

Before configuring any of the components in your subscriber management solution, verify that all the items on the following checklist have been successfully completed:

- The SCE platforms in your network are installed and configured as explained in the SCE 1000 and SCE 2000 User Guides, Chapter 3.
- The SM applications in your network are installed as explained in *Chapter 3, Installation and Upgrading* (on page 3-1).
- The Cisco Service Control Application for Broadband (SCA BB) is installed on all SCE platforms and SM systems in your network. Please refer to the SCA BB User Guide for an explanation of how to install the Service Control Application on the SCE platforms and SM systems.
- The subscriber integration concept has been determined, and an appropriate solution was designed for driving subscriber mappings and policy information into the SM. This can be implemented automatically using a LEG, or manually using the CLU.
- The subscriber introduction mode (push or pull) has been determined for each SM system, based on the number of subscribers that the relevant SCE platforms should be serving.
- The association between SCE platforms and the relevant SM systems has been determined.
- For each SM system, the association between the SCE platforms that it serves and the subscriber domains has been designed.

Step-by-Step Configuration Procedure

This configuration procedure applies to a single group, consisting of the following:

- A Subscriber Manager application
- A set of LEG applications or components that connect to this SM
- The SCE platforms that this SM serves

Every subscriber management solution can be divided into such groups, and this procedure can be applied to each of these groups.

This procedure consists of the following steps:

- *Step 1: Editing the SM Configuration File* (on page 4-4)
- *Step 2: Importing Subscribers to the SM from a CSV File* (on page 4-14)
- *Step 3: Configuring the SCE Platforms* (on page 4-14)

Note Step 2 and Step 3 are not always required.

Step 1: Editing the SM Configuration File

Edit the SM configuration file (*p3sm.cfg*) according to your system definition, and reload it using the CLU `p3sm --load-config` command.

For additional details about the SM configuration file, see *Appendix A, Configuration File Options* (on page A-1).

The following topics describe the sections in the SM configuration file and their parameters.

Configuring SCE Platform Repository

Use the **p3net** CLU command to verify the connection state of each SCE Platform that should be provisioned by the SM.

To configure the SCE platform repository:

Step 1 Configure the SCE.XXX sections to add the SCE platform to the repository.

Step 2 Use `p3sm --load-config` to load the SCE configuration to the SM.

Step 3 Use `p3net --show` to verify that the SCE platform was successfully connected.

Step 4 When finished, use `p3net --show-all` to verify your configuration before continuing to the next procedure - Configuring Domains.

SCE.XXX

After the physical installation of an SCE platform (by being connected to the management network), it must be explicitly added to the SM list, or repository, of existing SCE platforms before the SM will recognize it (see Configuring SCE Platform Repository). Conversely, after the removal of an SCE platform from that list, the SM will no longer recognize it, even though it is still physically connected.

Each SCE.XXX section defines the following configuration parameters that represent a single SCE platform:

- *ip*

  Defines the IP address of the SCE platform.

- *port*

  Defines the port through which to connect to the SCE platform. The default is 14374.

To view the SCE platforms, use the CLU `p3net` (on page 4-30) command.

For additional information, see *SCE.XXX Section* (on page A-14).

Configuring Domains

Use the p3domains command to verify the domain configuration and that the SCE platforms are set to these domains.

To configure the domain:

Step 1 Configure the Domain.XXX sections to add domains to the domain repository and to add SCE platforms to each domain.

Step 2 Use `p3sm --load-config` to load the domain configuration to the SM.

Step 3 When finished, use `p3domains --show-all` and `p3net --show-all --detail` to verify your configuration before you start editing the configuration file.

Domain.XXX

When a system has more than one SCE platform, they can be configured into groups, or domains (see *Configuring Domains* (on page 4-5)). A subscriber domain is one or more SCE platforms that share a specified group of subscribers. Before adding an SCE platform to a domain, you must add the SCE platform to the SCE platform repository.

Note The SM is preconfigured with a single subscriber domain called *subscribers*.

Each Domain.XXX section specifies the elements (SCE platforms), aliases, and properties for one domain. It contains the following parameters:

- *elements=<logical_name1[,logical_name2,...]>*

  Specifies the names of the SCE platforms that are part of the domain.

- *aliases=alias_name1[,alias_name2,...]*

  Defines domain aliases. When subscriber information is received from the LEG with one of the aliases (for example, `alias1`), the information is distributed to the domain that matches this alias (for example, `domain_name1`). A typical alias could be a network device IP address, where, for example, each string in the values can be the IP address of a NAS or a CMTS.

Note Each alias (for example, alias_name1) can only appear in *one* [Domain.XXX] section.

The specification `aliases=*` means that every subscriber that does not have a domain will be put in this domain.

Note Only one domain at any given time may specify this option (`aliases=*`).

- *property.<name1>=<value1>[,property.<name2>=<value2>,...]*

  Defines the default policy property values for a domain. Unless the LEG/API overrides these defaults when it introduces the subscriber to the SM, the subscriber policy is set according to the default policy property values of its domain. Property values must be integers.

To view the domains, use the CLU `p3domains` (on page 4-30) command.

For additional information, see *Domain.XXX Section* (on page A-6).

SM General

This section of the configuration file is relevant to any deployment topology. It addresses the following system-wide parameters:

- *introduction_mode*

  Defines whether the SM introduces the subscribers to the SCE platforms immediately after a login operation (Push mode), or only when the SCE requests the subscriber specifically (Pull mode).

- *application_subscriber_lock*

  Defines whether to lock subscriber-related operations (login, logout, etc.) at the application level. Set this flag to `true` in the cases when several LEG components can update subscribers simultaneously.

- *force_subscriber_on_one_sce*

  Defines whether the SM supports the solution where a Cisco 7600/6500 is used for load-balancing among several SCE platforms. In this solution, when one SCE platform fails, subscriber traffic is redistributed to a different SCE platform. The SM must remove these subscribers from the failed SCE platform and send the relevant subscriber information to the new SCE platform. This parameter is relevant only in the Pull mode.

- *logon_logging_enabled*

  Defines whether to enable the logging of subscriber logon events.

To view the SM settings, use the `p3sm` (on page 4-23) CLU command.

Note Setting `logon_logging_enabled` to `true` will cause performance degradation. For additional information, see *SM General Section* (on page A-2).

Data Repository

The Data Repository section defines the SM operation with the TimesTen In-Memory Database, via the following parameters:

- *support_ip_ranges*

  Defines whether the SM supports IP-Ranges.

Note Disabling this support provides better performance.

- *checkpoint_interval_in_seconds*

  Defines the interval, in seconds, for calling the TimesTen checkpoints. Reducing the value affects performance, increasing the value increases vulnerability to power-down.

- *max_range_size*

  Determines the maximum IP range size used in the system. This parameter is used for improving performance of the SM in Pull mode when the `[Data Repository]` section is configured with `support_ip_ranges=yes`.

Note Defining this parameter with too large a value may cause performance degradation in handling pull requests.

For additional information, see *Data Repository Section* (on page A-17).

High Availability Setup

The High Availability section defines in what kind of topology the SM should work, via the parameter:

- *topology*

  Defines in what kind of topology the SM should work (cluster or standalone).

For additional information, see *SM High Availability Setup Section* (on page A-4).

Subscriber State Persistency

The Subscriber State Persistency section defines the storage of subscriber information by the SCE:

- *SCE_subscriber_persistency*

  Defines whether the SCE stores subscriber information in persistent memory. Use this option *only* when there are just a few subscribers, which rarely change, in the system.

Note Setting the *SCE_subscriber_persistency* parameter to `true` may cause performance degradation and reduce the rate of subscriber updates.

For additional information, see *Subscriber State Persistency Section* (on page A-4)

SM-LEG Failure Handling

The SM-LEG Failure handling section deals with the configuration of the behavior of the SM upon a SM-LEG communication link failure. To activate the SM-LEG failure action, you must also define LEG-Domain associations (see *LEG-Domains Association* (on page 4-9)).

The following are the parameters for the SM-LEG Failure Handling configuration:

- *clear_all_mappings*

  Defines the behavior of the SM in case of an SM-LEG connection failure. This parameter is relevant only for cases where the SM and LEG are running on different machines.

- *timeout*

  Defines the time in seconds from an SM-LEG connection failure until the mappings are cleared in the SM database.

For additional information, see *SM-LEG Failure Handling Section* (on page A-4).

LEG-Domains Association

The LEG-Domains association section relates to the SM-LEG Failure Handling section. It associates LEGs and domains for the SM-LEG failure action. It contains pairs having the following format:

- *<LEG name>=domain_name1[,domain_name2,...]*

  Defines the domains whose subscriber mappings will be cleared on SM-LEG connection failure. The key is the <LEG name>. A <LEG name> of **"*"** specifies all LEG applications. The (comma-separated) values are the domains (`domain_name`) to be cleared in the event of a network link failure (connection failure) between the specified LEG and the SM. A value of **"*"** in the `domain_name` specifies all subscriber domains in the system.

EXAMPLES:

- All subscribers of domain **subscribers** will be cleared in case of connection failure between the SM and the specified LEG:

```
[LEG-Domains Association]
MYSERVER12.CNR.LEG=subscribers
```

- All subscribers of domains **domain1** and **domain2** will be cleared in case of connection failure between the SM and LEG *MYSERVER13.CNR.LEG*:

  In addition, all subscribers of domain **domain3** will be cleared in case of connection failure between the SM and LEG *MYSERVER14.CNR.LEG*:

```
[LEG-Domains Association]
MYSERVER13.CNR.LEG=domain1,domain2
MYSERVER14.CNR.LEG=domain3
```

- All domains of the specified LEG will be cleared in case of communication failure between this LEG and the SM:

```
[LEG-Domains Association]
MYSERVER15.CNR.LEG=*
```

- The specified domains will be cleared in case of communication failure between any LEG and the SM:

```
[LEG-Domains Association]
*=my_domain1,my_domain2
```

- All domains will be cleared in case of communication failure between any LEG and the SM:

```
[LEG-Domains Association]
*=*
```

For additional information, see *LEG-Domains Association Section* (on page A-5).

Default Domains Configuration

The Default Domains Configuration section defines the default policy properties values for all domains via the parameter:

- *property.<name1>=<value1>[,property.<name2>=<value2>,...]*

  Defines the default policy property values for a domain. Unless the LEG/API overrides these defaults when it introduces the subscriber to the SM, the subscriber policy is set according to the default policy property values of its domain. Property values must be integers.

All domains are assigned all the properties defined in this section, unless overridden in the `[Domain.XXX]` section. This section is particularly useful when all domains in the system have the same set of properties.

Use the following configuration, for example:

```
[Default Domains Configuration]
property.A=1
property.B=2

[Domain.d1]

[Domain.d2]
property.A=2
property.C=1

[Domain.d3]
property.A=3
property.B=4
```

The properties are assigned to the domains as follows:

Domain	Value of Property A	Value of Property B	Value of Property C
d1	1 (assigned by default)	2 (assigned by default)	
d2	2	2 (assigned by default)	1
d3	3	4	

To view the domains, use the CLU `p3domains` (on page 4-30) command.

For additional information, see *Default Domains Configuration Section* (on page A-7).

Auto Logout

The Auto Logout section deals with the configuration of the SM to automatically log out a subscriber IP address after a specified period of time has passed since the last login operation. This time is called the *lease-time* and the LEG sets it as part of the login operation. Auto Logout contains the following parameters:

- *`auto_logout_interval`*

  Configures the timeout value, in seconds, between consecutive runs of the SM auto-logout mechanism. When the mechanism identifies that the subscriber IP address lease-time has expired (and the grace_period has also expired), the subscriber IP address is logged out automatically and the IP mapping is removed.

- *`grace_period`*

  Defines the grace period, in seconds, for each subscriber IP logout. After the subscriber lease-time has expired, the subscriber is automatically logged out after the grace period has also expired.

- *`max_rate`*

  Defines the maximum rate (logouts per second) that the auto-logout task is allowed to perform logouts from the system. This limit spreads out the load of the logout operations over time, reducing the performance impact on other operations. It is recommended to calculate the value for this parameter so that it spreads the logouts over at least half of the *`auto_logout_interval`* time.

These parameters are applicable in automatic integrations where the LEG cannot provide the SM with logout operations. For additional information, see *Auto Logout Section* (on page A-8).

RADIUS Listener

The RADIUS Listener section contains the general configuration parameters for the RADIUS Listener:

- *start*

  Defines whether the SM should run the RADIUS Listener at startup.

- *accounting_port*

  Defines the RADIUS Listener accounting port number.

- *ip*

  Defines the IP address to where the Radius Listener should bind (optional). This parameter should be used in cluster setups or when the machine's local-host IP address is not the IP address to where the RADIUS messages are sent.

For additional information, see *Radius Listener Section* (on page A-9).

Radius.NAS.XXX

Each Radius.NAS.XXX section specifies the configuration parameters for a single Network Access System (NAS):

- *domain*

  Specifies the domain.

- *IP_address*

  Specifies the IP address in dotted notation (xxx.xxx.xxx.xxx).

- *NAS_identifier*

  Specifies the name of the NAS that exists in the *NAS-ID* attribute.

- *secret*

  Specifies a secret key defined in the NAS for this connection.

For additional information, see *Radius.NAS.XXX Section* (on page A-10).

Radius.Subscriber ID

The Radius.Subscriber ID section defines the RADIUS attribute on which to base the subscriber ID association, via the following parameters:

- *radius_attribute*

  Specifies the RADIUS protocol attribute number. Use the value of `26` for Vendor Specific Attributes (VSAs).

- *radius_attribute_type*

  Specifies the RADIUS attribute type.

- *radius_attribute_vendor_id*

  Used only if the association is based on Vendor Specific Attributes (VSAs).

- *radius_sub_attribute*

  Used only if the association is based on Vendor Specific Attributes (VSAs).

For additional information, see *Radius.Subscriber ID Section* (on page A-12).

Radius.Property.Package

The Radius.Property.Package section defines the RADIUS attribute from which to retrieve a subscriber package. It contains the following configuration parameters:

- *radius_attribute*

  Specifies the RADIUS protocol attribute number. Use the value of `26` for Vendor Specific Attributes (VSAs).

- *radius_attribute_type*

  Specifies the RADIUS attribute type.

- *use_default*

  Defines whether to use a default value if the attribute was not found (when not mandatory).

- *default*

  Defines the default value to use if the attribute was not found (when not mandatory).

- *radius_attribute_vendor_id*

  Used only if the association is based on Vendor Specific Attributes (VSAs).

- *radius_sub_attribute*

  Used only if the association is based on Vendor Specific Attributes (VSAs).

For additional information, see *Radius.Property.Package Section* (on page A-11).

RPC.Server

Each RPC.Server section defines the following configuration parameters of the PRPC server:

- *port*

  Defines the PRPC server port. The default is 14374.

- *security_level*

  Defines whether the PRPC server forces authentication on all connections, authenticates connections that support authentication while still accepting connections that do not, or does not enforce authentication.

  Possible values for this parameter are **`full`**, **`semi`**, and **`none`**. The default value is **`semi`**.

To view the PRPC settings, use the CLU `p3rpc` (on page 4-33) command.

For additional information, see *RPC.Server Section* (on page A-13).

FTP

The FTP section contains the FTP server configuration parameters:

- *start*

  Defines whether the SM should run the FTP server at startup.

- *port*

  Defines the FTP server port. The default is 21000.

For additional information, see *FTP Section* (on page A-14).

HTTP Tech-IF

The HTTP Tech-IF section contains the HTTP server configuration parameters:

- *start*

  Defines whether the SM should run the HTTP server at startup.

- *port*

  Defines the HTTP server port. The default is 8082.

For additional information, see *HTTP Tech-IF Section* (on page A-15).

RDR Server

The RDR Server section contains the RDR server configuration parameters:

- *start*

  Defines whether the SM should run the RDR server at startup.

- *port*

  Defines the RDR server port. The default is 33001.

- *max_connections*

  Defines the maximum number of connections accepted by the server. The default is 10.

For additional information, see *RDR Server Section* (on page A-16).

Cable Adapter

The Cable Adapter section defines whether to allow logins from unknown CMs, via the following parameter:

- *allow_dynamic_CM*

  Defines whether to allow logins from unknown CMs.

To view the cable settings, use the CLU `p3cable` (on page 4-31) command.

For additional information, see *Cable Adapter Section* (on page A-16).

Step 2: Importing Subscribers to the SM from a CSV File

This step is optional and should be performed only when using manual integration, or when performing a setup prior to the beginning of the automatic integration.

A csv file is a simple text file where each line consists of comma-separated values. Because each line may contain subscriber properties, which are application dependant, the documentation of the application that you installed on your system describes the format of a csv import file.

In most cases, when importing csv files, you should use the CLU **`p3subsdb --import`** command. When integrating with a cable AAA system and working in the CPE as Subscriber mode (see *Appendix C, CPE as Subscriber in Cable Environment* (on page C-1)), importing cable modems requires the CLU **`p3cable --import-cm`** command.

Step 3: Configuring the SCE Platforms

This step is optional and should be performed only when using the Pull mode to introduce subscribers, or when performing a special operation on SM-SCE connection failure. Use the SCE platform Command-Line Interface (CLI) to configure several configuration parameters, as discussed below.

Configuring these parameters ensures that the SCE platform correctly applies appropriate defaults to subscribers in the period between subscriber detection and pull response (for unmapped subscribers). For additional details, see the SCE 1000 or SCE 2000 User Guides, Chapter 8.

Anonymous Groups and Subscriber Templates

When the SCE platform encounters network traffic that is not classified to any introduced subscriber, it checks whether the mapping of the unfamiliar subscriber belongs to one of the anonymous groups. If the subscriber belongs to an anonymous group, a new anonymous subscriber is created, and a request for an updated subscriber record is sent to the SM. The properties of the anonymous subscriber are taken from the subscriber template that is assigned to the newly created subscriber anonymous group.

Anonymous Subscriber Groups

An anonymous group is a specified IP range, possibly assigned a subscriber template (defined in the next section). If a subscriber template has been assigned to the group, the anonymous subscribers generated have subscriber properties as defined by that template. If no subscriber template has been assigned, the default template is used.

Use the appropriate CLI commands to import anonymous group information from a csv file, or to create or edit these groups explicitly.

Subscriber Templates

Values for various subscriber properties for unmapped or anonymous subscriber groups are assigned in the system based on subscriber templates. A number from 0 to 199 identifies subscriber templates. CSV formatted subscriber template files define the subscriber templates 1 to 199. However, template 0 cannot change; it always contains the default values. If a template is not explicitly assigned to an anonymous group, the group uses template 0.

Use the appropriate CLI commands to import subscriber templates from a csv file, or edit these templates from the command line. Additionally, use the appropriate CLI commands to assign subscriber templates to the anonymous groups.

Subscriber Aging Parameters

To prevent SCE capacity problems in Pull mode, configure the aging of introduced subscribers. The aging parameter defines a timeout, and any subscriber that does not generate traffic during this timeout interval will be automatically logged out from the SCE.

SM-SCE Connection Failure

To prevent incorrect classification of a subscriber's traffic during a lengthy connection failure between the SM and the SCE, configure the SM connection failure parameters.

The SCE has several alternatives for connection failures handling:

- The SCE can clear the mappings of all of the subscribers
- The SCE can put the line in cut-off mode
- The SCE does nothing

The timeout between the connection detection and actually performing the operation is also configurable.

System Configuration Examples

This section presents and explains common subscriber management scenarios, including the correct configuration parameters for these scenarios. The following scenarios are described:

- Automatic introduction of subscribers, with Push mode and fail-over of SCE platforms
- Manual introduction of subscribers with Pull mode
- SM fail-over scenario

Example 1: Automatic Introduction of Subscribers, with Push Mode and Fail-Over of SCE Platforms

This scenario assumes the following:

- Automatic introduction of subscribers, that is, a provisioning system of an AAA system introduces the subscribers. This example assumes that integration with a DHCP server allows automatic introduction of subscribers to the SM.
- The SM is operating in Push mode.
- The application that is used includes states that should be preserved such as volume quotas states in the Service Control Application for Broadband (SCA BB).

Figure 4-1: Cable Topology with Automatic Integration with a DHCP Server, Push Mode, and Fail-Over of SCE Platforms

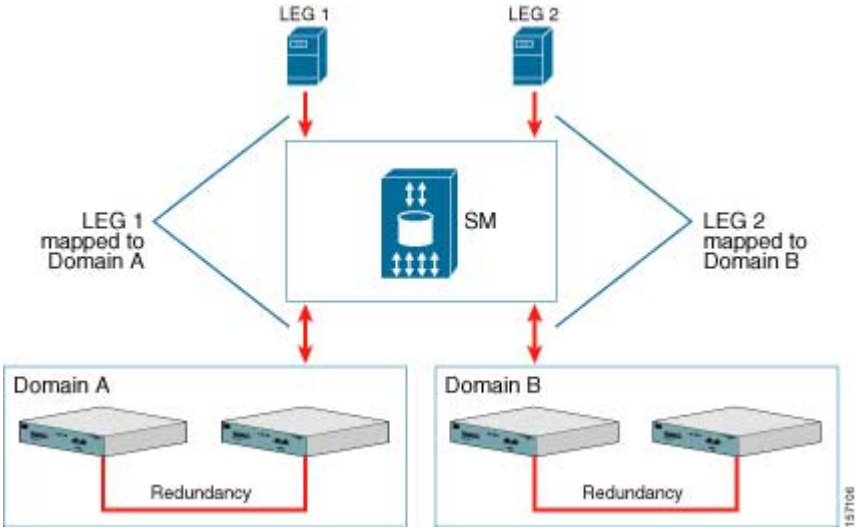

Note Ensure that everything is properly installed before proceeding with configuring the SM.

To configure the SM using this scenario:

Step 1 Edit the SM configuration file (*p3sm.cfg*) to add the SCE devices to the SCE device repository and group the SCE devices to domains, as depicted in the above figure.

Step 2 Edit the SM configuration file, as displayed in the table below.

Step 3 Reload the SM configuration file using the p3sm CLU.

Step 4 Import the cable modems to the SM database using the p3cable CLU.

This scenario does not need an SCE platform configuration.

Table 4-1 Configuration File Parameters for Automatic Integration with Push Mode in a Cable Environment

Section and Parameter	Value and Description
SM General	
introduction_mode	push
High Availability Setup	
topology	standalone The value should be set to *standalone* because the described scenario has just one SM.

Section and Parameter	Value and Description
Subscriber State Persistency	
save_subscriber_state	true The value must be set to *true* if working with an application that requires preserving the state. In any case, however, it is advisable to set the value to *true*.
SCE_subscriber_persistency	false The value should be set to *false* due to limited space on the SCE platform that probably will not suffice in an automatic integration scenario. Moreover, setting this value to *true* causes performance degradation.
SM-LEG Failure Handling	
timeout	300 seconds (more tolerance may be advisable for SM-LEG failures in actual configurations)
clear_all_mappings	true The value should be set to *true* because under the scenario conditions (automatic integration in cable environment), subscribers are likely to change their state or to logout from the SCE during an SM-LEG connection failure. You would therefore like to clean their mappings when the SM and LEG are connected again.
LEG-Domain Association	
<LEG name>	Define associations between LEGs and domains LEG-Domain associations should be defined in order that all subscriber mappings will be cleared on SM-LEG disconnection. If no association is defined, subscriber mappings will not be cleared (the `clear_all_mappings` value will be ignored).
Default Domains Configuration	
property	None Configure the policy here if all domains in the system have the same policy.
Domain.XXX	
aliases	Define aliases, to facilitate working in cable environment Define aliases if working with LEGs that are not aware of system domains. Alternatively, you can define domains with names that match values produced by LEGs.
property	None Configure the policy here only if you want it to be applied to all subscribers in the domain.
allow_dynamic_CM	no The value should be set to *no* in order to prevent uninstalled CMs from using the network.

Section and Parameter	Value and Description
Auto Logout	
auto_logout_interval	60 minutes Auto-logout should be activated assuming that the AAA system cannot provide logout events, which is true in cable environments. The value defined here should be smaller than the CPE lease time in this cable environment.
grace_period	10 seconds You should define a relatively high value to eliminate mistakes because of management network delays.
max_rate	100 logouts per second You should define a value similar to the max-login rate to the SM.

Example 2: Manual Introduction of Subscribers with Pull Mode

This scenario assumes the following:

- Manual introduction of subscribers
- Pull mode
- Application that does not require preserving state

Note Ensure everything is properly installed before proceeding with configuring the SM.

Figure 4-2: Topology with Manual Introduction of Subscribers and Pull Mode

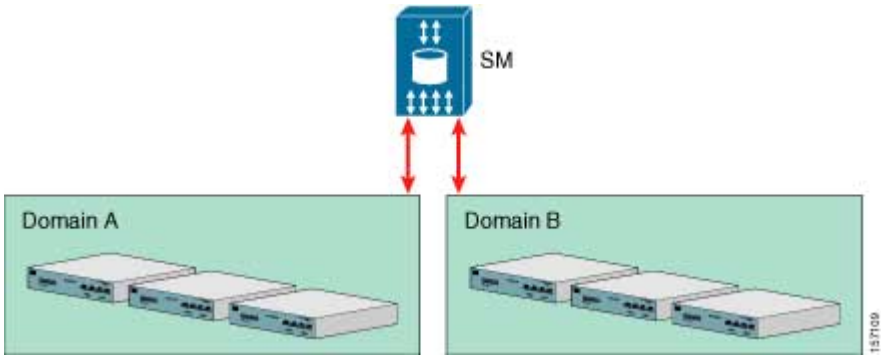

To configure the SM using this scenario:

Step 1 Edit the SM configuration file (*p3sm.cfg*) to add the SCE devices to the SCE device repository and group them to domains, as depicted in the above figure.

Step 2 Further edit the SM configuration file, as shown in the table below.

Step 3 Reload the SM configuration file using the p3sm CLU.

Step 4 Import the subscribers using the p3subsdb CLU (required for manual integration; there is no other way to bring subscribers into the SM).

Step 5 Use the SCE platform CLI to configure the system for Pull mode:

- Subscriber templates—In accordance with application
- Anonymous groups—In accordance with your network and subscribers
- Introduced subscriber aging—In accordance with your network and IP address allocation scheme

---

Table 4-2 Configuration File Parameters for Manual Integration with Pull Mode

Section and Parameter	Value
SM General	
introduction_mode	pull
High Availability Setup	
Topology	standalone The value should be set to *standalone* because this scenario involves just one SM.
Subscriber State Persistency	
save_subscriber_state	false
SCE_subscriber_persistency	false Defines whether the SCE stores subscriber information in persistent memory. Use this option only when there are just a few subscribers, which rarely change, in the system.
SM-LEG Failure Handling	
Timeout	Not applicable—No LEGs are involved in the scenario (use default—60 seconds)
clear_all_mappings	Not applicable—No LEGs are involved in the scenario (use default—no action)
LEG-Domain Association	
<LEG name>	Not applicable—No LEGs are involved in the scenario (Use default—no mappings)
Default Domains Configuration	
Default policy (property.XXX)	None Configure the policy here only if all domains in the system have the same policy.
Domain.XXX	
Aliases	Not applicable (use default—none)

Section and Parameter	Value
property.XXX	None Configure the policy here only if you want to apply it to all subscribers in the domain.
allow_dynamic_CM	Not applicable—valid for cable environment only (use default – no)
Auto Logout	
auto_logout_interval	Not applicable (use default—0)
grace_period	Not applicable (use default—10)
max_rate	Not applicable (use default—50)

Example 3: SM Fail-over Configuration (General)

When using a set-up with SM fail-over (based on two SM nodes connected in a cluster), the configuration is identical to the regular configuration, with one exception:

- The topology parameter in the *SM High Availability Setup Section* (on page A-4) should be set to *cluster* for both SM machines.

Other than that, SM fail-over configuration is performed normally via the `p3sm.cfg` configuration file. Manually copy the configuration file from the active SM to the standby SM.

For additional information of how to configure the VCS, see *Appendix E, Veritas Cluster Server Requirements and Configuration* (on page E-1).

Using the CLU

This section introduces the Command-Line Utilities (CLU), and describes how to use the CLU for viewing, subscriber management, and other tasks when working with the SM.

Note Some of the CLU operations and options can be specified by abbreviations. For this and additional information about the CLU, see *Appendix B, Command Line Utilities* (on page B-1).

This section contains the following subsections:

- *Informative Output* (on page 4-21)
- *Parsing CLU Operations and Options* (on page 4-21)
- *Reloading the SM Configuration (p3sm)* (on page 4-23)
- *Managing the SM (p3sm)* (on page 4-23)
- *Managing Subscribers, Mappings, and Properties (p3subs)* (on page 4-25)
- *Managing the Subscriber Database (p3subsdb)* (on page 4-28)
- *Viewing and Connecting Network Elements (p3net)* (on page 4-30)
- *Viewing Subscriber Domains (p3domains)* (on page 4-30)

- *Managing the Cable Support Module (p3cable)* (on page 4-31)
- *Installing an Application (p3inst)* (on page 4-32)
- *Viewing Information of the PRPC Interface Server (p3rpc)* (on page 4-33)
- *Managing a Cluster of Two SM Nodes (p3cluster)* (on page 4-34)
- *Managing the User Log (p3log)* (on page 4-34)
- *Viewing Statistics of the RADIUS Listener (p3radius)* (on page 4-35)
- *Utilities* (on page 4-35)

The procedures explained below invoke the following CLU commands:

- **p3batch**
- **p3cable**
- **p3clu**
- **p3cluster**
- **p3db**
- **p3domains**
- **p3inst**
- **p3log**
- **p3net**
- **p3radius**
- **p3rpc**
- **p3sm**
- **p3subs**
- **p3subsdb**

Informative Output

All CLU commands support the following operations for informative output:

Operation	Description
--help	Prints the help for the specified CLU command, then exits.
--version	Prints the SM program version number, then exits.

Parsing CLU Operations and Options

Place in quotation marks a command operation or option containing any of the following characters:

- A space character
- A separation sign (comma “,”; ampersand “&”; colon “:”)
- An escape character (backslash “\”)

A command operation or option that contains any of the following characters must have that character preceded by an escape character:

- An equal sign (=)
- A quotation mark (“ or ”)
- An escape character (backslash “\”)

Following are several examples of the above rules:

Operation/option contains the character	**Example of how operation/option should be written**
Space character	--property=“file name”
Escape character (backslash “\”)	--property=“good\\bad”
Equal sign (=)	--property=“x\=y”
Quotation marks (“ or ”)	--name=“\“myQuotedName\”” (in the above example, inner quotation marks are escaped)
Separation characters	
• comma (,)	--names=“x,y”
• ampersand (&)	--names=“x&y”
• colon (:)	--names=“myHost:myDomain”

One-letter abbreviations are available for some of the operations and options. For example, “-d” is an abbreviation for “--domain”. Note that only one hyphen (-), not two, precedes the letter for an abbreviation, and that if the operation or option takes a parameter, there is a space and not an equal sign before the parameter.

- Example of using full name --domain=subscribers
- Example of using abbreviated name -d subscribers

The abbreviations are useful if you want to specify an expression to be expanded by the UNIX shell, for example:

- p3subsdb --import –f ~pcube/file.csv ~pcube will be expanded by the UNIX shell

Reloading the SM Configuration (p3sm)

Use the p3sm utility to configure the SM by reloading the SM configuration file *p3sm.cfg*. Use any standard text editor to edit the configuration file.

To reload the SM configuration:

From the shell prompt, type the following command:

p3sm --load-config [--ignore-warnings] [--remote=*OTHER_SM_IP*[**:***port*]]

The configuration file is loaded, and the SM configuration updated accordingly.

The `--remote` option loads the configuration first to the local SM, and afterward to the remote SM (in High Availability setups).

Managing the SM (p3sm)

Use the p3sm utility to manage the SM on an ongoing basis. The p3sm utility enables you to start, stop, and resynchronize the SM.

To manage the SM:

From the shell prompt, enter a command using the following general format:

p3sm *operation* [**--ne-name=***SCE NAME*]

The following table lists the p3sm operations relevant to managing the SM.

Table 4-3 p3sm Management Operations

Operation	Description
--show	Displays the current SM configuration.
--resync	Resynchronizes the subscribers of the specified SCE with the SM database. The SCE is specified using the format `--ne-name=SCE_NAME`.
--resync-all	Resynchronizes the subscribers of all SCEs with the SM database.
--start [--wait]	Starts the server. The option `--wait` causes the CLU to return only after the SM is up.

Operation	Description
--stop	Stops the server. **Note:** When using fail-over, a simple shut-down of the SM does not work; the Veritas Cluster Server identifies that the SM is down and attempts to restart it. The correct procedure is: 1. Perform the manual fail-over. See *Subscriber Manager (SM) Fail-Over* (on page 2-11). 2. Use the Veritas Cluster Manager Application to stop the monitoring (probing) of the SM. 3. Use the SM CLU (`p3sm --stop`) to stop the SM.
--restart [--wait]	Stops the server operation and then restarts it. The option `--wait` causes the CLU to return only after the SM is up.
--sm-version	Displays the version of the currently installed SM.
--sm-status [--detail]	Displays the SM operational status: whether the SM is running or not and whether it is active or standby. If errors have occurred, a summary is displayed. To receive a detailed description, use the option `--detail`.
--reset-sm-status	Clears the errors displayed, enabling the user to see if new errors have occurred.
--extract-support-file	Retrieves the support file from the agent. This command extracts the SM support information to a defined file. The SM support information should be extracted and sent to Cisco Customer Support with each support request.
--output	Defines the name of the file to which the SM support information is extracted.
--show-stats	Displays statistics information regarding logon operations. The rate results are updated once every 10 seconds.
--reset-stats	Resets the statistics information.
--logging=[on/off]	Enables or disables the logging of user logon to the user log. **Note:** Enabling this flag may affect performance.

EXAMPLE 1:

The following example shows how to stop the server operation and then restart it.

```
p3sm --restart
```

EXAMPLE 2:

The following example shows how to resynchronize an SCE whose logical name is *SCE_1000A*.

```
p3sm --resync --ne-name=SCE_1000A
```

EXAMPLE 3:

The following example shows how to extract the SM support information to a file named *support.zip*.

```
p3sm --extract-support-file --output=support.zip
```

Managing Subscribers, Mappings, and Properties (p3subs)

Use the p3subs utility to manage subscribers. You can add or remove subscribers. You can also manage subscriber properties and mappings with this utility.

To manage subscribers:

From the shell prompt, enter a command using the following general format:

p3subs *operation* **--subscriber=***Subscriber-Name* **[--ip=***IP-address***] [--vlan-id=***VLAN***] [--mpls-vpn=VPN-ID@PE-IP] [--property=***property-name=value***] [--domain=***domain-name***] [--overwrite]**

The subscriber on whom the operation is to be performed is specified by using the format `--subscriber=`*subscriber-name*. A mapping (IP address, VLAN, or MPLS/VPN specification), property, or domain, if specified, uses the format displayed.

Note If a domain is not specified, the subscriber is added to the default *subscribers* domain.

This section describes the following:

- *Managing Subscribers* (on page 4-25)
- *Managing Mappings* (on page 4-26)
- *Mappings Specification* (on page 4-26)
- *Managing Properties* (on page 4-27)
- *Clearing the Subscriber Applicative State* (on page 4-27)

Managing Subscribers

The following table lists the p3subs operations relevant to managing subscribers.

Table 4-4 p3subs Subscriber Operations

Operation	Description
--add	Adds a subscriber or replaces the existing subscriber configuration.
--set	Updates mappings and/or properties for the specified subscriber. If the subscriber does not exist, this operation acts like the **add** operation.
--remove	Removes the specified subscriber.
--show	Displays information for the specified subscriber.

EXAMPLE 1:

The following example shows how to add a subscriber with the specified IP address.

```
p3subs --add --subscriber=jerry --ip=96.142.12.7
```

EXAMPLE 2:

The following example shows how to overwrite subscriber information. Because the subscriber named “jerry” already exists, the add operation would fail; however, the overwrite option allows the IP address to be overwritten.

```
p3subs --add --subscriber=jerry --ip=96.128.128.42 --overwrite
```

Managing Mappings

The following table lists all the p3subs operations relevant to managing mappings.

Table 4-5 p3subs Mapping Operations

Operation	Description
--set	Adds or updates a subscriber. If the mappings exist, it replaces the existing mappings, unless the **additive-mapping** option is used.
--remove-all-mappings	Removes all the mappings of the specified subscriber.
--remove-mappings	Removes specified mappings of the specified subscriber.
--show-all-mappings	Displays all the mappings for the specified subscriber.

Table 4-6 p3subs Mapping Options

Option	Description
--additive-mapping	Adds the specified mappings to the existing ones (instead of replacing the existing mappings when this option is not used). Used with the **set** operation.

EXAMPLE:

The following example shows how to removes all the mappings for a specified subscriber.

```
p3subs --remove-all-mappings --subscriber=jerry
```

Mappings Specification

You can specify the following mapping types for each subscriber:

- IP address or range—Use the `--ip-address` option. For an IP address, use the dotted notation. A range is used to specify several consecutive mappings, for example, the notation 1.1.1.0/30 is used to specify the IP addresses 1.1.1.0 to 1.1.1.3. You can specify multiple mappings by using a comma.
- VLAN—Use the `--vlan-id` option. An integer number specifies the VLAN. You can specify multiple mappings by using a comma.
- MPLS/VPN—Use the `--mpls-vpn` option. The notation of the MPLS/VPN mappings is `VPN-ID@PE-IP`. VPN-ID is the VPN identifier of the VPN site (it can be RT or RD). The PE-IP is the loopback IP address of the PE router attached to the VPN. For more information on the MPLS/VPN mappings, see the *Cisco SCMS SM MPLS/VPN BGP LEG Reference Guide*.

Note You cannot specify different types of mappings for the same subscriber.

Note When working with VLAN mapping types, the SCE must be configured using the following CLI:

```
SCE2000#> configure
SCE2000(config)#> in li 0
SCE2000(config if)#> VLAN symmetric classify
```

Managing Properties

The following table lists the p3subs operations relevant to managing subscriber properties. The property names depend on the application running on your system. To find descriptions of the application property names and values, see the documentation provided with the application installed on your system.

Table 4-7 p3subs Subscriber Property Operations

Operation	Description
--set	When specifying a property, the subscriber record is updated with the specified property and value.
--reset-property	Resets specified property of specified subscriber to its default value.
--show-property	Displays value of specified property of specified subscriber.
--show-all-properties	Displays all the properties of the specified subscriber.
--show-all-property-names	Displays all the property names and descriptions.
--remove-properties	Removes specified properties or custom properties from the specified subscriber record.

EXAMPLE:

The following example shows how to set a property value for a specified subscriber.

```
p3subs --set --property=packageId=1 --subscriber=jerry
```

Clearing the Subscriber Applicative State

The following table presents the **p3subs** operation that clears a subscriber's applicative state.

Table 4-8 p3subs Mapping Operations

Operation	Description
--clear-state	Clears the subscriber's applicative state. This command clears only the backup copy at the SM. It does not clear the applicative state record in the SCE platform.

EXAMPLE:

The following example shows how to clear the applicative state of the specified subscriber. Note that this command clears only the backup copy at the SM. It does *not* clear the applicative state record in the SCE platform.

```
p3subs –-clear-state --subscriber=jerry
```

Managing the Subscriber Database (p3subsdb)

Use the **p3subsdb** utility to manage the SM database. You can import subscriber information for a group of subscribers from a CSV file into the SM database. You can also export subscriber information from the SM database to a CSV file.

Note The format of the CSV file depends on the application. The documentation of a specific application specifies the CSV file format for that application.

To manage the SM database:

From the shell prompt, enter a command using the following general format:

`p3subsdb` *operation* **`[--domain=`***domain-name***`]`** **`[`***filename***`]`**

The following table lists the **p3subsdb** operations.

Table 4-9 p3subsdb Operations

Operation	Description
--clear-all	Removes all subscriber records from the database.
--clear-domain	Removes all subscriber records from the specified domain.
--show-num	Displays the number of subscribers in the database for the specified domain.
--show-all	Lists all subscriber names.
--show-domain	Lists all subscriber names in the specified domain.
--import	Imports subscribers to the database from a specified CSV file. The filename that is to be imported is specified using the format `--file=import-filename`. The results go to a results file, *`import-results.txt`*, which is created in the same directory as the CSV file.
--export	Exports subscribers from the database to a specified CSV file. The filename that is to be exported is specified using the format `--output=export-filename`. The results go to a results file, *`export-results.txt`*, which is created in the same directory as the CSV file.

Operation	Description
--remove	Removes subscribers from the database using a specified CSV file. The filename is specified using the format `--file=remove-filename`. The results go to a results file, `remove-results.txt`, which is created in the same directory as the CSV file.
--clear-all-states	Clears the state of all subscribers in the SM database.
--remove-property	Removes a specified property from all subscribers in the system. Note: After running this command, you should re-synchronize all SCE devices.
--remove-all-ip	Removes all the IP addresses of all subscribers.
--remove-all-vlan	Removes all the VLAN tags of all subscribers.
--remove-all-mpls-vpn	Removes all the MPLS/VPN mappings of all subscribers.

Table 4-10 p3subsdb Options

Operation	Description
--domain=DOMAIN	Performs the operation on the specified DOMAIN. If DOMAIN is **none**, then the operation refers to the subscribers who have no domain specified.
--prefix=NAME	Used in the export command for filtered export.
--property=PROPERTY	Used in conjunction with the `--remove-property` operation. You must specify this option for the `--remove-property` operation, otherwise the operation fails.

Table 4-11 p3subsdb File Options

Operation	Description
--file=FILE	One or more subscriber CSV FILEs to import from.
--output=FILE	Subscriber CSV FILE to export to.

EXAMPLE 1:

The following example shows how to list all subscribers in a specified domain.

```
p3subsdb --show-domain --domain=mainDomain
```

EXAMPLE 2:

The following example shows how to import subscribers from the specified CSV file.

```
p3subsdb --import --file=goldSubscriberFile.csv
```

EXAMPLE 3:

The following example shows how to export subscribers with filtering options to a specified CSV file.

```
p3subsdb --export --prefix=a --–output=silverSubscriberFile.csv
```

Viewing and Connecting Network Elements (p3net)

Use the **p3net** utility for viewing the connection status of network elements and trying to reconnect disconnected elements.

To view the connection status of a network element or to reconnect a disconnected element:

From the shell prompt, enter a command using the following general format:

p3net *operation* **--ne-name=***logical-name*

The following tables list **p3net** operations and options.

Table 4-12 p3net Operations

Operation	Description
--show-all	Shows all the configured network elements.
--show	Shows an element's connection status.
--connect	Tries to connect a disconnected element

Table 4-13 p3net Options

Network Element Option	Description
--ne-name=NAME	Specifies the logical NAME for the network element.
--detail	(Optional) Used with the **--show-all** operation for displaying additional information as a table.

EXAMPLE:

The following example shows how to display a network element's connection status.

```
p3net --show --ne-name=mainNE
```

Viewing Subscriber Domains (p3domains)

Use the **p3domains** utility for viewing the subscriber domains. As explained in *Introducing the Subscriber Manager*, subscriber domains are groups of SCE devices that serve the same subscribers.

To view subscriber domains:

From the shell prompt, enter a command using the following general format:

p3domains *operation* **--domain=***domain-name*

The following tables list **p3domains** operations and options.

Table 4-14 p3domains Operations

Operation	Description
--show-all	Shows all configured domains.
--show	Shows a domain and its associated network elements.

Table 4-15 p3domain Options

Domain/NE Option	Description	Notes
--domain=DOMAIN	DOMAIN specifies logical name	"**none**" cannot be used, it is a reserved word

Managing the Cable Support Module (p3cable)

In the cable environment, the SM supports two modes of operation: *CM as Subscriber* and *CPE as Subscriber*, as described in *CPE as Subscriber in Cable Environment* (on page C-1).

This section only discusses the support of the *CPE as Subscriber* mode. In this mode, the CPE is modeled as the subscriber, and it inherits its policy and domain from the cable modem (CM) through which it connects to the network. Each cable modem is linked with one or more CPEs. (For background information about special characteristics of the *CPE as Subscriber* mode in the cable environment, see *CPE as Subscriber in Cable Environment* (on page C-1).)

Use the **p3cable** utility commands to import cable modem information from a CSV file to the SM, and to export the cable modem information from the SM to a CSV file. You can also use this utility to clear the repository of all cable modems, and to show whether to allow or deny the login of CPEs that belong to unfamiliar cable modems; i.e., cable modems that do not exist in the SM database. However, for specifying whether to allow or deny such a login, use the *Cable Adapter* (on page 4-13) section of the `p3sm.cfg` configuration file.

To manage the cable modems:

From the shell prompt, enter a command using the following general format:

`p3cable` *operation* **`--cm=`***CM-name* [*filename*] [*other CM options*]

The following table lists the **p3cable** operations.

Table 4-16 p3cable Operations

Operation	Description
--import-cm	Imports cable modems from a CSV file. The filename that is to be imported is specified using the format `--file=import-filename`. The results go to a results file.

Operation	Description
--export-cm	Exports cable modems to a CSV file. The filename that is to be exported is specified using the format `--output=export-filename`. The results go to a results file.
--clear-all-cm	Clears the repository of all cable modems.
--show-dynamic-mode	Shows whether to allow or deny the login of CPEs that belong to unfamiliar cable modems; i.e., cable modems that do not exist in the SM database.

EXAMPLE 1:

The following example shows how to import cable modems from a CSV file.

```
p3cable -–import-cm --file=CMFile.csv
```

EXAMPLE 2:

The following example shows how to clear the repository of all cable modems:

```
p3cable --clear-all-cm
```

EXAMPLE 3:

The following example shows how to display login status (allow/deny) of CPEs that belong to cable modems that do not exist in the SM database:

```
p3cable --show-dynamic-mode
```

Installing an Application (p3inst)

Use the **p3inst** utility to install or uninstall an application (*pqi* file). You can install an application on the SM to customize the components. You can also upgrade an existing application to a new version, or uninstall a previously installed application.

To install or uninstall an application:

From the shell prompt, enter a command using the following general format:

`p3inst` *operation* [*file options*] [*installation/upgrade parameters*]

The following table lists the **p3inst** operations.

Table 4-17 p3inst Operations

Operation	Description
--install	Installs the specified application *pqi* file to the SM. It may be necessary to specify arguments for the installation procedure in the command line.
--uninstall	Uninstalls the specified application *pqi* file from the SM.
--upgrade	Upgrades an existing application using the specified application *pqi* file. It may be necessary to specify arguments for the upgrade procedure in the command line.

Operation	Description
--rollback	Returns the specified application to the previous version. Rollback is the reverse of an **upgrade** operation: it reverses the **upgrade**.
--describe	Displays the contents of the specified application *pqi* file.
--show-last	Lists the details of the last installed application *pqi* file

EXAMPLE 1:

The following example shows how to install the specified installation file.

```
p3inst --install --file=myInstallation.pqi
```

EXAMPLE 2:

The following example shows how to uninstall the specified installation file.

```
p3inst --uninstall –-file=oldInstallation.pqi
```

Viewing Information of the PRPC Interface Server (p3rpc)

Cisco provides a proprietary RPC (Remote Procedure Call) interface to the SM. Use the **p3rpc** utility to view the configuration and statistics of the PRPC server.

To display PRPC interface server information:

From the shell prompt, enter a command using the following general format:

`p3rpc` *operation*

The following table lists the **p3rpc** operations.

Table 4-18 p3rpc Operations

Operation	Description
--show	Displays the port number the PRPC server listens to, the maximum number of connections, the current number of active connections, and the host IP.
--show-statistics	Displays the PRPC server statistics. They contain information about the number of current PRPC sessions and statistics for PRPC server actions such as invocations and errors.
--reset-statistics	Clears the PRPC server statistics.
--show-client-names	Displays the names of the connected clients.

EXAMPLE:

The following example displays the statistics of the PRPC server.

```
p3rpc –-show-statistics
```

Managing a Cluster of Two SM Nodes (p3cluster)

Use the **p3cluster** utility to view the redundancy state of the SM and its components. This utility also supports operations that alter the redundancy state of the SM. These operations are used by the SM Cluster Agent and for administrative tasks.

To manage the cluster:

From the shell prompt, enter a command using the following general format:

`p3cluster` *operation*

The following table lists the **p3cluster** operations.

Table 4-19 p3cluster Operations

Operation	Description
--show	Displays the redundancy status of the SM and its components.
--active	Make the SM become the active SM in the cluster.
--standby	Make the SM become the standby SM in the cluster.

EXAMPLE:

The following example displays the redundancy status of the SM and it components.

```
p3cluster –-show
```

Managing the User Log (p3log)

Use the **p3log** utility to configure and manage the user log. All user-related events and errors are directed to the SM user log. You can extract the contents of the user log to a specified file in order to read and save its contents. You can also clear the user log.

To extract the user log to a specified file:

From the shell prompt, type:

`p3log –-extract --output=`*my_file*

To clear the user log:

From the shell prompt, type:

`p3log –-reset`

Viewing Statistics of the RADIUS Listener (p3radius)

Use the **p3radius** utility to view the statistics of the RADIUS Listener LEG. For information about this CLU, see the *Cisco SCMS SM RADIUS Listener LEG Reference Guide*.

Utilities

This section describes the following:

- Running a batch file (p3batch)
- Printing general help about CLU commands (p3clu)
- Database operations (p3db)

Running a Batch File (p3batch)

Use the **p3batch** utility to run a batch file and execute its commands. Using any text editor, you can create a batch file that contains a series of CLU commands, one command per line. This operation (**p3batch**) enables you to run multiple operations on a single connection to the SM.

To run a batch file:

From the shell prompt, enter a command using the following general format:

`p3batch` [*file-options*] [*error-options*]

Table 4-20 p3batch File Option

Operation	Description
--file=FILE	Runs a batch file, where FILE specifies the CLU script (batch) file to run.

Table 4-21 p3batch Error Option

Error Option	Description
--skip-errors	Specifies that the batch operation should *not* halt if an error occurs. If this flag is not used, the batch operation *will* halt if an error occurs.

EXAMPLE 1:

The following example shows how to run a batch file that will halt if an error occurs.

```
p3batch --file=mainBatchFile.txt
```

EXAMPLE 2:

The following example shows how to run a batch file that will not halt if an error occurs.

```
p3batch --file=mainBatchFile.txt --skip-errors
```

Printing General Help About CLU Commands (p3clu)

Use the **p3clu** utility to print a list of all supported CLU utilities and operations.

To print all CLU commands:

From the shell prompt, type:

`p3clu --help`

Database Operations

Use the **p3db** utility to manage and monitor the TimesTen database. The CLU exposes capabilities of the some of TimesTen CLUs with respect to specific needs of the SM.

To manage or monitor the TimesTen database:

From the shell prompt, enter a command using the following general format:

`p3db` *operation* [*options*]

The following tables list the **p3db** operations and options.

Note Use caution when activating commands that can affect the database. If used incorrectly, these commands may possibly **damage** the database.

Table 4-22 p3db Operations

Operation	Description
--status	Displays the database status.
--destroy-rep-db	Destroys the replicated data-store.
--destroy-local-db	Destroys the local data-store.
--keep-in-mem [SECS]	Indicates to the database daemon how many seconds to keep the database in the memory, after the last connection to the database is down. Use this option in large databases to reduce the SM restart time. **Note:** To prevent limitations in performing database destroy, do not use values above a few minutes (that is, above a few hundred seconds).
--rep-status	Displays status of the replication agent.
--rep-start	Starts the replication agent. **Note:** Use only for database recovery.

Operation	Description
--rep-stop	Stops the replication agent. **Note:** Use only for database recovery.
--duplicate	Copies the data-store from the *remote* machine (specified by "--remote=REMOTE_HOSTNAME") to the *local* machine (specified by "--local=LOCAL_HOSTNAME"). **Note:** This option is applicable only for a **cluster setup**. For additional information, see *Data Duplication Procedure* (on page 3-24).
--set-rep-scheme	Sets the database scheme to be replicated to the remote machine. **Note:** This option is applicable only for a **cluster setup**.
--drop-rep-scheme	Removes the database scheme that the remote machine uses for replication. **Note:** This option is applicable only for a **cluster setup**.

EXAMPLE

The following example shows how to request the status of the replication agent, and also lists a typical response:

```
bash-2.03$ p3db --rep-status
Peer name         Host name                 Port    State  Proto
---------------  ------------------------ ------  ------- -----
PCUBE_SM_REPOSITORY  SM_REP1          Auto   Start      11

Last Msg Sent Last Msg Recv Latency TPS     RecordsPS Logs
------------- ------------- ------- ------- --------- ----
00:00:02      00:00:00         1.15    2723      5447    1

the subscriber DB is ok
Command terminated successfully
```

Configuration File Options

This appendix describes in detail all the parameters that can be configured by using the Subscriber Manager (SM) configuration file. The shorter description of SM configuration given in *Configuration and Management* (on page 4-1) is more oriented toward the routine configuration tasks that can be performed online using the CLU.

Introduction

The SM can be configured only by using its configuration file (the CLU is used for displaying, not configuring). The SM has one configuration file, `p3sm.cfg`, which is located under `~pcube/sm/server/root/config/p3sm.cfg`. To change any configuration parameter, edit the configuration file using a standard text editor, and then use the CLU to reload it (see *Reloading the SM Configuration (p3sm)* (on page 4-23)).

You can use the `p3sm.cfg` configuration file for setting parameters for the following:

- SM configuration
- Radius Listener configuration
- Redundancy (cluster/standalone) configuration
- Domain configuration
- SCE configuration
- Cable adapter configuration
- PRPC port configuration
- FTP port configuration
- HTTP port configuration
- Database configuration

Description of the Configuration File Options

The following sections correspond to the sections of the SM configuration file, *p3sm.cfg*.

- *SM General* (on page A-2)
- *SM High Availability Setup* (on page A-4)

- *Subscriber State Persistency* (on page A-4)
- *SM-LEG Failure Handling* (on page A-4)
- *LEG-Domains Association* (on page A-5)
- *Domain.XXX* (on page A-6)
- *Default Domains Configuration* (on page A-7)
- *Auto Logout* (on page A-8)
- *Radius Listener* (on page A-9)
- *Radius.NAS.XXX* (on page A-10)
- *Radius.Property.Package* (on page A-11)
- *Radius.Subscriber ID* (on page A-12)
- *RPC.Server* (on page A-13)
- *SCE.XXX* (on page A-14)
- *FTP* (on page A-14)
- *HTTP Tech-If* (on page A-15)
- *Cable Adapter* (on page A-16)
- *Data Repository* (on page A-17)

For an explanation of the terms and concepts used in these sections, see *Configuring a Subscriber Management Solution* (on page 4-3).

SM General Section

The `[SM General]` section contains the following parameters:

- *`introduction_mode`*

  Defines whether the SM introduces the subscribers to the SCE devices immediately after a login operation (Push mode), or only when the SCE requests the subscriber specifically (Pull mode).

  Possible values for this parameter are **`push`** and **`pull`**. The default value is **`push`**.

  The following is an example of assigning a value to this parameter:

  ```
  introduction_mode=push
  ```

- *`application_subscriber_lock`*

  Defines whether to lock subscriber-related operations (login, logout, etc.) at the application level. Set this flag to `true` only if several LEG applications can update simultaneously the same parameters of a subscriber.

  Possible values for this parameter are **`true`** and **`false`**. The default value is **`true`**.

  The following is an example of assigning a value to this parameter:

  ```
  application_subscriber_lock=true
  ```

- *force_subscriber_on_one_sce*

  Defines whether the SM supports the solution when a Cisco 7600/6500 Router is used for load-balancing among several SCE platforms. In this solution when one SCE platform fails, subscriber traffic is redistributed to a different SCE platform. The SM must remove subscribers from the failed SCE platform and send the relevant subscriber information to the new SCE platform. This parameter is relevant only in the pull introduction mode.

  Possible values for this parameter are **true** and **false**. The default value is **false**.

Note Changing this value requires a restart of the SM process.

  The following is an example of assigning a value to this parameter:

  ```
  force_subscriber_on_one_sce=false
  ```

- logon_logging_enabled

  Defines whether to enable the logging of subscriber logon events.

Note Setting this flag to **true** might cause performance degradation.

  Possible values for this parameter are **true** and **false**. The default value is **false**.

  The following is an example of assigning a value to this parameter:

  ```
  logon_logging_enabled=false
  ```

- *subscriber_id_case_sensitivity*

  Optional parameter that defines whether subscriber IDs are case sensitive or not. When this flag is set to **no**, all subscriber IDs in the subscriber database are set to be lower case. For example, 'JohnSmith' is converted to 'johnsmith'.

Note Setting this flag to **no** when the subscriber database is not empty is not allowed. An error message will be shown and the configuration will not be loaded. To overcome this limitation, you can do the following:
1. Export the subscriber database to an external file
2. Clear the subscriber database
3. Change the configuration by setting the flag to **no** and load the new configuration
4. Import the subscriber database from the external file

  Possible values for this parameter are **yes** and **no**. The default value is **yes**.

  The following is an example of assigning a value to this parameter:

  ```
  subscriber_id_case_sensitivity=yes
  ```

SM High Availability Setup Section

The `[SM High Availability Setup]` section contains the following parameter:

- *`topology`*

  Defines in what kind of topology the SM should work.

  Possible values for this parameter are **`standalone`** and **`cluster`**. The default value is **`standalone`**.

  The following is an example of assigning a value to this parameter:

  ```
  topology=standalone
  ```

Subscriber State Persistency Section

The `[Subscriber State Persistency]` section contains the following parameters:

- SCE_subscriber_persistency

  Defines whether the SCE stores subscriber information in persistent memory. Use this option **only** when there are just a few subscribers, which rarely change, in the system.

Note Setting this parameter to **`true`** causes performance degradation, reducing the rate of subscriber updates.

  Possible values for this parameter are **`true`** and **`false`**. The default value is **`false`**.

  The following is an example of assigning a value to this parameter:

  ```
  SCE_subscriber_persistency=false
  ```

SM-LEG Failure Handling Section

The `[SM-LEG Failure Handling]` section contains parameters that affect the discovery of an SM-LEG connection failure event and the actions taken by SM upon a connection failure event. A network problem or a severe failure (such as reboot) of the SM or the LEG can cause an SM-LEG connection failure event.

The `[SM-LEG Failure Handling]` section contains the following parameters:

- *`clear_all_mappings`*

  Defines the behavior of the SM in case of LEG-SM connection failure. This parameter is relevant only for cases where the SM and LEG are running on **different** machines.

Note This parameter defines a behavior that is similar for **all** LEG applications connected to the SM.

  If this parameter is set to **`true`** and a SM-LEG connection failure occurs that is not recovered within the defined *`timeout`*, the mappings of all subscribers in the domains defined in the *[LEG-Domains Association]* (on page A-5) section for the LEG that was disconnected will be removed.

Note **Important:** If you set the **clear_all_mappings** parameter to **`true`**, you must also set the **LEG-Domains Association** parameter to clear the mappings in the SM if an SM-LEG connection failure occurs.

Possible values for this parameter are **`true`** and **`false`**. The default value is **`false`**.

The following is an example of assigning a value to this parameter:

```
clear_all_mappings=false
```

- *`timeout`*

  Defines the time in seconds from a SM-LEG connection failure until clearing the mappings in the SM database.

Note It takes several seconds for the SM to detect an SM–LEG connection failure. You must add this time to the value of the *`timeout`* parameter when calculating how long it will take the SM to react in case of an SM–LEG connection failure. For example, if *`timeout`* is set to 80, it will take 80 seconds **from the failure detection time** until the SM clears the subscribers database.

The default value for this parameter is 60 (seconds).

The following is an example of assigning a value to this parameter:

```
timeout=60
```

LEG-Domains Association Section

The *`[LEG Domains Association]`* section defines the domains for which the mapping of all subscribers that belong to them will be cleared on SM-LEG connection failure. This section is relevant only if the **clear_all_mappings** parameter has been set in the *[SM-LEG Failure Handling]* (on page A-4) section.

Note Even though you set the **LEG-Domains Association** parameter, you must also set the **clear_all_mappings** parameter to **`true`** to actually clear the mappings in the SM if an SM-LEG connection failure occurs.

This section contains a list of LEG-Domain associations, each item in a separate line. Each LEG-Domain association is specified as shown for the following parameter:

- *`<LEG name>=domain_name1[,domain_name2,...]`*

  Defines the domains whose subscriber mapping will be cleared on an SM-LEG connection failure. The key is the *`<LEG name>`*.

To determine which value or values to use for the <LEG name> key, consult the documentation of the LEG that you use. The *<LEG name>* usually is divided into two parts: *<hostname>.<common suffix>*. The first part is a general LEG identifier. The second part is extracted from the machine on which the LEG is running. Alternatively, you can use the CLU command `p3rpc --show-client-names`.

A *<LEG name>* of **"*"** specifies all LEGs. The (comma-separated) values are the domains (*domain_name*) to be cleared in the event of a network link failure (connection failure) between the specified LEG and the SM. A value of **"*"** for the *domain_name* specifies all subscriber domains in the system.

Note The LEG name is **case sensitive**.

By default, there are no LEG domain mappings.

The following is an example of assigning values to this parameter:

```
10.1.12.76.NB.SM-API.J=home_users
10.1.12.77.B.SM-API.C=office_users
```

The following is an example of specifying all subscriber domains:

```
10.1.12.31.CNR.LEG=*
```

The following is an example of specifying all LEGs and all subscriber domains:

```
*=*
```

Domain.XXX Section

Each `[Domain.XXX]` section specifies one domain, where *XXX* represents the domain name.

This section contains the following parameters:

- *`elements=logical_name1[,logical_name2,...]`*

  Specifies the name or names of the SCE platforms that are part of the domain.

Note Each name must be the **exact** "XXX" name (**case sensitive**) that appears in the *SCE.XXX Section* (on page A-14).

Following is an example of assigning a value to this parameter:

```
elements=se0,se1
```

- *aliases=alias_name1[,alias_name2,...]*

Defines domain aliases. When subscriber information is received from the LEG with one of the aliases (for example, *alias1*), the information is distributed to the domain that matches this alias (for example, *domain_name1*). A typical alias could be a network access device IP address, where, for example, each string in the values can be the IP address of a NAS or a CMTS.

Note Each alias can appear in only one domain section.

By default, there are no domain aliases.

The following is an example of assigning a value to this parameter:

```
aliases=10.10.88.99,10.10.88.98
```

- *property.name1=value1[,property.name2=value2,...]*

Defines the default policy properties values for a domain. Unless the LEG/API overrides these defaults when it logs in the subscriber to the SM, the subscriber policy is set according to the default policy properties values of its domain.

The policy format is a comma-separated list of `property_name=property_value`, where each property value is an integer.

Note To learn more about policy specification, see the *Service Control Application for Broadband (SCA BB) User Guide*.

The following is an example of assigning a value to this parameter:

```
property.packageId=1
```

Default Domains Configuration Section

The `[Default Domains Configuration]` section defines default policy for the domain. It is used for those domain properties that were not defined in the domain policy configuration (see `[Domain.XXX]` (on page A-6)), and for properties of domains for which no policy whatsoever was defined.

This section contains the following parameter:

- *property.name1=value1[,property.name2=value2,...]*

Defines the default policy properties values for all domains. Unless the LEG/API overrides these defaults when it logs in the subscriber to the SM, or unless they are overwritten by the default policy property values of the subscriber domain, the subscriber policy is set according to the global default policy property values defined in this section.

The policy format is a comma-separated list of `property_name=property_value`.

Note To learn more about policy specification, see the *Cisco Service Control Application for Broadband (SCA BB) User Guide*.

The following is an example of assigning a value to this parameter:

```
property.packageId=1
```

Auto Logout Section

The `[Auto Logout]` section defines the parameters for the **Auto Logout** feature, which is relevant mainly for cable environments. This feature is relevant to automatic integrations where the LEG/API cannot provide logout indications. In such a case, you can turn on the automatic logout mechanism, which instructs the SM to log out a subscriber automatically after a certain period of time. Note that a login event for a subscriber resets the subscriber logout timer.

Note *Not* using the **Auto Logout** feature in the scenario described above (a provisioning system that can provide subscriber login events to the SM but cannot provide subscriber logout events) might result in exhausting the SCE resources, because subscribers are logged in but are never logged out.

This section contains the following parameters:

- *`auto_logout_interval`*

  Configures the interval value, in seconds, of the SM auto-logout mechanism. Every interval, the SM checks for which subscriber IP addresses the lease time has expired, and begins to automatically remove these IP addresses from the system.

  Lease time is the timeout defined by the LEG during the login operation per each IP address. All subscriber login events will start a timer of lease-time seconds. When the timer expires and the grace period (see below) has also expired, the subscriber IP addresses are removed, causing the subscriber to be removed from the SCE platform database. Any login event by the subscriber with an existing IP address during the timer countdown period resets the timer, causing it to restart.

  Setting the interval value to zero (0) disables the SM auto-logout mechanism.

  Setting the interval to a value greater than zero enables the SM auto-logout mechanism.

Note The interval should be smaller, but of the order of the lease time used in the system. It is recommended that during a lease time the auto-logout task will run several times.

The default value for this parameter is 0 (seconds), meaning the auto-logout mechanism is disabled.

The following is an example of assigning a value to this parameter:

```
auto_logout_interval=600
```

The following is an example of deactivating the Auto Logout feature:

```
auto_logout_interval=0
```

- *grace_period*

  Defines the grace period, in seconds, for each subscriber. After a subscriber auto-logout timeout has expired, the subscriber IP address is logged out automatically after the grace period has also expired.

  The default value for this parameter is 10 (seconds).

  The following is an example of assigning a value to this parameter:

  ```
  grace_period=10
  ```

- *max_rate*

  Defines the maximum rate (logouts per second) that the auto-logout task is allowed to perform logouts from the system. This limit spreads out the load of the logout operations over time, reducing the performance impact on other operations.

  Calculate the value for this parameter to spread the logouts over at least half of the **auto_logout_interval** time. The default value is 50.

  The following is an example of assigning a value to this parameter:

  ```
  max_rate=50
  ```

Note Use the lowest rate possible to reduce the influence of the auto-logout process on other operations. However, as a guideline, calculate the value so that the auto-logout process takes about half of the **auto_logout_interval** and be similar to the maximum login rate to the SM.

Radius Listener Section

Use the `[Radius Listener]` section for integrating with the RADIUS Listener LEG.

Note For additional RADIUS Listener configuration information, see the *Cisco SCMS SM RADIUS Listener LEG Reference Guide*.

This section contains the following parameters:

- *start*

  Defines whether the SM should run the RADIUS Listener at startup.

  Possible values for this parameter are **yes** and **no**. The default value is **no**.

  The following is an example of assigning a value to this parameter:

  ```
  start=no
  ```

- *accounting_port*

  Defines the RADIUS Listener's accounting port number.

  The default value is 1813.

  The following is an example of assigning a value to this parameter:

  ```
  accounting_port=1813
  ```

- *ip*

  (Optional) Defines the IP address to where the RADIUS Listener should bind. Use this parameter in cluster setups or when the machine local-host IP is not the IP to where the RADIUS messages are sent.

  By default, this parameter is not set.

  The following is an example of assigning a value to this parameter:

  ```
  ip=192.56.21.200
  ```

Radius.NAS.XXX Section

Each `[Radius.NAS.XXX]` section specifies a single Network Access System (NAS), where *XXX* represents the NAS name.

Note The RADIUS Listener LEG refers to all of its RADIUS clients as NAS devices, even though they might be RADIUS servers.

This section contains the following parameters:

- *domain*

  Specifies the Cisco Service Control subscriber domain name.

  The following is an example of assigning a value to this parameter:

  ```
  domain=my_domain
  ```

- *IP_address*

  Specifies the IP address in dotted notation (xxx.xxx.xxx.xxx).

  The following is an example of assigning a value to this parameter:

  ```
  IP_address= 1.1.1.1
  ```

- *NAS_identifier*

  Specifies the name of the NAS that exists in the **NAS-ID** attribute. For information about the use of this parameter, see the *Cisco SCMS SM RADIUS Listener LEG Reference Guide*.

  The following is an example of assigning a value to this parameter:

  ```
  NAS_identifier=RedHat37
  ```

- *secret*

  Specifies a secret key defined in the NAS for this connection.

  The following is an example of assigning a value to this parameter:

  ```
  secret=mysecret
  ```

Radius.Property.Package Section

The `[Radius.Property.Package]` section defines the *RADIUS attribute* from which a subscriber package is retrieved.

Note This section is commented out in the configuration file. If you want to retrieve a subscriber package from an attribute other than a *RADIUS attribute*, you should uncomment this section.

This section contains the following parameters:

- *radius_attribute*

  Specifies the RADIUS protocol attribute number. Use the value of **26** for Vendor Specific Attributes (VSA). Use **-1** if you do not want to extract from any attribute.

  The following is an example of assigning a value to this parameter:

  ```
  radius_attribute=26
  ```

- *radius_attribute_type*

  Specifies the RADIUS attribute type.

  Possible values for this parameter are **integer** and **string**. The default value is **integer**.

  The following is an example of assigning a value to this parameter:

  ```
  radius_attribute_type=integer
  ```

- *use_default*

  Defines whether to use a default value if the attribute was not found.

  Possible values for this parameter are **true** and **false**. The default value is **true**.

  The following is an example of assigning a value to this parameter:

  ```
  use_default=true
  ```

- *default*

  Defines the default value to use if the attribute was not found.

  There is no default value. (If this parameter is used, it should have non-empty value.)

  The following is an example of assigning a value to this parameter:

  ```
  default=0
  ```

- Additional parameters for VSA association

  Use the following additional parameters only if the association is based on Vendor Specific Attributes (VSA):

  ```
  radius_attribute_vendor_id=<insert the vendor identifier number here>
  radius_sub_attribute=<insert the attribute number here>
  ```

Radius.Subscriber ID Section

The `[Radius.Subscriber ID]` section defines the RADIUS attribute on which to base the subscriber ID association. Default association is based on the **user name** attribute.

Note This section is commented out in the configuration file. If you want to base subscriber ID association on an attribute other than **user name**, you should uncomment this section.

This section contains the following parameters:

- *`radius_attribute`*

  Specifies the RADIUS protocol attribute number. Use the value of **`26`** for Vendor Specific Attributes (VSA). Use the value of **`-1`** to use the default.

  The following is an example of assigning a value to this parameter:

  ```
  radius_attribute=26
  ```

- *`radius_attribute_type`*

  Specifies the RADIUS attribute type.

  Possible values for this parameter are **`integer`** and **`string`**. The default value is **`integer`**.

  The following is an example of assigning a value to this parameter:

  ```
  radius_attribute_type=string
  ```

- Additional parameters for VSA association

  Use the following additional parameters only when basing the association on Vendor Specific Attributes (VSA). For example, to use the 3GPP_IMSI attribute, set:

  ```
  radius_attribute_vendor_id=10415
  radius_sub_attribute=1
  ```

RPC.Server Section

The `[RPC.Server]` section represents the PRPC server configuration.

This section contains the following parameters:

- *`port`*

  Defines the PRPC server port. The default is 14374.

  The following is an example of assigning a value to this parameter:

  ```
  port=14374
  ```

- *`security_level`*

  Defines whether the PRPC server forces authentication on all connections (**`full`**), authenticates connections that support authentication while still accepting connections that do not (**`semi`**), or not to enforce authentication (**`none`**). When clients attempt to connect to the SM, they are authenticated if configured correctly.

  Possible values for this parameter are **`full`**, **`semi`**, and **`none`**.

  The default value is **`semi`**.

Note SCA BB Console 3.0.5 supports authentication with the SM PRPC Server; therefore, it can be used in conjunction with all *`security_level`* values.

Note In version 3.0.5, the SM Java API, SM C/C++ API, and CNR LEG do not support authentication with the SM PRPC Server; therefore, if installed, the security level cannot be configured to **`full`**.

The following is an example of assigning a value to this parameter:

```
security_level=semi
```

MPLS-VPN Section

The `[MPLS-VPN]` section contains configuration parameters that are relevant to MPLS/VPN installations. See the *Cisco SCMS SM MPLS/VPN BGP LEG Reference Guide* for a description of Subscriber Management in MPLS/VPN networks.

This section contains the following parameters:

- *`vpn_id`*

  Defines the BGP attribute to use to identify the VPN subscribers.

  Possible values for this parameter are **`RD`** or **`RT`**. The default value is **`RT`**.

  The following is an example of assigning a value to this parameter:

  ```
  vpn_id=RD
  ```

- *log_all*

  Defines the logging level of the BGP LEG.

  Possible values for this parameter are **true** or **false**. The default value is **false**. If set to **true** the SM logs all BGP packets that were received by it and is useful during the integration and testing phase.

  The following is an example of assigning a value to this parameter:

  ```
  log_all=false
  ```

SCE.XXX Section

Each [SCE.XXX] section represents a single SCE platform, where *XXX* represents the SCE logical name.

This section contains the following parameters:

- *ip*

  Defines the IP address of the SCE device.

  The following is an example of assigning a value to this parameter:

  ```
  ip=11.12.13.1
  ```

- *port*

  Defines the port through which to connect to the SCE platform. The default is 14374.

  The following is an example of assigning a value to this parameter:

  ```
  port=14374
  ```

FTP Section

The SM manages an internal FTP server for various purposes.

The [FTP] section contains the following parameters:

- *start*

  Defines whether the SM should run the FTP server at startup.

  Possible values are **yes** and **no**. The default is **no**.

  The following is an example of assigning a value to this parameter:

  ```
  start=yes
  ```

- *port*

  Defines the FTP server port. The default is 21000.

  The following is an example of assigning a value to this parameter:

  ```
  port=21000
  ```

The following parameters define the TCP ports range when the FTP server works in passive mode:

```
#first_passive_port = 21001
#last_passive_port = 21100
```

Note Uncomment these parameters when working with the FTP server via a firewall.

HTTP Tech-IF Section

The SM manages an internal HTTP adapter server that is a technician interface.

The `[HTTP Tech-IF]` section contains the following parameters:

- *start*

  Defines whether the SM should run the HTTP server at startup.

Note This is a technician interface and normally should not be used.

  Possible values are **yes** and **no**. The default is **no**.

  The following is an example of assigning a value to this parameter:

  ```
  start=yes
  ```

- *port*

  Defines the HTTP server port. The default is 8082.

  The following is an example of assigning a value to this parameter:

  ```
  port=8082
  ```

RDR Server Section

The SM manages an internal RDR server that is used to receive RDR messages from the SCE.

The `[RDR Server]` section contains the following parameters:

- *`start`*

  Defines whether the SM should run the RDR server at startup.

Note This interface should be used when installing the SCE-Sniffer LEGs on the SM.

  Possible values are **`yes`** and **no**. The default is **no**.

  The following is an example of assigning a value to this parameter:

  ```
  start=yes
  ```

- *`port`*

  Defines the RDR server port. The default is 33001.

  The following is an example of assigning a value to this parameter:

  ```
  port=33001
  ```

- *`max_connections`*

  Defines the maximum number of connections accepted by the server. The default is 10.

  The following is an example of assigning a value to this parameter:

  ```
  max_connections=10
  ```

Cable Adapter Section

The SM manages a Cable Adapter, a special cable support module that is a translator between the cable world (DHCP events) and the SM. For additional information, see *CPE as Subscriber in Cable Environment* (on page C-1).

The `[Cable Adapter]` section contains the following parameter:

- *`allow_dynamic_CM`*

  Defines whether to allow logins from cable modems (CM) that are not in the SM database.

  Possible values are **`yes`** and **no**. The default is **no**.

  The following is an example of assigning a value to this parameter:

  ```
  allow_dynamic_CM=no
  ```

Data Repository Section

The `[Data Repository]` section defines the SM operation with the TimesTen In-Memory Database.

Note Some of the parameters in this section are discarded on regular configuration loading, and resetting them requires restarting the SM.

The `[Data Repository]` section contains the following parameters:

- *`support_ip_ranges`*

  Defines whether the SM supports IP-Ranges. Disabling this support provides better performance.

Note Resetting this parameter requires restarting the SM. This parameter is discarded on regular configuration loading (using CLU).

  Possible values are **`yes`** and **`no`**. The default is **`no`**.

  The following is an example of assigning a value to this parameter:

  ```
  support_ip_ranges=yes
  ```

- *`checkpoint_interval_in_seconds`*

  Defines the interval, in seconds, for calling the TimesTen checkpoints. Reducing the value affects performance, increasing the value increases vulnerability to power-down.

  The default value is 60 (seconds).

  The following is an example of assigning a value to this parameter:

  ```
  checkpoint_interval_in_seconds=60
  ```

- *`max_range_size`*

  Determines the maximum IP range size used in the system.

  This parameter is used for improving performance of the SM in Pull mode when the Data Repository section is configured with *`support_ip_ranges=yes`*.

Note Defining this parameter with too low a value may cause incorrect operation in handling pull requests.

  The default value is 256.

  The following is an example of assigning a value to this parameter:

  ```
  max_range_size=256
  ```

Command-Line Utilities

This appendix describes the Command-Line Utilities (CLU) that is distributed with the Subscriber Manager (SM) application.

Introduction

The SM provides a set of Command-Line Utilities (CLU). The CLU is designed mainly for viewing SM operations and statistics, and subscriber management, therefore only subscriber-related CLU changes are persistent. The CLU is used for configuration only indirectly, in that it loads the edited configuration file to the SM.

This appendix describes in detail all the CLU commands, their operations and options. The shorter description of the CLU commands given in the *Configuration and Management* (on page 4-1) chapter is more oriented toward the performance of routine management and configuration tasks.

CLU commands are executable only when the user is logged in to the machine using the **pcube** account, which is always installed (see *Installation and Upgrading* (on page 3-1)). In general, the CLU runs as a separate process to the configured entity and communicates with it via a predefined communication port and interface. Therefore, the configured entity must keep a certain communication port open at all times, at least locally on the configured machine.

Description of the CLU Commands

This section describes the Command-Line Utilities commands, their operations and options, in detail.

This section contains the following topics:

- *Informative Output* (on page B-2)
- *Parsing CLU Operations and Options* (on page B-2)
- *p3batch Utility* (on page B-3)
- *p3cable Utility* (on page B-4)
- *p3clu Utility* (on page B-5)
- *p3cluster Utility* (on page B-5)
- *p3db Utility* (on page B-6)

- *p3domains Utility* (on page B-7)
- *p3ftp Utility* (on page B-7)
- *p3http Utility* (on page B-8)
- *p3inst Utility* (on page B-8)
- *p3log Utility* (on page B-10)
- *p3net Utility* (on page B-11)
- *p3radius Utility* (on page B-11)
- *p3rpc Utility* (on page B-12)
- *p3sm Utility* (on page B-13)
- *p3subs Utility* (on page B-16)
- *p3subsdb Utility* (on page B-19)

Informative Output

All CLU commands support the following operations for informative output:

Operation	Description
--help	Prints the help for the specified CLU command, then exits.
--version	Prints the SM program version number, then exits.

Parsing CLU Operations and Options

Place in quotation marks a command operation or option containing any of the following characters:

- A space character
- A separation sign (comma “,”; ampersand “&”; colon “:”)
- An escape character (backslash “\”)

A command operation or option that contains any of the following characters must have that character preceded by an escape character:

- An equal sign (=)
- A quotation mark (“ or ”)
- An escape character (backslash “\”)

Following are several examples of the above rules:

Operation/option contains the character	Example of how operation/option should be written
Space character	--property=“file name”
Escape character (backslash “\”)	--property=“good\\bad”
Equal sign (=)	--property=“x\=y”
Quotation marks (“ or ”)	--name=“\“myQuotedName\”” (in above example, inner quotation marks are escaped)

Separation character

- comma (,) --names="x,y"
- ampersand (&) --names="x&y"
- colon (:) --names="myHost:myDomain"

One-letter abbreviations are available for some of the operations and options. For example, "-d" is an abbreviation for "--domain". Note that only one hyphen (-), not two, precedes the letter for an abbreviation, and that if the operation or option takes a parameter, there is a space and not an equal sign before the parameter.

- Example of using full name --domain=subscribers
- Example of using abbreviated name -d subscribers

p3batch Utility

The **p3batch** utility enables you to run many operations on a single connection with the SM. You can use any text editor to create a batch file that contains a series of CLU commands, one command per line (terminated by a new-line sign). Use the **p3batch** utility to run this file and execute the commands, where empty lines are skipped.

All batch file command line operations use the same connection option. The **p3batch** utility ignores any connection options in the script file commands. While processing the operations in the batch file, a progress indicator is displayed.

The command format is: `p3batch [FILE-OPTION] [ERROR-OPTION]`

The following tables list the **p3batch** options.

Table B-1 p3batch File Option

File Operation	Abbreviation	Description	Notes
--file=FILE	-f	Runs a batch file, where FILE specifies the CLU script (batch) file to run.	A progress indicator is displayed.

Table B-2 p3batch Error Option

Error Option	Description
--skip-errors	Specifies that the batch operation should *not* halt if an error occurs. If this flag is not used, the batch operation *will* halt if an error occurs.

Examples of using the p3batch utility

- To run a batch file that will halt if an error occurs.

  ```
  p3batch --file=mainBatchFile.txt
  ```

- To run a batch file that will *not* halt if an error occurs.

  ```
  p3batch --file=mainBatchFile.txt --skip-errors
  ```

p3cable Utility

In a cable environment, the CPE is modeled as the subscriber, and inherits its policy and domain from the cable modem (CM) through which it connects to the network. Each cable modem is linked with one or more CPEs. For background information about special characteristics of the cable environment, see *CPE as Subscriber in Cable Environment* (on page C-1).

You can use the **p3cable** utility commands to import cable modem information from a CSV file to the SM and to export the cable modem information from the SM to a CSV file. You can also use this utility to clear the repository of all cable modems, and to show whether to allow or deny the login of CPEs that belong to unfamiliar cable modems; i.e., cable modems that do not exist in the SM database. However, to specify whether to allow or deny such a login, use the *Cable Adapter* (on page 4-13) section of the configuration file *p3sm.cfg*.

The command format is: `p3cable OPERATION [FILE-OPTIONS]`

The following tables list the **p3cable** operations and options.

Table B-3 p3cable Operations

Operation	Description
--import-cm	Imports cable modems from a CSV file. The filename that is to be imported is specified using the format `--file=import-filename`. The results go to a results file.
--export-cm	Exports cable modems to a CSV file. The filename that is to be exported is specified using the format `--output=export-filename`. The results go to a results file.
--clear-all-cm	Clears the repository of all cable modems.
--show-dynamic-mode	Shows whether to allow or deny the login of CPEs that belong to unfamiliar cable modems; i.e., cable modems that do not exist in the SM database.

Table B-4 p3cable File Options

File Option	Abbreviation	Description
--file=FILE	-f FILE	Specifies the CSV FILE to import from.
--output=FILE	-o FILE	Specifies the subscriber CSV FILE to export to.

Examples of using the p3cable utility

- To import cable modem information from the specified csv file:

  ```
  p3cable –-import-cm –f CMFile.csv
  ```

- To export cable modem information to the specified csv file:

  ```
  p3cable –-export-cm –-outfile=myCMFile.csv
  ```

- To clear the repository of all cable modems:

  ```
  p3cable --clear-all-cm
  ```

- To display whether the login of CPEs that belong to unfamiliar cable modems (cable modems that do not exist in SM database) is allowed or denied:

  ```
  p3cable --show-dynamic-mode
  ```

p3clu Utility

The **p3clu** utility prints a list of all supported CLU utilities and options.

The command format is: `p3clu OPERATION`

The following table lists the **p3clu** operations.

Table B-5 p3clu Operations

Operation	Description
--help	Prints the currently supported CLU commands.

Example of using the p3clu utility

- To display a listing of all supported CLU utilities and operations:

  ```
  p3clu --help
  ```

p3cluster Utility

The **p3cluster** utility displays the redundancy state of a cluster of two SM nodes and its components. This utility also supports operations that alter the redundancy state of the SM. These operations are used by the SM Cluster Agent and for administrative tasks.

The command format is: `p3cluster OPERATION`

The following table lists the **p3cluster** operations.

Table B-6 p3cluster Operations

Operation	Description
--show	Displays the redundancy status of the SM and its components.
--active	Make the SM become the active SM in the cluster.
--standby	Make the SM become the standby SM in the cluster.

Example of using the p3cluster utility

- To display the redundancy status of the SM and its components:

  ```
  p3cluster –-show
  ```

p3db Utility

The **p3db** utility manages and monitors the TimesTen database. The CLU exposes capabilities of some of the TimesTen CLUs with respect to specific needs of the SM.

The command format is: `p3db OPERATION [OPTIONS]`

The following tables list the **p3db** operations and options.

Caution Use caution when activating commands that can affect the database. If used incorrectly, these commands can possibly damage the database.

Table B-7 p3db Operations

Operation	Description
--rep-status	Displays status of the replication agent.
--rep-start	Starts the replication agent. **Note:** Use only for database recovery.
--rep-stop	Stops the replication agent. **Note:** Use only for database recovery.
--status	Displays the database status.
--destroy-rep-db	Destroys the replicated data-store.
--destroy-local-db	Destroys the local data-store.
--duplicate	Copies the data-store from the *remote* machine to the *local* machine. **Note:** This option is applicable only for a **cluster setup**. For additional information, see *Data Duplication Procedure* (on page 3-24).
--keep-in-mem [SECS]	Indicates to the database daemon how many seconds to keep the database in the memory, after the last connection to the database is down. Use this option with large databases to reduce the SM restart time. **Note:** To prevent limitations in performing a database destroy, do not use values above a few minutes (that is, above a few hundred seconds).

Table B-8 p3db Options

Option	Description
--local=LOCAL_HOSTNAME	Specifies the local machine.
--remote=REMOTE_HOSTNAME	Specifies the remote machine.

Examples of using the p3db Utility

- To display the status of the replication agent:

  ```
  p3db –-rep-status
  ```

p3domains Utility

The **p3domains** utility displays the subscriber domains. When a system has more than one SCE platform, you can configure the platforms into groups or domains. A subscriber domain is one or more SCE platforms that share a specified group of subscribers. You must add the SCE platform to the network and create the domain before you can add an SCE platform to a domain.

The command format is: `p3domains OPERATION [OPTIONS]`

The following tables list the **p3domains** operations and options.

Table B-9 p3domains Operations

Operation	Description
--show-all	Displays all configured domains.
--show	Displays a domain and its associated network elements.

Table B-10 p3domains Domain/Network Element (NE) Options

Domain/NE Option	Abbreviation	Description	Notes
--domain=DOMAIN	-d DOMAIN	DOMAIN specifies logical name.	"**none**" cannot be used, it is a reserved word.

Examples of using the p3domains utility

- To display all configured domains:

  ```
  p3domains --show-all
  ```

- To display the specified domain and its associated network elements:

  ```
  p3domains --show --domain=myDomain
  ```

p3ftp Utility

The **p3ftp** utility monitors the SM internal FTP server.

The command format is: `p3ftp OPERATION`

The following table lists the **p3ftp** operations and options.

Table B-11 p3ftp Operations

Operation	Description	Notes
--show	Displays the port number the FTP server listens to, the passive FTP port range the server uses, the current number of open sessions, the maximum number of sessions supported, and the state (ONLINE/OFFLINE) of the FTP server.	

Examples of using the p3ftp utility

- To display the port number that the FTP server listens to, the passive FTP port range that the server uses, the current number of open sessions, the maximum number of sessions supported, and the state (ONLINE/OFFLINE) of the FTP server.

  ```
  p3ftp --show
  ```

p3http Utility

The **p3http** utility monitors the HTTP adapter server.

Note The HTTP adapter server is a technician interface and normally should not be used.

The command format is: `p3http OPERATION`

The following table lists the **p3http** operations:

Table B-12 p3http Operations

Operation	Description	Notes
--show	Displays the port number that the server listens to, the state of the server, and the current number of open sessions.	

Examples of using the p3http utility

- To display the port number to which the server listens, the state of the server, and the current number of open sessions:

  ```
  p3http --show
  ```

p3inst Utility

The **p3inst** utility installs or uninstalls an application (*pqi* file).

Note Before using p3inst to install an application *pqi* file, read the application installation instructions that came with the application you are using.

The command format is: `p3inst OPERATION [FILE-OPTION] [ARGUMENT-OPTION]`

The following tables list the **p3inst** operations and options.

Table B-13 p3inst Operations

Operation	Abbreviation	Description	Notes
--install	-i	Installs the specified application *pqi* file to the SM. It may be necessary to specify arguments for the installation procedure in the command line. Requires a file option.	Progress indicator
--uninstall		Uninstalls the specified application *pqi* file from the SM. Requires a file option.	Progress indicator
--upgrade		Upgrades an existing application using the specified application *pqi* file. It may be necessary to specify arguments for the upgrade procedure in the command line. Requires a file option.	Progress indicator
--rollback		Returns the specified application to the previous version. Rollback is the opposite of an **upgrade** operation: it reverses the **upgrade**.	Progress indicator
--describe	-d	Displays the contents of the specified application *pqi* file.	
--show-last		Lists the details of the last installed application *pqi* file.	

Table B-14 p3inst File Options

File Option	Abbreviation	Description
--file=FILE[;FILE...]	-f FILE[;FILE...]	Specifies one or more installation FILEs to use. If there is more than one FILE, semicolons should separate them.

Table B-15 p3inst Argument Options

Argument Option	Description
--arg=ARG1[,ARG2…]	Specifies one or more arguments for the install and upgrade procedures.

Examples of using the p3inst utility

- To install the specified installation file:
  `p3inst --install --file=myInstallation.pqi`
- To uninstall the specified installation file:
  `p3inst --uninstall –f oldInstallation.pqi`
- To upgrade an existing application using the specified application *pqi* file:
  `p3inst –-upgrade –-file=newInstallation.pqi`

- To upgrade an existing application using the specified application *pqi* file, using arguments in the command line:

  ```
  p3inst –-upgrade –f newInstallation.pqi
  ```

- To return the specified application to the previous version:

  ```
  p3inst --rollback
  ```

- To display the contents of the specified application *pqi* file:

  ```
  p3inst –-describe –-file=myInstallation.pqi
  ```

- To list the details of the last installed application *pqi* file:

  ```
  p3inst --show-last
  ```

p3log Utility

The **p3log** utility configures and manages the SM user log. The user log contains all user-related events and errors. Use the user log to view the history of the system events and errors.

The command format is: `p3log OPERATION [FILE-OPTION]`

The following tables list the **p3log** operations and options.

Table B-16 p3log Operations

Operation	Description	Notes
--extract	Retrieves the user log from the agent.	Progress indicator
--reset	Clears the user log.	

Table B-17 p3log File Option

File Option	Abbreviation	Description
--output=FILE	-o FILE	Specifies to where the SM user log file should be extracted.

Examples of using the p3log utility

- To extract the SM user log to the specified file:

  ```
  p3log –-extract –o aug20.log
  ```

- To clear the SM user log:

  ```
  p3log –-reset
  ```

p3net Utility

The **p3net** utility shows the connection status of network elements and tries to reconnect disconnected elements.

The command format is: `p3net OPERATION [NETWORK-ELEMENT-OPTION]`

The following tables list the **p3net** operations and options.

Table B-18 p3net Operations

Operation	Description
--show-all	Shows all the configured network elements.
--show	Shows the element connection status/general information.
--connect	Tries to connect a disconnected element.

Table B-19 p3net Network Element Options

Network Element Option	Abbreviation	Description
--ne-name=NAME	-n NAME	Specifies the logical NAME for the network element.
--detail		(Optional) Used with the **--show-all** operation for displaying additional information as a table.

Examples of using the p3net utility

- To connect a disconnected element to the network:

  ```
  p3net –-connect –n mainNE
  ```

- To display the names of all configured network elements:

  ```
  p3net –-show-all
  ```

- To display all configured network elements details (as a table):

  ```
  p3net –show-all --detail
  ```

- To display the connection status of the specified network element:

  ```
  p3net –-show –-ne-name=mainNE
  ```

p3radius Utility

The **p3radius** utility displays the statistics of the RADIUS Listener LEG. For information about this CLU, see the *Cisco SCMS SM RADIUS Listener LEG Reference Guide*.

p3rpc Utility

The **p3rpc** utility displays the information of the proprietary Cisco RPC (Remote Procedure Call) server interface to the SM. It also authenticates users.

The command format is: `p3rpc OPERATION [OPTIONS]`

The following tables list the **p3rpc** operations and options.

Table B-20 p3rpc Operations

Operation	Description
--show	Displays the port number to which the PRPC server listens, the maximum number of connections, the current number of active connections, and the host IP to which the server listens.
--show-client-names	Displays the names of the connected clients. Can be used for extracting the LEG_NAME key (see *LEG-Domains Association Section* (on page A-5)).
--show-statistics	Displays the PRPC server statistics. They contain information about the number of current PRPC sessions and statistics for PRPC server actions such as invocations and errors.
--reset-statistics	Clears the PRPC server statistics.
--set-user	Adds or updates the username and password.
--validate-password	Validates the username and password.
--delete-user	Deletes a user configuration.
--show-users	Displays all configured users.

Table B-21 p3rpc User Options

User Option	Abbreviation	Description
--username=USER-NAME	-u	Specifies the name of the user. Used with **--set-user**, **--validate-password**, and **--delete-user** operations.
--password=USER-PASSWORD	-P	Specifies the password of the user. Used with **--set-user**, **--validate-password**, and **--delete-user** operations.

Table B-22 p3rpc Miscellaneous Options

Option	Abbreviation	Description
--remote=IP[:port]	-r	(Optional) Used with **--set-user**, **--validate-password**, and **--delete-user** for users operations on the remote SM in High Availablity setups. The port option should be used if the PRPC Server port on the remote SM machine differs from the default value (14374).

Examples of using the p3rpc utility

- To display the port number to which the PRPC server listens, the maximum number of connections, the current number of active connections, the host IP to which the server listens, and the name of the configuration file used by the server:

  `p3rpc –-show`

- To display the statistics of the PRPC server:

  `p3rpc –-show-statistics`

- To clear the statistics of the PRPC server:

  `p3rpc –-reset-statistics`

- To show all the users configured at the PRPC server:

  `p3rpc –-show-users`

p3sm Utility

The **p3sm** utility performs general configuration and management of the SM.

The command format is: `p3sm OPERATION [OPTIONS]`

The following tables list the **p3sm** operations and options.

Table B-23 p3sm Operations

Operation	Description	Notes
--show	Displays the current SM configuration and statistics.	
--load-config	Reloads the SM configuration file. If the **–f** option is not used, file *p3sm.cfg* is loaded.	
--resync	Resynchronizes subscribers of specified SCE with the SM database. The SCE is specified using the option `--ne-name=SCE_NAME`.	Progress indicator
--resync-all	Resynchronizes all subscribers of all SCEs with the SM database.	Progress indicator
--start [--wait]	Starts the server. The option `--wait` causes the CLU to return only after the SM is up.	Default: started

Operation	Description	Notes
--stop	Stops the server. **Note:** When using fail-over, a simple shut-down of the SM does *not* work, the Veritas Cluster Server identifies that the SM is down and attempts to restart it. The correct procedure is: 1. Perform the manual fail-over. See *Subscriber Manager (SM) Fail-Over* (on page 2-11) 2. Use the Veritas Cluster Manager Application to stop the monitoring (probing) of the SM 3. Use the SM CLU (the `p3sm` command) to stop the SM	
--restart [--wait]	Stops the server operation and then restarts it. The option `--wait` causes the CLU to return only after the SM is up.	
--sm-version	Displays the currently installed SM version.	
--sm-status [--detail]	Displays the SM operational status: whether the SM is running or not, and whether it is Active or Standby. If errors have occurred, it also displays their summary. To receive a detailed description, use the option `--detail`.	
--extract-support-file	Retrieves the support file from the agent. This command extracts the SM support information to a defined file, which is defined using the option `--output=FILE`. SM support information should be extracted and sent to Cisco customer support with each support request.	
--reset-sm-status	Clears errors and warnings that were displayed to the user.	
--logging=[on/off]	Enables/disables the logging of user logon to the UserLog. **Note:** Enabling this flag may affect performance.	
--show-stats	Displays statistics information regarding logon operations. The rate results are updated once every 10 seconds.	
--reset-stats	Resets the statistics information.	

Table B-24 p3sm SM Options

SM Option	Abbreviation	Description
--ne-name=NAME	-n NAME	Specifies logical NAME of the SCE platform to resynchronize.

Table B-25 p3sm File Options

File Option	Abbreviation	Description
--output=FILE	-o FILE	Where to extract the support information file, relative to the SM root directory
--file=FILE	-f FILE	File to load the configuration from, relative to the SM configuration directory.

Table B-26 p3sm Miscellaneous Options

File Option	Abbreviation	Description
--ignore-warnings	-i	Ignore configuration validation warnings while loading the configuration file.
--remote=IP[:port]	-r	Used with **--load-config** to load the local configuration file to both the local SM and the remote SM.
--detail		Displays a detailed view of the SM status.
--wait		Used with **--start** or **--restart** to signal the CLU to return only when the SM is up.

Examples of using the p3sm utility

- To start the server:

  ```
  p3sm --start
  ```

- To stop the server:

  ```
  p3sm --stop
  ```

Note When using fail-over, a simple shut-down of the SM does **not** work, the Veritas Cluster Server identifies that the SM is down and attempts to restart it. The correct procedure is:
1. Perform the manual fail-over. See *Subscriber Manager (SM) Fail-Over* (on page 2-11).
2. Use the Veritas Cluster Manager Application to stop the monitoring (probing) of the SM.
3. Use the SM CLU (p3sm --stop) to stop the SM.

- To display the SM configuration:

  ```
  p3sm --show
  ```

- To resynchronize the subscribers of the specified SCE with the SM database:

  ```
  p3sm --resync --ne-name=my_SCE_100
  ```

- To stop the server operation and then restart it:

  ```
  p3sm --restart
  ```

- To reload the SM configuration file, *p3sm.cfg*:

  ```
  p3sm --load-config
  ```

- To display the SM operational status (active or inactive):

  ```
  p3sm --sm-status
  ```

- To extract the SM support information to the specified file:

  ```
  p3sm --extract-support-file --output=support.zip
  ```

p3subs Utility

The **p3subs** utility manipulates individual subscriber information in the SM database. Regarding properties, the property names depend on the application running on your system. To find descriptions of the application property names and values, see the documentation provided with the application installed on your system.

The command format is: `p3subs OPERATION [SUBSCRIBER-OPTIONS]`

The following tables list the **p3subs** operations and options.

Table B-27 p3subs Operations

Operation	Description
--add	Adds/updates a subscriber. The operation fails if the subscriber exists, unless the **--overwrite** option is used.
--set	Adds/updates mappings and/or properties for the specified subscriber. A new mapping overwrites all existing mappings, unless the **--additive-mapping** option is used. A property is overwritten only when a new value is assigned to it, but not when a *different* property has a new value assigned to it.
--show	Displays information for the specified subscriber.
--remove	Removes the specified subscriber.
--show-all-mappings	Displays all the mappings for the specified subscriber.
--remove-mappings	Removes the specified mapping of the specified subscriber.
--remove-all-mappings	Removes all the mappings of the specified subscriber.
--show-property	Displays the value of the specified property of the specified subscriber.
--show-all-properties	Displays the values of all the properties of the specified subscriber.
--show-all-property-names	Displays all the property names and descriptions.
--reset-property	Resets the specified property of the specified subscriber to its default value.
--remove-properties	Removes all properties and custom properties from the subscriber record.
--clear-state	Clears applicative state of specified subscriber. This command clears **only** the backup copy at the SM; it does **not** clear the applicative state record in the SCE platform.

Table B-28 p3subs Subscriber Options

Subscriber Option	Abbreviation	Description
--overwrite		Used in add operations to replace the existing subscriber configuration, instead of failing.
--subscriber=NAME	-s NAME	Performs operation using specified subscriber NAME.

Subscriber Option	Abbreviation	Description
--additive-mappings		Adds the new mapping(s) to any existing ones. (Without this option, any existing mappings are overwritten.)
--ip= IP1[/RANGE][,…]		Performs operation using specified IP mapping(s). IP is in dotted notation. "/RANGE" is used for specifying several consecutive mappings, by specifying the number of consecutive set bits in the mask. For example, **1.1.1.0/30** means **1.1.1.0 to 1.1.1.3**, or **1.1.1.0 with mask 255.255.255.252**.
--vlan-id=VLAN1[,…]		Performs operation using specified VLAN mapping(s).
--mpls-vpn=VPN-ID@PE-IP[...]		Performs operation using specified MPLS/VPN specification. The notation of the MPLS/VPN mappings is VPN-ID@PE-IP. VPN-ID is the VPN identifier of the VPN site (can be RT or RD). The PE-IP is the loopback IP address of the PE router attached to the VPN. For more information on the MPLS/VPN mappings, see the *Cisco SCMS SM MPLS/VPN BGP LEG Reference Guide*.
--property= KEY1[=VAL1][;...]	-p KEY1[=VAL1][;...]	Performs operation using the specified KEY=VAL property/properties. These properties are defined by the application and influence the subscriber service in the SCE.
--custom-property= KEY1[=VAL1][;...]		Performs operation using the specified KEY=VAL custom property/properties. These properties are user defined and have no influence on the service the subscriber receives.
--domain=DOMAIN	-d DOMAIN	Performs operation on specified DOMAIN. If DOMAIN is **none**, the operation refers to subscribers who have no domain specified.

Note When working with VLAN mapping types, the SCE must be configured using the following CLI:

```
SCE2000#> configure
SCE2000(config)#> in li 0
SCE2000(config if)#> VLAN symmetric classify
```

Examples of using the p3subs utility

- To add a subscriber with the specified IP address:

```
p3subs --add --subscriber=jerry --ip=96.142.12.7
```

- To overwrite subscriber information (because the subscriber **jerry** already exists, this operation would fail, but the **overwrite** option allows the IP address to be overwritten):

```
p3subs --add --subscriber=jerry --ip=96.128.128.42 --overwrite
```

- To set a property value for the specified subscriber:

```
p3subs --set --subscriber=jerry --property=packageId=1
```

- To add new mappings for the specified subscriber; any existing ones are overwritten:

```
p3subs --set --subscriber=jerry –-vlan-id=1
```

- To add new mappings to the existing ones for the specified subscriber:

```
p3subs --set --subscriber=jerry –-vlan-id=4,2 --additive-mappings
```

- To display information for the specified subscriber:

```
p3subs --show --subscriber=jerry
```

- To remove the specified subscriber:

```
p3subs --remove --subscriber=jerry
```

- To display all the mappings for the specified subscriber:

```
p3subs --show-all-mappings --subscriber=jerry
```

- To remove the specified mappings for the specified subscriber:

```
p3subs --remove-mappings --subscriber=jerry --ip=96.142.12.7,96.128.128.42
```

- To remove a range of consecutive mappings for the specified subscriber:

```
p3subs --remove-mappings --subscriber=jerry --ip=1.1.1.0/30
```

- To remove all the mappings for the specified subscriber:

```
p3subs --remove-all-mappings --subscriber=jerry
```

- To display the value of the specified property for the specified subscriber:

```
p3subs --show-property --subscriber=jerry –-property=reporting
```

- To display the values of all the properties for the specified subscriber:

```
p3subs --show-all-properties --subscriber=jerry
```

- To display all the property names and descriptions:

```
p3subs --show-all-property-names
```

- To reset specified property of specified subscriber to its default value:

```
p3subs –-reset-property --subscriber=jerry
--property=rdr.transaction.generate
```

- To clear the applicative state of the specified subscriber. This command clears only the backup copy at the SM, it does **not** clear the applicative state record in the SCE platform:

```
p3subs –-clear-state --subscriber=jerry
```

p3subsdb Utility

The **p3subsdb** utility manages the subscriber database and performs operations on groups of subscribers.

The command format is: `p3subsdb OPERATION [OPTIONS] [FILE-OPTIONS]`

The following tables list the **p3subsdb** operations and options.

Table B-29 p3subsdb Operations

Operation	Description	Notes
--clear-all	Removes all subscriber records from the SM database.	Progress indicator
--clear-domain	Removes all subscriber records from the specified domain.	Progress indicator
--show-num	Displays number of subscribers in database for the specified domain.	
--show-all	Lists all the subscriber names.	
--show-domain	Lists all the subscriber names in the specified domain.	
--import	Imports subscribers to the database from a specified CSV file. The filename that is to be imported is specified using the format "--file=import-filename". The results go to a result file, *import-results.txt*, which is created in the same directory as the CSV file.	Progress indicator
--export	Exports subscribers from the database to a specified CSV file. The filename that is to be exported is specified using the format "--output=export-filename". The results go to a result file, *export-results.txt*, which is created in the same directory as the CSV file.	Progress indicator
--clear-all-states	Clears the state of all subscribers in the SM database.	
--remove-property	Removes a specified property from all subscribers in the system. **Note:** After running this command you should re-synchronize all SCE devices.	
--remove-all-ip	Removes all the IP addresses of all subscribers.	
--remove-all-vlan	Removes all the VLAN tags of all subscribers.	
--remove-all-mpls-vpn	Removes all the MPLS/VPN mappings of all subscribers.	

Table B-30 p3subsdb Options

Option	Abbreviation	Description
--prefix=NAME		Used in the export operation for filtering the export.
--property=PROP		Used in removing of property PROP from all of the subscribers.

Option	Abbreviation	Description
--domain=DOMAIN	-d DOMAIN	Performs the operation on the specified DOMAIN. If DOMAIN is **none**, the operation refers to the subscribers who have no domain specified.

Table B-31 p3subsdb File Options

File Option	Abbreviation	Description
--file=FILE	-f FILE	Specifies the subscriber CSV FILEs to import from.
--output=FILE	-o FILE	Specifies the subscriber CSV FILE to export to.

Examples of using the p3subsdb utility

- To import subscribers from a specified CSV file:
  `p3subsdb --import --file=mySubscriberFile.csv`
- To export subscribers to a specified CSV file:
  `p3subsdb --export -o mySubscriberFile.csv`
- To export subscribers to a specified CSV file, using filtering options:
  `p3subsdb --export --prefix=a –-output=mySubscriberFile.csv`
- To export subscribers to a specified CSV file, using filtering options:
  `p3subsdb --export --prefix=a –o a.csv`
- To remove all subscriber records from the SM database:
  `p3subsdb –-clear-all`
- To remove all subscriber records from the specified domain:
  `p3subsdb –-clear-domain --domain=myDomain`
- To list all the subscribers:
  `p3subsdb --show-all`
- To list all subscribers in a specified domain:
  `p3subsdb --show-domain --domain=myDomain`
- To show the number of subscribers in a specified domain:
  `p3subsdb --show-num --domain=myDomain`
- To list all subscribers who have no domain specified:
  `p3subsdb –-show-domain --domain=none`
- To clear the state of all subscribers in the SM database:
  `p3subsdb –-clear-all-state`
- To remove a property from all subscriber records:
  `p3subsdb –-remove-property –-property=monitor`

APPENDIX C

CPE as Subscriber in Cable Environment

The cable market presents special issues in terms of subscribers, in addition to the normal subscriber management issues that exist in other markets, such as DSL and Wireless.

This appendix deals with the special case when the CPE is considered as the subscriber in the Cisco Service Control Solution for a cable environment. *This appendix is* **not** *relevant for the more common case where the cable modem with all the CPEs behind it is considered the subscriber.*

This appendix makes use of the following cable terms:

- Customer Premises Equipment (CPE)—The CPE is an individual computer. The cable support module considers the CPE as the subscriber.
- Cable Modem (CM)—There may be multiple CPEs connected through a single CM. Although the CPE is the subscriber, the policy is assigned to the CM, and the CPEs inherit the policy from the CM. All CPEs connected through the same CM always have the same policy.
- Cable Modem Termination System (CMTS)—The CMTS groups several thousand CMs and connects them to the network. Several CMTS devices can be served by the same SCE platform and thus belong to the same domain.
- DHCP Server—Assigns IP addresses and provides the login/logout event information to the system.

Cable Support Module

The Subscriber Manager (SM) includes a special cable support module (**p3cable**) for dealing with the special case where the CPE is considered as a subscriber in a cable environment. The cable support module functions as a translator between the cable world (DHCP events) and the Cisco SM, for this special case. It provides an API on top of the basic SM API functionality. This API is accessible using the Java/C/C++ APIs by calling the `cableLogin` and `cableLogout` methods.

To ensure the correct behavior of the cable support module, certain configuration steps are necessary, such as the correct domain configuration and the static/dynamic CM configuration.

The cable support module, which translates between the SM and the DHCP events in the cable world, performs the following functions:

- Associates between CPEs and CMs
- Makes CPEs inherit application policy from their CM
- Allows/denies the introduction of CPEs whose CM is unfamiliar to the SM

For additional information regarding the functions of the cable support module, see *p3cable Utility* (on page B-4).

The cable support module uses the hardware (MAC) addresses of the CM as its subscriber name. The subscriber name of the CPE is the hardware address of its CM followed by the hardware address of the CPE.

CM and CPE in the SM

In the special case when CPEs are considered as subscribers, cable modems are not delegated to the SCE in any way, and are not considered as subscribers in the Cisco Service Control Solution. However, for ease of integration and for the sake of simplicity, CMs are saved as subscribers in the SM only (but are **never** introduced to the SCE).

Cable modem SM subscribers are saved in special hidden subscriber domains called *CM domains*. These CM domains do not contain any SCE and are created automatically upon an insertion of a CM. For a CPE in a given subscriber domain, its CM will reside in a CM domain having the same name as the CPE domain but with the prefix **CM_**.

Because CM domains are hidden, they cannot be configured by the configuration file. However, it is possible to run subscriber-related commands (**p3subs** and **p3subsdb**) on these domains.

A CM subscriber name has the following form: <CM MAC> (the MAC of the CM as sent in the DHCP protocol).

A CPE subscriber name (for such a CM) has the following form: <CM MAC>__<CPE MAC> (the MAC of the CM, followed by **two** underscore signs, followed by the MAC of the CPE).

The **p3cable** command imports and exports cable modems, similar to importing and exporting subscribers, except that it is unnecessary to import the CM with an IP address.

Note When importing cable modems, the **full** CM domain name (**CM_** plus the domain name of its CPEs) **must** be provided.

EXAMPLE:

In the configuration of this example, the SM has a domain called DomainA. We want CPEs arriving from CMTS with IP 1.2.3.4 to reach this domain; therefore, we have configured 1.2.3.4 as an alias of DomainA.

During operation, because of a DHCP request-response, the DHCP LEG event sends a login event of a cable modem with MAC 0X0Y0Z from CMTS 1.2.3.4.

In the login event, the alias sent was 1.2.3.4 (the alias of domain DomainA), so the cable modem subscriber will be entered into domain CM_DomainA with the name 0X0Y0Z.

When a login event of its CPE with MAC 0A0B0C is sent with the same alias (as the CPE that arrived from the same CMTS), the CPE subscriber will be entered into domain DomainA with the name 0X0Y0Z__0A0B0C.

Static and Dynamic CMs

Login and logout events of CPEs whose CM does not exist in the subscriber database will be ignored, since no subscriber will be created in the SM and aggregated to the SCE. This CPE traffic will be treated as default subscriber.

The SM supports two modes of integrating with cable modems. Editing and loading the *p3sm.cfg* configuration file controls these modes. (Configuring dynamic CM support is described in *Configuration File Options* (on page A-1)). Use the CLU **p3cable** to view the current status.

- **Deny dynamic CM**—In this mode, login/logout events of cable modems that were **not** imported using the **p3cable** command will be ignored. Consequently, the CPE traffic of these CMs will be treated as default subscriber.
- **Allow Dynamic CM**—In this mode, login/logout events of cable modems that were **not** imported using the **p3cable** command will result in automatic addition of the cable modem to the subscriber database. These cable modems will receive the application tuneables that were defined in the domain tunable template section of the configuration file. For a description of application tuneables, see the *Cisco Service Control Application for Broadband User Guide*.

Troubleshooting

This appendix describes how to troubleshoot the SM installation and daily operation.

Using the Troubleshooting Chapter

Each entry in this appendix consists of an error message, probable cause(s), and solution. Note that the same error message may appear in more than one section of this appendix.

When an unexpected error occurs during the system's installation or daily operation, search for the error message throughout this appendix (the message may appear in more than one place). When you find the error message, read the section below the message and try the recommended solution. If the message appears more than once, try to correct the most probable cause first.

General Errors

SM Not Running

Error message:

The following sequence of output appears (in response to the command `p3sm --sm-status`):

```
> p3sm –-sm-status
Could not connect to SM.
```

Probable cause:

The SM server has not been started.

Solution:

Use the following command to start/restart the SM server:

```
> p3sm –-start
```

SM in Failure Mode

Error message:

The following sequence of output appears (in response to the command `p3sm --sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is Failure
Command terminated successfully
```

Probable cause:

The SM server restarted three times in 30 minutes due to an internal error.

Solution:

This error can happen only in a cluster setup. Check the pcube user log and the Veritas Cluster Server log for the reason for the failure that caused the reboots. Act according to the problem in the logs.

Additional operations that can be taken are:

- Use the following command to extract a support file:

  ```
  > p3sm –-extract-support-file –f ../support.zip
  ```

  Send the support file to Cisco's customer support (see *Technical Support* (on page xiii))
- Use the following command to start/restart the SM server and get out of the Failure state:

  ```
  > p3sm –-restart
  ```

General Setup Errors

install-sm.sh Script–User is not Root

Error message:

The following sequence of output appears (in response to the command `./install–sm.sh`):

```
> ./install-sm.sh
install-sm.sh: Starting SM installation sequence
install-sm.sh: Error - this script must be run by root - exiting.
```

Probable cause:

You started the installation sequence as **user** and not as **superuser**.

Solution:

Run the *`install-sm.sh`* script as **superuser**.

install-sm.sh Script–User pcube Exists

Error message:

The following sequence of output appears (in response to the command `./install-sm.sh`):

```
# ./install-sm.sh
install-sm.sh: Starting SM installation sequence
install-sm.sh: Error - pcube user exists and has home /export/home/pcube,
not /opt/pcube - remove it or use -o - exiting.
#
```

Probable cause:

Your machine already has the user **pcube**.

Solution:

Run the installation using the **–o** option (overwrite), as follows:

```
# ./install-sm.sh -o
```

install-tt.sh Script

Note A minimum of 1.5 GB of free hard disk space is required to install the TimesTen database.

Error message:

The following sequence of output appears (in response to the command `install-tt.sh/export/home/pcube/lib/tt`):

```
> install-tt.sh /export/home/pcube/lib/tt
install-tt.sh: Starting TimesTen P-Cube installation sequence
install-tt.sh: Error - This script must be run by root - exiting.
```

Probable cause:

You started the installation sequence as **user** and not as **superuser**.

Solution:

Run the *`install-tt.sh`* script as **superuser**.

Note that the TimesTen directory name given, */lib/tt/*, **must** be relative to the (pcube) user directory. For example, if the user directory is pcube, install TimesTen in */export/home/pcube/lib/tt/*.

install-dsn.sh Script

Note that the TimesTen directory name given, */var/tt/*, **must** be relative to the (pcube) user directory. For example, if the user directory is pcube, install TimesTen in */export/home/pcube/var/tt/*.

TimesTen Database Setup Errors

Introduction

The TimesTen configuration consists of several configuration files. This section explains the purpose and scope of each of these files. When troubleshooting the TimesTen, you will be requested to edit these configuration files and reboot the machine or restart the SM. In most cases, the defaults applied by the SM installation procedure are satisfactory.

Note Changing the TimesTen configuration files should be done with **extreme care**, and it is best to **consult Cisco technical support** prior to making any changes. See *Obtaining Technical Assistance* (on page xiii) for more information.

System (Kernel) Configuration File

The kernel configuration file is a system configuration file, which affects system-wide configuration parameters:

- For Solaris, it is file */etc/system*.
- For RedHat, it is file */etc/sysctl.conf*.

The Subscriber Manager installation procedure configures this file to add extra semaphores and shared memory to the system. After editing this file, you have to reboot the machine for the changes to take effect.

If you are running other applications that require changes in this file's semaphore and shared memory values, **take care** that the TimesTen configuration does not override the other application's configuration, or vice versa. **In particular**, when installing the Cisco Collection Manager (CM) and SM on the same machine, you should consult with the Cisco technical support for the proper values to use for the file configuration parameters.

Configuration File /var/TimesTen/sys.odbc.ini

The file */var/TimesTen/sys.odbc.ini* is a TimesTen configuration file that configures system DSNs. Any user on the machine on which the system DSN is defined can use this file. The SM DSNs are system DSNs that are named PCube_SM_Repository and PCube_SM_Local_Repository, and which have the following system DSN configuration parameters:

- LogFileSize—The size of the TimesTen log file, in megabytes.
- PermSize—The size of the permanent memory region for the data store, in megabytes. You may increase PermSize but not decrease it.

  The data stored in the permanent memory region includes tables and indexes that make up a TimesTen data store. The permanent data partition is written to the disk periodically.
- TempSize—The size of the memory allocated to the temporary region, in megabytes.

Temporary data includes locks, cursors, compiled commands, and other structures needed for command execution and query evaluation. The temporary data partition is created when a data store is loaded into memory and is destroyed when the data store is unloaded.

For additional information, see the recommended database parameter configuration table in RAM and Configuration Parameters Versus Number of Subscribers.

- SMPOptLevel—Optimizes the database operation on multi-processor machines. If the machine is a multi-processor platform, set parameter SMPOptLevel to 1 (default is 0).

Configuration File ~pcube/.odbc.ini

The file *~pcube/.odbc.ini* is a TimesTen configuration file that configures user DSNs.

TimesTen DSN Configuration—Cannot Find Requested DSN

Error message:

The following sequence of output appears (in response to the command `p3sm –-sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 CLIENT]Cannot find the requested DSN
(PCube_SM_Repository_CS) in ODBCINI
```

Probable cause:

The TimesTen Client DSN is not configured correctly in file ~pcube/.odbc.ini.

Solution:

Ensure that file *~pcube/.odbc.ini* contains the following:

```
[ODBC Data Sources]
PCube_SM_Repository_CS=TimesTen 5.0 Client Driver
```

TimesTen DSN Configuration—Data Source Name Not Found

Error message:

The following sequence of output appears (in response to the command `p3sm –-sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 ODBC Driver]Data source name not found and no
default driver specified
```

Probable cause:

The TimesTen Client DSN is not configured correctly in file ~pcube/.odbc.ini.

Solution:

Ensure that file *~pcube/.odbc.ini* contains the following:

```
[PCube_SM_Repository_CS]
TTC_SERVER_DSN=PCube_SM_Repository
```

Ensure that file */var/TimesTen/sys.odbc.ini* contains the following:

```
[ODBC Data Sources]
PCube_SM_Repository=TimesTen 5.0 Driver
[PCube_SM_Repository]
Driver=__TTDIR__/TimesTen/pcubesm22/lib/libtten.so
DataStore=__VARDIR__/pcube_SM_Repository
```

TimesTen Database Settings—Cannot Connect to Data Source

Error message:

The following sequence of output appears (in response to the command `p3sm --sm-status`):

```
> p3sm --sm-status
SM is running.
SM operational state is XXX
Error - Times-Ten DB is not set up correctly:
[TimesTen][TimesTen 5.0.35 CLIENT]Unable to connect to data source
(DSN: pcube_SM_Repository_CS; Network Address: X.X.X.X; Port Number: XXX):
This operation has Timed Out. Try increasing your ODBC timeout attribute
or check to make sure the target TimesTen Server is running
```

Probable cause:

The following causes are possible:

- The address of the Server DS is incorrect.
- The port of the Server DS is incorrect.
- TimesTen is not active.

Solution:

The Service Control solutions for the above causes are:

- (The address of the Server DS is incorrect.) Ensure that file *~pcube/.odbc.ini* contains the following:

  ```
  TTC_SERVER=127.0.0.1
  ```

- (The port of the Server DS is incorrect.) On a default installation, ensure that file *~pcube/.odbc.ini* does *not* contain *TTC_SERVER_PORT*.

  On a non-default installation, ensure that file *~pcube/.odbc.ini does* contain *TTC_SERVER_PORT=Non-default-port*.

- (TimesTen is not active.) Run the following command:

```
~pcube/lib/tt/TimesTen/pcubesm22/bin/ttStatus
```

If TimesTen is not working, re-install TimesTen.

If the above solutions do not work, please refer to the TimesTen manual.

TimesTen Configuration Error—Not Enough Memory

Error message:

The following sequence of output appears (in response to the command `p3sm –-sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 ODBC Driver][TimesTen]TT0836:
Cannot create data store shared-memory segment,
error 1455 -- file "db.c", lineno 6289, procedure "sbDbConnect()"
```

Probable cause:

There is not enough memory for creating TimesTen's in-memory database.

Solution:

Do **all** of the following:

- Ensure that the Unix machine has at least 1024 MB of memory installed.
- Ensure that the configured memory size parameters (*PermSize and TempSize*) specified in file */var/TimesTen/sys.odbc.ini* are less than the total amount of memory installed in the machine.
- For **Solaris**, ensure that the maximum shared memory (parameter *shmsys:shminfo_shmmax*) specified in file */etc/system* is less than the total amount of memory installed in the machine.
- For **Red Hat**, ensure that the maximum shared memory (parameter *kernel.shmmax*) specified in file */etc/sysctl.conf* is less than the total amount of memory installed in the Linux Machine.

TimesTen Configuration Error—Incorrect Memory Definitions

Error message:

The following sequence of output appears (in response to the command `p3sm --sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX.
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 ODBC Driver]Overflow in converting data store or
log file size from megabytes to bytes,
or in converting log buffer size from kilobytes to bytes
```

Probable cause:

The memory definitions of DSN are incorrect.

Solution:

Ensure that the configured permanent memory size and log file size (parameters *PermSize and LogSize*) are less than the total amount of memory and of disk space specified in file */var/TimesTen/sys.odbc.ini*.

TimesTen Configuration Error—Cannot Create Semaphores

Error message:

The following sequence of output appears (in response to the command `p3sm –-sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 ODBC Driver][TimesTen]TT0925: Cannot create data
store semaphores (Invalid argument) –
file "db.c", lineno 5124, procedure "sbDbCreate()", sqlState: 08001,
errorCode: 925
```

Probable cause:

TimesTen was unable to create the data store semaphores that are defined in the kernel configuration file (**Solaris**: */etc/system* for Solaris; for **Red Hat**: */etc/sysctl.conf* for Red Hat).

Solution:

Do **all** of the following:

- Ensure that the machine has at least 1024 MB of memory installed.
- Reboot the machine after the first time that TimesTen is installed.
- Verify the contents of the system (kernel) configuration file:
  - For **Solaris**, ensure that file */etc/system* contains the following:

```
semsys:seminfo_semmni = 20
semsys:seminfo_semmsl = 100
semsys:seminfo_semmns = 2000
semsys:seminfo_semmnu = 2000
```

  - For **Red Hat**, ensure that file */etc/sysctl.conf* contains the following:

```
*---- Begin settings for TimesTen
kernel.sem = “SEMMSL_250 SEMMNS_32000 SEMOPM_100 SEMMNI_100
*---- End of settings for TimesTen
```

TimesTen Configuration Error—Cannot Read Data Store File

Error message:

The following sequence of output appears (in response to the command `p3sm –-sm-status`):

```
> p3sm –-sm-status
SM is running.
SM operational state is XXX.
Error - Times-Ten DB is not setup correctly:
[TimesTen][TimesTen 5.0.35 ODBC Driver][TimesTen]TT0845: Cannot read data
store file.
OS-detected error: Error 0 -- file "db.c", lineno 6320, procedure
"sbDbConnect()"
```

Probable cause:

TimesTen was unable to read the data store file, probably due to an error during the installation. This error occurs when installing a TimesTen application on top of an existing TimesTen, without having first uninstalled the old TimesTen database.

Solution:

Perform the following procedure:

Step 1 Remove the database, by using the SM p3db CLU with the following commands:

```
>p3db –destroy-rep-db
>p3db –destroy-local-db
```

Step 2 Uninstall TimesTen with the following commands:

```
> su
Password:
# ~pcube/lib/tt/TimesTen4.5/32/bin/setup.sh -uninstall
```

Step 3 Re-install TimesTen, either by running the SM *`install-tt.sh`* script or by using the installation files supplied by TimesTen.

TimesTen Configuration Error—Data Store Space Exhausted

Error message:

The following sequence of output appears in the SM log (while using the SM APIs):

```
java.io.IOException: Failure in putting subscriber
45977166__00:50:bf:97:c1:b2 : [TimesTen][TimesTen 5.0.35 ODBC
Driver][TimesTen]TT0802: Data store space exhausted -- file "blk.c", lineno
1571, procedure "sbBlkAlloc"
```

Probable cause:

The TimesTen database has already reached its maximum capacity, which caused the operation of adding a new subscriber to the database to fail.

Solution:

Usually, doing just one of the following is sufficient:

- Reduce the number of the subscribers handled by the SM (of course, this solution is not always possible).
- Configure the system to support a larger number of subscribers. Note that this solution may require editing one or more of the TimesTen configuration files discussed in *Introduction* (on page D-4), as well as rebooting the machine.
- Move the SM to a more powerful machine; this could be a faster CPU (or more CPUs), a larger disk, more RAM, etc.

For help and guidance in implementing the last two solutions, please contact Cisco *Technical Support* (on page xiii).

Network Management Command Line Utility (p3net) Errors

First Connection—Operation Timed Out

Error message:

The following sequence of output appears (in response to the command `p3net --connect`):

```
> p3net --connect --ne-name=YYYY
Error - failed to connect to element 'YYYY'
Operation timed out: connect
```

Probable cause:

The following causes are possible:

- The IP address is incorrect.
- The element YYYY is down.

Solution:

The Service Control solutions for the above causes are:

- Ensure that the IP address is correct.
- Ensure that the element YYYY is online and is connected via its management port.

Status Error—Connection Down

Error message:

The following sequence of output appears (in response to the command `p3net --show-ne`):

```
>p3net --show --ne-name=se0
Network Element Information:
============================
Name:        YYY
Description: testing element
Host:        X.X.X.X
Ip:          X.X.X.X
Port:        14374
Status:      Connection down (Failure in connecting to agent on host,
Connection refused: connect)
Type:        SCE1000
Domain:        smartNET.policy.unitTestSubscribers
Subscriber Management: Not Active
```

Probable cause:

The following causes are possible:

- The IP address is incorrect.
- The element YYY is down.

Solution:

The Service Control solutions for the above causes are:

- Ensure that the IP address is correct.
- Ensure that the element YYY is online and is connected via its management port.
- Ensure that the PRPC adapter is online on the port that the status indicates.

Status Error—Subscriber Management Down

Error message:

The following sequence of output appears (in response to the command `p3net --show-ne`):

```
>p3net --show --ne-name=se0
Network Element Information:
============================
Name:        YYY
Description: testing element
Host:        X.X.X.X
Ip:          X.X.X.X
Port:        14374
Status:      Connection ready
Type:        SCE1000
Domain:        smartNET.policy.unitTestSubscribers
Subscriber Management: Not Active
```

Probable cause:

The Subscriber Management field indicates whether the SM successfully performed SM-SCE subscriber synchronization. If the value of the field is *Not Active*, it is possible that the SM failed to synchronize the SCE.

Solution:

One possible solution is to force SM-SCE resynchronization by using the CLU command `p3sm --resync`.

Subscriber Database Command Line Utility (p3subsdb) Errors

CSV File Validation Error

Error message:

The following sequence of output appears (in response to the command p3subsdb `--import`):

```
> p3subsdb --import --file=/export/home/pcube/XXX.csv
Error - Failed to validate the file XXX.csv

See import-results.txt for detailed errors description.

> cat import-results.txt
x.csv:1: expected 2 items but got 4 items.
1 subscribers, 1 errors.
NO APPLICATION INSTALLED, MAKE SURE TO INSTALL PQI BEFORE IMPORTING CM WITH
TUNEABLES.
```

Probable cause:

You tried to import a four-field *csv* file to the SM, but no application (*SCA BB*) was installed.

For example, the following *csv* file for a *SCA BB* application contains four fields:

```
# CSV line format: subscriber-id, domain, mappings, package-id
JerryS,subscribers,80.179.152.159,0
ElainB,subscribers,194.90.12.2,3
```

However, the default definition file that defines *csv* file parsing rules contains only two fields: **name** and **ip mapping**.

Solution:

Do one of the following:

- Install an application (*SCA BB*) on the SM (for details, see the *Cisco SCA BB User Guide*).
- Import a *csv* file that has just two fields.

Cable Support Command Line Utility (p3cable) Errors

CSV File Import Error

Error message:

The following sequence of output appears (in response to command `p3cable –-import-cm`):

```
> p3cable –-import-cm --file=/export/home/pcube/XXX.csv
Importing cable modems ... 0%
Importing cable modems ... 100%
Error - Errors during import from
'H:\work\Mng\dev\install\ems\bin\win32\x.csv':
Imported 1 CM(s).  1 Error(s).
See cm-import-results.txt for detailed errors description.

> cat cm-import-results.txt
x.csv:1: expected 2 items but got 4 items.
1 cable modem(s); 1 error(s).
NO APPLICATION INSTALLED, MAKE SURE TO INSTALL PQI BEFORE IMPORTING CM WITH
TUNEABLES.
```

Probable cause:

You tried to import a four-field *csv* file to the SM, but no application (*SCA BB*) was installed.

For example, the following *csv* file for a *SCA BB* application contains four fields:

```
# CSV line format: subscriber-id, domain, mappings, package-id
JerryS,subscribers,80.179.152.159,0
ElainB,subscribers,194.90.12.2,3
```

However, the default definition file that defines *csv* file parsing rules contains only two fields: **name** and **ip mapping**.

Solution:

Perform one of the following:

- Install an application (*SCA BB*) on the SM (for details, see the *Cisco SCA BB User Guide*).
- Import a *csv* file that has just two fields.

Configuration Errors

The errors in the following sections may appear after the CLU command `p3sm –-load-config` or in a user log during the SM startup procedure.

Network Management Errors

Error message (1):

```
Error section  <section name>: cannot contain white spaces.
```

Probable cause:

[SCE.XXX] section cannot contain white spaces (SCE name cannot contain white space).

Solution:

Remove the white spaces.

Error message (2):

```
Error in section <section name>: host  <ip address> already exists in
section <section name>
```

Probable cause:

Configuration cannot contain two SCEs with the same IP address.

Solution:

Change the IP address of one of the SCEs.

Error message (3):

```
Unknown NE <name> found in domain <domain name> section: it does not have
[SCE.<name>] section
```

Probable cause:

The section <domain name> includes, under the `elements` property, an SCE that is not defined in an [SCE.xxx] section.

Solution:

Add the missing [SCE.xxx] section to the file.

Error message (4):

```
Duplicate NE <name> found in domain <domain name> section: it already
appears in < domain name > domain section.
```

Probable cause:

Same SCE cannot belong to more than one domain.

Solution:

Remove the SCE from all but one of the domains.

Domain Errors

Error message (1):

```
Error section <section name>: cannot contain white spaces.
```

Probable cause:

[Domain.XXX] section cannot contain white spaces (domain name cannot contain white spaces).

Solution:

Remove the white spaces.

Error message (2):

```
Error  <alias name> value - alias name should not start with ‘CM_’.
```

Probable cause:

The alias name cannot start with `CM_` because this is the prefix of hidden domain generated by the SM when working with CMs (see *CPE as Subscriber in Cable Environment* (on page C-1)).

Solution:

Use a different prefix for the alias.

Error message (3):

```
Alias <alias name> already exists in [<domain name>] section
```

Probable cause:

Each alias can appear in only one [Domain.XXX] section.

Solution:

Alias mentioned in the error message should be removed from all but one [Domain.XXX] section.

Error message (4):

```
Unknown domain <domain name> found in [LEG-Domains Association]. It does not
appear as a section.
```

Probable cause:

Domain mentioned in given section does not have a [Domain.XXX] section.

Solution:

Domain mentioned in error message should be given a [Domain.XXX] section.

Error message (5):

```
Invalid non-integer value: <value> for property '<property name>' in section
[section name].
```

Probable cause:

Properties in [Domain.XXX] sections do not have integer values.

Solution:

Properties mentioned in error message should be given integer values.

Error message (6):

```
Error in section <domain name>: Property - <name> not found: <list of
application properties>
```

Probable cause:

Property defined in [Domain.XXX] section is not found in properties list defined by installed application.

Solution:

Delete the properties mentioned in the error message (or define them in the installed application).

Error message (7):

```
New configuration was not applied due to the following warnings:
Warning - Cannot remove domain <domain name> with <num of SCEs> SCEs. Note
that all subscribers will be removed from domain db.
Please use '--ignore-warnings' option to complete the operation.
```

Probable cause:

This is actually a warning: its purpose is to warn that the user removed from the p3sm.cfg file a domain that contained SCEs (which probably contained subscribers, etc.), and that the user will lose all subscriber data relevant for that domain. This warning appears only after the CLU command `p3sm –-load-config` is activated.

Solution:

To avoid this warning, use the `–-ignore-warnings` option.

PRPC Errors

Error message:

```
New configuration was not applied due to the following warnings:
Warning - PRPC configuration was changed. Note: Reloading may take up to 5
seconds.
Please use '--ignore-warnings' option to complete the operation.
```

Probable cause:

This is actually a warning: it is displayed after the CLU command `p3sm –load-config` is activated when the PRPC configuration in the p3sm.cfg file has been changed.

Solution:

Use the `–-ignore-warnings` option to complete the operation.

RADIUS Listener Errors

Error message:

```
Duplicate NAS identifier <nasID> found in section [NAS name] : already
exists in <other NAS name>
```

Probable cause:

<nasID> is not unique.

Solution:

Change one of the <nasID> so that both are unique.

Common Validation Errors

The following configuration errors are relevant for all sections/parameters of the *`p3sm.cfg`* file.

Error message (1):

```
Unknown property <property name> found in section [<section name>] in
configuration file <file name>.
```

Probable cause:

Property written in the *p3sm.cfg* file is unknown to the SM. Maybe the name is misspelled or the property belongs in a different section.

Solution:

Ensure that the name is spelled correctly and that the property resides in the correct section.

Error message (2):

```
Unknown section [<section name>] found in configuration file {2}.
```

Probable cause:

The section written in the p3sm.cfg file is unknown to the SM. Maybe the name is misspelled.

Solution:

Ensure that the section name is spelled correctly.

Error message (3):

```
Error value <value> for property <property name> in section [<section
name>]. Optional values: [<values range>]
```

Probable cause:

Value of the property is invalid. The <values range> field contains the valid values.

Solution:

Specify any valid value for the property.

Error message (4):

```
Missing mandatory property <property name> in section [<section name>].
```

Probable cause:

The property <property name> is **mandatory** and must appear in the section <section name>.

Solution:

Set a value for the requested property in the specified section.

Error message (5):

```
Error value <property value> for property <property name> in [<section
name>] section. Valid format: [0..255].[0..255].[0..255].[0..255]
```

Probable cause:

The value is an invalid IP address.

Solution:

Specify a valid IP address.

Error message (6):

```
Error empty value for <property name> property in [<section name>] section -
must have at least one character
```

Probable cause:

Value of the <property name> is empty; for example, `prop=`

Solution:

Specify a non-empty value for the property.

Error message (7):

```
Section <section name> added when already exists
```

Probable cause:

The section with <section name> appears more than once. This error is most likely to occur for the [SCE.XXX] and [Domain.XXX] sections.

Solution:

Use the specified section name only once.

Veritas Cluster Server Requirements and Configuration

This appendix provides basic guidelines for the Veritas Cluster Server (hereafter VCS) configuration in an SM cluster installation. It assumes basic knowledge of the VCS environment; it does not come in place of the VCS user guide. This appendix does not cover installation of the cluster or of the SM cluster agents.

This appendix lists the software and hardware system requirements for the VCS. It also gives a systematic explanation on how to configure the SM cluster using the VCS configuration tools. Most of the examples are taken from the use of the Java Veritas Manager GUI, although the operations can also be done via the Veritas command-line utilities.

SM 3.0.3 supports Veritas Cluster Server version 3.5 and 4.1 on Solaris machines and Veritas Cluster Server version 2.2 and 4.1 on Red-Hat Linux machines.

Veritas Software was acquired by Symantec. Currently, the cluster solution is still called the Veritas Cluster Server.

Veritas Cluster Server System Requirements

For your convenience, the following Veritas Cluster Server System Requirements have been taken from the Veritas site:

http://www.veritas.com

- Supported Platforms:

  Sun Solaris 8, 9, 10

  Red Hat Linux 3.0, 4.0

- Networking:

  Public Network: 10 MB/100 MB/Gigabit Ethernet

  Private Network: 10 MB/100 MB/Gigabit Ethernet

- Ethernet Controllers:

  Requires at least three independent Ethernet connections per system

- Memory:

  Each VERITAS Cluster Server system requires at least 128 MB of RAM (256 MB of RAM is recommended)

- Supported Server Hardware:

  Please refer to http://support.veritas.com or contact your VERITAS sales representative for the latest list of certified server hardware.

  - Sun Solaris 8, 9, 10
  - Red Hat Linux 2.1, 3.0, 4.0

- Supported Storage Hardware:

  Please refer to http://support.veritas.com or contact your VERITAS sales representative for the latest list of certified storage hardware.

  - Sun
  - Red Hat Linux

Veritas Cluster Server Nodes on Remote Sites

The heartbeat links use the Low Latency Transport (LLT) Ethernet/dlpi protocol. It uses Ethernet broadcasts and must be on the same broadcast network. A separate layer-2 switch per heartbeat link is supported. The distance limitation is based on performance. A number of factors govern cluster distance. The primary factors for LLT are network connectivity and latency. Direct L2 low latency connections must be provided for LLT with a maximum round-trip time of 500 milliseconds. Large campus clusters or metropolitan area clusters must be very carefully designed to provide two completely separate paths for heartbeat to prevent a single fiber optic or fiber bundle failure from removing the heartbeat links.

Although the database replication network uses IP as its transport, it must have two separate connection paths between the nodes that provide at least 10Mbps for the subscriber data replication.

When planning an SM Cluster setup where the nodes are at a distance from each other, please consult the Veritas support.

Replication Configuration Guidelines

Replication Scheme Setup

After the replication network has been setup (as described in the *Replication Network Configuration* (on page E-3) section), the replication scheme needs to be set to the database. Setting the replication scheme is performed using the p3db CLU:

```
> p3db --set-rep-scheme
```

This operation configures the database to send every subscriber-data update to the peer machine.

If the setup operation fails because there might be an existing replication scheme already set, run the following CLU to drop the previous replication scheme and then set the new scheme:

```
> p3db --drop-rep-scheme
```

By configuring and running the VCS agent of the replication agent or by running the p3db **–rep-start** CLU, the replication agent starts its work.

Replication Network Configuration

The configuration of the replication private network between the two SMs of the cluster must be carefully planned. This section discusses some of the guidelines for performing the configuration.

The TimesTen replication agent uses hostnames to implement fail-over between the two replication NICs. The agent uses the first IP address of the hostname supplied to the agent to connect to the other agent. If the connection fails and cannot be reconstructed on the first IP, the replication agent tries the next IP addresses assigned to this hostname, and so on.

Editing the /etc/hosts file to assign hostnames to IP addresses.

You **must** use the predefined hostnames *SM_REP1* and *SM_REP2* as the hostnames for replication.

Note Verify in the OS configuration files that the /etc/hosts file is used before using a name server.

EXAMPLE:

The following figure shows an example of a replication network.

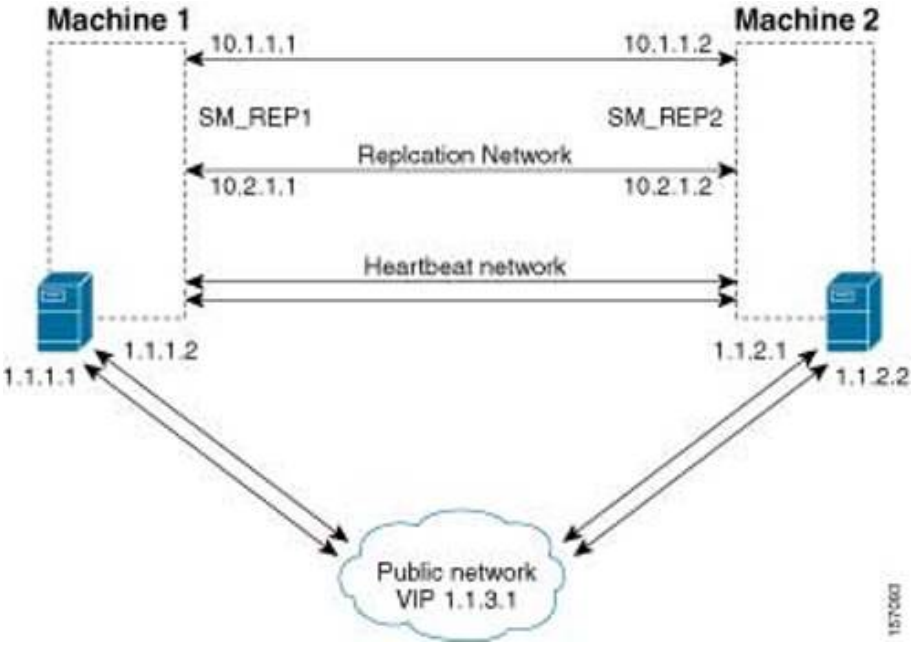

To configure the replication network shown by the above figure, do the following:

- Configure the IP addresses of each of the replication NICs, each one in a different network. In this example, the IP addresses of the Machine1 replication NICs are 10.1.1.1 and 10.2.1.1.
- Assign a hostname *SM_REP1* to both of the local replication NIC IP addresses. In this example, the hostname *SM_REP1* is assigned to the IP addresses of the replication NICs on Machine1. In the */etc/hosts* file, be sure to also assign the local hostname (Machine1) to the local replication NICs. Ensure that there are no empty lines between the lines containing the local hostname.
- Assign a hostname *SM_REP2* to both of the remote replication NIC IP addresses. In this example, the hostname *SM_REP2* is assigned to the IP addresses of the replication NICs of Machine2
- The /etc/hosts file on Machine1 should appear as follows:

```
127.0.0.1   localhost
1.1.1.1     Machine1    loghost
10.1.1.1    Machine1    SM_REP1    REP_1_NIC_1
10.2.1.1    Machine1    SM_REP1    REP_1_NIC_2
10.1.1.2    SM_REP2       REP_2_NIC_1
10.2.1.2    SM_REP2       REP_2_NIC_2
```

- The /etc/hosts file on Machine2 should appear as follows:

```
127.0.0.1   localhost
1.1.2.1     Machine2    loghost
10.1.1.2    Machine2    SM_REP2       REP_2_NIC_1
10.2.1.2    Machine2    SM_REP2       REP_2_NIC_2
10.1.1.1    SM_REP1    REP_1_NIC_1
10.2.1.1    SM_REP1    REP_1_NIC_2
```

Note Make sure the system uses the /etc/hosts file before it performs DNS or any other name service operation.

Veritas Cluster Server Configuration Guidelines

The sections below assume that the following operations were performed before starting the VCS configuration:

- Installation of the VCS on both machines. As part of this installation:
  - Each machine was given a hostname, which is used as the *system* name for the VCS configuration.
  - The machines are configured to recognize each other's hostname.
- An IP address was allocated for the cluster (hereafter, *the cluster's IP*).
- The SM and the TimesTen database were installed on both machines.
- The SM VCS agents were installed on each machine.
- The VCS manager Java console was installed on the administrator PC.

Note that in an SM cluster, two SM machines are connected to each other in a fully redundant way. The connection uses four cables: two for the VCS heartbeat mechanism, and two for the TimesTen database replication mechanism. Each machine is connected to the network via one of two redundant NICs. To access the cluster, you should use the *cluster IP address*, which is a virtual IP managed by the VCS. For management operations, you should use the local IP address of each machine.

SM Cluster Resources Configuration

To configure the VCS with the SM cluster resources, perform the procedures described in the following sections

Adding Clusters

To add a cluster:

Step 1 Open the VCS cluster manager Java console by using START | PROGRAMS | VERITAS CLUSTER MANAGER | CLUSTER MANAGER (JAVA CONSOLE).

Step 2 Add a new cluster by selecting from the menu FILE | NEW CLUSTER.

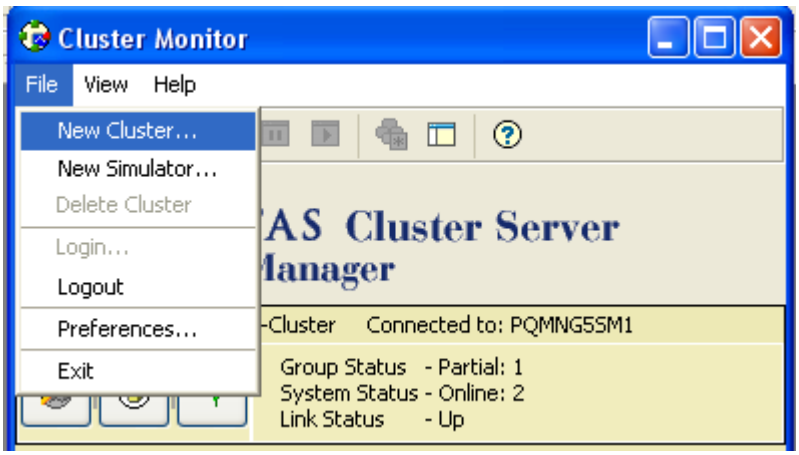

Step 3 Configure the cluster (see following figure).

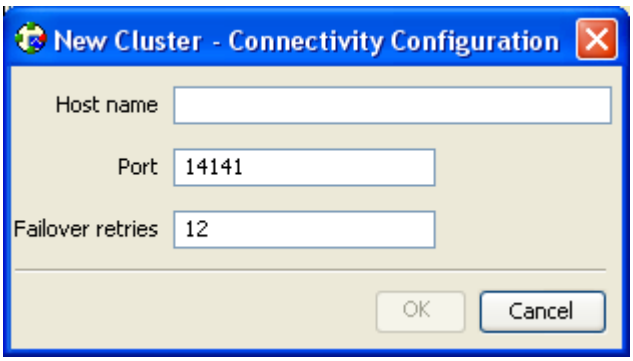

- Cluster Alias—cluster name.
- Host Name—one of the machine's IP addresses or hostname.

Step 4 Log in to the cluster.

Adding Services Groups

To add a service group to the cluster:

Step 1 In the cluster explorer, from the service group tab, right-click the cluster, and select ADD SERVICE GROUP.

Step 2 The Add Service Group window appears.

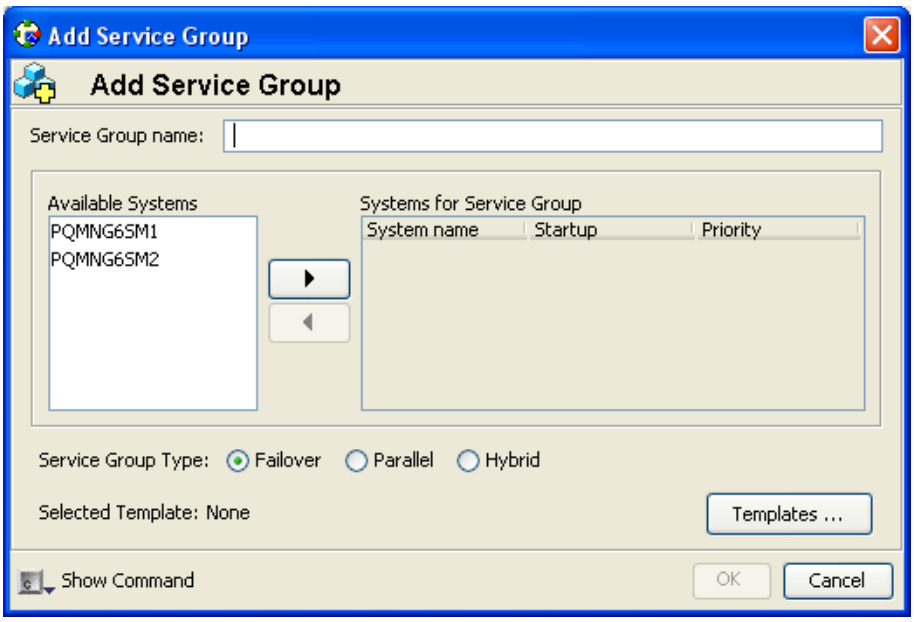

Step 3 Enter a name for the service group.

Step 4 Add the two machines as part of the service group and define their priority in the cluster.

Step 5 Click OK.

Setting the Auto-start

This section describes how to set the auto-start parameters that define which machine will start after a boot of both nodes. If these parameters are not set, then at boot of both nodes the cluster will stay offline.

To set the auto-start parameters:

Step 1 From the service group display, click SHOW ALL ATTRIBUTES.

Step 2 Make sure that both nodes are defined in the AutoStartList (see the following figure).

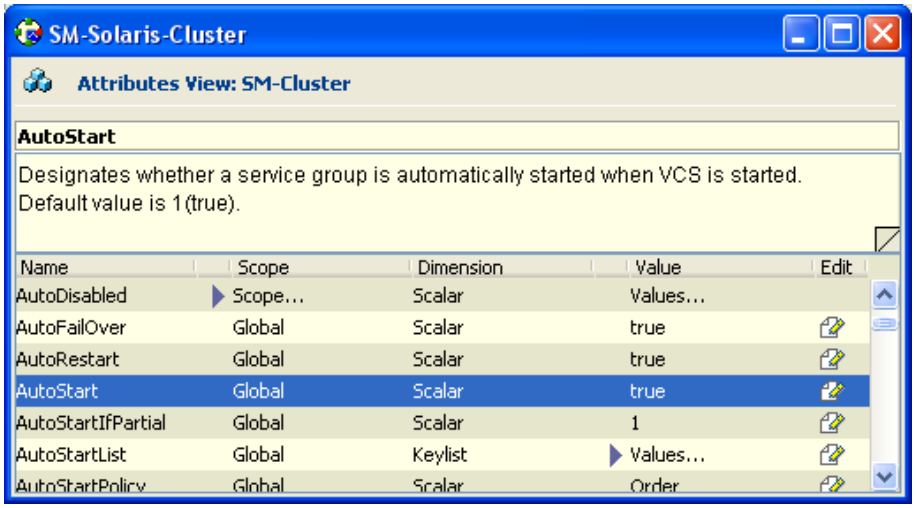

Step 3 Specify which node will start by defining the *AutoStartPolicy* parameter.

Step 4 Make sure that the *AutoStart* parameter is set to `true` (if `false`, both nodes will come up as standby).

Adding SM Cluster Resources

This section describes how to add the various SM cluster resources.

Adding Resources - General Guidelines

To add a resource (general guidelines):

Step 1 From the right-click menu of the service group, select ADD RESOURCE.

Step 2 When the following screen is displayed, select the RESOURCE TYPE, and give the resource a name.

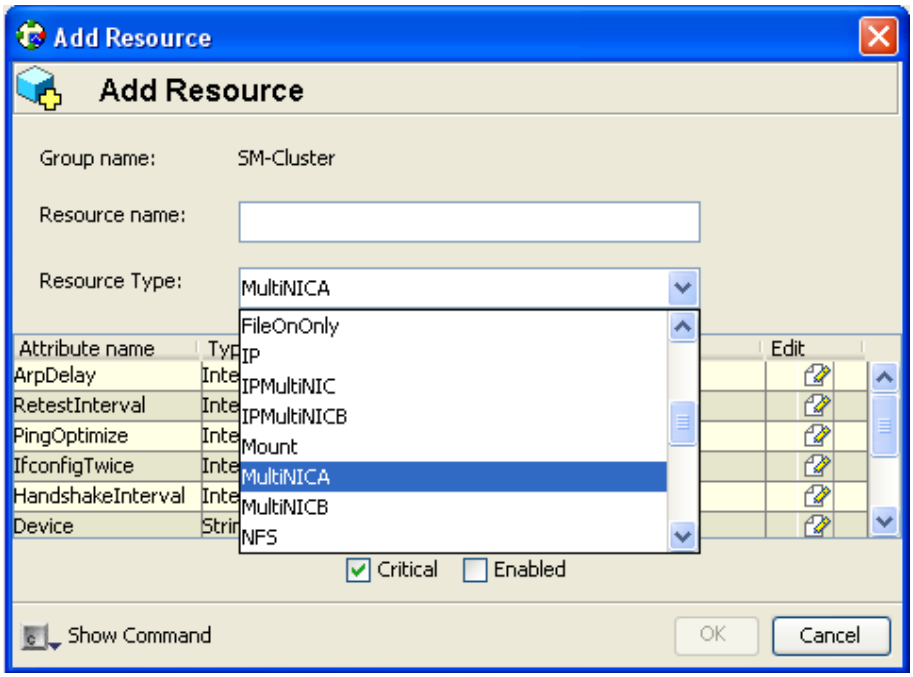

Step 3 Scroll through the attributes, and configure the ones that you need.

Step 4 When you are done, click OK.

Adding Network NICs

To add a Network NIC:

Step 1 Decide which two network interfaces to use for the network connection.

Step 2 Add a MultiNICA resource called *Network-NICs* to the service group.

Step 3 Define the following parameter:

- Device—Write the names of the Network NICs in the KEY column and their corresponding IP addresses in the VALUE column.

  Note: Use the LOCAL option and configure each machine separately, because the IP addresses are different in each machine.

  In the following example, `bge0` and `bge3` are the network NICs.

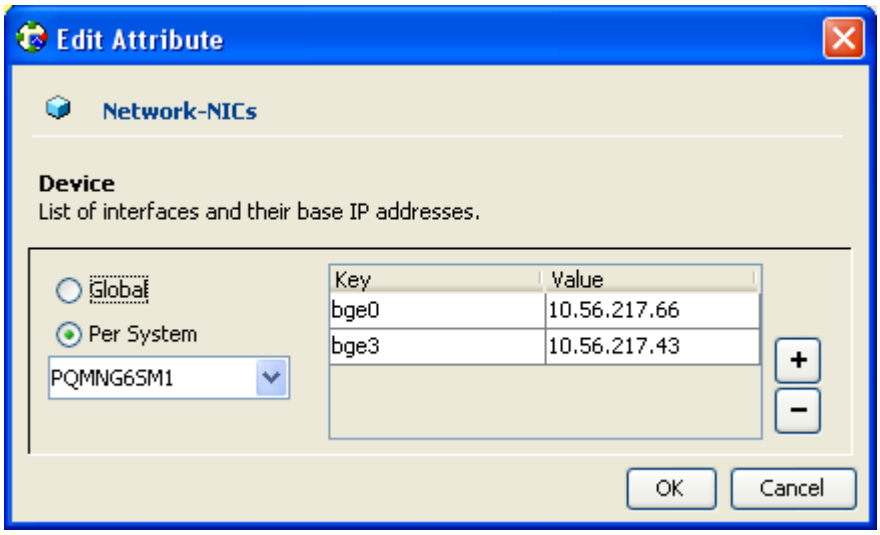

- NetMask—Assign a relevant network mask. For example, 255.255.255.255 can be defined as a network mask.

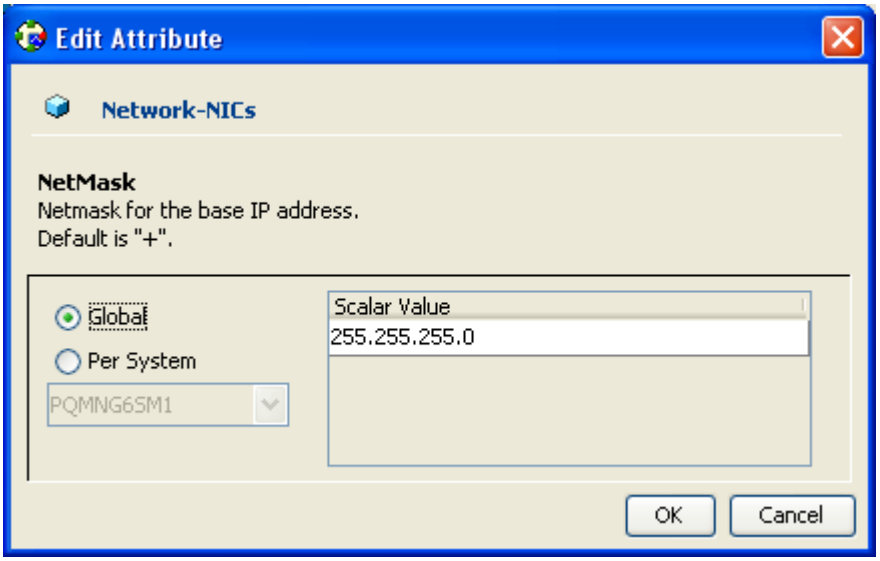

Step 4 Check the ENABLED and CRITICAL attributes of the resource (see the following figure):

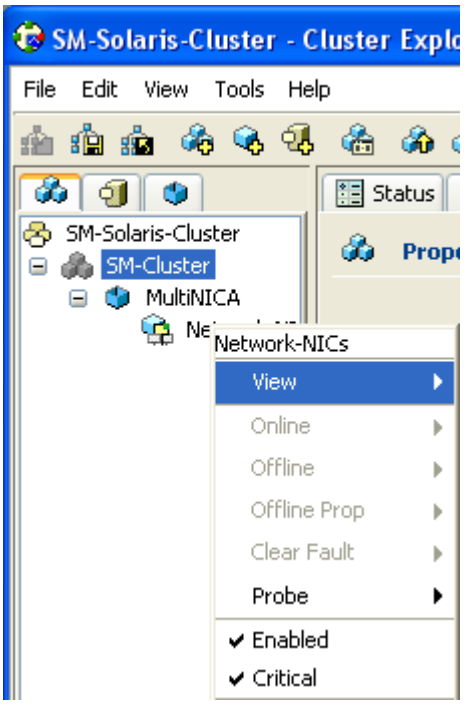

Adding Network VIPs

To add a Network VIP:

Step 1 Decide on the IP address of the cluster.

Step 2 Add to the service group an IPMultiNIC resource called *Network-VIP*.

Step 3 Define the following parameters:

- Address—Type the Cluster IP address.
- Net-mask—Type the network-mask you want to use for this IP.
- MultiNICAResName—Type *Network-NICs* to specify the relevant NICs.

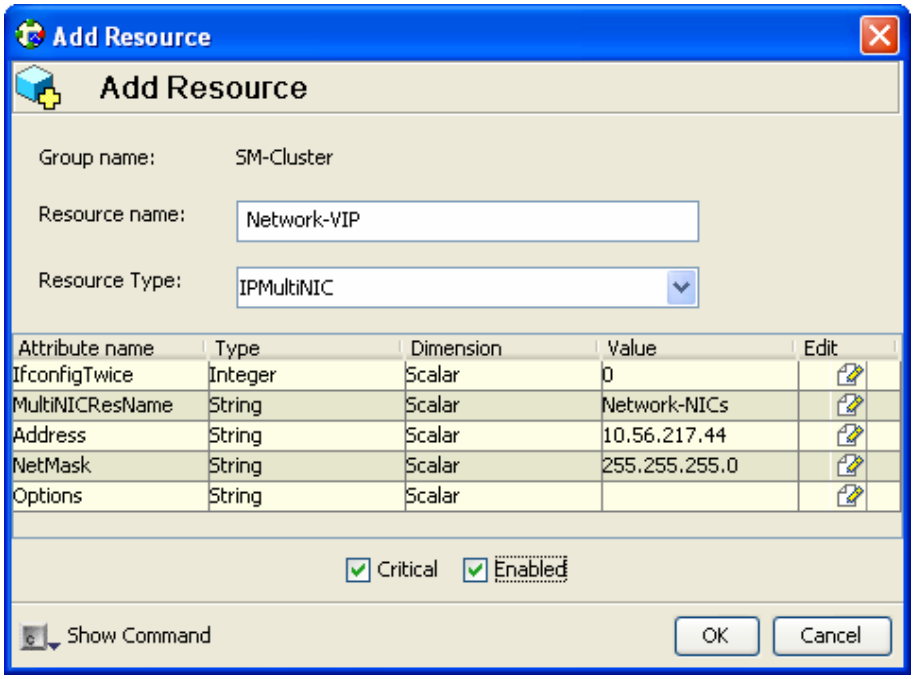

Step 4 Set the resource to be ENABLED and CRITICAL (see *Adding Network NICs* (on page E-10)).

Adding SM Resources

To add an SM resource:

Step 1 Import the SubscriberManager agent's type from file */opt/VRTSvcs/bin/SubscriberManager/SubscriberManager.cf*.

Step 2 Add to the service group a SubscriberManager resource called *SM*.

Step 3 Define the following parameters:

- SmBinPathName—Type the path to the *bin* directory under the SM installation directory; for example, */opt/pcube/sm/server/bin/*.
- SmDebugLevel—Type a number between 1 and 4 to view debug messages, type 0 to disable debug messages.

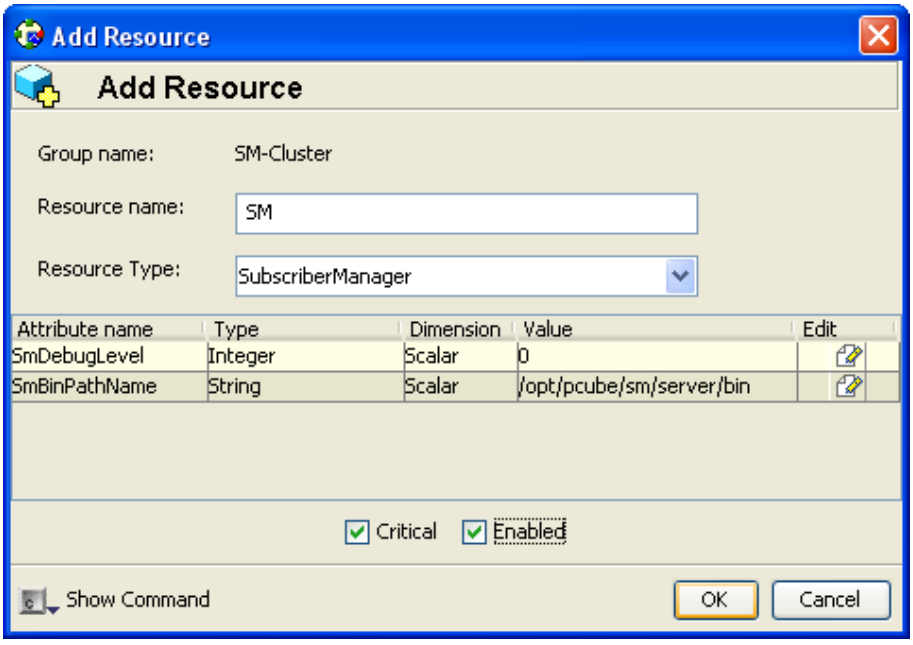

Adding TimesTen Daemon Resources

To add a TimesTen Daemon Resource:

Step 1 Import the OnOnlyProcess agent's type from file */opt/VRTSvcs/bin/OnOnlyProcess/OnOnlyProcess.cf*.

Step 2 Add to the service group an OnOnlyProcess resource called *"TimesTenDaemon*".

Step 3 Define the following parameters (see following figure):

- OnlineCmd—Type the TimesTen Daemon start command: */etc/init.d/tt_pcubesm22 start*.
- PathName—Type the TimesTen Daemon process path; for example, */opt/pcube/lib/tt/TimesTen/pcubesm22/bin/timestend*.
- Arguments—To view the arguments, run the following command of the machine: **ps –eaf | grep timestend**

  For example, the arguments can be: -initfd 13

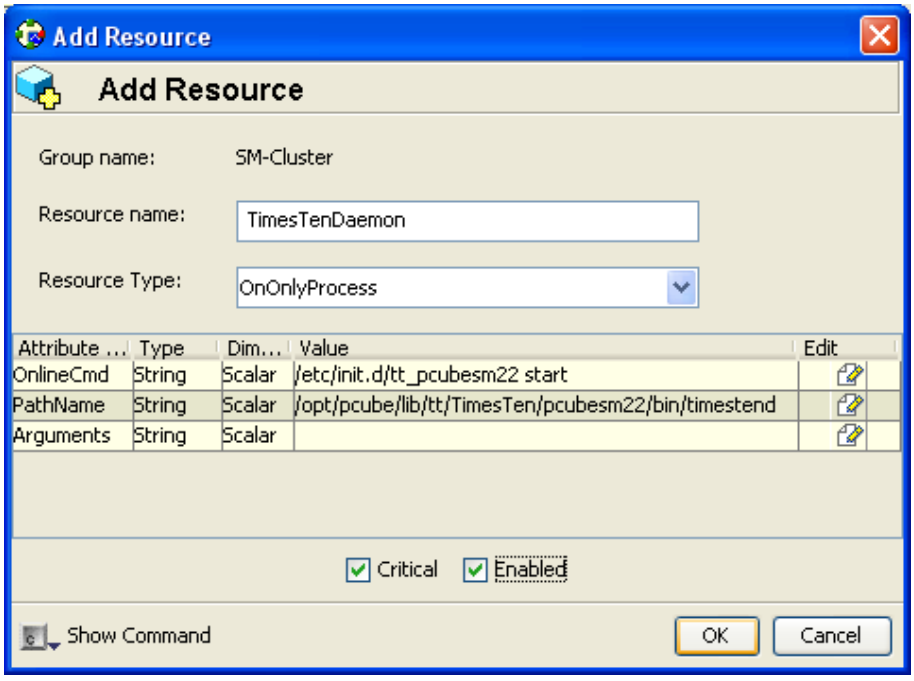

Adding TimesTen Replication Agent Resources

To add a TimesTen replication agent resource:

Step 1 Import the TimesTenRep agent's type from file */opt/VRTSvcs/bin/TimesTenRep/TimesTenRep.cf*.

Step 2 Add to the service group a TimesTenRep resource called *ReplicationAgent*.

Step 3 Define the following parameters (see following figure):

- TtBinPathName—Type the TimesTen bin directory path; for example, */opt/pcube/lib/tt/TimesTen/pcubesm22/bin*.
- TtDebugLevel—Type a number in the range of 1-4 for viewing debug messages; enter 0 for disabling the debug messages.

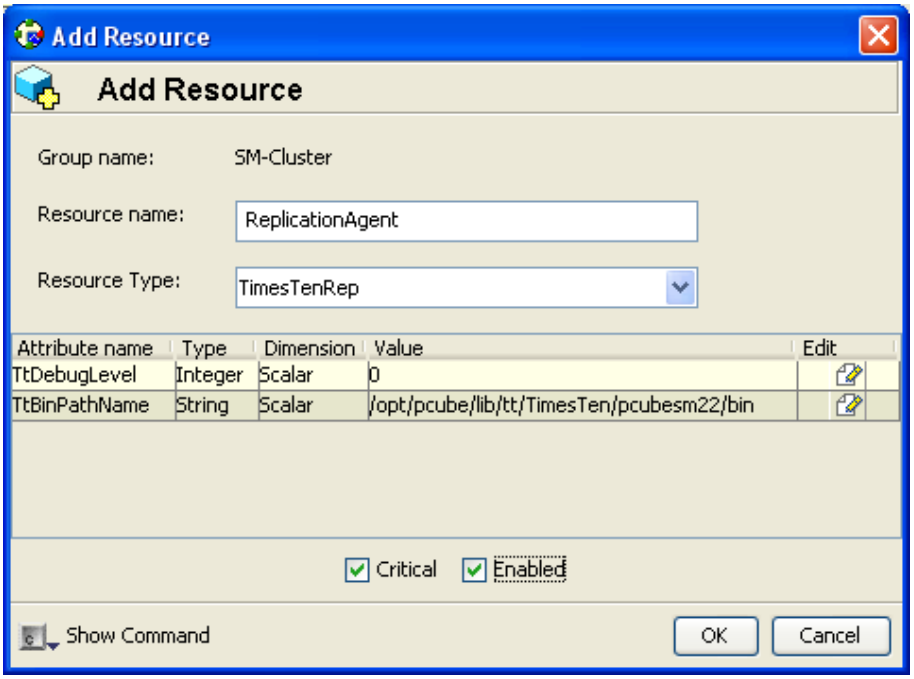

Useful Operations

The following sections are useful operations for the management of the VCS.

Logging into the Cluster

After you have added the cluster, click the icon: a login window appears.

Use the initial user and password (user: *admin*; password: *password*) for logging in.

Saving the Configuration

Click the icon, or from the menu select FILE | SAVE CONFIGURATION.

Before exiting the VCS make sure to save your configuration; otherwise, your configuration will be lost.

Closing the Configuration

Click the icon, or from the menu select FILE | CLOSE CONFIGURATION.

Before exiting the VCS, make sure your configuration is closed. Some operations (like rebooting the system) could fail or cause a configuration conflict if performed while the configuration is in read/write mode.

Importing Types

To configure the SM Veritas agents, you first have to import the type file of these agents. This operation is performed by selecting from the menu FILE | IMPORT TYPES. A navigation window appears:

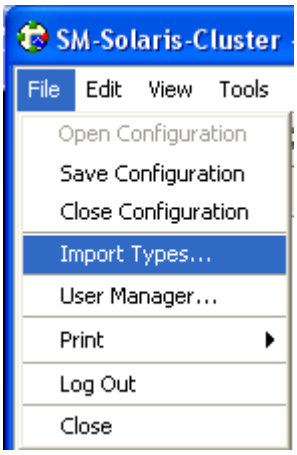

The window navigates through one of the cluster-system's file system. Go to the agent directory under */opt/VRTSvcs/bin/<agent-dir>* where there is a file whose name extension is *.cf*. Select this file (see following figure).

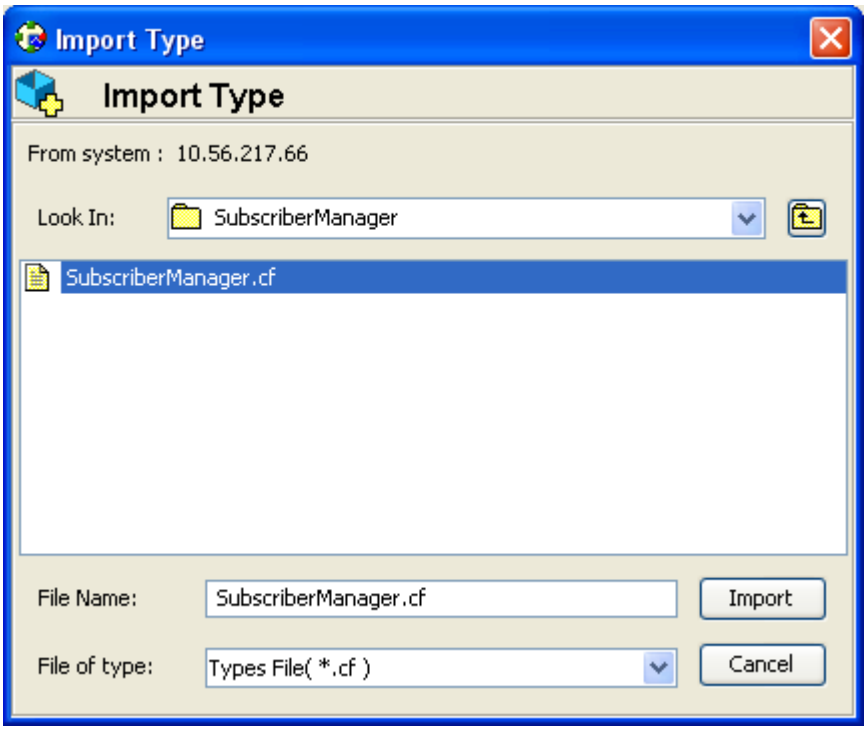

You can see the resource's parameters in the following window:

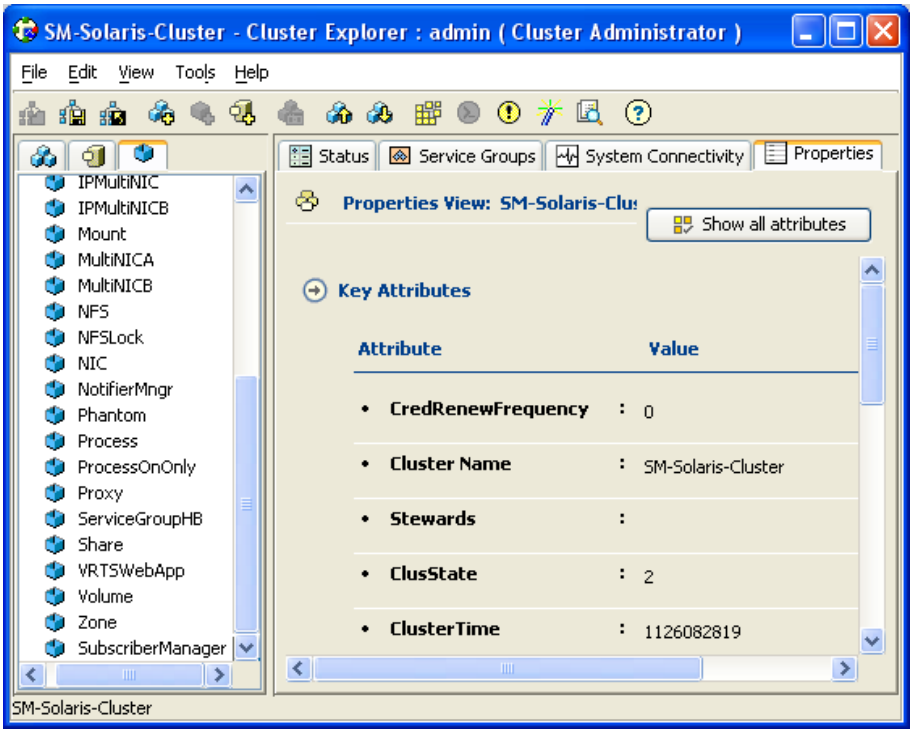

Linking the Resources

Linking the resources defines the order of becoming online and going offline.

To link the resources:

Step 1 Select the service group and enter the Resources tab.

Step 2 To link two resources, click once on one resource, pull the line to the second resource, and click once over the icon of the second resource.

The final links should look like the ones in the following figure:

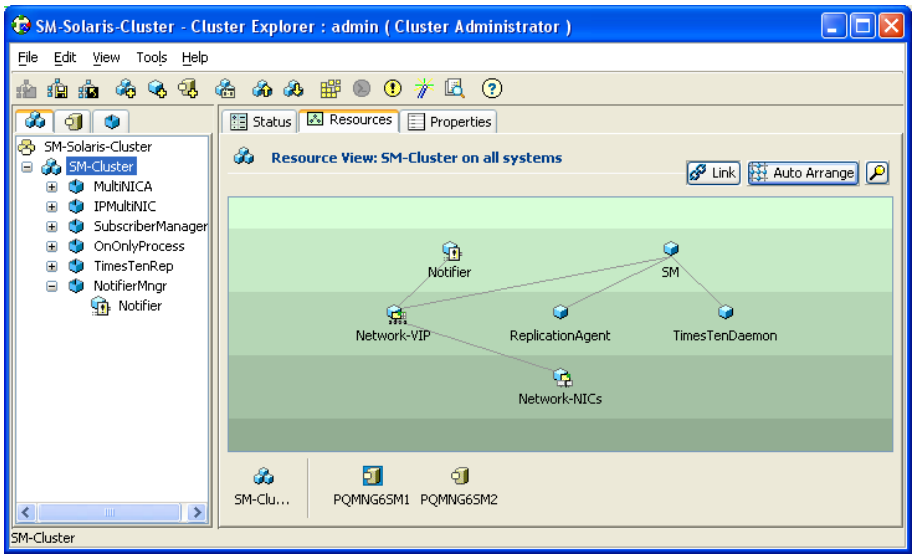

Saving and Closing Your Cluster Configuration

Remember to save your work at the end of the configuration process (see *Saving the Configuration* (on page E-16) and *Closing the Configuration* (on page E-16)).

Verifying That Service Group is Online

Check that all of the resources are online/offline, according to the system and the resource type.

Note The TimesTen Daemon, the NICs, and the TimesTen Replication Agent should all be online on all of the systems.

The state should be similar to the following figure:

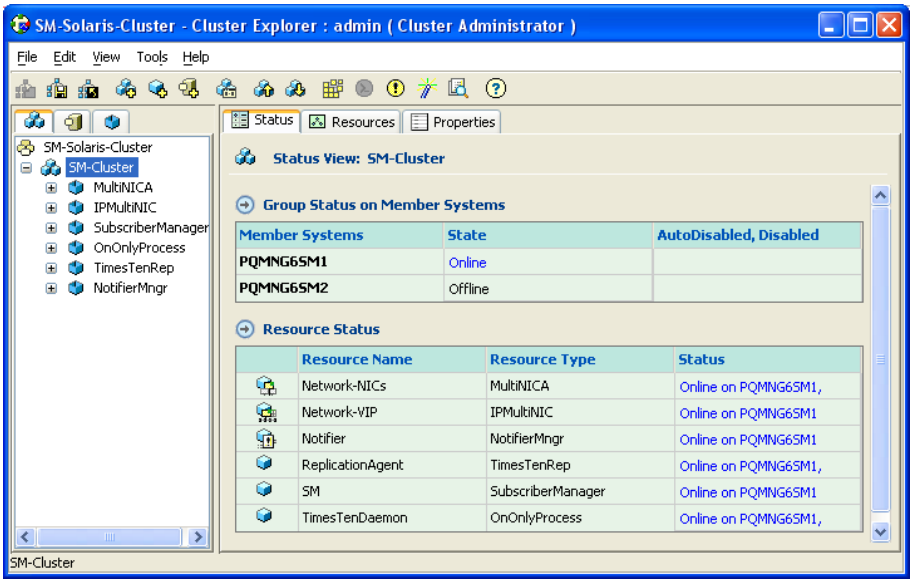

SNMP Support

VCS provides a method for notifying the user of important events such as a resource or system fault. For this purpose, VCS supplies a NotifierMngr agent that enables the reception of messages from VCS and the delivery of those messages to SNMP consoles. This section describes configuring NotifierMngr in order to enable SNMP support.

Configuring NotifierMngr

Add and configure NotifierMngr using either the command line or the Cluster Manager Java Console.

When started from the command line, Notifier is a process that VCS does not control.

For best results, use the NotifierMngr agent bundled with VCS to configure Notifier as part of a highly available service group, which can then be monitored, brought online, and taken offline.

The following sections describe the configuration process using the Cluster Manager Java Console.

Adding NotifierMngr Resource

Add NotifierMngr as a resource to the SM Cluster service group.

To add the NotifierMngr Resource:

Step 1 Add to the service group a NotifierMngr resource called Notifier.

Step 2 When the following screen is displayed, choose NotifierMngr as a ResourceType.

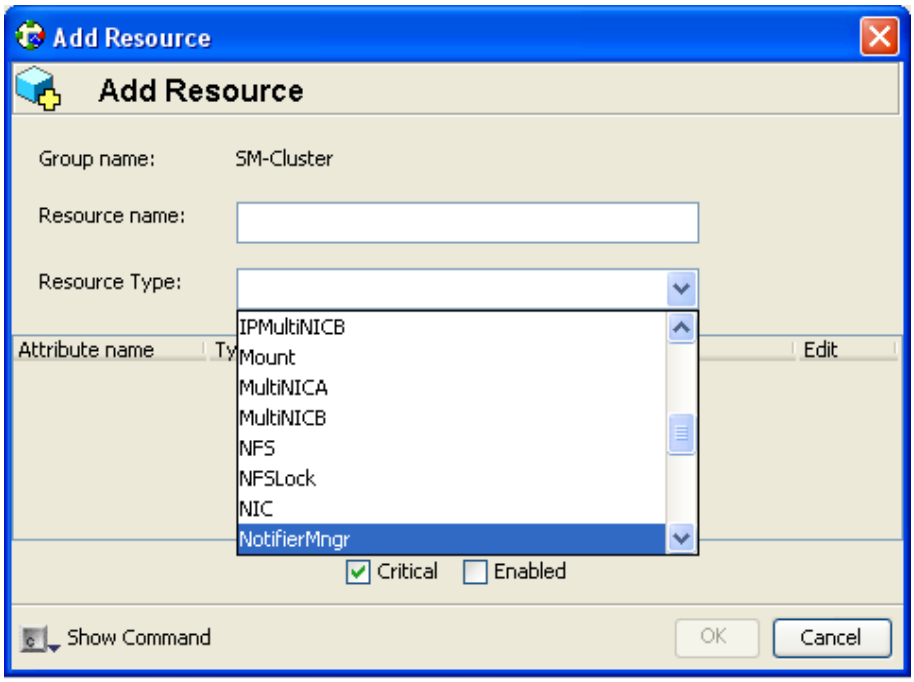

Configuring the NotifierMngr Attributes

After adding the NotifierMngr resource, configure its attributes.

To configure the NotifierMngr attributes:

Step 1 Select NotifierMngr as a Resource Type; the following screen appears:

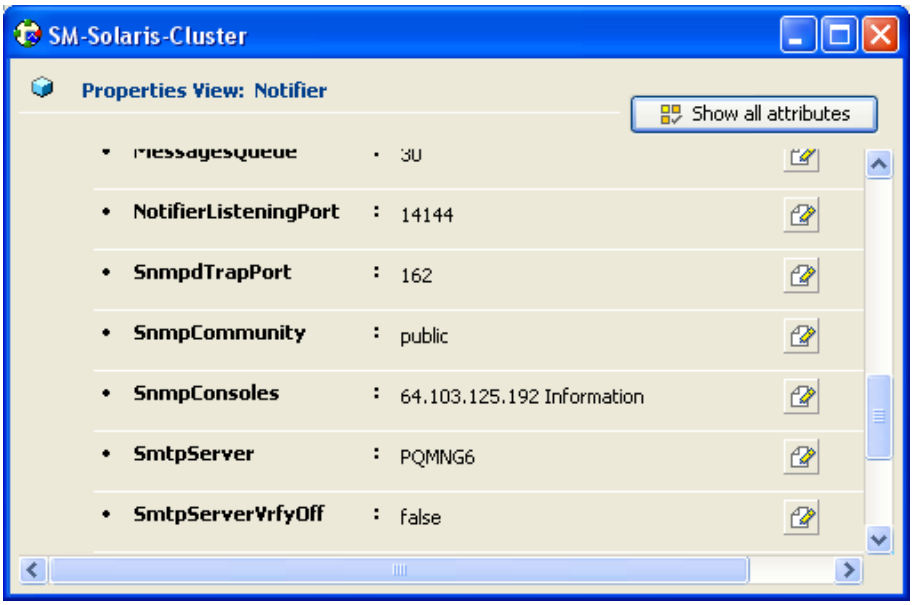

Step 2 Define the following parameters (see the above figure):

- SnmpConsoles—Specify the machine name of the SNMP manager and the severity level of messages to be delivered to the SNMP manager. The severity levels of messages are Information, Warning, Error, and SevereError. Specifying a given severity level for messages generates delivery of all messages of equal or higher severity.
- SnmpdTrapPort—Specify the port to which the SNMP traps are sent. The value specified for this attribute is used for all consoles if more than one SNMP console is specified. The default is `162`.
- SnmpCommunity—Specify the community ID (a string scalar) for the SNMP manager. The default is `public`.

Configuring the SnmpConsole Attribute

The SnmpConsole attribute specifies the IP addresses to which you want the SNMP traps to be sent. You can specify different trap severity for each IP address:

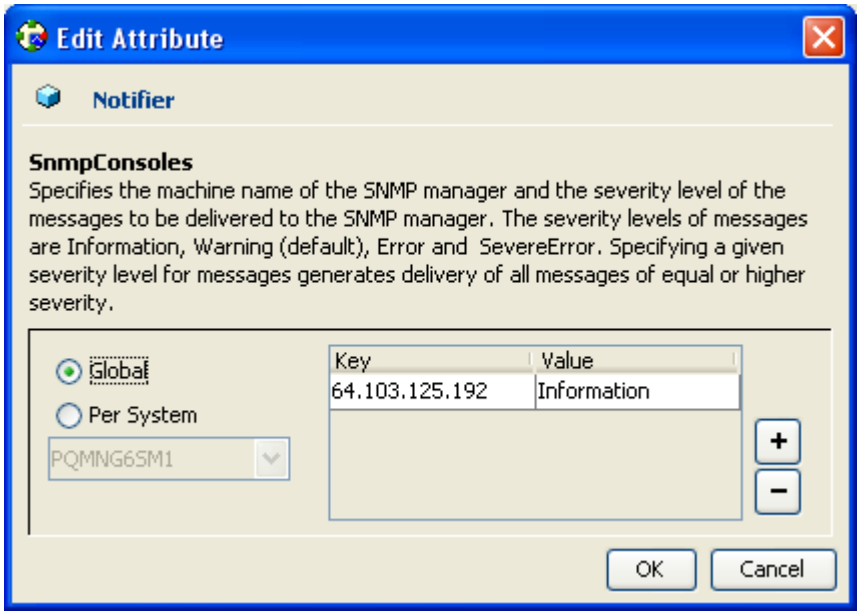

Linking to IPMultiNIC

To link to IPMultiNIC:

Using the resources viewer, connect the *Notifier* to the *Network-VIP* resource so that it will be online after the VIP.

For more information, see the figure in *Adding TimesTen Replication Agent Resources* (on page E-15).

Viewing Traps

After adding and configuring NotifierMngr, it will send traps according to the configured severity to the destinations configured by the SnmpConsole Attribute.

View these traps using SNMP trap viewer/MIB Browser (for example, AdventNet MibBrowser).

For a complete list of traps/severities, please see Chapter 10 of the VERITAS Cluster Server User Guide.

Glossary of Terms

A

Access Control List (ACL)

Permits or denies incoming connections on any of the management interfaces. It is an ordered list of entries, each consisting of an IP address and an optional wildcard “mask” defining an IP address range, and a permit/deny field.

Active subscriber

An online subscriber who is actually generating IP traffic.

Anonymous subscriber ID

The SCE device assigns each unknown subscriber IP address with a temporary subscriber ID until it receives the real Subscriber ID from the Policy Server. This occurs in a Pull Model integration in the Subscriber Manager.

For more information on the Pull Model integration, see the *Cisco SCMS SCE Subscriber API Programmer Guide*.

Anonymous subscriber mode

A mode of the Service Control solution in which the system monitors traffic and automatically assigns a Service Configuration based on to the individual's IP address used on its subscriber-side. This mode can be used to control a subscriber's traffic anonymously, without integrating the system with an OSS system. In this mode, the subscribers defined in the system are anonymous and are distinguished only by their IP address or VLAN ID.

Automatic integration

Subscribers are introduced into the SM automatically upon login, by the LEG.

B

Blocking API

In a Blocking API, which is the common type, every method returns *after* its operation has been performed.

C

Cable/Satellite modem

A data modem that provides Internet access over cable and satellite networks. The modem usually corresponds to a single subscriber of the Internet Service Provider (ISP).

CLI

One of the management interfaces to the SCE platform. It is accessed through a Telnet session or directly via the console port on the front panel of the SCE platform.

Cluster Topology

The SM is implemented based on the Veritas Cluster Server (VCS) technology. The cluster includes two machines, each of them running SM TimesTen and Veritas software. The Veritas Cluster Server software consolidates the two SM machines and exposes a single entity by providing a single virtual IP address for the entire cluster.

The added functionality when operating in a cluster topology provides powerful new features such as fail-over and high availability.

The cluster software distinguishes an active and a standby machine. The active machine ‘owns’ the virtual IP address and all network connections, while the standby machine is passive until a fail-over occurs. At fail-over, the IP address is passed from the failing server to the backup server, which becomes activated and re-establishes all network connections.

Command-Line Interface

See CLI.

CPE

Any type of equipment that an end-user can connect to the network through a modem. The end-user can connect multiple CPE devices to the Internet through a single modem.

Customer premise equipment (CPE)

See CPE.

D

DHCP Forwarder

An application that acts as a bridge between the DHCP Lease Query LEG and the DHCP servers. The request sent from the LEG is received by the forwarder and is forwarded to the appropriate DHCP server. The replies from the DHCP servers are forwarded to the LEG. The LEG signals the forwarder which server should receive each request.

DHCP Lease Query Transaction

The DHCP Lease Query transaction is a DHCP transaction with special message types that enable, among other things, clients to query DHCP servers regarding the owner and the lease-expiration-time of an IP address.

The DHCP Lease-Query transaction is defined in an Internet Engineering Task Force (IETF) draft. For more information, see the *IETF website*.

DHCP Sniffer

The software logic inside the SCE device that analyzes DHCP traffic and sends the information to the SCE-Sniffer DHCP LEG using the RDR protocol.

Domains

A group of SCE platforms that share a group of subscribers. The subscriber traffic can pass through any SCE platform in the domain.

A subscriber and an SCE can belong to only one domain.

F

Fail-Over

In a fail-over solution, the two Subscriber Manager platforms exchange keep alive messages via the cascade ports. This keep alive mechanism enables fast detection of failures between the Subscriber Manager platforms and fast fail-over to the standby SM platform when required.

H

High Availability

In a high availability topology, the Subscriber Manager software runs on two machines, designated as the active machine and the standby machine. Subscriber data is continuously replicated from the active to the standby machine, ensuring there is minimal data loss in case of active SM failure. When the active machine fails, the standby machine discovers the failure and becomes active.

L

LEG

Software components that use the Subscriber Manager API to generate subscriber-record update messages such as login/logout, and send the messages to the SM.

LEGs are usually installed on AAA/OSS platforms, and serve as translators from events generated by the customer system to Cisco subscriber update events.

Login Event Generator

See LEG.

M

Manual integration

Subscribers are introduced into the SM by importing them in the form of a *csv* file, or by using the appropriate SM CLU.

N

NAS

A network device that serves as an access point for a remote user. It initiates RADIUS transactions to the RADIUS server to authenticate a remote user.

Network Access System (NAS)

See NAS.

Network ID Mappings

A network identifier that the SCE device can relate to a specific subscriber record. A typical example of a network ID mapping (or simply mapping) is an IP address. Currently, the Cisco Service Control solution supports IP address, IP range, and VLAN mappings.

For additional information, see the *Cisco SCMS Subscriber Manager User Guide*.

Non-blocking API

Non-blocking methods return immediately, even before their operation has been completed. The operation results are either returned to an Observer object (Listener) or not returned at all.

The non-blocking method is advantageous when the operation is lengthy and involves I/O. Performing the operation in a separate thread allows the caller to continue doing other tasks and improves overall system performance.

P

Pull mode

Subscriber information is not distributed automatically. Only when an SCE platform cannot associate the IP traffic with a subscriber, will it request (pull) information from the SM.

Pull-request

A message sent from an SCE device to the SM/LEG when it identifies a use of a new subscriber IP address in the network. The SM uses the IP address provided in this message to query the database to retrieve the subscriber data of the subscriber associated with this address and to send its data to the SCE.

Push mode

The SM automatically distributes (pushes) subscriber information to the relevant SCE platform immediately after the record is added or changed. Therefore, when the subscriber starts producing traffic through the SCE platform, it is ready with the required subscriber information.

Q

Quota

A (subscriber's) limit for a specific metric, such as bandwidth or volume.

S

SCE Platform

The Service Control purpose-built hardware service component. This hardware device is capable of performing smart analysis of the packets at wire speed. It monitors the traffic on the line, producing raw data to be provided to the loaded application, which processes the data for functions such as reporting, policy management, subscription management, and implementation of tiered service subscriber aware traffic policies.

The SCE platform comes in the following models: SCE 1000 2xGBE, SCE 2000 4xGBE, and SCE 2000 4/8xFE. There may be one or more SCE platforms on the provider network.

Service Control Engine platform

See SCE platform.

Service Control Management Suite Subscriber Manager

See SM.

SM

A middleware software component used in cases where dynamic binding of subscriber information and service configurations is required. The SM manages subscriber information and provisions it in real time to multiple SCE platforms. It can store subscriber service configurations information internally, and act as a state-full bridge between the AAA system (for example, RADIUS and DHCP) and the SCE platforms.

Subscriber

The generic term used to refer to the managed entity for which a Service Configuration is enforced, and usage is monitored, by a Service Control solution. A subscriber can be defined as an individual IP address, or ranges of IP addresses or VLANs. The system supports different modes of operations including: subscriber-less mode (all control is performed globally), anonymous subscriber mode, and dynamic and static subscriber-aware modes.

Subscriber domain

The Subscriber Manager (SM) provides the option of partitioning SCE platforms and subscribers into subscriber domains. A subscriber domain is a group of SCE platforms that share a group of subscribers. Subscriber domains can be configured using the SM configuration file and can be viewed using the SM Command-Line Utility (CLU).

It is also possible to configure domain aliases. A domain alias is a synonym for the actual domain name in the SM. Domain aliases are configured in the SM configuration file.

Subscriber mappings

The SCE platform requires mappings between the network IDs (IP addresses) of the flows it encounters and the subscriber IDs. The SM database contains the network IDs that map to the subscriber IDs. The SCE network-ID-to-subscriber mappings are constantly updated from the SM database.

For more information about the SCE platforms, see the *SCE 1000* and the *SCE 2000 User Guides.*

Subscriber package

A policy enforced by Cisco SCMS solutions on each subscriber, which is usually defined by a policy package ID. The various Subscriber Manager Login Event Generator (LEG) component software can handle the package ID in any of the following ways:

- Sets the package ID according to configurable options of the DHCP initial logon or lease extension transactions
- Sets the package ID using a constant default value
- Leave the package ID unset

Index

A

audience • x

C

cable support module
  managing • 4-31, See p3cable
Cisco Service Control • 1-1
  broadband service providers • 1-2
  capabilities • 1-2
cisco.com • xiv
CLU • B-1
  output • 4-21, B-2
  parsing • 4-21, B-2
  using • 4-20
commands, CLU
  p3batch • B-3
  p3cable • B-4
  p3clu • B-5
  p3cluster • B-5
  p3db • B-6
  p3domains • B-7
  p3ftp • B-7
  p3http • B-8
  p3inst • B-8
  p3log • B-10
  p3net • B-11
  p3radius • B-11
  p3rpc • B-12
  p3sm • B-13
  p3subs • B-16
  p3subsdb • B-19
communication failures • 2-9
  failure detection • 2-9
  failure handling • 2-9
  failure recovery • 2-9
configuration • 3-12, 4-1
  configuring the SM • 3-16
  reloading the SM configuration • 4-23, See p3sm
connection failure • 4-15
conventions, document • xi
customer premises equipment (CPE) • C-1, C-2
  cable modem • C-3
  cable support module • C-1

D

documentation, obtaining • xiii

F

fail-over • 2-11, 2-12, 2-15
  normal operation • 2-12
  operation • 2-15
  recovery • See fail-over recovery
  topology • 2-13
fail-over recovery • 2-15
  database duplication recovery • 2-17
  database duplication recovery management • 2-17
  machine reboot • 2-16
  server replacement • 2-16

I

installation • 3-1, 3-2
  additional installation procedures • 3-17, 3-18, 3-19
  changes made by installation scripts • 3-21
  distribution file contents • 3-2
  documentation • 3-5
  installating the SM • See typical installation
  installing an application • 3-20, 4-32, See p3inst
  procedures • 3-8

system requirements • 3-5
troubleshooting • 3-19
verification • 3-16
introduction mode • 2-6
pull mode • 2-7
push • 2-7

M

management • 4-1
management and collection • 1-4
collection • 1-6
network management • 1-5
subscriber management • 1-5
management, subscribers, mappings, and properties • 4-25, See p3subs
clearing subscriber applicative state • 4-27
mappings • 4-26
properties • 4-27
subscribers • 4-25

O

obtaining documentation • xiii
obtaining technical assistance • xiii

P

preface • ix

S

SCE Platform • 1-3
configuring • 4-14
service configuration management • 1-4
SM configuration file • 4-4
Auto Logout • 4-10, A-8
Cable Adapter • 4-13, A-16
Data Repository • 4-7, A-17
Default Domains Configuration • 4-9
Domain.XXX • 4-5, A-6
Domains • 4-5, A-7
FTP • 4-13, A-14
High Availability • 4-8, A-4
HTTP Tech-IF • 4-13, A-15
LEG-Domains Association • 4-9, A-5
MPLS-VPN • A-13
RADIUS Listener • 4-11, A-9
RADIUS.NAS.XXX • 4-11, A-10
RADIUS.Property.Package • 4-12, A-11
RADIUS.Subscriber ID • 4-11, A-12
RDR Server • 4-13, A-16
RPC.Server • A-13
RPC.XXX • 4-12
SCE platform repository • 4-4
SCE.XXX • 4-5, A-14
SM General • 4-6, A-2
SM-LEG Failure handling • 4-8, A-4
Subscriber State Persistency • 4-8, A-4
SM configuration file options • A-1, See SM configuration file
usage • A-1
SM management and configuration • 4-1
command-line utilities • 4-2
configuration file • 4-1, See SM configuration file
importing subscribers • 4-14
subscriber management solution • 4-3
SM node cluster
managing • 4-34, See p3cluster
subscriber • 2-2, See introduction mode
aging parameters • 4-15
groups • 4-14
handling • 2-2
information flow • 2-3
numbers • 2-4
SM database • 2-5
templates • 4-15
subscriber manager • 2-1
API • 2-6
cluster • 2-10
communication failure • See communication failures
database • 2-5, 4-28, See p3subsdb
domains • 2-8
fail-over • See fail-over
fundamentals • 2-6
login event generators • 2-6
management • 2-10, 4-23, See p3sm
system configuration
examples • 4-15

T

Technical Assistance Center (TAC) • xiv
technical assistance, obtaining • xiii
troubleshooting • D-1
configuration • D-13
general • D-1
general setup • D-2
p3cable • D-13
p3net • D-10
p3subsdb • D-12
TimesTen database setup errors • D-4

typical installation • 3-8
distribution file extraction • 3-9
engage package • 3-15
install-sm.sh script • 3-9
maximum number of subscribers • 3-12
pcube user password • 3-10
reboot • 3-15
system memory configuration • 3-10
TimesTen • 3-12, 3-14, 3-15

U

uninstalling • 3-27
uninstalling VCS agents • 3-29
uninstall-sm.sh script • 3-28
upgrading • 3-22
data duplication procedure • 3-25
from cluster setup version 2.x • 3-27
from cluster setup version 3.x • 3-27
from non-cluster setup to cluster setup • 3-25
upgrade-sm.sh script • 3-23
upgrading from • 3-25
upgrading the SM • 3-22
user log
managing • 4-34, See p3log
utilities • 4-35
CLU help • 4-36, See p3clu
database operations • 4-36, See p3db
running batch files • 4-35, See p3batch

V

VCS configuration
adding clusters • E-6
adding network NICs • E-10
adding network VIPs • E-12
adding services groups • E-7
adding SM cluster resources • E-8
adding SM resources • E-13
adding TimesTen Daemon resources • E-14
adding TimesTen replication agent resources • E-15
closing • E-16
guidelines • E-5
importing types • E-16
linking resources • E-18
logging into the cluster • E-15
NotifierMngr • E-19, E-20, E-21
saving • E-16
setting the auto-start • E-8
SM cluster resources • E-5
SNMP support • E-19
SnmpConsole • E-22
Veritas cluster server • E-1
remote site nodes • E-2
replication configuration guidelines • E-3
requirements • E-1
viewing
RADIUS listener statistics • 4-35, See p3radius
RPC interface server information • 4-33, See p3rpc
subscriber domains • 4-30, See p3domains
viewing and connecting network elements • 4-30, See p3net